

Verification of Information Flow and Access Control Policies with Dependent Types

IMDEA-SW-TR-2010.0; Version of February 26, 2011

Aleksandar Nanevski Anindya Banerjee
IMDEA Software Institute
Madrid, Spain
{aleks.nanevski, anindya.banerjee}@imdea.org

Deepak Garg
Carnegie Mellon University
Pittsburgh, PA, USA
dg@cs.cmu.edu

Abstract

We present Relational Hoare Type Theory (RHTT), a novel language and verification system capable of expressing and verifying rich information flow and access control policies via dependent types. We show that a number of security policies which have been formalized separately in the literature can all be expressed in RHTT using only standard type-theoretic constructions such as monads, higher-order functions, abstract types, abstract predicates, and modules. Example security policies include conditional declassification, information erasure, and state-dependent information flow and access control. RHTT can reason about such policies in the presence of dynamic memory allocation, deallocation, pointer aliasing and arithmetic. The system, theorems and examples have all been formalized in Coq.

1 Introduction

Several challenges persist in existing work on specification and enforcement of confidentiality policies. First, many practical applications require a combination of a number of different classes of policies: authentication, authorization, conditional declassification, erasure, etc. Yet, most existing systems are tailored for enforcing specific classes of policies in isolation. Second, where policy combinations have been considered (e.g. [6, 9, 15]), policy conformance is typically formalized for simple languages without important programming features such as dynamic allocation, mutable state and pointer aliasing, or without modern modularity mechanisms that aid programming in the large. There has been little work on confidentiality policies pertaining to linked data structures (lists, trees, graphs, etc.), and even less work exists for structures that are heterogeneous; that is, structures that contain mixed secret and public data as well as mixed secret and public links. Third, despite their efficiency, enforcement mechanisms are often imprecise in their handling of implicit information flow (that arises due to program control structures such as conditionals or procedure calls) and reject perfectly secure programs.

In this paper we revisit the foundations of information flow — its specification as well as its static enforcement — and address the above challenges of policy specificity, language expressiveness and precision, simultaneously. The key insight of our work is that all the three problems can

be addressed using *standard* linguistic features from dependent type theory [30]: (a) higher-order functions, abstract data types and modules, that provide for software engineering concepts such as abstraction and information hiding, and (b) a logic for higher-order assertions, including quantification over predicates, that serves as the foundation for a rich policy specification language. We additionally consider an extension of dependent types with (c) general recursion, mutable state, dynamic allocation, and pointer aliasing. We use the dependent types as a policy specification language, and typechecking (i.e., program verification) to enforce conformance of programs to policy. As is standard in type theory, we assume that programs are typechecked before they are executed.

As our first contribution, we show that a number of security policies which have been previously considered in isolation, such as declassification [17, 44], information erasure [18, 19], state-dependent access control [14, 15] and state-dependent information flow policies [9], can be combined in the same system using the mentioned type-theoretic abstractions. We explain this point further below, and illustrate it through several verified examples in the paper.

As our second contribution, we show that these policies can be enforced in the presence of dynamic allocation, deallocation, and pointer aliasing, and in particular, over programs involving linked, heterogeneous data structures. To achieve this, we employ a *semantic* definition of what constitutes confidential (high) vs. public (low) data, in contrast to most related work where variables are *syntactically* labeled with a desired security level [34, 51]. The semantic characterization allows the same variable or pointer to contain data of different security levels at different points in program execution, which gives us the needed flexibility of enforcement. The semantic characterization also facilitates precise specification of programs with implicit information flow such as procedure calls or (possibly nested) conditionals.

Our third and technically central contribution is a novel verification system, Relational Hoare Type Theory (RHTT), that integrates a programming language and a logic into a common substrate underlying all of (a)–(c) above. In more detail, RHTT provides (a) and (b) by including the type theory of the Calculus of Inductive Constructions (CiC) [31, Chapter 4], as implemented in the Coq proof assistant. To provide for (c), RHTT introduces a new type constructor `STsec`, which classifies side-effectful computations similar to Haskell monads [38], except that the `STsec` monad is indexed with a precondition and a postcondition, as in a Hoare triple. `STsec` types separate the imperative from the purely functional fragment of the type theory, ensuring soundness of their combination.

RHTT’s preconditions specify constraints on the environment under which it is safe to run a program, and can be used to enforce authentication and authorization policies, even when they depend on state. RHTT’s postconditions are *relational* assertions; they specify the behavior of *two runs* of a program [3]. The relational formulation directly captures in the types the notion of *noninterference* [24], a prominent semantic characterization of confidentiality. Together with higher-order type theory, this provides an architecture for uniform treatment of all the policies mentioned above.

For example, we show that the fundamental linguistic abstractions required to specify and implement declassification are `STsec` types, modules and abstract predicates. A module can be used to delimit the scope in which data is considered public, by hiding the publicity of the data from module clients via existential type abstraction [32]. Then declassification amounts to breaking the abstraction barrier by an exported interface method that reveals this in-module publicity. This is orthogonal to revealing the data itself. The latter can always be done even without declassification, but the clients will have to use such data as if it were confidential. Declassification may be uncon-

ditional or conditional [9], where the condition might be stateful and involve, e.g., authentication.

In information erasure policies [18, 19] confidential data may be released within a delimited scope, provided there is a guarantee that such data will be erased upon exit from the scope. We show that such policies can be specified using a combination of higher-order functions with local state, modules and abstract predicates. The key facilitating component here is that `STsec` types may appear in argument positions in function types, which is similar to having Hoare logic where one can reason about Hoare triples hypothetically wrt. the truth of other Hoare triples. A similar combination of features can be used to grant a method access to data only if the method provably conforms with some desired confidentiality policy.

Finally, state-dependent information flow and access control policies require abstract predicates combined with mutable state. This allows expressing security policies that can change with time due to state updates [48].

Our fourth contribution is the development of a logic for relational reasoning about RHTT programs. Inference rules of the logic have been verified sound against a semantic model, and are formally implemented as lemmas in RHTT. The distinguishing characteristic is that we can reason relationally about structurally dissimilar programs. This is in contrast to previous relational logics [11, 52], which support reasoning only about programs with similar control flow, and provide inference rules for conditionals and loops with only low boolean guards.

Our development of RHTT overcomes a number of technical challenges. First, for relational reasoning to be applicable at all, the type system must give special status to instantiations of a program e with high values. The special status is needed so that the *same* postcondition of e can relate e 's *different* instantiations. Our solution is to introduce new typing and programming primitives for abstraction and instantiation wrt. a number of variables, simultaneously (Section 2). This illustrates why our type system had to be developed hand-in-hand with the associated relational verification logic (Section 4), as each must possess the requisite constructs to facilitate the other. The second challenge concerns the semantic treatment of allocation and deallocation, pertaining to dynamic data structures. Existing techniques [3, 8] for modelling allocation in the relational security setting cannot cope with deallocation; hence the need for two different allocators — one for low and another for high addresses (Section 3).

The soundness of our program logic, the domain theoretic implementation of our semantic model, as well as all of our examples, have been fully and formally verified in Coq. Additional technical difficulties arise in this process, but we elide them here for readability. The interested reader is invited to look at our Coq proofs, which are available at <http://software.imdea.org/~aleks/rhtt/>.

2 RHTT by examples

Overview As suggested by the introduction, this paper assumes understanding of the following aspects of type theory: (1) *Dependent function types*, used to specify how the body of a function depends on the input arguments. To illustrate, consider the type `vector(n)`, of integer-storing arrays. This type is dependent on the size parameter n . A function computing the inner product of two vectors can be typed as

$$\prod n:\text{nat}. \text{vector}(n) \times \text{vector}(n) \rightarrow \text{nat}$$

capturing the invariant that that the argument vectors must be of equal size. In RHTT, dependent function types naturally arise when specifying any kind of program behavior. (2) *Module systems (including abstract types and predicates)*, for information hiding, and as we show, declassification. (3) *Inductive types*, for specifications of programs that manipulate (possibly heterogeneous) data structures such as lists, trees, etc.

To use RHTT in practice, it is further important to be familiar with some implementation of type theory (our chosen one is Coq [31], but others exist too), as one needs to interact with the system to discharge verification conditions. Our presentation in this paper does not include such interaction aspects, and hence does not assume familiarity with Coq.

RHTT basics: types, specifications, opaque sealing To begin with, our types must be able to express at least noninterference: that low outputs of a computation are independent of high inputs. To illustrate, assume a function $f:A^2\rightarrow A^2$, where $A^2 = A \times A$. Also, let $e.1$ and $e.2$ denote resp. the first and the second component of the ordered pair e . Then, mathematically, f 's first output is independent of f 's second argument iff

$$\forall x_1 x_2 y_1 y_2. x_1 = x_2 \rightarrow f(x_1, y_1).1 = f(x_2, y_2).1$$

In other words, in two runs of f , equal x inputs, lead to equal $f(x, y).1$ outputs. This relational statement of independence can be viewed as a definition of noninterference in terms of f alone [3, 11], without recourse to outside concepts such as security lattices [10, 20]. Consequently, inputs and outputs related by equality in the two runs of f are considered low (x and $f(x, y).1$ above), and the unconstrained values (y and $f(x, y).2$) are by default considered high. So defined, the notions of low and high security are intrinsic to the considered specification, rather than to the code itself; one is free to consider statements about f in which the inputs and outputs take other security levels.

In RHTT, program specifications are stated using a *monadic* type $\text{STsec } A(p, q)$, which classifies heap-manipulating, potentially diverging computations e whose return value has type A . e 's *precondition* p is a predicate over heaps, i.e., function of type $\text{heap} \rightarrow \text{prop}$. The reader can roughly think of prop as type bool which in addition to the usual logical operations supports quantifiers as well. The precondition selects a set of heaps from which e 's execution will be memory-safe (e.g., there will be no dangling-pointer dereferences or run-time type errors). This automatically provides a mechanism for controlling access to heap locations, in a manner identical to that of separation logic [40]: e may only access those locations that are provably in all the heaps satisfying e 's precondition, or that e allocated itself. We will illustrate access control via preconditions in subsequent examples (see, e.g., Example 4).

The *postcondition* q relates the output values, input heaps and output heaps of any *two terminating* executions of e . Thus q has the type $A^2 \rightarrow \text{heap}^2 \rightarrow \text{heap}^2 \rightarrow \text{prop}$. The postcondition does not apply if one or both of the executions of e are diverging. In that respect, our type system is *termination insensitive* [43]. While p controls access to locations x , we use q to implement information flow policies about x . This is why q is a predicate over two runs. For example, q may specify that x is low, so that e may freely propagate x 's value. Or x may be high, requiring that all x -dependent outputs of e must be high too. Or x may be high but q may require all of e 's final heap to be low, in which case e must deallocate or rewrite any portion of its final heap that depends on x .

RHTT is implemented via shallow-embedding into Coq, which it extends with STsec types. In the implementation of STsec types in Coq, we rely on the ability of Coq modules to perform *opaque sealing* [25, 28]; that is, hiding the implementations of various values within a module,

while only exposing their types, thus forcing the clients of the module to be generic with respect to implementations of the module. Moreover, the actual implementations of opaquely-sealed functions, types and propositions cannot be recovered by clients, because RHTT does not contain constructs for pattern-matching (i.e., making observations) on the structures of such values.

We point out that our types can only describe the properties of the input and output states of the program (via pre- and postconditions), but not of intermediate states. Although this is not a significant limitation for a sequential, non-reactive language like RHTT, further work in this direction is left for future work.

Syntax, heaps, implicit flow Consider the following program, P_1 , adapted from Terauchi and Aiken [49], and presented here in a Haskell-like notation. We use side-effecting primitives such as `write x y` , which stores the value y into the location x ; `read x` , which returns the contents of x ; and `$x \leftarrow e_1; e_2$` , which sequentially composes e_1 and e_2 , binding the return value of e_1 to x . In future examples, we will also use `alloc x` , which returns a fresh memory location initialized with x ; and `dealloc x` , which deallocates the location x from the heap. Additionally, we use `do` to delimit the scope of the side-effectful computations. Our actual syntax implemented in Coq differs somewhat from the one here in the treatment of variable binding, an issue we ignore for the time being but to which we return in Section 3. Further, we freely use all the constructors inherited from CiC and Coq, such as for example, functions (`fun`), and dependent function type constructor (Π).

$$P_1 \hat{=} \text{fun } x \ y \ z \ lo \ hi : \text{ptr.} \\ \text{do (write } z \ 1; b \leftarrow \text{read } hi; \\ \text{if } b \text{ then write } x \ 1 \text{ else } (w \leftarrow \text{read } z; \text{write } x \ w); \\ u \leftarrow \text{read } x; v \leftarrow \text{read } y; \text{write } lo \ (u + (v \text{ mod } 10)))$$

Pointers x, y, z, lo store integers, and hi stores a boolean. The policy is: contents of lo and y are low at program input and output, while contents of x, z, hi are high. P_1 satisfies the policy because: (1) the value of y is not modified, and (2) the value of lo is modified to store the sum of the contents of x and the contents of y modulo 10, but this sum is independent of high data: at the time of writing lo , x has been rewritten by 1 in both branches of the conditional. Thus, P_1 can be ascribed the following dependent type, U .

$$U \hat{=} \Pi x \ y \ z \ lo \ hi : \text{ptr. STsec unit} \\ (\text{fun } i. \exists u \ v \ w \ c : \text{nat. } b : \text{bool. } j : \text{heap.} \\ i = x \mapsto u \bullet y \mapsto v \bullet z \mapsto w \bullet hi \mapsto b \bullet lo \mapsto c \bullet j, \\ \text{fun } rr \ ii \ mm. \\ (ii.1 \ lo = ii.2 \ lo) \rightarrow (ii.1 \ y = ii.2 \ y) \rightarrow \\ (mm.1 \ lo = mm.2 \ lo \wedge mm.1 \ y = mm.2 \ y))$$

The precondition states that P_1 must start in an initial heap i containing the five pointers x, y, z, lo, hi , with appropriately-typed contents. The heap i may be larger still; this is stated by existentially quantifying over the heap variable j . Heaps are (finite) maps from pointers to values; $x \mapsto u$ is a singleton heap containing only the location x storing value u ; and \bullet is *disjoint* heap union. The precondition insists that i be a disjoint union of smaller singleton heaps; hence there be no aliasing between the five pointers. The postcondition binds over three variables $rr : \text{unit}^2$, $ii, mm : \text{heap}^2$ which are, respectively, the pair of return values, the pair of initial heaps and the pair

of ending heaps for the two runs of P_1 . The postcondition states that if the contents of lo and y in the two initial heaps are equal (hence low), then they are low in the output heaps too.

Other types for P_1 are possible too. For example, we may specify that only the last digit of y is low, by replacing $ii.1\ y$ with $(ii.1\ y) \bmod 10$ in the postcondition, and similarly with $ii.2$, $mm.1$ and $mm.2$. Or, the postcondition may state that the contents of x and z are low at the end of P_1 , though not at the beginning. RHTT (like [3]) can deem arbitrary expressions as low, even though they may have high subparts. The only requirement is that the values of the expressions in two runs are the same. Because we are considering full functional verification, which STsec type a program should have is a matter of programmer's choice. The system merely issues a proof obligation that the desired type is indeed valid, to be discharged interactively, using the logic we outline in Section 4. This proof obligation may not only be about security but also may concern full functional correctness.

Opaque sealing The ascription of STsec types in RHTT is opaque, as mentioned earlier in this section. Even if program execution makes more values low, this knowledge cannot be utilized by clients if it is not exposed in the postcondition. For example, using P_1 's type U , program

$$P_2 \hat{=} \text{fun } x\ y\ z\ lo\ hi. \text{do } (P_1\ x\ y\ z\ lo\ hi; t \leftarrow \text{read } x; \text{return } t),$$

cannot be given a type in which t is low, because the postcondition in U does not expose the property that x is low at the end of P_1 .

Local contexts While the STsec type of P_1 classifies the security of the *contents* of x , y , z , lo , hi , it cannot classify the pointer addresses themselves, as the latter requires discerning the address names in the two different runs (e.g., $x.1$ and $x.2$). We therefore extend the STsec constructor with a *local context*, which is a list of types of the variables we consider local to the computation. For example, the type for P_1 in which the five pointer addresses are high, even though the contents of lo and y are low, can be written as follows, using the list `[ptr, ptr, ptr, ptr, ptr]` as the local context.

$$\begin{aligned} & \text{STsec } [\text{ptr}, \text{ptr}, \text{ptr}, \text{ptr}, \text{ptr}] \text{ unit} \\ & (\text{fun } x\ y\ z\ lo\ hi:\text{ptr}. i:\text{heap}. \\ & \quad \exists u\ v\ w\ c:\text{nat}. b:\text{bool}. j:\text{heap}. \\ & \quad \quad i = x \mapsto u \bullet y \mapsto v \bullet z \mapsto w \bullet hi \mapsto b \bullet lo \mapsto c \bullet j, \\ & \quad \text{fun } xx\ yy\ zz\ llo\ hhi:\text{ptr}^2. rr:\text{unit}^2. ii\ mm:\text{heap}^2. \\ & \quad ii.1\ llo.1 = ii.2\ llo.2 \rightarrow ii.1\ yy.1 = ii.2\ yy.2 \rightarrow \\ & \quad mm.1\ llo.1 = mm.2\ llo.2 \wedge mm.1\ yy.1 = mm.2\ yy.2) \end{aligned}$$

The type of the precondition (and similarly for postconditions) now changes to $\text{ptr}^5 \rightarrow \text{heap} \rightarrow \text{prop}$, so that we can bind additional names for the pointers x , y , ... in the precondition, and *pairs* of pointers xx , yy , ... in the postcondition. The program syntax changes too, as the local variables now have to be bound within the scope of `do`. In other words, our program now looks like

$$P_3 \hat{=} \text{do } (\text{fun } x\ y\ z\ lo\ hi. \text{write } z\ 1; \dots).$$

Remark 1. Ordinary function arguments, corresponding to the \rightarrow and Π -types, can be viewed as a special kind of STsec-local arguments, where the security level is low by default. Indeed, any

function $f:\Pi x:A. \text{STsec } \Gamma (p \ x, q \ x)$ can be transformed into

$$\begin{aligned} & \text{do } (\text{fun } x \ \gamma_1 \dots \gamma_n. f \ x \ \gamma_1 \dots \gamma_n) : \\ & \text{STsec } (A::\Gamma) \ B \\ & \quad (\text{fun } x \ \gamma_1 \dots \gamma_n. p \ x \ \gamma_1 \dots \gamma_n, \\ & \quad \text{fun } xx \ \gamma\gamma_1 \dots \gamma\gamma_n \ yy \ ii \ mm. \\ & \quad \quad xx.1 = xx.2 \rightarrow q \ xx \ \gamma\gamma_1 \dots \ yy \ ii \ mm) \end{aligned}$$

Here, the variables $\gamma_1, \dots, \gamma_n$ are typed with types from the local context Γ , and the postcondition explicitly declares x to be low, by inserting the hypothesis $xx.1 = xx.2$. To summarize, function arguments are always low, whereas variables in local contexts may be low, high, or subject to a more precise security specification, depending on the postcondition.

Example 1 (Nested conditionals). The following program is adapted from Simonet [45]. It uses low arguments a, b, c, u, v , and a high argument x which is declared in the local context but is unrestricted by the postcondition. It nests two conditionals to compute the final result, but the result is independent of x , and hence is low. Owing to the non-trivial implicit control flow, however, most security type systems will not be able to establish this independence and typecheck the example accordingly. Simonet’s type system for sum types *can* typecheck the example using types annotated with matrices containing security levels. In contrast, in RHTT the type can precisely describe the final result, y , as a function of the inputs: $y \hat{=} (a = c) \parallel (b = c) \parallel u \parallel v$. Clearly y does not depend on x and we prove that $yy.1 = yy.2$ in the postcondition.

$$\begin{aligned} P_4 : & \ \Pi a \ b \ c \ u \ v:\text{bool}. \text{STsec } [\text{bool}] \ \text{bool} \\ & \quad (\text{fun } x \ i. \ \text{True}, \\ & \quad \quad \text{fun } xx \ yy \ ii \ mm. \ yy.1 = yy.2 \wedge mm = ii \wedge \\ & \quad \quad \quad yy.1 = (a = c) \parallel (b = c) \parallel u \parallel v) \hat{=} \\ & \text{fun } a \ b \ c \ u \ v. \\ & \quad \text{do}(\text{fun } x. \ t \leftarrow \text{if } u \ \text{then} \\ & \quad \quad \text{if } x \ \text{then return } a \ \text{else return } b \\ & \quad \quad \text{else} \\ & \quad \quad \text{if } v \ \text{then return } a \ \text{else return } c; \\ & \quad \text{return } ((t = a) \parallel (t = b))) \end{aligned}$$

Example 2 (Access control through abstraction). What if we want to allow read access, but not write access to some data (or vice-versa), or that access should be made conditional upon successful authentication? To enforce this kind of access control, we employ the standard abstraction mechanisms of type theory, such as abstract types, predicates and modules. The data to be protected can be hidden behind module boundaries, so that it can be accessed only via dedicated methods that enforce access control. For example, let Alice be a module storing some integer data, say salary, whose integrity should be enforced: Alice allows the salary to be readable globally, but only Alice herself can update it, to keep the value coherent with promotions at work. Thus, she exports unconstrained functions for creating new instances and for reading the salary, but the function for writing requires a check against a password that is also stored locally. The signature, `AliceSig` in Figure 1, presents the specifications that Alice wants to export, and a possible module implementing `AliceSig` by keeping two local pointers – one for the salary, one for the password – is given in Figure 2. Referring to Figure 2, the method `new` takes a `nat` salary and a `string` password, and

```

alice : Type
sshape : alice2 → nat2 → string2 → heap2 → prop
shape ≐ fun a s p h. sshape (a, a) (s, s) (p, p) (h, h)
srefl : ∀aa ss pp ii. sshape aa ss pp ii →
    shape aa.1 ss.1 pp.1 ii.1 ∧ shape aa.2 ss.2 pp.2 ii.2
new : nat → string → STsec nil alice
    (fun i. True,
     fun aa ii mm.
       ∃ss pp hh. mm = ii •• hh ∧ sshape aa ss pp hh)
read_salary : STsec [alice] nat
    (fun a i. ∃s p j h. i = j • h ∧ shape a s p j,
     fun aa yy ii mm. ∀ss pp jj hh.
       ii = jj •• hh → sshape aa ss pp jj →
       mm = ii ∧ yy = ss)
write_salary : nat → string → STsec [alice] unit
    (fun a i. ∃s p j h. i = j • h ∧ shape a s p j,
     fun aa qq yy ii mm. ∀ss pp jj hh.
       ii = jj •• hh → sshape aa ss pp jj →
       ∃jj' ss'. mm = jj' •• hh ∧ sshape aa ss' pp jj')

```

Figure 1: AliceSig: access control via abstract predicates.

```

type alice ≐ ptr × ptr
salary (a : alice) ≐ a.1
passwd (a : alice) ≐ a.2
sshape (aa : alice2) (ss : nat2) (pp : string2) (ii : heap2) ≐
    ii.1 = salary aa.1 ↦ ss.1 • passwd aa.1 ↦ pp.1 ∧
    ii.2 = salary aa.2 ↦ ss.2 • passwd aa.2 ↦ pp.2 ∧
    ss.1 = ss.2 ∧ pp.1 = pp.2
new s p ≐ do (x ← alloc s; y ← alloc p; return (x, y))
read_salary ≐ do (fun a. read (salary a))
write_salary s p ≐
    do (fun a. x ← read (passwd a);
       if x = p then write (salary a) s else return ())

```

Figure 2: Implementation of AliceSig.

generates a new instance of Alice, initialized with this data; `read_salary` takes a local-context argument representing Alice, and returns her current salary; `write_salary` takes a new salary, a password, and an `alice` argument in the local context, and updates the salary only if the supplied password matches the password stored in the `alice` argument. Referring back to Figure 1, `AliceSig` specifies an abstract type `alice`, abstract predicates `sshape` and `shape`, a relation, `srefl`, between `shape` and `sshape`, and the types of the methods. Although these figures may look complicated, the reader should bear in mind that they are intended for full functional verification. Also, the definitions of the various abstract predicates and types such as `sshape`, `shape` and `alice`, will be hidden from the clients, and do not contribute to the complexity.

The `sshape` predicate is a relational invariant of the module’s local state (i.e., invariant over two runs). It is parametrized over pairs of alices, `nat` salaries, `string` passwords, and heaps that are current during execution. The parametrization by all these values captures that different instances of Alice that may be allocated at run time all have different local states, which can potentially store different salaries and passwords. If we were not interested in tracking the changes to salaries and passwords, but only in restricting write access, then these can be omitted from `sshape`, resulting in fewer quantifiers and hence simpler `STsec` types for the methods.

For use in preconditions for access control, we employ the non-relational variant `shape` which is a diagonal of `sshape`, as constrained by `srefl`. Recall that a computation in `RHTT` can access locations only in those heaps that provably satisfy its precondition. Correspondingly, a method that wants to access Alice’s local state, has to describe the desired parts of that state in its own precondition. This is why `AliceSig` keeps `sshape` and `shape` abstract. The abstraction hides the layout of Alice’s local state from the clients, thus preventing them from describing the layout in their preconditions and forcing them to access Alice’s local state exclusively via the exported methods. Apart from giving code for the methods, the implementation also provides a *proof* of `srefl` (elided here, but present in the Coq scripts).

The `STsec` types in Alice’s methods describe several additional properties. For example, that the local state of each instance of Alice is disjoint from that of another instance. For `new`, this is achieved by stating that the pair of ending heaps `mm` extends the initial ones `ii` by newly allocated sections `hh` ($mm = ii \bullet\bullet hh$). Here $\bullet\bullet$ generalizes the disjointness operator \bullet to pairs of heaps, that is, $(i_1, i_2) \bullet\bullet (h_1, h_2) = (i_1 \bullet h_1, i_2 \bullet h_2)$. For `read_salary`, we allow that the state in which the function executes be larger than the module’s local state by allowing $ii = jj \bullet\bullet hh$ where `jj` names the local state and `hh` is the potential global part. For `write_salary` we require that the global part, `hh`, remain invariant, but the local part may be changed by storing the new salary.

Finally, the specifications expose that `read_salary` does not change Alice’s local state ($mm = ii$ in the postcondition). On the other hand, `write_salary` may change the salary field, but not the password field, as the `sshape` predicate changes from using the salary `ss` to using `ss'`, but `pp` persists.

Notice that the salary and password arguments in `new` and `write_salary` are ordinary function arguments, whereas `alice` is in the local context of `STsec` in `read_salary` and `write_salary`. Thus, within the scope of Alice’s methods, the salary and the password are low (c.f. Remark 1) whereas the `alice` argument is high because it is unconstrained by the methods’ pre- and postconditions. Of course, as far as clients of `AliceSig` are concerned, all three of these are high: the abstraction over `sshape` hides all relations between the stored values.

Example 3 (Declassification). One consequence of making salary and password internally low is that whenever a new instance of Alice is allocated, or a salary of an existing instance is changed, the

salary and password have to be computed only out of low arguments – it is not possible for Alice to store confidential data into her local fields. Additionally, the specifications of `new` and `write_salary` must hide that the stored salary and password are equal to the supplied ones. The latter are internally low, while the former are to be externally high. The hiding is achieved by existential quantification over ss and pp in the postcondition of `new`, and over ss' in the postcondition of `write_salary`.

Alice can use the internal knowledge that salary and password are low, to implement and export an additional function which declassifies her salary – that is, reveals the internal knowledge that the salary is low. This declassification can be based on arbitrary conditions – say, it is only granted if a correct password has been supplied.

```

declassify :  $\Pi p:\text{string}.$  STsec [alice] bool
  (fun a i.  $\exists s q j h. i = j \bullet h \wedge \text{shape } a s q j,$ 
   fun aa yy ii mm.  $\forall ss qq jj hh.$ 
     ii = jj  $\bullet\bullet$  hh  $\rightarrow$  sshape aa ss qq jj  $\rightarrow$ 
      $\exists jj'. mm = jj' \bullet\bullet$  hh  $\wedge$  sshape aa ss qq jj'  $\wedge$ 
     yy.1 = yy.2  $\wedge$  yy = (p = qq.1, p = qq.2)  $\wedge$ 
     yy.1  $\rightarrow$  ss.1 = ss.2)  $\hat{=}$ 
  fun p. do (fun a. x  $\leftarrow$  read (passwd a); return (p = x))

```

The code of `declassify` checks if the supplied password equals the stored one, and returns the corresponding boolean. `declassify` does not return the value of the salary; for that, one has to use `read_salary`, but the specification of `declassify` shows that the salary is low if `declassify` returned `true` ($yy.1 \rightarrow ss.1 = ss.2$). This is possible because the low-status of the salary has been hardwired into the implementation of `sshape`, and Alice can reveal it at will.

Example 4 (State-based policies). Alice can implement policies that change depending on her local state. For example, she may control the granting of read access with functions `grant` and `revoke`, as specified in Figure 3. These enable and disable reading by, respectively, adding and removing a new abstract predicate – `readable` – from the knowledge exposed about Alice’s local state. Typically, such functions require authentication, but for simplicity, we forgo that aspect. The postcondition of `grant` exposes that the newly obtained state jj' is readable, while `revoke` omits this property, thus revoking the read access. To associate the predicate with reading, the specification of `read_salary` has to require a proof of `readable`.

The signature keeps `readable` abstract, so that the only way `readable` can be derived is if `readable` has been placed into the proof context by a previous call to `grant`, without an intervening `revoke`. The signature can be implemented by extending Alice’s state with an additional boolean pointer that is set and reset by `grant` and `revoke`: `readable` is in force once the boolean is set `true`. Our Coq scripts provide several different implementations of this interface.

Example 5 (Conditional access and erasure policies). Suppose Alice wants to download a program from Bob for computing tax returns. Alice is willing to let Bob access her local state and read her salary directly using `read_salary`, but wants to prevent Bob from stealing her secret by copying it into his own local state. Alice may insist that Bob not keep any local state, or that he deallocate all of it before termination. But this is too restrictive, for Bob may want to keep in his local state a count of how many times his program has executed. Such local state should be allowed to escape the function call as it is independent of Alice’s salary. In RHTT, Alice can formulate such a permissive policy.

```

rreadable : alice2 → heap2 → prop
readable ≡ fun a h. rreadable (a, a) (h, h)
rrefl : ∀aa ii. rreadable aa ii →
        readable aa.1 ii.1 ∧ readable aa.2 ii.2
grant : STsec [alice] unit
  (fun a i. ∃s p j h. i = j • h ∧ shape a s p j,
   fun aa yy ii mm. ∀ss pp jj hh.
    ii = jj •• hh → sshape aa ss pp jj →
    ∃jj'. mm = jj' •• hh ∧ sshape aa ss pp jj' ∧
    rreadable aa jj')
revoke : STsec [alice] unit
  (fun a i. ∃s p j h.
   i = j • h ∧ shape a s p j ∧ readable a j,
   fun aa yy ii mm. ∀ss pp jj hh.
    ii = jj •• hh → sshape aa ss pp jj →
    ∃jj'. mm = jj' •• hh ∧ sshape aa ss pp jj')
read_salary : STsec [alice] nat
  (fun a i. ∃s p j h.
   i = j • h ∧ shape a s p j ∧ readable a j,
   fun aa yy ii mm. ∀ss pp jj hh.
    ii = jj •• hh → sshape aa ss pp jj →
    mm = ii ∧ yy = ss)

```

Figure 3: Extension of AliceSig with state-based read access.

```

bbshape : bob2 → G2 → heap2 → prop
bshape ≡ fun b k t. bbshape (b, b) (k, k) (t, t)
brefl : ... (* similar to srefl *)

epre (a : alice) (b : bob) (j i : heap) ≡
  ∃ s p k t h. i = j • t • h ∧ shape a s p j ∧ bshape b k t
epost aa bb jj yy ii mm ≡
  ∀ ss pp kk tt hh. ii = jj •• tt •• hh →
    sshape aa ss pp jj → bshape bb kk tt →
  ∃ jj' tt'. mm = jj' •• tt' •• hh ∧ sshape aa ss pp jj' ∧
    bshape bb (bcmp kk.1, bcmp kk.2) tt' ∧
    (tt.1 = tt.2 → tt'.1 = tt'.2)

```

Figure 4: Some definitions for conditional access and erasure policies.

We start the description of Alice’s specification by assuming that `bob` is an abstract type representing Bob’s local state (usually implemented as the set of root pointers for Bob’s local state). G is another abstract type representing the values that Bob keeps in his local state. For example, if Bob wants to count how many times his program has been invoked, then $G = \text{nat}$. Further, we assume `bbshape`: $\text{bob}^2 \rightarrow G^2 \rightarrow \text{heap}^2 \rightarrow \text{prop}$ is an abstract predicate describing Bob’s local state, and `bcmp`: $G \rightarrow G$ describes how Bob’s local state changes within one function call. In the counting example, `bcmp` will be the program Bob needs to run over both Alice’s and Bob’s local heap. The initial heaps i for his program can therefore be split in three ways: j belonging to Alice, t belonging to Bob, and the remainder h that is untouched. Predicate `epre` in Figure 4 describes this situation. On the other hand, `epost` states that Bob’s local state tt' at the end stores the correct statistics (`bcmp` of $kk.1$ and $kk.2$), and if Bob’s initial local state tt is assumed low, then tt' is low as well. In other words, Bob did not copy into tt' any of the high values that he may have read from Alice. A program that requests read access to Alice’s local store, and respects the described policy has the type

```

T ≡ STsec [alice, bob] nat
  (fun a b i. ∃ j. readable a j ∧ epre a b j i,
   fun aa bb yy ii mm. ∀ jj.
     rreadable aa jj → epost aa bb jj yy ii mm)

```

Alice now wants to ratify programs with type T by granting them read access to her salary. She can do so by exporting from her module a function `ratify` which removes `readable` from T , much like the `grant` program would do. After that, Bob’s program can execute without needing special reading privileges. In this respect, `ratify` is a *higher-order* function because in its type, `STsec` appears in a negative (argument) position. `ratify` can be said to implement a conditional access policy, because it grants access only after Bob supplies a proof that his program satisfies the type T , i.e., the

program does not leak Alice's salary.

```

ratify : T →
  STsec [alice, bob] nat
    (fun a b i. ∃j. epre a b j i,
     fun aa bb yy ii mm.
       ∀jj. epost aa bb jj yy ii mm) ≐
  fun e : T. do (fun a b. e a b)

```

This specification can be instantiated in several ways, by choosing different values for `bob`, `G`, `bbshape` and `bcmp`. For example, if `bob ≐ unit`, `G ≐ unit`, and `bbshape` states that Bob's local heap is empty, then Bob's program does not keep any local state, and Alice, correctly, ratifies it. Alternatively, if `bob ≐ ptr`, `G ≐ nat`, `bcmp ≐ succ` and `bbshape bb kk tt ≐ (tt.1=bb.1 ↦ kk.1 ∧ tt.2 = bb.2 ↦ kk.2)` then Bob's program keeps a single pointer whose content is incremented by 1 after every execution.

Bob's program which computes the tax of 24% of Alice's salary, while also keeping its invocation count, can be implemented and then immediately ratified by the following function call. Notice that by the type of `ratify`, the return value of Bob's program is high as there is no requirement `yy.1 = yy.2` in `epost`. Hence, the fact that this value is a function of Alice's salary, is not a security leak.

```

ratify (do (fun a:alice b:bob.
            x ← read_salary a; k ← read b;
            write b (k + 1); return (x * 24%)))

```

Suppose Bob keeps the count with two `nat` pointers, whose contents `p` and `q` are both increased at every call, so that the overall count is the difference between the two. This is represented by taking `bob ≐ ptr × ptr` (one `ptr` for `p` and one for `q`), `G ≐ nat`, `bcmp ≐ succ` and `bbshape bb kk tt` is

$$\begin{aligned}
&\exists pp \ qq: \text{nat}^2. tt.1 = \text{fst } (bb.1) \mapsto pp.1 \bullet \text{snd } (bb.1) \mapsto qq.1 \wedge \\
&\quad tt.2 = \text{fst } (bb.2) \mapsto pp.2 \bullet \text{snd } (bb.2) \mapsto qq.2 \wedge \\
&\quad kk = (pp.1 - qq.1, pp.2 - qq.2)
\end{aligned}$$

Bob's program can read Alice's salary, then increment `p` and `q` by amount of the salary, and additionally, increment `p` by 1. In terms of required specifications for Bob's local state, the program still keeps the invocation count. However, the program is actually stealing Alice's salary, because the salary can be inferred by deducting the old value of `q` from the new one. Bob will fail to get such a program ratified by Alice, if he calls `ratify` with the argument

```

do (fun a:alice b:bob.
    x ← read_salary a; p ← read (fst b); q ← read (snd b);
    write (fst b) (p + x + 1); write (snd b) (q + x);
    return (x * 24%))

```

`ratify` forces Bob to prove that his ending state is low ($tt'.1 = tt'.2$) as defined in `epost`, but this is not provable if `tt'` stores Alice's salary `x`. Indeed, as `x` is high, Bob lacks the information that `x` is equal in the two runs, so he cannot prove that his pointers store equal values in two runs. For ratification, Bob will have to *erase* Alice's salary from his state, perhaps by mutating his pointers to store `p + 1` and `q` instead of `p + x + 1` and `q + x`. `ratify` may thus be said to implement an *erasure policy*, similar to those of Chong and Myers [18, 19]. Alternatively, Bob may try to *declassify* Alice's salary, using the function from Example 2, but then he has to provide the correct

3 Typing rules and Semantic Model

Each command of the stateful fragment of RHTT comes with a dependent `STsec` type that captures the command’s specification using pre- and post-conditions. We start our description with the types of the basic commands; descriptions of the other commands appear later in the section.

```

return : STsec [A] A
        (fun x i. True,
         fun xx yy ii mm. mm = ii ∧ yy = xx)

read   : STsec [ptr] A
        (fun ℓ i. ∃h:heap v:A. i = ℓ ↦ v • h,
         fun ℓℓ yy ii mm. mm = ii ∧ ∀hh vv.
           ii = (ℓℓ.1 ↦ vv.1, ℓℓ.2 ↦ vv.2) •• hh → yy = vv)

write  : STsec [ptr, A] unit
        (fun ℓ v i. ∃h B:type w:B. i = ℓ ↦ w • h,
         fun ℓℓ vv yy ii mm. ∀hh B1 B2 w1:B1 w2:B2.
           ii = (ℓℓ.1 ↦ w1, ℓℓ.2 ↦ w2) •• hh →
           mm = (ℓℓ.1 ↦ vv.1, ℓℓ.2 ↦ vv.2) •• hh)

dealloc : STsec [ptr] unit
        (fun ℓ i. ∃h B:type w:B. i = ℓ ↦ w • h,
         fun ℓℓ yy ii mm. ∀hh B1 B2 w1:B1 w2:B2.
           ii = (ℓℓ.1 ↦ w1, ℓℓ.2 ↦ w2) •• hh → mm = hh)

```

`return` immediately terminates with the value that was supplied as a local argument. Its `STsec` constructor records the argument type in the local context, and the type of the returned value (here, both types are A). The precondition states that `return` can execute in any heap, as it performs no heap operations. The postcondition states that `return` does not change the input heaps ($mm = ii$) and passes the input argument to the output ($yy = xx$). The precondition of `read` `write` and `dealloc` all require that the initial heaps contain at least the pointer ℓ to be read from, written to or deallocated. In the case of `read`, the contents of the pointer must have the expected type A . For `write` and `dealloc`, this type is irrelevant and is hence existentially quantified. The postconditions of all three commands explicitly describe the layout of the new heap and, in particular, state that parts of the input heaps that are disjoint from ℓ (hh above) remain invariant.

Allocation presents the following challenge. If under a high guard, a pointer is allocated in one branch of a conditional, but not in the other, this may constitute a leak of the high guard, if the pointer itself is of low security. Such “unmatched” allocations should therefore always produce high pointers. This is why we provide two allocation primitives: `lalloc` for allocating low pointer

addresses, and `alloc`, for allocating high ones.

```

lalloc : STsec [A] ptr
  (fun v i. True, fun vv yy ii mm.
    mm = (yy.1 ↦ vv.1, yy.2 ↦ vv.2) •• ii ∧
    (ii.1 ≅ ii.2 → yy.1 = yy.2 ∧ mm.1 ≅ mm.2))

alloc : STsec [A] ptr
  (fun v i. True, fun vv yy ii mm.
    mm = (yy.1 ↦ vv.1, yy.2 ↦ vv.2) •• ii ∧
    even yy.1 ∧ even yy.2)

```

Both commands take a local argument $v:A$, and return a *fresh* pointer initialized with v . The freshness is captured in the postcondition by demanding that the initial heaps ii be *disjoint* from the returned pointers yy in the equation for the ending heaps mm . However, `alloc` chooses the returned location non-deterministically, while `lalloc` is *deterministic*; that is, it returns *equal* (and hence low) pointers, when invoked under appropriate conditions. We make the two allocators operate on disjoint pools of locations: `alloc` always returns an *even* pointer (albeit, a randomly chosen one), while `lalloc` returns the next unallocated *odd* pointer. Here we rely on the property that type `ptr` is isomorphic to `nat` in our model.

Definition 1. Heaps h_1 and h_2 are *low-equivalent*, written $h_1 \cong h_2$, iff their domains contain the same odd pointers. The content of the pointers is irrelevant.

The postconditions of `lalloc` and `alloc` further capture the behavior of the two commands with respect to the \cong relation. In the case of `lalloc`, we expose that if invoked in low-equivalent input heaps ($ii.1 \cong ii.2$), the command returns equal pointers ($yy.1 = yy.2$), and low-equivalent output heaps ($mm.1 \cong mm.2$). In the case of `alloc`, we expose the evenness of $yy.1$ and $yy.2$, and provide a number of lemmas, that can be used to relate evenness with \cong . For example, the lemma

$$\forall x:\text{ptr}. \text{even } x \rightarrow (x \mapsto v \bullet h_1 \cong h_2) \leftrightarrow (h_1 \cong h_2)$$

when iterated, can show that low equivalence of h_1 and h_2 is preserved after arbitrary number of high allocations. Other related lemmas are present in our Coq scripts.

Example 6 (The need for both allocators). The following program can be given a type in which the returned pointer y is low, no matter what the boolean h is.

```

do (fun h. if h then y ← lalloc 2; return y else
    x ← alloc 1; y ← lalloc 2; dealloc x; return y) :
STsec [bool] ptr
  (fun h i. True,
   fun hh yy ii mm. mm = ii •• (yy.1 ↦ 2, yy.2 ↦ 2) ∧
    (ii.1 ≅ ii.2 → yy.1 = yy.2))

```

The program does not typecheck if the high allocation of x is replaced by `lalloc`. In that case, it is possible that the two executions of the program select different branches of the conditional (depending on h). If we started with low-equivalent heaps $i_1 \cong i_2$, then at the point of allocation of y , the heaps will not be low equivalent anymore, since one of them has been extended with an odd location x , while the other has not. Thus, we cannot conclude that the returned pointer is low ($yy.1 = yy.2$).

Remark 2. Deterministic allocation forces `STsec` to use *large-footprint specifications*, whereby specifications describe the full heaps in which commands operate. This is in contrast to separation logic, where specifications describe only those heap parts that commands touch, and implicitly assume invariance of the remaining heap. The latter style is more succinct, but cannot support deterministic allocation [53]. With large footprints, we can specify `lalloc` (specifically, the antecedent $ii.1 \cong ii.2$ in the postcondition), but the invariance of untouched parts of heaps has to be stated explicitly for every program, as witnessed by the quantification over hh in the postconditions of `write` and `dealloc`. Note that *the concrete layouts of untouched parts of heaps do not need to appear in the specifications* — thus alleviating concerns of scalability of specifications. Moreover, the overhead between large and small footprint specifications is *constant*, as we discuss in Section 6. The two styles also lead to similar proofs. What matters in proofs is the ability to effectively reason about heap disjointness, and we can do that equally well in both styles by relying on the operator \bullet [36].

Another way of treating allocation in the relational setting is to model its non-determinism by means of partial bijections between pointers [3, 8]. Then one can avoid using two different allocators, albeit at a price of increasing the complexity of reasoning. Such proposals, however, only work in the *absence of deallocation*. For example, the definition of noninterference of Amtoft et al. [3] allows that the input heaps to the computation are related by some bijection between pointers, and requires that the ending heaps are also related by a bijection. However, the ending bijection has to be an *extension* of the initial one. Obviously, such a definition cannot support deallocation, as deallocation produces smaller, not larger heaps. Alternatively, one can omit the extension requirement; but that leads to counterexamples which satisfy the weakened requirement even though they actually leak information. For example, consider a 2-node, acyclic, singly-linked list, with the usual implementation: each node has a *data* field and a *nxt* pointer that points to the next node in the list. Let p be the list header, and let $p.data = 0$, $p' = p.nxt$ and $p'.data = 0$. Suppose also that all *data* are low and all *nxt* pointers are low. Now consider the conditional

if high then *flip* else skip

where *flip* reverses the list, so that now the list header is p' . Suppose the first run of the program takes the *then* branch and the second run takes the *else* branch. In this case there exists a bijection $\{(p', p), (p'.nxt, p.nxt), (p'.data, p.data)\}$. Yet, by knowing that p and p' are different addresses, information is leaked about the guard.

We proceed to describe our constructor for sequential composition, but first we need some notational conventions. Let Γ be a list of types. We denote by $\bar{\Gamma}$ the product of all the types in Γ , e.g., $\bar{\text{nil}} = \text{unit}$ and $\overline{[A, B, C]} = A \times B \times C$. We further conflate the function types $\bar{\Gamma} \rightarrow T$ and $\Gamma_1 \rightarrow \Gamma_2 \rightarrow \dots \rightarrow T$, and their corresponding terms. For example, we freely interchange `fun γ : $\overline{[A, B, C]}$. . .`, or `fun γ . . .` if the types are clear from the context, with `fun $x:A$ $y:B$ $z:C$. . .`. Similarly, we interchange $e(x, y, z)$ with $e\ x\ y\ z$. We hope that no confusion arises due to this abuse of notation; all of our exposition has been checked in Coq, where the notation is formally resolved.

For sequential composition $e_1; e_2$, let $e_1 : \text{STsec } \Gamma\ A\ (p_1, q_1)$ and $e_2 : \text{STsec } (A::\Gamma)\ B\ (p_2, q_2)$. Then $e_1; e_2$ first executes e_1 , passing the returned value as the first local argument to e_2 . Assuming $\gamma:\bar{\Gamma}$

and $\gamma\gamma:\bar{\Gamma}^2$, the STsec type for $e_1; e_2$ is

```

STsec  $\Gamma B$ 
(fun  $\gamma i. p_1 \gamma i \wedge$ 
     $\forall y m. q_1 (\gamma, \gamma) (y, y) (i, i)(m, m) \rightarrow p_2 (y, \gamma) m,$ 
    fun  $\gamma\gamma yy ii mm.$ 
     $\exists vv:A^2. hh:heap^2. q_1 \gamma\gamma vv ii hh \wedge$ 
     $q_2 ((vv.1, \gamma\gamma.1), (vv.2, \gamma\gamma.2)) yy hh mm)$ 

```

In English, the precondition requires e_1 to be safe in the initial heap of the sequential composition, and that any value y and heap m obtained as output of e_1 – and which thus satisfy e_1 ’s “squared” postcondition – make e_2 safe. The postcondition states that intermediate values vv and heaps hh exist, obtained after running e_1 but before running e_2 .

As e_1 ’s output is bound in the local context of e_2 , we cannot treat this output as an ordinary functional variable, despite our suggestive notation in Section 2. Indeed, as discussed previously in Section 2, ordinary variables are always low, whereas the ones in the local context may be high, depending on the specification. Thus, we must rely on variable-free representation via *combinators*, as described next.

Our first combinator is for changing the local context of an STsec type. Given $\Gamma_1, \Gamma_2, f:\bar{\Gamma}_1 \rightarrow \bar{\Gamma}_2$, and $e:\text{STsec } \Gamma_2 A (p, q)$, we can *instantiate* the local variables of e according to f , to produce a computation with context Γ_1 .

```

 $e @ f : \text{STsec } \Gamma_1 B$ 
    (fun  $\gamma. p (f \gamma), \text{fun } \gamma\gamma. q (f \gamma\gamma.1, f \gamma\gamma.2))$ 

```

We denote by $e @_0 \gamma$ the special instance of $@$, when $\Gamma_1 = \text{nil}$ and hence, f is isomorphic to a tuple $\gamma:\bar{\Gamma}_2$. We refer to e in $e @ f$ or $e @_0 \gamma$ as the *head* of the instantiation, and to f as the *explicit substitution*.

Example 7. In Example 2, we implemented `declassify` as

```

fun  $p. \text{do } (\text{fun } a. x \leftarrow \text{read } (\text{passwd } a); \text{return } (p = x))$ 

```

The actual implementation using combinators is

```

fun  $p:\text{string}. \text{do } (\text{read } @ (\text{fun } a. \text{passwd } a);$ 
     $\text{return } @ (\text{fun } x a. (p = x)))$ 

```

Programs thus become lists of commands instantiated with explicit substitutions, where the domains of substitutions grow with each command to provide names for the results of previous commands. In the above example, the domain of the substitution for `read` includes only the variable $a:\text{alice}$, but the substitution for `return` also includes $x:\text{string}$, which names the result of the previous `read` (Alice’s stored password). Similarly, the functions `new` and `read_salary` are reimplemented as

```

new  $s p \hat{=} \text{do } (\text{alloc } @_0 s; \text{alloc } @ (\text{fun } x. p);$ 
     $\text{return } @ (\text{fun } y x. (x, y)))$ 
read_salary  $\hat{=} \text{do } (\text{read } @ (\text{fun } a. \text{salary } a))$ 

```

This brings us to the combinator `do`, which mediates between the semantics of types, and the logic of assertions. The role of `do` is to change the type of e : $\text{STsec } \Gamma A (p_1, q_1)$ into $\text{STsec } \Gamma A (p_2, q_2)$, if a proof can be constructed, possibly interactively, that e can be ascribed the precondition p_2 and postcondition q_2 . This ascription can be made if two properties are proved.

First, e should be safe to execute in any heap satisfying p_2 ; that is, it should be possible to strengthen the precondition and prove $\forall \gamma: \bar{\Gamma}. i: \text{heap}. p_2 \ \gamma \ i \rightarrow p_1 \ \gamma \ i$. We will use a somewhat different proposition, which exposes that safety can be proved by syntax-directed reasoning on the structure of the involved program – a property we exploit in Section 4.

$$\begin{aligned} C_1(e) &\hat{=} \forall \gamma: \bar{\Gamma}. i: \text{heap}. p_2 \ \gamma \ i \rightarrow \text{safe } (e \ @_0 \ \gamma) \ i \\ \text{safe} &\hat{=} \text{fun } e: \text{STsec nil } (p, q). i: \text{heap}. p \ i \end{aligned}$$

Second, it should be possible to change the postcondition of e into q_2 . Semantically, we should prove that two executions of e result in heaps and values satisfying q_2 . We capture this property via the following predicate, which formalizes the reasoning in a relational Hoare logic over *two programs* [11, 52]. Assuming $e_1: \text{STsec nil } A (r_1, t_1)$ and $e_2: \text{STsec nil } A (r_2, t_2)$, a pair of input heaps ii , and a predicate $q: A^2 \rightarrow \text{heap}^2 \rightarrow \text{prop}$, we define

$$\begin{aligned} \text{verify2 } ii \ e_1 \ e_2 \ q &\hat{=} \\ &\forall yy: A^2. mm: \text{heap}^2. \\ &\quad (ii.1, yy.1, mm.1) \in \text{runs_of } e_1 \rightarrow \\ &\quad (ii.2, yy.2, mm.2) \in \text{runs_of } e_2 \rightarrow q \ yy \ mm \end{aligned}$$

Here, `runs_of` coerces programs into relations between input heaps, output values and output heaps. We will define it further below, for our particular model. `verify2` is almost like a Hoare quadruple from relational Hoare logic except that it lacks the preconditions. We add the preconditions by the following definition.

$$\begin{aligned} C_2(e) &\hat{=} \forall \gamma \gamma: \bar{\Gamma}^2. ii: \text{heap}^2. \\ &\quad p_2 \ \gamma \gamma.1 \ ii.1 \rightarrow p_2 \ \gamma \gamma.2 \ ii.2 \rightarrow \\ &\quad \text{verify2 } ii \ (e \ @_0 \ \gamma \gamma.1) \ (e \ @_0 \ \gamma \gamma.2) \\ &\quad (\text{fun } yy \ mm. q_2 \ \gamma \gamma \ yy \ ii \ mm) \end{aligned}$$

In other words, if the input heaps are assumed to satisfy the new precondition p_2 , then the two instantiations of e satisfy the new postcondition q_2 . Notice how the two instantiations of e are actually different programs, which is why `verify2` had to take two program arguments.

We now allow changing the type of e only if proofs of both $C_1(e)$ and $C_2(e)$ are provided as arguments to the constructor `do`, corresponding to the typing rule:

$$\begin{aligned} \text{do} &: \Pi e: \text{STsec } \Gamma A (p_1, q_1). \\ &\quad C_1(e) \rightarrow C_2(e) \rightarrow \text{STsec } \Gamma A (p_2, q_2). \end{aligned}$$

Remark 3. In practice, our implementation in Coq allows that the proofs of $C_1(e)$ and $C_2(e)$ can be left out when using `do`. In such cases, the system emits the appropriate proof obligations, to be discharged at some later point, possibly interactively. For this reason, we consider `do` to have only one explicit argument e ; the arguments standing for proofs of $C_1(e)$ and $C_2(e)$ can be ignored as they merely serve to guide the generation of verification conditions for e .

Finally, a development similar to the one for `do` has to be carried out for conditionals as well. Given programs $e_i : \text{STsec } \Gamma \ A \ (p_i, q_i)$ for $i=1,2$, corresponding to branches of a conditional, and a boolean guard $b : \bar{\Gamma} \rightarrow \text{bool}$ (here parametrized over a context), which type should we ascribe to the conditional? We would like to be precise, and ascribe the weakest precondition sufficient for the safety, and the strongest postcondition sound wrt. the expected semantics. Unfortunately, computing that postcondition seems impossible in the case when the boolean guard is high. Indeed, we know that q_1 (resp. q_2) relates the output heaps if both runs of the conditional choose the same branch e_1 (resp. e_2), but nothing can be said if the branches chosen in the two runs are different. Since the principal specification cannot be computed, the best we can do is ask the programmer for the desired precondition p and postcondition q , and emit proof obligations for checking that (p, q) is valid for the conditional.

$$\begin{aligned} \text{cond} & : \Pi b : \bar{\Gamma} \rightarrow \text{bool}. \\ & \Pi e_1 : \text{STsec } \Gamma \ A \ (p_1, q_1). \Pi e_2 : \text{STsec } \Gamma \ A \ (p_2, q_2). \\ & D_1(b, e_1, e_2) \rightarrow D_2(b, e_1, e_2) \rightarrow \text{STsec } \Gamma \ A \ (p, q). \end{aligned}$$

Here D_1 captures the safety of the conditional, and D_2 the Hoare-style correctness.

$$\begin{aligned} D_1(b, e_1, e_2) & \hat{=} \forall \gamma \ i. p \ \gamma \ i \rightarrow \\ & \quad \text{safe (if } b \ \gamma \ \text{then } e_1 \ @_0 \gamma \ \text{else } e_2 \ @_0 \gamma) \ i \\ D_2(b, e_1, e_2) & \hat{=} \\ & \forall \gamma \gamma' \ ii. p \ \gamma \gamma'.1 \ ii.1 \rightarrow p \ \gamma \gamma'.2 \ ii.2 \rightarrow \\ & \quad \text{verify2 } ii \ (\text{if } b \ \gamma \gamma'.1 \ \text{then } e_1 \ @_0 \ \gamma \gamma'.1 \ \text{else } e_2 \ @_0 \ \gamma \gamma'.1) \\ & \quad \quad (\text{if } b \ \gamma \gamma'.2 \ \text{then } e_1 \ @_0 \ \gamma \gamma'.2 \ \text{else } e_2 \ @_0 \ \gamma \gamma'.2) \\ & \quad \quad (\text{fun } yy \ mm. q \ \gamma \gamma' \ yy \ ii \ mm) \end{aligned}$$

The definitions of D_1 and D_2 make use of the purely-functional conditional `if` to define when each of the branches is taken. In this paper, we conflate `cond` and `if` and use `if` for both. Note that, in contrast to other relational Hoare logics [11, 52], we do not restrict the reasoning to only the situation where the same branch of the conditional is taken in both runs; nor do we need side conditions, as in Amtoft et al. [3], that prohibit updates of low variables under a high guard (which would prevent verification of P_1 in Section 2).

Example 8. The function `write_salary` from Example 2 is implemented with combinators (omitting annotations and proofs) as follows. Notice that the guard of the conditional is a term with a local context consisting of $a:\text{alice}$ and $x:\text{string}$.

$$\begin{aligned} \text{write_salary } s \ p & \hat{=} \\ & \text{do (read @ (fun } a. \text{passwd } a); \\ & \quad \text{if (fun } x \ a. x = p) \ \text{then} \\ & \quad \quad \text{write @ (fun } x \ a. (\text{salary } a, s)) \\ & \quad \text{else return @ (fun } x \ a. ()))} \end{aligned}$$

Semantic model. We justify the soundness of our type system by building a denotational model for `STsec` types. This development can be fully carried out as a shallow embedding in CiC that we have formalized in Coq. The model is based on predicate transformers [21]. In the first step, we temporarily ignore the local context Γ and the postcondition q , and build the type `prog A p`

of stateful programs that return values of type A and are safe to execute in heaps satisfying p . In order to obtain the right denotational structure and information ordering, we make each element of $\mathbf{prog} A p$ be a function taking a *set of heaps* satisfying p (not just a single heap), and producing a set of possible output values and output heaps. However, we impose an additional condition of *coherence* on such an extension. Formally,

$$\begin{aligned} r \sqsubseteq p &\hat{=} \forall x:\mathbf{heap}. r x \rightarrow p x \\ \mathbf{ideal} p &\hat{=} \{r:\mathbf{heap} \rightarrow \mathbf{prop} \mid r \sqsubseteq p\} \\ \mathbf{prog} A p &\hat{=} \{f:\mathbf{ideal} p \rightarrow A \rightarrow \mathbf{heap} \rightarrow \mathbf{prop} \mid \mathbf{coherent} f\} \end{aligned}$$

Here $f:\mathbf{ideal} p \rightarrow A \rightarrow \mathbf{heap} \rightarrow \mathbf{prop}$ is coherent iff f 's action on a set of heaps is a union of f 's actions on the elements of the set:

$$\forall r:\mathbf{ideal} p. x:A. m:\mathbf{heap}. m \in f r x \leftrightarrow m \in \bigcup_{i:\mathbf{heap}, r i} f \{i\} x$$

By making $f:\mathbf{prog} A p$ apply only to (sets of) heaps satisfying p , as opposed to all heaps, we avoid the need to model programs that go wrong during their execution, perhaps because of memory errors (reading a dangling pointer) or type errors (reading a pointer storing a boolean, when an integer is expected). The typechecker will reject such applications as ill-typed expressions which do not need to be given semantics.

We now define the relation $\mathbf{runs_of} f:\mathbf{heap} \times A \times \mathbf{heap} \rightarrow \mathbf{prop}$, mentioned previously in this section, which relates input heaps, output values and output heaps produced by executions of f .

$$\mathbf{runs_of} f \hat{=} \{(i, y, m) \mid i \in p \wedge m \in f \{i\} y\}$$

It is easy to show that $\mathbf{prog} A p$, with the pointwise ordering, is a poset with least upper bounds for all sets of programs.

Next, we restore the type list Γ . Given parametrized precondition $p:\bar{\Gamma} \rightarrow \mathbf{heap} \rightarrow \mathbf{prop}$, parametrized programs are defined inductively on the structure of Γ as shown below. The poset structure of \mathbf{prog} trivially lifts to \mathbf{pprog} as well.

$$\mathbf{pprog} \Gamma A p \hat{=} \begin{cases} \mathbf{prog} A p & \text{if } \Gamma = \mathbf{nil} \\ \prod x:B. \mathbf{pprog} \Gamma' A (p x) & \text{if } \Gamma = B::\Gamma' \end{cases}$$

Now, $\mathbf{STsec} \Gamma A (p, q)$ is defined as a subset of $\mathbf{pprog} \Gamma A p$, consisting of those programs whose two executions satisfy the postcondition q .

$$\begin{aligned} \mathbf{STsec} \Gamma A (p, q) &\hat{=} \\ &\{c : \mathbf{pprog} \Gamma A p \mid \forall \gamma \gamma \ ii \ yy \ mm. \\ &\quad (ii.1, yy.1, mm.1) \in \mathbf{runs_of} (c \ \gamma \gamma.1) \rightarrow \\ &\quad (ii.2, yy.2, mm.2) \in \mathbf{runs_of} (c \ \gamma \gamma.2) \rightarrow q \ \gamma \gamma \ yy \ ii \ mm\} \end{aligned}$$

Because satisfaction of q is part of the definition of the \mathbf{STsec} type, any program that can be ascribed an \mathbf{STsec} type must automatically satisfy q . In particular, if q asserts noninterference, then it is immediate the program is noninterferent; there is no need for a separate proof of this fact as in a type system such as DCC [1, Theorem 4.2], Fine [46, Theorem 2], or Fable [47, Theorem 25 in the full version].

We have further shown in our Coq scripts that $\text{STsec } \Gamma A (p, q)$ is a complete partial order, thus supporting a combinator for least fixed points of continuous functions between monadic types.

$$\begin{aligned} \text{fix} &: \Pi f : \text{STsec } \Gamma A (p, q) \rightarrow \text{STsec } \Gamma A (p, q). \\ &\text{continuous } f \rightarrow \text{STsec } \Gamma A (p, q) \end{aligned}$$

We have also provided a number of lemmas showing continuity of all the combinators; these are useful for discharging the continuity obligations required by `fix`.

Finally, we note that each of the combinators presented in this section can be given a denotation (also carried out in the Coq scripts). For example, `fix` is modeled using the usual Kleene construction, sequential composition is realized as functional composition of predicate transformers, `do` is the identity function, etc.

4 Logic for discharging verification conditions

As illustrated in Section 3, our program semantics relies on two predicates: `safe e i` establishing that e is safe in the heap i , and `verify2 ii e1 e2 r` establishing that r holds after the execution of e_1 and e_2 in heaps $ii.1$ and $ii.2$, respectively. The proof obligations that RHTT generates are mostly of this form. For example, in the case of a program with conditionals, the system will issue proof obligations in the form of appropriate D_1 and D_2 predicates, to be discharged in interaction with the system. When proving programs involving recursion, the system also issues obligations about domain-theoretic continuity, but we do not discuss those here (our Coq scripts, however, contain a number of useful lemmas for reasoning about continuity). In this section, we illustrate how we reason about obligations with `safe` and `verify2`.

Proving obligations is largely directed by the syntax of the programs that appear as arguments to `safe` and `verify2`. The first case to consider is when the programs in the proof obligation are instantiations $e @ f$. Our strategy in such situations is to push `@` further into the program, to eventually reveal some primitive command at the top-level of e . In the case of `safe` (and `verify2` is similar), we have the following lemmas for dealing with `@`. Application of any of the lemmas replaces the goal in the consequent with an antecedent in which `@` has been pushed inside.

$$\begin{aligned} \text{sval_inst} &: \text{safe } (e @_0 (f \ \gamma)) \ i \rightarrow \text{safe } ((e @ f) @_0 \ \gamma) \ i \\ \text{sbnd_push} &: \text{safe } (e_1 @_0 \ \gamma; e_2 @ (\text{fun } y. (y, \gamma))) \ i \rightarrow \\ &\quad \text{safe } ((e_1; e_2) @_0 \ \gamma) \ i \\ \text{sbnd_inst} &: \text{safe } (e_1 @_0 (f \ \gamma); e_2) \ i \rightarrow \\ &\quad \text{safe } ((e_1 @ f) @_0 \ \gamma; e_2) \ i \end{aligned}$$

The second case is when e has been revealed as the head of the first (or the sole) command in the program, and e is not itself an instantiation. If the type of e is $e : \text{STsec } \Gamma A (p, q)$, we reason using p and q .

$$\begin{aligned} \text{sval_do} &: p \ \gamma \ i \rightarrow \text{safe } (e @_0 \ \gamma) \ i \\ \text{sbnd_do} &: p \ \gamma \ i \rightarrow (\forall y \ m. q \ (\gamma, \gamma) \ (y, y) \ (i, i) \ (m, m) \rightarrow \\ &\quad \text{safe } (e_2 @_0 \ y) \ m) \rightarrow \\ &\quad \text{safe } (e @_0 \ \gamma; e_2) \ i \end{aligned}$$

The above two lemmas will typically be applied if e is a call to a previously verified program, a program variable, a conditional or a fixed point.

The third case is when e is one of the primitive commands. Then `sval_do` and `sbnd_do` can be specialized to utilize the knowledge of p and q of the particular commands. For example, we show below the specialization of `sval_do` to `return` and of `sbnd_do` to `alloc`.

$$\begin{aligned} \text{sval_ret} &: \text{safe } (\text{return } @_0 x) i \\ \text{sbnd_alloc} &: (\forall x:\text{ptr}. \text{safe } (e @_0 x) (x \mapsto v \bullet i)) \rightarrow \\ &\quad \text{safe } (\text{alloc } @_0 v; e) i \end{aligned}$$

In `sval_ret`, x is the value that is being immediately returned (see the type of `return` in Section 3). As the precondition of `return` is always `True`, `return` is always trivially safe. In `sbnd_alloc`, allocating a fresh pointer x initialized with v is always safe. Thus, if e is a continuation of `alloc`, the whole composition is safe if e can be proved safe in the heap in which the existing heap i has been extended with $x \mapsto v$.

Sometimes, even though the top primitive command is revealed, the corresponding lemma does not apply because the heap expression is not in the expected form. For example, `sval_read` applies to the heap $x \mapsto v \bullet i$, but not to $i \bullet x \mapsto v$. In those situations, we have to first rearrange the heap expressions using the commutativity and associativity of \bullet .

$$\begin{aligned} \text{unC} &: h_1 \bullet h_2 = h_2 \bullet h_1 \\ \text{unA} &: h_1 \bullet (h_2 \bullet h_3) = (h_1 \bullet h_2) \bullet h_3 \end{aligned}$$

A similar strategy applies when reasoning about `verify2 ii e1 e2 r` except that now we have two kinds of lemmas: those that apply to both programs simultaneously, and those that apply to only one of them. The first kind is used when reasoning relationally. Typically, our programs start with the same command, say $e:\text{STsec } \Gamma A (p, q)$, which is instantiated with two different explicit substitutions. Since q relates two runs of e in different heaps, we can advance the verification in both *ii.1* and *ii.2* simultaneously, and strip e from the residual goal. In case e is a `return`, or `alloc`, the appropriate lemmas are as follows.

$$\begin{aligned} \text{val_ret} &: r (v_1, v_2) (i_1, i_2) \rightarrow \\ &\quad \text{verify2 } (i_1, i_2) (\text{return } @_0 v_1) (\text{return } @_0 v_2) r \\ \text{bnd_alloc} &: \\ &(\forall x_1 x_2:\text{ptr}. \\ &\quad \text{verify2 } (x_1 \mapsto v_1 \bullet i_1, x_2 \mapsto v_2 \bullet i_2) (e_1 @_0 x_1) (e_2 @_0 x_2) \\ &\quad (\text{fun } yy \text{ mm}. \text{even } x_1 \rightarrow \text{even } x_2 \rightarrow r \text{ yy mm})) \rightarrow \\ &\text{verify2 } (i_1, i_2) (\text{alloc } @_0 v_1; e_1) (\text{alloc } @_0 v_2; e_2) r \end{aligned}$$

In `val_ret`, for example, when both programs are returns of (possibly different) values v_1 and v_2 , a proof of `verify2` proceeds by proving that the postcondition r holds of (v_1, v_2) . In `bnd_alloc`, we change the goal about allocation to a goal about continuation, where the considered heaps are extended appropriately with fresh, even, locations.

The second kind of lemmas is used when the two programs in `verify2` do not start with the same command. In that case, we advance the verification of the programs separately, until some common structure is revealed, or until the verification is over. For example, advancing over the left program, when it starts with a `return` is done by the following lemma.

$$\begin{aligned} \text{bnd_retL} &: \text{verify2 } (i_1, i_2) (e_1 @_0 v) e_2 r \rightarrow \\ &\quad \text{verify2 } (i_1, i_2) (\text{return } @_0 v; e_1) e_2 r \end{aligned}$$

When advancing on the right is required, we can switch sides using the following structural rule.

$$\begin{aligned} \text{vrfC} : \text{verify2 } ii \ e_1 \ e_2 \ r \leftrightarrow \\ \text{verify2 } (ii.2, ii.1) \ e_2 \ e_1 \\ (\text{fun } yy \ mm. r \ (yy.2, yy.1) \ (mm.2, mm.1)) \end{aligned}$$

Just like `verify2` corresponds to a Hoare specification, each of the above lemmas corresponds to an inference rule in a relational Hoare logic. We have proved a number of other lemmas as well, such as, for example, one that corresponds to the standard Hoare rule of conjunction, but we omit them here for lack of space. In proof scripts, we employ a few simple tactics to automate parts of the reasoning in this logic. These tactics come in two flavors. The first one, automates reasoning about `safe` and `verify2` by inspecting the structure of the verified program, and automatically chooses which of the above lemmas to apply in order to advance the verification. The second one automates reasoning about equalities between heap expressions, exploiting commutativity, associativity and cancellativity of `•`. The inference rules and tactics appear in files `stlog.v` and `heaps.v` of the Coq scripts.

Example 9 (A worked-out example using the lemmas.). We consider proving that the implementation of `new` in Figure 2 matches its specification from Figure 1. Here is the specification of `new`:

$$\begin{aligned} \text{new} : \text{nat} \rightarrow \text{string} \rightarrow \\ \text{STsec nil alice } (\text{fun } i. \text{True}, \text{fun } aa \ ii \ mm. \exists ss \ pp \ hh. mm = ii \bullet\bullet \ hh \wedge \text{sshape } aa \ ss \ pp \ hh) \end{aligned}$$

Here is its implementation using combinators:

$$\begin{aligned} \text{new } s \ p = \text{do } (\text{alloc } @_0 \ s; \\ \text{alloc } @ \ (\text{fun } x. \ p); \\ \text{return } @ \ (\text{fun } y \ x. (x, y))) \end{aligned}$$

The code says: allocate new locations, x, y containing s, p respectively. The allocations are done by the high allocator, so x, y are high pointers. Return the pair (x, y) .

When the type system encounters the `do` commands, it issues a proof obligation which consists of two conjuncts. The first conjunct requires showing that the program is safe, and has the form:

$$\begin{aligned} \forall i : \text{heap}. \text{safe}((\text{alloc } @_0 \ s; \\ \text{alloc } @ \ (\text{fun } x. \ p); \\ \text{return } @ \ (\text{fun } y \ x. (x, y))) @_0 \ ()) \ i \end{aligned}$$

The second conjunct requires showing that the program satisfies the relational postcondition, and has the form:

$$\begin{aligned} \forall i_1 \ i_2 : \text{heap}. \\ \text{verify2 } i_1 \ i_2 \\ ((\text{alloc } @_0 \ s; \text{alloc } @ \ (\text{fun } x. \ p); \text{return } @ \ (\text{fun } y \ x. (x, y))) @_0 \ ()) \\ ((\text{alloc } @_0 \ s; \text{alloc } @ \ (\text{fun } x. \ p); \text{return } @ \ (\text{fun } y \ x. (x, y))) @_0 \ ()) \\ (\text{fun } (yy : \text{alice}^2) \ (m : \text{heap}^2). \\ \exists ss : \text{nat}^2. \exists pp : \text{nat}^2. \exists hh : \text{heap}^2. m = (i_1, i_2) \bullet\bullet \ hh \wedge \text{sshape } yy \ ss \ pp \ hh) \end{aligned}$$

For both conjuncts, the obligation requires an instantiation with a unit element $()$, which signifies an empty explicit substitution. This instantiation corresponds to the fact that the `STsec` type of `new` has a local context which is `nil`. If this type had a non-trivial local context, the explicit substitution would also be non-trivial. Be it as it may, even the empty explicit substitution has to be pushed inward. In the case of `safe`, we achieve this by applying the lemma `sbnd_push`, transforming the subgoal corresponding to the safety conjunct into the following.

$$\begin{aligned} & \text{safe } ((\text{alloc } @_0 s) @_0 ()); \\ & \quad (\text{alloc } @ (\text{fun } x. p)); \\ & \quad \text{return } @ (\text{fun } y x. (x, y)) @ (\text{fun } z. (z, ())) i \end{aligned}$$

Now `sbnd_inst` applies to reassociate the instantiations in the first `alloc` command, causing the first command to be transformed into `alloc @0 (s @0 ())`. In our notational conventions, s is isomorphic to `fun x:unit. s`. Hence, the above application $s @_0 ()$ simply transforms into s , and the subgoal becomes

$$\begin{aligned} & \text{safe } (\text{alloc } @_0 s; \\ & \quad (\text{alloc } @ (\text{fun } x. p)); \\ & \quad \text{return } @ (\text{fun } y x. (x, y)) @ (\text{fun } z. (z, ())) i \end{aligned}$$

Next, we apply `sbnd_alloc`, to remove the first command, and reduce the subgoal to

$$\begin{aligned} & \forall l : \text{ptr. safe } (((\text{alloc } @ (\text{fun } x. p); \\ & \quad \text{return } @ (\text{fun } y x. (x, y)) @ (\text{fun } z. (z, ())) @_0 l) (l \mapsto s \bullet i) \end{aligned}$$

Picking a fresh l for the quantified variable, we can apply `sval_inst` to push in the instantiation with l . This changes the term `fun z. (z, ())` into `(fun z. (z, ())) l`, which itself beta reduces to $(l, ())$, transforming our goal into

$$\begin{aligned} & \text{safe}((\text{alloc } @ (\text{fun } x. p); \\ & \quad \text{return } @ (\text{fun } y x. (x, y)) @ (l, ())) (l \mapsto s \bullet i) \end{aligned}$$

Next to apply will be `sbnd_push`, in order to push the substitution with $(l, ())$ into the sequential composition, etc. We can continue applying such lemmas till the proof obligation for safety is discharged. The actual proof in Coq of the safety conjunct looks like

```

move ⇒ i.
apply: sbnd_push; apply: sbnd_inst; apply: sbnd_alloc ⇒ l.
apply: sval_inst; apply: sbnd_push; apply: sbnd_inst; apply: sbnd_alloc ⇒ l'.
apply: sval_inst; apply: sval_inst; apply: sval_ret.

```

At each point in the above proof, there is only one lemma that can apply. Hence the proof can be constructed fully automatically. We have implemented a tactic `sauto` to do just that, simplifying the above proof to merely:

$$\text{move } \Rightarrow i; \text{ do } ![\text{sauto } \Rightarrow ?].$$

To discharge the subgoal corresponding to the conjunct involving `verify2`, we must show that with a pair of input heaps i_1, i_2 , and two copies of the code (inside `do`), the postcondition is satisfied. First we massage the subgoal using applications of several lemmas for pushing explicit substitutions into the programs. These are lemmas are similar to the previously shown `sval_inst`, `sbnd_inst` and

`sbnd_push`, so we omit them here. Suffices to say, we have another tactic, called `vauto`, which applies the lemmas as appropriate. At the end of the execution of `vauto`, we will have variables l_1 and l_2 for the pointers storing the salaries in the two runs, and l'_1 and l'_2 for the pointers storing the passwords in the two runs. The subgoal looks as follows.

$$\begin{aligned} & \exists ss:\text{nat}^2. \exists pp:\text{nat}^2. \exists hh:\text{heap}^2. \\ & (l'_1 \mapsto p \bullet (l_1 \mapsto s \bullet i_1), l'_2 \mapsto p \bullet (l_2 \mapsto s \bullet i_2)) = (i_1, i_2) \bullet \bullet hh \wedge \\ & \text{sshape } ((l_1, l'_1), (l_2, l'_2)) \text{ } ss \text{ } pp \text{ } hh \end{aligned}$$

Now we instantiate the existentially bound variables ss , pp with (s, s) and (p, p) respectively, so that the proof obligation becomes

$$\begin{aligned} & \exists hh:\text{heap}^2. \\ & (l'_1 \mapsto p \bullet (l_1 \mapsto s \bullet i_1), l'_2 \mapsto p \bullet (l_2 \mapsto s \bullet i_2)) = (i_1, i_2) \bullet \bullet hh \wedge \\ & \text{sshape } ((l_1, l'_1), (l_2, l'_2)) (s, s) (p, p) hh \end{aligned}$$

Now we use associative and commutative laws to rewrite the left-hand side of the equality above so that the proof obligation becomes

$$\begin{aligned} & \exists hh:\text{heap}^2. \\ & (i_1 \bullet (l'_1 \mapsto p \bullet l_1 \mapsto s), i_2 \bullet (l'_2 \mapsto p \bullet l_2 \mapsto s)) = (i_1, i_2) \bullet \bullet hh \wedge \\ & \text{sshape } ((l_1, l'_1), (l_2, l'_2)) (s, s) (p, p) hh \end{aligned}$$

Next, we let the system pick logical variables to instantiate the existential with; let us call these variables α and β – they will be instantiated with concrete values further below. Now there are two subgoals that need to be discharged corresponding to the two conjuncts. They are:

$$\begin{aligned} & (i_1 \bullet (l'_1 \mapsto p \bullet l_1 \mapsto s), i_2 \bullet (l'_2 \mapsto p \bullet l_2 \mapsto s)) = (i_1, i_2) \bullet \bullet (\alpha, \beta) \text{ and} \\ & \text{sshape } ((l_1, l'_1), (l_2, l'_2)) (s, s) (p, p) (\alpha, \beta) \end{aligned}$$

The first conjunct is discharged by noting that

$$(i_1, i_2) \bullet \bullet (\alpha, \beta) = i_1 \bullet \alpha, i_2 \bullet \beta$$

so that there are unique solutions for α and β :

$$\alpha = l'_1 \mapsto p \bullet l_1 \mapsto s \text{ and } \beta = l'_2 \mapsto p \bullet l_2 \mapsto s.$$

The second subgoal is easy to discharge by appealing to the definition of `sshape` (see Figure 2). Our actual proof in Coq looks like

```

move => i1 i2; vauto => l1 l2; vauto => l'1 l'2; vauto.
exists (s, s); exists (p, p); rewrite - (unC i1) - (unCA i1) - (unC i2) - (unCA i2).
exists (-, -); split; first by rewrite /plus2 /= .
by rewrite /sshape - toPredE /=; split => //; rewrite unC.

```

```

linked_list : type
shape : linked_list → list T → heap → prop
sshape (pp : linked_list2) (vvs : (list T)2) (ii : heap2) ≐
  shape pp.1 vvs.1 ii.1 ∧ shape pp.2 vvs.2 ii.2
low_links : linked_list2 → heap2 → prop

new : STsec nil linked_list
  (fun i. True,
   fun pp ii mm. ∃jj. mm = jj •• ii ∧
    sshape pp (nil, nil) jj ∧ (ii.1 ≅ ii.2 → low_links pp jj))

insert : STsec [linked_list, T] unit
  (fun p v i. ∃h j vs. i = j • h ∧ shape p vs j,
   fun pp vv yy ii mm. ∀hh jj vvs.
    ii = jj •• hh → sshape pp vvs jj →
    ∃jj'. mm = jj' •• hh ∧
    sshape pp (vv.1::vvs.1, vv.2::vvs.2) jj' ∧
    (low_links pp jj → hh.1 ≅ hh.2 → low_links pp jj'))

remove : STsec [linked_list] (option T)
  (fun p i. ∃h j vs. i = j • h ∧ shape p vs j,
   fun pp yy ii mm. ∀hh jj vvs.
    ii = jj •• hh → sshape pp vvs jj →
    ∃jj'. mm = jj' •• hh ∧
    sshape pp (tail vvs.1, tail vvs.2) jj' ∧
    yy = (if vvs.1 is v1::_ then some v1 else none,
          if vvs.2 is v2::_ then some v2 else none) ∧
    (low_links pp jj → low_links pp jj'))

```

Figure 5: ListSig: signature for linked lists (excerpts).

5 Linked data structures

In this section we develop a small library for linked lists to illustrate RHTT’s support for stateful abstract data types (ADTs), and their interaction with information flow. Working with ADTs essentially requires a number of higher-order features. For example, to support linked lists in a reasonable way, it has to be possible to: (1) describe the layout of the lists in the heap (is the list singly-linked, doubly-linked, etc?). This requires quantification in the assertion logic, definition of predicates by recursion, and inductive definitions of types; (2) abstract the definition of the heap layout from the specification of the ADT, so that the ADT clients can freely interchange implementations with different layouts (hence the need for abstract predicates); (3) parametrize the ADT with respect to the type of list elements (hence the need for type polymorphism in both programs and the assertion logic). All of these features are present in RHTT, and used in the Figures 5 and 6, which show one possible interface, ListSig, and a module, List, implementing ListSig. The interface exports methods that create a new empty list, insert an element to the head of a list, and remove the head element, should one exist.

Both ListSig and List are parametrized in the type of list elements T . The interface declares

the abstract predicate `shape p vs i`, capturing that the heap i stores a valid *singly-linked list* whose content is the mathematical (i.e., purely-functional) sequence vs of type `list T`. The pointer p stores the address of the list head, so that adding new elements at the head can be done by updating p . The linkage between the elements is described by the predicate `lseq x vs` which recurses over the contents vs and states that each node, starting from the head x , contains a single pointer z to the next node in the linked list. The interface hides the details of `shape`, however, and can thus be ascribed to other implementations of `shape`, such as ones describing doubly-linked lists.

The interface in Figure 5 contains one more abstract predicate `low_links pp ii`, which we use in combination with `sshape pp vvs ii`, to describe that the linkage of the list stored in the heap ii is of low security, no matter the security levels of the contents vvs . The latter may be heterogeneous; that is, some elements of vvs may be of low security, while others are high. Similar to `lseq`, `low_links` recurses over the linked lists, declaring that each node is stored at a low address; that is, an address which is equal in the two heap instances, $ii.1$ and $ii.2$. (The formal definition of `low_links` is elided here but appears in file `llist3.v` of the Coq scripts.)

The types of the methods declare how the methods modify the contents of the list as well as the linkage. For example, the `shape` predicate in the preconditions of `insert` and `remove` requires that the initial heaps of these methods store valid linked lists. The `sshape` predicate in the postconditions guarantees that valid linked lists are produced at the end. The postconditions additionally contain conjuncts describing that the methods preserve the low security level of the linkage. For example, `new` will allocate a fresh pointer p , and initialize it with `null`. If the deterministic allocator is used to obtain p , then p will be low *only if the allocator is executed in low-equivalent initial heaps*. Thus, in order to get `low_links pp jj`, we require an antecedent $ii.1 \cong ii.2$. Similarly, `insert` specifies that `low_links pp jj \rightarrow hh.1 \cong hh.2 \rightarrow low_links pp jj'`. In other words, if the initial lists have low linkage, and the remainders of the global heaps are low equivalent, then we can allocate a list node with low linkage. This is so, because the initial heaps must be low equivalent under the described conditions.

The implementations of the methods are standard (Figure 6), but due to the combinator syntax, we describe them in prose. `new` returns a fresh pointer, initialized with `null`. This will be the pointer p in the `shape` predicate. `insert` takes the pointer p to the list, and a value v to insert. It reads the address of the first element (bound to variable hd), and allocates a node x whose contents field is v and next pointer field is hd . Finally, x is written to p . `remove` reads the address of the first element of the list p into the variable hd . If hd is `null`, then the list is empty, and the function terminates. Otherwise, it reads the contents of the node at hd , binding it to the variable v . p is made to point to next v , before v is deallocated.

To establish that the implementation satisfies the signature, we need a number of helper lemmas about `lseq` and `linked_list`, which are kept local to the module. For example, for `lseq`, we need properties that describe the behavior of `lseq x vs i`, in case x is `null` (then the whole list is empty), and non-`null` (then x points to the head).

$$\begin{aligned} \text{lseq_null} &: \forall vs\ i. \text{lseq null } vs\ i \rightarrow vs = \text{nil} \wedge i = \text{empty_heap} \\ \text{lseq_pos} &: \forall vs\ x\ i. x \neq \text{null} \rightarrow \text{lseq } x\ vs\ i \rightarrow \\ &\quad \exists z\ j. i = x \mapsto \text{node (head } vs) z \bullet j \wedge \text{lseq } z\ (\text{tail } vs)\ j \end{aligned}$$

For `low_links`, we show that if two heaps store lists with low linkage *and equal contents*, then the heaps themselves are equal.

$$\begin{aligned} \text{low_linkR} &: \forall vs\ pp\ ii. \text{sshape } pp\ (vs, vs)\ ii \rightarrow \\ &\quad \text{low_links } pp\ ii \rightarrow ii.1 = ii.2 \end{aligned}$$

```

linked_list  $\hat{=}$  ptr
node : type  $\hat{=}$  node of ( $T \times$  ptr)
elem ( $e$  : node)  $\hat{=}$   $e.1$ 
next ( $e$  : node)  $\hat{=}$   $e.2$ 

lseq ( $x$  : ptr) ( $vs$  : list  $T$ ) : heap  $\rightarrow$  prop  $\hat{=}$ 
  if  $vs$  is  $v::vt$  then
    fun  $i. \exists z:ptr j:heap.$ 
       $i = x \mapsto$  node  $v z \bullet j \wedge$  lseq  $z vt j$ 
    else fun  $i. x = \text{null} \wedge i = \text{empty\_heap}$ 
shape ( $p$  : linked_list) ( $vs$  : list  $T$ ) ( $i$  : heap)  $\hat{=}$ 
   $\exists x:ptr. j:heap. i = p \mapsto x \bullet j \wedge$  lseq  $x vs j$ 

new  $\hat{=}$  do (lalloc @0 null)
insert  $\hat{=}$ 
  do (read @ (fun  $p v. p$ );
      lalloc @ (fun  $hd p v. \text{node } v hd$ );
      write @ (fun  $x hd p v. (p, x)$ ))
remove  $\hat{=}$ 
  do (read @ (fun  $p. p$ );
      if (fun  $hd p. hd = \text{null}$ ) then
        return @ (fun  $hd p. \text{none}$ )
      else
        read @ (fun  $hd p. hd$ );
        write @ (fun  $v hd p. (p, \text{next } v)$ );
        dealloc @ (fun  $_ v hd p. hd$ );
        return @ (fun  $_ _ v hd p. \text{some } (elem v)$ ))

```

Figure 6: Module List: implementation of singly-linked lists (excerpts).

```

P5 : STsec [bool] linked_list
  (fun b i. True,
   fun bb yy ii mm. ∃jj. mm = jj •• ii ∧
    sshape yy ([if bb.1 then 1 else 2, 0],
               [if bb.2 then 1 else 2, 0]) jj ∧
    (ii.1 ≅ ii.2 → low_links yy jj)) ≐
do (new @ (fun b. ());
   insert @ (fun p b. (p, 0));
   if (fun _ p b. b) then insert @ (fun _ p b. (p, 1))
   else insert @ (fun _ p b. (p, 2)) fi;
   return @ (fun _ _ p b. p))

P6 : STsec [bool] linked_list
  (fun b i. True,
   fun bb yy ii mm. ∃jj. mm = ii •• jj ∧
    sshape yy (if bb.1 then [0] else [1, 0],
               if bb.2 then [0] else [1, 0]) jj) ≐
do (new @ (fun b. ());
   insert @ (fun p b. (p, 0));
   insert @ (fun _ p b. (p, 1));
   if (fun _ _ p b. b) then
     remove @ (fun _ _ p b. p);
   return @ (fun _ _ _ p b. p)
   else return @ (fun _ _ p b. p))

```

Figure 7: Programs with heterogeneous lists.

Example 10. The program P_5 in Figure 7 illustrates heterogeneous lists, i.e., lists that contain both high and low values. It takes a high boolean argument b , creates a new linked list, and inserts 0 (a constant, hence low) at the head. Then, depending on b , it inserts either 1 or 2, resulting in a heterogeneous list with a high first element and low second element. This is described in the postcondition by conditionals over the values of b in the two different runs ($bb.1$ and $bb.2$). Irrespective of the contents, the ending *linkage* is low, assuming we started with low-equivalent input heaps.

Example 11. The program P_6 in Figure 7 is similar to P_5 , but branches on b to decide whether to remove the head element. Therefore, the length of the resulting list may differ in the two runs, depending on b . We can specify it with the type shown in the Figure. Notice however that we cannot prove that $\text{low_links } yy \text{ } jj$ holds at the end of P_6 . The length of the produced list is dependent on b , which implies that the resulting linkage may differ in two runs of P_6 , and hence cannot be low itself.

Our Coq scripts implement other interfaces for linked list, where the *sshape* predicates are parametrized by the linkage as well. This exposes more implementation details (e.g., that the list is singly-linked), but allows more precise reasoning about linkage. For example, we may prove that executing one more conditional over b , with a call to `remove` in the `else` branch, will restore the low

linkage.

We are not aware of any other system in literature that can reason statically about heterogeneous structures. In the dynamic setting, a recent example is the work of Russo et al. [42], which tracks information-flow through DOM trees, with the goal of preventing information leakage via node deletion or navigation. The system works by assigning to each node two security labels: one for the contents, and another for the existence of the node. These annotations are very specific to DOM trees, however, and it seems that the label assignment would have to be designed differently for different data structures and enforced properties. Thus, if one wants to work with a number of structures simultaneously, one must employ a very rich specification logic, just as we do.

We close with an example which combines linked lists with the Alice module from Section 2.

Example 12. In Example 5, Alice ratifies Bob’s tax function, which may keep local state, as long as Bob can prove that his final state does not steal Alice’s salary. Here we instantiate Bob’s local state to a linked list, which dynamically grows as various instances of Alice execute Bob’s program, but the values stored in the linked list are always independent of any instance’s salary and the list’s linkage is always low. Observe from the specifications of `new` and `insert` that Bob’s newly allocated nodes will be low only if he can generate them in low-equivalent heaps. To express this low equivalence the specification of `epost` used in `ratify`’s specification must change as emphasized below.

$$\begin{aligned}
\text{epost } aa \ bb \ jj \ yy \ ii \ mm &\hat{=} \\
\forall ss \ pp \ kk \ tt \ hh. & \\
ii = jj \bullet \bullet \ tt \bullet \bullet \ hh \rightarrow \text{sshape } aa \ ss \ pp \ jj \rightarrow & \\
\text{bbshape } bb \ kk \ tt \rightarrow & \\
\exists jj' \ tt'. mm = jj' \bullet \bullet \ tt' \bullet \bullet \ hh \wedge \text{sshape } aa \ ss \ pp \ jj' \wedge & \\
\text{bbshape } bb \ (\text{bcmp } kk.1, \text{bcmp } kk.2) \ tt' \wedge & \\
\boxed{jj.1 \bullet hh.1 \cong jj.2 \bullet hh.2} \rightarrow & \\
tt.1 = tt.2 \rightarrow tt'.1 = tt'.2) &
\end{aligned}$$

Bob can now be granted access to Alice’s salary and can keep the count in a linked list. For example, the implementation below defines Bob’s local state as a linked list which counts the number of times Bob’s program has been called by linking in new nodes into Bob’s list. The nodes are filled with 1 for simplicity, but arbitrary values would do, including dynamically computed ones, as long as they are independent of Alice’s salary.

$$\begin{aligned}
\text{bob} &\hat{=} \text{linked_list} \\
G &\hat{=} \text{list nat} \\
\text{bbshape } (bb : \text{bob}^2) \ (kk : G^2) \ (ii : \text{heap}^2) &\hat{=} \\
\text{List.sshape } bb \ kk \ ii \ \wedge \ bb.1 = bb.2 & \\
\text{bcmp} : G \rightarrow G &\hat{=} \text{fun } k.1 :: k
\end{aligned}$$

Bob’s program, which reads Alice’s salary, allocates a new node in his list, and then returns the computed tax for the salary, can then be created and ratified as follows.

$$\begin{aligned}
\text{linked_client} &\hat{=} \\
\text{ratify } (\text{do } (\text{read_salary } @ \ (\text{fun } a \ b. a); & \\
\text{insert } @ \ (\text{fun } x \ a \ b. (b, 1)); & \\
\text{return } @ \ (\text{fun } _ \ x \ a \ b. x * 24\%))) &
\end{aligned}$$

6 Discussion

Small vs. large footprint specifications In one of the previous iterations of the system, we did have small footprint specifications, before we realized that deterministic allocation is important. In that iteration of the system, the method `new` from Figure 1, had the type

$$\text{nat} \rightarrow \text{string} \rightarrow \text{STsec nil alice } (\text{fun } i. \text{emp } i, \text{fun } aa \ ii \ mm. \exists ss \ pp. \ \text{sshape } aa \ ss \ pp \ mm)$$

When we switched to large footprint specification, this became:

$$\begin{aligned} \text{nat} \rightarrow \text{string} \rightarrow \text{STsec nil alice} \\ (\text{fun } i. \text{True}, \text{fun } aa \ ii \ mm. \exists ss \ pp \ hh. \ mm = ii \bullet\bullet \ hh \wedge \text{sshape } aa \ ss \ pp \ hh) \end{aligned}$$

In other words, one more existentially quantified variable (hh) and one more assertion expressing disjointness between the initial and allocated heap ($mm = ii \bullet\bullet hh$).

In other examples, where the preconditions were not so simple as here, we needed one more variable and one more assertion for the precondition too. As parts of the precondition sometimes have to appear in the postcondition, for the purposes of correct scoping, this may add one more variable and one more assertion in the postcondition as well. Altogether, at most three new variables, and three more disjointness assertions between those variables. This pattern applied throughout the development and caused minimal refactoring on our already written proofs.

Completeness We have informally justified the completeness of our system through several examples, covering a wide range of security relevant policies including access control, information flow, declassification, erasure, and their combinations. Unfortunately we are not aware of a clear and exhaustive formal definition of what constitutes, say, an erasure, or access-control policy, or a combination thereof. Therefore, we do not know how to state a formal completeness result.

If we focus on Cook completeness for RHTT, then, as we have argued in Section 3, our specifications for all of the primitive effectful combinators compute weakest preconditions and strongest postconditions using the specifications of the components. The exception are the conditionals, for which this cannot be done when the boolean guard is high. However, RHTT is still capable of checking high conditionals against programmer-supplied postconditions. The lack of Cook completeness therefore results in an increase in code annotations that the programmer has to supply, but does not decrease the reasoning power of the logic. In particular, the programmer-supplied annotations are also in higher-order logic, and therefore can be as expressive as needed. For example, we implemented a denotational model for our language in the assertion logic itself (presented in the TR), and the programmer’s annotations can talk about arbitrary properties of this model.

Noninterference for finite security lattices The standard notion of noninterference when locations are classified into elements of a security lattice is compatible with RHTT and can be expressed in postconditions of RHTT programs. If each variable of a program is classified at some level of a finite lattice L , then the program is noninterfering if for each $\ell \in L$, the following holds: $(xx_1.1 = xx_1.2 \wedge \dots \wedge xx_n.1 = xx_n.2) \rightarrow (yy_1.1 = yy_1.2 \wedge \dots \wedge yy_m.1 = yy_m.2)$ where xx_1, \dots, xx_n are pairs of values of variables at or below security level ℓ in the two initial heaps and yy_1, \dots, yy_m are pairs of values of variables at security level ℓ in the two final heaps. If the lattice L is finite, then the noninterference property can be represented in the postcondition of a program as the *conjunction* of such requirements for each level $\ell \in L$.

This representation can be combinatorially explosive, if one uses one conjunct for each variable. But, in our higher-order assertion logic we can introduce predicates that abstract over a number of such conjuncts at once, and hence avoid the explosion. For example, if the variables in question all stand for the contents of some linked list, we can define a predicate that conjoins equations of the above form for each value reachable from the head of the list. Note that reachability can be expressed in higher-order but not in first-order logic. Thus, we expect that once we move into investigating security lattices, there will not be one unique definition which captures the semantics of the multiple levels. Rather, we expect to introduce abstractions that exploit the specifics of particular programs and data layouts in order to avoid combinatorial explosion. But, at this point, we do need more experience with the system in order to validate the above conjecture.

On proof sizes We have found that the size of interactive proofs is not too overwhelming in general. However, the amount of interaction varies with programs. Programs with complex loop invariants usually require large proofs, whereas simpler programs can be verified in just a number of lines proportional to the size of the program.

Programs that branch on high boolean guards invariably have larger proofs than programs that branch only on low: the latter always choose the same branches of conditionals in two runs, so the verification of the two runs proceeds in lockstep. High-branching programs can choose branches asymmetrically, thus doubling the number of proof obligations. In addition, when branches are chosen asymmetrically, the proofs usually require some mathematical insight from the programmer (for example, algebraic simplification of expressions) in order to argue that the high secret has not been leaked. The latter, however, seems unavoidable, and inherent to the nature of programs branching on high guards.

To substantiate, consider the programs from Examples 2 and 3, our first examples that do not branch on high. We have the following statistics given as the pair (*code+spec size*, *proof size*). For `new`, we have (7, 5); for `read_salary` (7, 4); for `write_salary` (18, 15) and for `declassify` (11, 5). The above proofs share common definitions and lemmas which are altogether 10 lines long.

The program P_2 in Example 1, which contains nested conditionals and branching on high, is implemented using 36 lines of code, most of which are inlined user-supplied annotations. The corresponding proof is 44 lines long.

We have also implemented examples that iterate over linked data structures (not presented in the paper, but available in the accompanying Coq scripts). In a program for in-place list reversal, in which the linkage of the list is high, the code and annotations together take 43 lines. The proof is 94 lines long, because there is a high conditional branching on a null-pointer check.

7 Related work

Banerjee et al. [9] specify expressive declassification policies using Hoare style specifications (termed flowspecs); preconditions thereof are conjunctions of ordinary state conditions based on first-order logic (for specifying conditions *when* declassification can happen) as well as relational predicates (that specify *what* is being declassified) [44]. We extend the ideas in [9] and consider a higher-order imperative language and also a policy specification language based on higher-order logic, where Hoare-style specifications may appear in negative (i.e., argument) positions, which is required for conditional access and erasure policies.

A recent line of work [29, 41] uses type-theoretic technology, namely Haskell, to specify and enforce information-flow properties in a non-dependently-typed setting. While Haskell already provides the important higher-order constructs for abstraction and modularity, non-dependent types by definition cannot specify behaviors that are dependent on some condition such as authorization, conformance to a policy, or local state. Thus, we do not think they can be used directly to enforce involved security policies such as the ones considered in this paper.

Some other recent languages, with somewhat similar high-level goals to ours, and which use some form of dependent types are Fine [46], Fable [47], FX [13], Aglet [33], F7 [14] and Aura [26]. They all support some, but not all features that we provide in RHTT.

In Swamy et al.’s purely functional programming language Fine [46], access and information flow policies can mention attributes like high and low, that statically label data. The type system enforces these policies by tracking flows of attributes. Unlike RHTT, Fine’s type system does not track changes to the state (heap), so the effect of state in policies must be simulated through ghost variables, whose (static) updates are governed by specifications of primitive functions. A token passing mechanism based on affine kinds ensures that at most one static state is valid at each program point, but it makes programming in Fine inconvenient. Fine includes a simple module system which allows a programmer to hide type definitions, but does not allow abstraction over predicates as RHTT does. In an earlier language, Fable [47], data can be statically labelled with attributes that can be used to enforce both access control and information flow policies. However, Fable’s type system lacks the affine kinds of Fine as well as Fine’s logic-based sublanguage for policies and, therefore, cannot be used to reason about state-dependent policies.

The language FX [13] succeeds Fine with the purpose of verifying stateful programs that permit object allocation, mutation and deallocation. The type system of FX admits computation (Hoare) types and caters to the verification of safety properties of FX programs by translating into Fine programs and typechecking the latter. The translation is a simulation under strong bisimilarity, rather than the stronger property that well-typed FX programs are translated into well-typed Fine programs. The verification of security policies, particularly, of non-safety properties such as noninterference, is not the overarching goal of FX’s type system, although a lattice of labels can be encoded and used to prove, e.g., an integrity property that untrusted data does not get consumed at trusted sinks. A proof of noninterference is not supplied; as in most label-based security type systems, such a proof cannot be carried out in FX’s (or Fine’s) type system directly (in contrast to our work) but rather must be established as a metatheorem of the type system by reasoning about two runs of programs. As regards reasoning about stateful higher-order programs, the formalization is left for future work and we expect that it will elucidate how the type system reasons about (security properties of) unbounded dynamic data structures e.g., linked lists, trees with back pointers etc., that contain significant use of aliased mutable objects. In particular, because FX proposes to reason about aliasing using a library of permissions the above formalization might be delicate.

Morgenstern and Licata have recently proposed a type system called Aglet [33], for enforcement of state-dependent access control policies. Aglet is an extension of Agda [37] with a computation monad similar to our `STsec` types. However, Aglet’s computation monad lacks semantics and, consequently, the soundness of its inference rules has to be taken on faith (in contrast, the RHTT model is formalized in Coq). Moreover, the pre- and post-conditions of Aglet’s computation monad can only mention a restricted form of state, namely, a mutable list of authorization-relevant credentials, which can be used to discharge authorization obligations at various program points. Due to this

restriction, Aglet cannot be used to reason about data structures written in Agda. Also, Aglet’s postconditions do not consider simultaneous runs of programs. As a result of these limitations, Aglet cannot be used to represent many of our examples. On the other hand, we believe that examples from the paper on Aglet can be expressed in RHTT easily.

Borgström et al. [14] reason about access control behavior of programs in an extension of F7 that has a state monad with pre- and post-conditions. Although the state monads in their work and ours are technically similar, that work differs from ours in two significant ways. First, the goals are different: whereas we consider enforcement of information flow properties and declassification in addition to access control properties, Borgström et al. consider access control and show how the state monad can be used to enforce different flavors of it, *viz.* role-based, stack-based, and history-based. Second, in common with other work based in F7, a priori evidence for discharging verification conditions in Borgström et al’s work is programmer specified assumptions that are not necessarily semantically grounded, and verification is correct only to the extent that these assumptions are correct. In contrast to their axiomatic approach, we verify the soundness of our type theory on a semantic model. Nonetheless, due to the common state-monad based approach, and RHTT’s more general type system, we believe that Borgström et al’s work can be encoded in RHTT without much change. As a first step in this direction, our Coq scripts contain an example that shows how RHTT supports reasoning about principals and roles.

The languages Aura [26], PCML₅ [7], and PCAL [16], based on the proof-as-authorization paradigm [4], enforce logic-represented access policies by statically ensuring that each call to a protected interface is accompanied by proper authorization. Although work in the context of Aura shows that noninterference can be encoded [27], Aura currently does not handle state in the form that we consider in this paper. However, it is conceivable that mutable state can be added to Aura along the lines of the STsec monad.

The Paralocks language [15] also allows logic-based access control policies that are enforced statically in the type system. Information flow policies can be encoded as a specific mode of access control as, for instance, is demonstrated through an encoding of Myers’ and Liskov’s Decentralized Label Model. Like Fine, Paralocks includes two kinds of state, of which, one, called *locks*, is tracked through the type system, while the other is not. Locks are boolean variables that can be used to encode a wide range of policies. The semantics of Paralocks is trace-based and, like gradual release [5], uses a knowledge-based definition of information leaks. A meta-theorem guarantees that access policies of a well-typed program are respected at all program points during the program’s execution.

Finally, RHTT extends the work on Hoare Type Theory and Ynot [35] with the ability to reason relationally about security. The addition caused significant changes in the semantics and the usage of the language. For example, because of the reasoning about two-runs, the modeling of STsec deals with CPO’s, whereas for HTT complete lattices sufficed. The logic for discharging verification conditions in RHTT has to reason separately about safety and correctness, whereas in HTT, safety and correctness could be captured in one judgment, corresponding to the separation logic triple. In RHTT we use large footprint specification and combinator syntax, whereas HTT used small footprints for programming monadically. On the other hand, we were able to reuse from HTT the library for reasoning about heap disjointness, to keep RHTT proofs relatively short.

8 Conclusion

We have presented RHTT, a system implemented in Coq that is targeted for full interactive verification of state-based access control and information flow policies via dependent types. Examples of such security policies include declassification, information erasure and state-based access control and information flow. We have presented typing rules for the stateful fragment of RHTT and implemented a semantic model that provides a denotation to every well-typed RHTT program. We have also developed a logic for discharging verification conditions that arise in the verification process.

Beyond what has been achieved in this paper much remains to be done. While we can specify expressive program properties, and verify that programs comply with them, it is currently impossible to reason about specifications themselves. For example, we cannot reason that indirect inference (say using aggregation operators such as average) do not lead to unwanted leaks. For that, one might need to reason about quantitative information flow and apply knowledge-based reasoning [23].

Currently, RHTT does not support reasoning about trace-based, temporal properties. For example, while it is intuitively clear that our specification of functions `grant`, `revoke`, `read_salary` (Example 4) indeed encodes a temporal discipline on the usage of `read_salary` (e.g., “no reads occur unless a grant has occurred and no revoke has occurred after the grant”) this cannot be formally proved in our logic itself. We note that very little is known on how enforcement of trace-based properties, in security or other areas such as concurrency, interacts with type theoretic constructions such as higher-order functions, abstract types or modules. We intend to investigate this in the future, in the context of reactive, non-deterministic and concurrent higher-order languages.

A related property is parametricity of the type system of Coq – that opaque sealing does not divulge the actual implementation of the sealed data. We have assumed this property without a proof here, and we are not aware of such a proof in the literature. (We believe though that this is a reasonable property to assume; it would be surprising if Coq were not parametric.) The closest related proofs of which we are aware in the imperative world, are the recent ones for ML with references [2] and for separation logic [50]. In the world of dependent types, recently Bernardy et al. [12] have established parametricity for a class of pure type systems (including the calculus of constructions). We intend to investigate in the future if and how this proof can be extended to mutable state and our `STsec` types.

We also intend to investigate the use of relations other than equality, such as distance metrics and continuity (small changes to the distance of inputs cause small changes to the distance of outputs), in the postconditions of `STsec` types. Such definitions may be particularly useful in the context of differential privacy guarantees [22, 39].

Acknowledgements Thanks to Gilles Barthe, Alexey Gotsman, Boris Köpf, Jamie Morgenstern, Greg Morrisett and David Naumann for their comments on earlier drafts. We thank Trent Jaeger for his advice on improving the presentation of the paper.

This research was partially supported by Madrid Regional Government Project S2009TIC-1465 Prometidos; MICINN Projects TIN2009-14599-C03-02 Desafios and TIN2010-20639 Paran10; EU Projects GA-231620 Hats and NoE-256980 Nessos; Ramon y Cajal grant RYC-2010-07433; AMAROUT grant PCOFUND-GA-2008-229599; U.S. NSF Trustworthy Computing grant 1018061 “Compositional End-to-End Security for Systems”, and the U.S. AFOSR MURI “Collaborative

Policies and Assured Information Sharing”.

References

- [1] M. Abadi, A. Banerjee, N. Heintze, and J. G. Riecke, “A core calculus of dependency,” in *POPL*, 1999.
- [2] A. Ahmed, D. Dreyer, and A. Rossberg, “State-dependent representation independence,” in *POPL*, 2009.
- [3] T. Amtoft, S. Bandhakavi, and A. Banerjee, “A logic for information flow in object-oriented programs,” in *POPL*, 2006.
- [4] A. W. Appel and E. W. Felten, “Proof-carrying authentication,” in *ACM CCS*, 1999.
- [5] A. Askarov and A. Sabelfeld, “Gradual release: Unifying declassification, encryption and key release policies,” in *IEEE Symp. Security and Privacy*, 2007.
- [6] A. Askarov and A. Myers, “A semantic framework for declassification and endorsement,” in *ESOP*, 2010.
- [7] K. Avijit, A. Datta, and R. Harper, “Distributed programming with distributed authorization,” in *TLDI*, 2010.
- [8] A. Banerjee and D. A. Naumann, “Stack-based access control and secure information flow,” *JFP*, vol. 15, no. 2, pp. 131–177, 2005.
- [9] A. Banerjee, D. A. Naumann, and S. Rosenberg, “Expressive declassification policies and their modular static enforcement,” in *IEEE Symp. Security and Privacy*, 2008.
- [10] D. Bell and L. LaPadula, “Secure computer systems: Mathematical foundations,” MITRE Corp., Tech. Rep. MTR-2547, 1973.
- [11] N. Benton, “Simple relational correctness proofs for static analyses and program transformations,” in *POPL*, 2004.
- [12] J.-P. Bernardy, P. Jansson, and R. Paterson, “Parametricity and dependent types,” in *ICFP*, 2010.
- [13] J. Borgstrom, J. Chen, and N. Swamy, “Verifying stateful programs with substructural state and Hoare types,” in *PLPV*, 2011.
- [14] J. Borgström, A. D. Gordon, and R. Pucella, “Roles, stacks, histories: A triple for Hoare,” *JFP*, forthcoming.
- [15] N. Broberg and D. Sands, “Paralocks: role-based information flow control and beyond,” in *POPL*, 2010.
- [16] A. Chaudhuri and D. Garg, “PCAL: Language support for proof-carrying authorization systems,” in *ESORICS*, 2009.

- [17] S. Chong and A. C. Myers, “Security policies for downgrading,” in *ACM CCS*, 2004.
- [18] —, “Language-based information erasure,” in *CSFW*, 2005.
- [19] —, “End-to-end enforcement of erasure and declassification,” in *CSF*, 2008.
- [20] D. Denning, “A lattice model of secure information flow,” *CACM*, vol. 19, no. 5, pp. 236–242, 1976.
- [21] E. W. Dijkstra, “Guarded commands, nondeterminacy and formal derivation of program,” *Communications of the ACM*, vol. 18, no. 8, pp. 453–457, August 1975.
- [22] C. Dwork, F. McSherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *TCC*, 2006.
- [23] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi, *Reasoning About Knowledge*. MIT Press, 1995.
- [24] J. Goguen and J. Meseguer, “Security policies and security models,” in *IEEE Symp. Security and Privacy*, 1982, pp. 11–20.
- [25] R. Harper and M. Lillibridge, “A type-theoretic approach to higher-order modules with sharing,” in *POPL*, 1994.
- [26] L. Jia, J. A. Vaughan, K. Mazurak, J. Zhao, L. Zarko, J. Schorr, and S. Zdancewic, “AURA: A programming language for authorization and audit,” in *ICFP*, 2008.
- [27] L. Jia and S. Zdancewic, “Encoding information flow in Aura,” in *PLAS*, 2009.
- [28] X. Leroy, “Manifest types, modules, and separate compilation,” in *POPL*, 1994.
- [29] P. Li and S. Zdancewic, “Arrows for secure information flow,” *Theoretical Comput. Sci.*, vol. 411, no. 19, pp. 1974–1994, 2010.
- [30] P. Martin-Löf, *Intuitionistic Type Theory*. Bibliopolis, 1984.
- [31] The Coq development team, *The Coq proof assistant reference manual*, LogiCal project, INRIA, 2009, version 8.2. [Online]. Available: <http://coq.inria.fr>
- [32] J. C. Mitchell and G. D. Plotkin, “Abstract types have existential type,” *ACM Trans. Prog. Lang. Syst.*, vol. 10, no. 3, pp. 470–502, 1988.
- [33] J. Morgenstern and D. Licata, “Security-typed programming within dependently-typed programming,” in *ICFP*, 2010.
- [34] A. C. Myers, “JFlow: Practical mostly-static information flow control,” in *POPL*, 1999.
- [35] A. Nanevski, J. G. Morrisett, and L. Birkedal, “Hoare type theory, polymorphism and separation,” *JFP*, vol. 18, no. 5-6, pp. 865–911, 2008.
- [36] A. Nanevski, V. Vafeiadis, and J. Berdine, “Structuring the verification of heap-manipulating programs,” in *POPL*, 2010.

- [37] U. Norell, “Towards a practical programming language based on dependent type theory,” Ph.D. dissertation, Chalmers University of Technology, 2007.
- [38] S. L. Peyton Jones and P. Wadler, “Imperative functional programming,” in *POPL*, 1993.
- [39] J. Reed and B. C. Pierce, “Distance makes the types grow stronger,” in *ICFP*, 2010.
- [40] J. C. Reynolds, “Separation logic: a logic for shared mutable data structures,” in *LICS*, 2002.
- [41] A. Russo, K. Claessen, and J. Hughes, “A library for light-weight information-flow security in Haskell,” in *Haskell Symposium*, 2008.
- [42] A. Russo, A. Sabelfeld, and A. Chudnov, “Tracking information flow in dynamic tree structures,” in *ESORICS*, 2009.
- [43] A. Sabelfeld and D. Sands, “A PER model of secure information flow in sequential programs,” in *ESOP*, 1999.
- [44] —, “Declassification: Dimensions and principles,” *Journal of Computer Security*, vol. 17, no. 5, pp. 517–548, 2009.
- [45] V. Simonet, “Fine-grained information flow analysis for a λ -calculus with sum types,” in *CSFW*, 2002.
- [46] N. Swamy, J. Chen, and R. Chugh, “Enforcing stateful authorization and information flow policies in Fine,” in *ESOP*, 2010.
- [47] N. Swamy, B. J. Corcoran, and M. Hicks, “Fable: A language for enforcing user-defined security policies,” in *IEEE Symp. Security and Privacy*, 2008, full version: Technical report CS-TR-4895, Univ. Maryland.
- [48] N. Swamy, M. Hicks, S. Tse, and S. Zdancewic, “Managing policy updates in security-typed languages,” in *CSFW*, 2006.
- [49] T. Terauchi and A. Aiken, “Secure information flow as a safety problem,” in *SAS*, 2005.
- [50] J. Thamsborg, L. Birkedal, and H. Yang, “Two for the price of one: Lifting separation logic assertions,” 2010, manuscript. [Online]. Available: <http://www.itu.dk/~birkedal/papers/relational-lifting.pdf>
- [51] D. M. Volpano, C. E. Irvine, and G. Smith, “A sound type system for secure flow analysis,” *Journal of Computer Security*, vol. 4, no. 2/3, pp. 167–188, 1996.
- [52] H. Yang, “Relational separation logic,” *Theoretical Comput. Sci.*, vol. 375, pp. 308–334, 2007.
- [53] H. Yang and P. W. O’Hearn, “A semantic basis for local reasoning,” in *FoSSaCS*, 2002.