

Verification Condition Generation for Conditional Information Flow

Torben Amtoft

Kansas State University, Manhattan, KS, USA
tamtoft@cis.ksu.edu

Anindya Banerjee*

IBM T. J. Watson Research Center,
Hawthorne, NY, USA
ab@cis.ksu.edu

Abstract

We formulate an intraprocedural information flow analysis algorithm for sequential, heap manipulating programs. We prove correctness of the algorithm, and argue that it can be used to verify some naturally occurring examples in which information flow is conditional on some Hoare-like state predicates being satisfied. Because the correctness of information flow analysis is typically formulated in terms of noninterference of pairs of computations, the algorithm takes as input a program together with two-state assertions as postcondition, and generates two-state preconditions together with verification conditions. To process heap manipulations and while loops, the algorithm must additionally be supplied “object flow invariants” as well as “loop flow invariants” which are themselves two-state, and possibly conditional.

1. Introduction

Information flow analyses are used to ensure that programs satisfy confidentiality policies. Such policies are expressed by labeling variables with security levels, e.g., H for secrets/classified and L for public/observable/unclassified. For a given policy, a program P satisfies *noninterference* (NI) [18] provided that for any *two* runs of P , if P is executed from two

input states that are L -indistinguishable (i.e., the input states agree on the values of L -variables) then it yields output states that are also L -indistinguishable. A sound information flow analysis guarantees that the programs it accepts are noninterferent.

This paper formulates a sound intraprocedural information flow analysis *algorithm* — rather than a type-based or logic-based *specification* — for heap manipulating programs. We assume that such programs are more or less decorated with assertion statements and loop/object invariants; those can be automatically checked by tools such as BLAST [20], ESC/Java [14] or Spec# [8]. A novel aspect of the algorithm is that it reasons about possibly *conditional* information flow, and also handles while loops and common data structures when armed with *flow invariants* (introduced in the sequel). We leave the automatic *inference* of flow invariants for future work.

Given a variable x labeled L , the formulation of noninterference entails that we restrict our attention to pair of states σ_1, σ_2 where $\sigma_1(x) = \sigma_2(x)$. This observation inspired Amtoft et al. [2, 1] to a logical rendition of NI which uses *agreement* assertions of the form $x \bowtie$, where two states σ_1, σ_2 satisfy $x \bowtie$ when $\sigma_1(x) = \sigma_2(x)$. If a program P has observable input variables x_1, \dots, x_n , and observable output variables y_1, \dots, y_m , then NI can be recast as

$$\{x_1 \bowtie \wedge \dots \wedge x_n \bowtie\} P \{y_1 \bowtie \wedge \dots \wedge y_m \bowtie\}$$

The meaning (partial correctness) of the above triple is that for any two states σ_1, σ_2 that agree on the values of x_1, \dots, x_n (as asserted by the precondition), if one run of P transforms σ_1 to σ'_1 and another run of P transforms σ_2 to σ'_2 , then the values of y_1, \dots, y_m

* On sabbatical leave from Kansas State University, Manhattan, KS, USA.

agree in the final states, σ'_1, σ'_2 (as asserted by the postcondition).¹

Amtoft et al.[1] specify, in logical form, a modular information flow analysis for sequential, heap-manipulating programs. If a triple is derivable for a program then NI holds for the program. The specification is flow sensitive (unlike most type-based approaches), can check information leaks caused by aliasing, and can be used for analyzing observational purity. Moreover, the specification can be used to check compliance with delimited release policies [23] in a technically straightforward manner: extend agreements over variables to agreements over “escape-hatch” expressions that syntactically specify such policies. More recently, the specification has been proposed as a crucial component for the verification of state-dependent declassification policies [6].

The logical specification of [1] comes with an analysis algorithm which, however, has some shortcomings: it needs to know the shape of the heap, and it does not integrate well with programmer assertions. Also, the specification itself does not capture *conditional* information flows. These shortcomings make it difficult to analyze information flow in non-trivial programs, especially ones that involve reasoning about common data structures. (A similar situation prevails with extant security type systems [25, 5, 21]).

Contributions. This paper shows how to reason about information flow that may be conditional, and how to compute it for programs that may manipulate common data structures. The algorithm (Sect. 4) takes as input a program and a (possibly conditional) agreement assertion as postcondition, and as output generates preconditions and verification conditions (VCs). Currently, the algorithm expects the user to provide loop invariants and object invariants that are themselves (conditional) agreement assertions; we call such invariants *flow invariants*. The algorithm always terminates, but the VCs may be unsatisfiable; this will happen if the flow invariants are not strong enough. We prove the correctness of the algorithm, and use it to verify some naturally oc-

¹Two remarks: (a) The connection with NI based on security labels [25] is that for any well-labeled program, P , if l_1, \dots, l_n are all the L -variables in P then $l_1 \times \wedge \dots \wedge l_n \times$ is an invariant. (b) To model security lattices with more than two elements, say $L \leq M \leq H$, multiple specifications are needed, like “if input states agree on L then output states agree on L ” and “if input states agree on L, M then output states agree on L, M ”.

curing examples. A prototype implementation² is currently being developed by Jonathan Hoag.

An example loop flow invariant is $x \times$, with the following informal semantics: if two states, σ_1 and σ_2 , agree on the value of x , and one iteration of the loop transforms σ_1 into σ'_1 and σ_2 into σ'_2 , then also σ'_1 and σ'_2 agree on the value of x . If the invariant is conditional, like $i > n \Rightarrow x \times$, then σ'_1 and σ'_2 are required to agree on x only if they both assert $i > n$, whereas σ_1 and σ_2 can be assumed to agree on x only if they both assert $i > n$. (We defer examples of object flow invariants to Sect. 2.) A second contribution of the paper is the underlying semantic framework (Sect. 3) for such conditional assertions that mixes ordinary, Hoare-logic style predicates with two-state agreement assertions.

A third contribution is the smooth integration with standard assertions, the presence of which can help the algorithm to increase precision. A simple example of this is the program

```
if  $w$  then  $x := 7$  else  $x := 7$ ; assert( $x = 7$ )
```

Given the postcondition $x \times$, the algorithm will compute $x = 7 \Rightarrow x \times$ as the precondition of the assertion statement; this is justified in all contexts because we employ a correctness criterion which considers only executions that terminate successfully, and the assertion will abort if $x \neq 7$ (which of course cannot happen in the given context). Since $x = 7 \Rightarrow x \times$ always holds, it can be simplified to *true*, which, when given as postcondition to the conditional is also returned as the precondition. Without the ability to use and/or derive/infer the assertion statement, however, the precondition would need to include $w \times$. The inference of such “standard” assertions can be done by, e.g., BLAST, but will not be our concern in this paper.

2. Examples

We now illustrate, by way of examples in Figs. 1 and 2, the issues involved in verifying information flow policies for while loops, as well as for programs that manipulate the heap using field update, field access and object allocation.

Loop flow invariants. Consider the program P in Fig. 1(a), and the policy specification $\{x \times\} - \{result \times\}$.

²Available at <http://people.cis.ksu.edu/~jch5588/securityflow/SecurityFlow.html>. It requires Java 1.5.11. As of writing, it handles assignments, conditionals, and while loops.

Does P satisfy this specification? That is, will two runs of P for which the values of x agree in the initial states also yield final states in which the values of $result$ agree? Note that the precondition does not make any commitments about $v \times$ and $h \times$.

To answer the above question, observe that since the program updates $result$ (line 4), for $result \times$ to hold at the end, $v \times$ must also hold. Alas, $v \times$ holds only at the beginning of every *odd* iteration of the loop — but fortunately, this is exactly when v is used to update $result$. It turns out that to verify the program we need the loop flow invariant $odd(i) \Rightarrow v \times$ which testifies to *conditionally secure information flow* within the loop.³ Furthermore, after $result$ is updated, the assignment to v (line 5) *invalidates* the invariant because $h \times$ need not hold. But because i is incremented by 1 (line 8), $odd(i)$ is falsified and the invariant is reestablished, vacuously, at the beginning of the next (even) iteration of the loop. Our algorithm, applied to the program in Fig. 1(a) and equipped with the above loop flow invariant, generates valid verification conditions (VCs) together with a precondition that includes $x \times$ but *not* $h \times$. Thus the program is deemed secure.

Note that standard security type systems do not take *conditional* loop flow invariants like the one above into account and therefore, given that $result$ has type L and h has type H , reject the program as insecure. (The security type given to a while loop can be interpreted as an unconditional loop flow invariant, which in this case is not precise enough.) For, well-typedness demands v to have type L , due to the assignment to $result$ (line 4), and also to have type H , due to the assignment to v (line 5).

Object flow invariants. The next example is motivated by an actual program, used in hardware verification of operational amplifiers, that was provided by our industrial collaborators, Rockwell-Collins. The example also serves to introduce the heap manipulating fragment of the language we analyze. We are given a collection of objects where each object has three fields: val containing its “value”, src containing the “source” object whose value will be used to update the val field, and idx containing the object’s index in the collection. The overall policy specification is that *odd* elements should be public; formally, we need to specify

$$\begin{aligned} odd(o.idx) &\Rightarrow (o.val) \times \text{ and} \\ odd(o.idx) &\Rightarrow (o.src) \times. \end{aligned}$$

Given this *object flow invariant*, we now ask whether the program

```
y := x.src; i := x.idx;
q := y.val; x.val := q; result := x.val
```

satisfies the policy $\{x \times\} - \{odd(i) \Rightarrow result \times\}$.

Intuitively, for this to hold we must demand that if the val field of an object with odd index is updated with a value q then the source object whose val field contains q must be one with odd index. We therefore assert an implication based on the above intuition:

```
y := x.src; i := x.idx;
assert (odd(i) → odd(y.idx));
q := y.val; x.val := q; result := x.val
```

It is well-known that standard Hoare logic does not handle heaps very well, a key issue being “pointer swing” that leads to aliasing. An update of $u.f$ may affect $w.f$ if u and w may alias. Rather than employ a may-alias analysis, we demand that all field accesses and updates be *scoped*. For example, a field access, $y := x.f$, occurs as **open** x **in** $y := .f$; **close**. A field update, $x.f := y$, occurs as **open** x **in** $.f := y$; **close**.

Fig. 1(b) shows the program that corresponds to the one above. It also exemplifies the syntax of the language that we analyze: it is a simple imperative language, extended with assertions and scoped heap manipulating commands (field accesses, field updates, object allocation). A formal BNF appears in Sect. 3.

Because of scoped field accesses and updates, we no longer need a prefix for a field as this is clear from the scope. In general, to compare claims about two different scopes, as in $\text{assert}(odd(x.idx) \rightarrow odd(y.idx))$, we need to save the result of $x.idx$ into a variable i . Finally, it turns out that we must assist our analysis by explicitly asserting (line 10) that when x is opened the second time, the index is still i .

The task of each scope is now to maintain the object flow invariant. To see that reasoning about aliasing is not a problem, observe that it is possible that updating the object pointed to by x also updates the object pointed to by y . However, this is permissible as long as the new object state satisfies the object flow invariant.

Note that the assertions used in the program (lines 6, 10) can be eliminated by theorem proving tools used

³Note that we do not want $odd(i)$ in the precondition along with $x \times$; i can be any integer, odd or even.

| | |
|---|---|
| <pre> 1. $i := 0; result := 0;$ 2. while ($i < 7$) do 3. if $odd(i)$ 4. then $result := result + v;$ 5. $v := v + h;$ 6. else $v := x;$ 7. fi; 8. $i := i + 1;$ 9. od </pre> <p style="text-align: center;">(a)</p> | <pre> 1. open x in 2. $y := .src;$ 3. $i := .idx;$ 4. close; 5. open y in 6. assert ($odd(i) \rightarrow odd(.idx)$); 7. $q := .val;$ 8. close; 9. open x in 10. assert ($.idx = i$); 11. $.val := q;$ 12. $result := .val;$ 13. close; </pre> <p style="text-align: center;">(b)</p> |
|---|---|

Figure 1. Two examples that illustrate (a) loop flow invariants, and (b) object flow invariants and scoped heap operations. $odd(i)$ is expressible as $(i \bmod 2 = 1)$ in our language.

in conjunction with other static analyses. In particular, the first assertion (line 6) could be eliminated in case we can prove, say, that for all objects o we have $o.src.idx = o.idx + 2$.

Our algorithm for verification condition generation, when given as input the program in Fig. 1(b) with postcondition $odd(i) \Rightarrow result \times$ and object flow invariant $\{odd(.idx) \Rightarrow .val \times, odd(.idx) \Rightarrow .src \times\}$, generates (as sketched in Sect. 5) valid VCs, and the precondition $true \Rightarrow x \times$ (equivalent to $x \times$).

Combining loop flow invariants, object flow invariants, and allocation. Next, we consider the example in Fig. 2, featuring a heterogeneous list pointed to by x and represented as a node chain, where one node can be reached from another by traversing $next$ links. The val field of each node contains either a high (H) value or a low (L) value, where the protocol is that a value is L provided it is less than 10. Informally, the list satisfies an object flow invariant $.val < 10 \Rightarrow val \times$.

We wish to split the list pointed to by x and output two homogeneous lists, pointed to by y and z ; here y will point to a list containing all the nodes of x with val fields that are L , i.e., less than 10, whereas z will point to a list containing the other nodes of x . Since the final value of $result$ is taken from the list pointed to by y , the overall policy specification is $\{x \times\} - \{result \times\}$. Our algorithm verifies that the program in Fig. 2 satisfies

this specification, in that from postcondition $result \times$ it generates precondition $x \times$ and some valid VCs.

For the verification process, object flow invariants are needed; one might think that we need one invariant for each kind of node but those can be combined into a “universal” object flow invariant, using a field t which tags the lists x , y and z with 1, 2, 3 respectively.

$$\begin{aligned}
 (.t = 1 \wedge .val < 10) &\Rightarrow .val \times \\
 .t = 1 &\Rightarrow .next \times \quad .t = 1 \Rightarrow (.val < 10) \times \\
 .t = 2 &\Rightarrow .val \times \quad .t = 2 \Rightarrow .next \times
 \end{aligned}$$

Here $(.val < 10) \times$ is satisfied by a pair of states if they agree on the value of the comparison (but not necessarily on the value of $.val$).

The example also shows a scoped object allocation, where new objects (pointed to by y_1 and z_1) are allocated in the heap and their fields initialized as shown. Once all fields are initialized, the object flow invariant must have been established so that when the scope **new** . . . **close** is exited the object is in a “steady state”.

Readers familiar with the Boogie methodology [7] might notice some similarity between **open** . . . **close** and Boogie’s **unpack** and **pack**, where the object invariant must be reestablished at the end of every field update. Boogie requires object invariants to be associated with every object of a class. Our language seems impoverished in comparison to Boogie’s in that we have the equivalent of a single universal class, but as

| | |
|--|--|
| <pre> 1. $y := nil; z := nil;$ 2. while $x \neq nil$ do 3. open x in assert($.t = 1$); $v := .val$; $n := .next$; close; 4. $x := n$; 5. if $v < 10$ 6. then new y_1 in $.val := v$; $.t := 2$; $.next := y$; close; 7. $y := y_1$; 8. else new z_1 in $.val := v$; $.t := 3$; $.next := z$; close; 9. $z := z_1$; 10. fi; 11. od; </pre> | <pre> 12. $result := nil$; 13. while $y \neq nil$ do 14. open y in 15. assert($.t = 2$); 16. $result := .val$; 17. $y := .next$; 18. close; 19. od </pre> |
|--|--|

Figure 2. List splitting

the above object flow invariant shows, the use of tags enables us to encode multiple invariants.

3. Syntax and Semantics

Expression syntax. An expression $E \in \mathbf{Exp}$ is either an arithmetic expression $A \in \mathbf{AExp}$ or a boolean expression $B \in \mathbf{BExp}$, given by the syntax

$$\begin{aligned}
A &::= x \mid .f \mid c \mid nil \mid A \text{ op } A \\
B &::= A \text{ bop } A
\end{aligned}$$

where we use x, y, \dots to range over variables in \mathbf{Var} , and f, g, \dots to range over field names in \mathbf{Fld} , and c to range over integer constants, and op to range over arithmetic operators in $\{+, \times, \text{mod}, \dots\}$, and bop to range over comparison operators in $\{=, <, \dots\}$.

We write $\text{fv}(E)$ (or $\text{ff}(E)$) for the variables (field names) occurring free in E . We write $E[A/x]$ for the result of substituting all occurrences of x in E by A ; similarly we write $E[A/.f]$. We say that E is *field-free* if E contains no field names, and that E is an *object expression* if E contains no variables.

We assume that each variable and each field is either for integers or for pointers (to objects), as prescribed by a function type mapping $\mathbf{Var} \cup \mathbf{Fld}$ into $\{\text{int}, \text{obj}\}$. We shall only consider programs that are “well-typed” in that respect. In particular, we disallow pointer arithmetic; the only operation allowed on pointers is pointer equality. Therefore, if $\text{type}(x) = \text{obj}$ then $x \in \text{fv}(A)$ implies $A = x$, and $x \in \text{fv}(B)$ implies that B is either $x = x$ or $A = x$ or $x = A$ with $x \notin \text{fv}(A)$.

Semantic domains. A value ($v \in \text{Val}$) is an integer n , a location $l \in \text{Loc}$, or nil ; default values are defined as $\text{deflt}(\text{int}) = 0$ and $\text{deflt}(\text{obj}) = nil$, and we

write $\text{deflt}(f)$ for $\text{deflt}(\text{type}(f))$. A *store* $s \in \text{Store}$ maps variables to values, an *object state* r maps field names to values, and a *heap* $h \in \text{Heap}$ maps locations to object states; the notions of $\text{dom}(\cdot)$ and $\text{ran}(\cdot)$ are as usual except that (with misuse of notation) we write $\text{ran}(h) = \{v \mid \exists l \in \text{dom}(h), f \in \mathbf{Fld} \bullet v = h(l)(f)\}$. We write $[s \mid x \mapsto v]$ for the store that is like s except that it maps x into v ; similarly we write $[r \mid f \mapsto v]$ and $[h \mid l \mapsto r]$.

Expression semantics. The semantics of an arithmetic (boolean) expression is a function from stores and object states into values (booleans). If an expression E is field-free (an object expression), the “ r ” component (the “ s ” component) can be omitted.

$$\begin{aligned}
\llbracket x \rrbracket_r^s &= s(x), \quad \llbracket .f \rrbracket_r^s = r(f), \quad \llbracket c \rrbracket_r^s = c, \\
\llbracket nil \rrbracket_r^s &= nil \\
\llbracket A_1 + A_2 \rrbracket_r^s &= \llbracket A_1 \rrbracket_r^s + \llbracket A_2 \rrbracket_r^s, \text{ etc.} \\
\llbracket A_1 < A_2 \rrbracket_r^s &= \text{True iff } \llbracket A_1 \rrbracket_r^s < \llbracket A_2 \rrbracket_r^s, \text{ etc.}
\end{aligned}$$

One-state assertions. We use $\phi \in \mathbf{1Assert}$ to range over “standard” assertions, given by the syntax

$$\phi ::= B \mid \phi \wedge \phi \mid \phi \vee \phi \mid \neg \phi$$

We may define *true* as $0 = 0$, and *false* as $0 = 1$; as usual, we define $\phi_1 \rightarrow \phi_2$ as $\neg \phi_1 \vee \phi_2$. We write $\phi[A/x]$ for the result of substituting all occurrences of x in ϕ by A ; similarly we define $\phi[A/.f]$.

The satisfaction relation for assertions reads $s, r \models \phi$ and denotes that ϕ holds in the *one state* comprised by the store s and the object state r . The definition is standard and elided. We say that ϕ is field-free if ϕ contains no field names, in which case the r component

can be omitted; we say that ϕ is an *object assertion* if ϕ contains no variables, in which case the s component can be omitted.

Command syntax. A command $S \in \mathbf{Cmd}$ is either a *top-level command* $TS \in \mathbf{TCmd}$ or a *record command* $RS \in \mathbf{RCmd}$; the latter is executed within the scope of a *single* object and is thus allowed to reference its fields. The syntax is given in Fig. 3, where in the grammar for TS we demand that all instances of A , B , and ϕ are field-free.

Command semantics. A record command transforms the store, and the state of the object being manipulated, into another store and another object state; hence its semantics is given in relational style, in the form $s, r \llbracket RS \rrbracket s', r'$. A top-level command transforms a store and a heap into another store and another heap; thus its semantics is given in the form $s, h \llbracket TS \rrbracket s', h'$. The semantics is defined inductively on RS and TS ; some key clauses are given in Fig. 6 in Appendix. Note that for some TS and s, h , there may not exist any s', h' such that $s, h \llbracket TS \rrbracket s', h'$ (modulo the choice of fresh location for object allocation, there exists at most one s', h'); this can happen if a **while** loop does not terminate, or an **assert** fails.

Two-state assertions. We shall use $\theta \in \mathbf{2Assert}$ to range over conditional agreement assertions, also called *2-assertions*; they are of the form $\phi \Rightarrow E \times$ which intuitively is satisfied by a pair of states if either at least one of them does not satisfy ϕ , or they agree on the value of E . As we cannot expect two runs to choose the same fresh location for object allocation, we employ a bijection β between locations; we extend β so that $c \beta c$ for all integers c , $nil \beta nil$, $\text{True} \beta \text{True}$, and $\text{False} \beta \text{False}$.

Then we define $s, r \& s_1, r_1 \models_{\beta} \theta$, the satisfaction relation for 2-assertions, by

$$s, r \& s_1, r_1 \models_{\beta} \phi \Rightarrow E \times \text{ iff whenever } s, r \models \phi \\ \text{ and } s_1, r_1 \models \phi \text{ then } \llbracket E \rrbracket_r^s \beta \llbracket E \rrbracket_{r_1}^{s_1}.$$

For $\theta = (\phi \Rightarrow E \times)$, we call ϕ the antecedent of θ and write $\phi = \text{ant}(\theta)$, and we call E the consequent of θ and write $E = \text{con}(\theta)$. We say that θ is field-free if it contains no field names, in which case the r and r_1 can be omitted, and say that θ is an object assertion if it contains no variables, in which case the s and s_1 can be omitted.

We use $\Theta \in \mathcal{P}(\mathbf{2Assert})$ to range over sets of 2-assertions, with conjunction implicit. Thus

$$s, r \& s_1, r_1 \models_{\beta} \Theta \text{ iff } \forall \theta \in \Theta \bullet s, r \& s_1, r_1 \models_{\beta} \theta.$$

Object flow invariants. We assume that there exists an object assertion \mathcal{I} that serves as a flow invariant for every object (cf. the discussion at the end of Sect. 2). We shall demand that for two runs of the program, the heap part obeys this invariant (except when an object is being manipulated within a scoped construct), and thus define

$$h \& h_1 \models_{\beta} \mathcal{I} \text{ iff for all } l, l_1 \text{ with } l \beta l_1: \\ h(l) \& h_1(l_1) \models_{\beta} \mathcal{I}.$$

4. Algorithm

We shall define, as done in Figs. 4 & 5, an algorithm VCgen for inferring preconditions, and verification conditions, from postconditions. We write

$$[VC]\{\Theta\}(R) \Leftarrow S \{\Theta'\}$$

if from input S and Θ' , VCgen returns output Θ , R , and VC . Here S is a command, Θ' is the desired postcondition for S , and Θ is a precondition for S that is designed so as to be sufficient to establish Θ' ; if S is a top-level command then VCgen requires Θ' to be field-free and ensures that Θ is field-free. We shall shortly explain the role of the verification conditions VC , but shall first explain the R component which captures how 2-assertions in Θ relate to 2-assertions in Θ' . More precisely, we have $R \subseteq \Theta \times \{m, u\} \times \Theta'$ where tags m, u are mnemonics for “modified” and “unmodified”; we use γ to range over $\{m, u\}$. We write $\text{dom}(R) = \{\theta \mid \exists(\theta, -, -) \in R\}$ and $\text{ran}(R) = \{\theta' \mid \exists(-, -, \theta') \in R\}$. Intuitively, if $(\theta, -, \theta') \in R$ then θ is in the precondition because θ' is in the postcondition (θ' is an origin of θ); moreover, if $(\theta, u, \theta') \in R$ then additionally it holds that S modifies no “relevant” variable or field name, where a “relevant” variable is one occurring in the consequent of θ' . For example, if S is $x := w$ then R might contain the triplets $(q > 4 \Rightarrow w \times, m, q > 4 \Rightarrow x \times)$ and $(w > 3 \Rightarrow z \times, u, x > 3 \Rightarrow z \times)$.

Verification conditions. These are either of the form $\phi \triangleright^1 \phi'$, meaning that ϕ logically implies ϕ' , or of the form $\Theta \triangleright^2 \theta$, again meaning that Θ logically implies θ but now for 2-assertions. Thus $\models \phi \triangleright^1 \phi'$ iff for all s, r : whenever $s, r \models \phi$ then also $s, r \models \phi'$; and $\models \Theta \triangleright^2 \theta$ iff for all s, r, s_1, r_1, β : whenever

| | | | |
|----------------------|--|--|--|
| | $RS ::= \text{skip}$ | | $TS ::= \text{skip}$ |
| assertion | $\text{assert}(\phi)$ | | $\text{assert}(\phi)$ |
| sequential execution | $RS ; RS$ | | $TS ; TS$ |
| conditional | $\text{if } B \text{ then } RS \text{ else } RS$ | | $\text{if } B \text{ then } TS \text{ else } TS$ |
| iteration | $\text{while } B \text{ do } RS$ | | $\text{while } B \text{ do } TS$ |
| variable assignment | $x := A$ | | $x := A$ |
| field update | $.f := A$ | | |
| object allocation | | | $\text{new } x \text{ in } RS \text{ close}$ |
| object manipulation | | | $\text{open } x \text{ in } RS \text{ close}$ |

Figure 3. Command syntax

$s, r \& s_1, r_1 \models_{\beta} \Theta$ then also $s, r \& s_1, r_1 \models_{\beta} \theta$. We use VC to range over sets of verification conditions, and write $\models VC$ iff $\models vc$ holds for all $vc \in VC$.

Now assume that some vc in the output of $VCgen$ cannot be satisfied. (This is the only way that $VCgen$ can “fail” on a well-typed program.) Looking at the clauses, we see that vc must have been generated by either **open** or **while**. The former case would reflect the failure to prove that \mathcal{I} is indeed a flow invariant for objects in the heap; the user would then need to propose another object flow invariant. The latter case would reflect the failure to prove that the given postcondition is indeed a loop flow invariant; the user would then need to strengthen it. The above situations are the only places where $VCgen$ needs user assistance.

Correctness results. Ultimately, we must express that if $[VC]\{\Theta\} (-) \Leftarrow S \{\Theta'\}$ with $\models VC$ then Θ is indeed a precondition that is strong enough to establish Θ' . (Θ may not be the *weakest* such precondition, however.) For record commands, this is stated as:

PROPOSITION 4.1 (Correctness of record commands).

Assume that

1. $[VC]\{\Theta\} (-) \Leftarrow RS \{\Theta'\}$ and that $\models VC$
2. $s, r \llbracket RS \rrbracket s', r'$ and $s_1, r_1 \llbracket RS \rrbracket s'_1, r'_1$
3. $s, r \& s_1, r_1 \models_{\beta} \Theta$.

Then $s', r' \& s'_1, r'_1 \models_{\beta} \Theta'$.

Note that Proposition 4.1 is termination-*insensitive*, as is also Theorem 4.2; this is not surprising given our choice of a relational semantics (but see [3] for a logic-based approach that is termination-sensitive).

Proposition 4.1 is used to prove correctness of top-level commands, for which the correctness statement is slightly more complex:

THEOREM 4.2 (Correctness). Assume that

1. $[VC]\{\Theta\} (-) \Leftarrow TS \{\Theta'\}$ and that $\models VC$
2. $s, h \llbracket TS \rrbracket s', h'$ and that $s_1, h_1 \llbracket TS \rrbracket s'_1, h'_1$
3. $s \& s_1 \models_{\beta} \Theta$ and $h \& h_1 \models_{\beta} \mathcal{I}$.
4. There exists $\theta'_0 \in \Theta'$ such that $s' \models \text{ant}(\theta'_0)$ and $s'_1 \models \text{ant}(\theta'_0)$.

Then there exists β' extending β such that $s' \& s'_1 \models_{\beta'} \Theta'$ and $h' \& h'_1 \models_{\beta'} \mathcal{I}$.

If TS contains no **new** commands, we may choose $\beta' = \beta$, but otherwise β' may be a proper extension of β so as to model that new heap locations have been allocated. Condition 4 is a bit nonintuitive, but it is (at least currently) needed for the proofs to carry through, and it is non-restrictive as it can be fulfilled by adding to Θ' a trivial 2-assertion $\text{true} \Rightarrow 0 \times$.

Theorem 4.2 is proved in a forthcoming⁴ technical report [4], by establishing a number of auxiliary properties. These properties have largely determined the design of $VCgen$ and will thus guide us as we later explain the various clauses of Figs. 4 & 5.

The first such property is a variant of the “*-property” by Bell and La Padula [10], also called “write confinement” [5], which is used to preclude, e.g., “low writes under high guards”. In our setting, it captures the role of the R component and reads as follows:

LEMMA 4.3 (Write Confinement).

Assume $[VC]\{\Theta\} (R) \Leftarrow S \{\Theta'\}$. Then $\text{dom}(R) = \Theta$ and $\text{ran}(R) = \Theta'$. Given $\theta' \in \Theta'$, there exists at most one θ such that $(\theta, u, \theta') \in R$. If there exists such θ , then $\text{con}(\theta) = \text{con}(\theta')$, and with $E = \text{con}(\theta)$ we have

⁴Note to referees: all proofs have been done; we are about to convert them into \LaTeX .

- if $s, _ \llbracket S \rrbracket s', _$ then s agrees with s' on $\text{fv}(E)$;
- if $s, r \llbracket S \rrbracket s', r'$ (thus S is of form RS) then also r agrees with r' on $\text{ff}(E)$.

Lemma 4.3 is needed in the proof of Theorem 4.2 (and Prop. 4.1) to handle the case where the two runs in question follow *different branches* in a conditional, as we must then ensure that neither run modifies a variable (field name) on which we want the two runs to agree afterwards.

We now embark on explaining the various clauses of VCgen in Figs. 4 and 5. For an assignment $x := A$, each 2-assertion $\phi \Rightarrow E \times$ in Θ' produces exactly one 2-assertion in Θ , given by substituting A for x (as in standard Hoare logic) in ϕ as well as in E ; the connection is tagged m when x occurs in E . The treatment of field update is similar, and of **skip** even simpler. The rule for $S_1 ; S_2$ works backwards, first computing the precondition for S_2 which is then used to compute the precondition for S_1 ; the tags express that a consequent is modified iff it has been modified in either S_1 or S_2 . The rule for **assert** allows us to weaken 2-assertions, by strengthening their antecedents; this is sound since execution will abort from states not satisfying the new antecedents.

To motivate the treatment (Fig. 4) of a conditional **if** B **then** S_1 **else** S_2 , assume that $\phi \Rightarrow E \times$ occurs in Θ' . If $(\phi \Rightarrow E \times) \in \Theta'_u$, we can assume from Lemma 4.3 that neither S_1 nor S_2 has modified E , and that the precondition of each S_i will contain a 2-assertion of the form $\phi_i \Rightarrow E \times$; these can now be combined by R_0 into one single precondition. On the other hand, if $(\phi \Rightarrow E \times) \in \Theta'_m$ then E has been modified by at least one branch; therefore, we should not allow two runs to take *different branches* if they both satisfy ϕ afterwards. This is ensured by R'_0 , while R'_1 (R'_2) caters for the case where both runs choose S_1 (S_2).

EXAMPLE 4.4. Consider the result of applying VCgen to the body of the **while** loop in Fig 1(a), with postcondition $\{x \times, \text{odd}(i) \Rightarrow v \times\}$. (We write $x \times$ for $\text{true} \Rightarrow x \times$.) Working backwards, the assignment to i transforms $\text{odd}(i) \Rightarrow v \times$ to $\text{odd}(i+1) \Rightarrow v \times$, which amounts to $\neg \text{odd}(i) \Rightarrow v \times$, but keeps $x \times$ unchanged. To process the conditional, we apply VCgen to the branches; the **else** branch produces R_2 given by

$$\begin{aligned} &(x \times, u, x \times), \\ &(\neg \text{odd}(i) \Rightarrow x \times, m, \neg \text{odd}(i) \Rightarrow v \times) \end{aligned}$$

while the **then** branch produces R_1 given by

$$\begin{aligned} &(x \times, u, x \times), \\ &(\neg \text{odd}(i) \Rightarrow (v + h) \times, m, \neg \text{odd}(i) \Rightarrow v \times) \end{aligned}$$

Referring to the clause for **if** in Fig. 4, we have $\Theta'_u = \{x \times\}$ and $\Theta'_m = \{\neg \text{odd}(i) \Rightarrow v \times\}$. The former contributes, by R_0 , the precondition $(\text{odd}(i) \vee \neg \text{odd}(i)) \Rightarrow x \times$ which amounts to $x \times$. The latter contributes by R'_1 the precondition $(\neg \text{odd}(i) \wedge \text{odd}(i)) \Rightarrow (v + h) \times$ which is vacuously true, by R'_2 the precondition $(\neg \text{odd}(i) \wedge \neg \text{odd}(i)) \Rightarrow x \times$ which amounts to $\neg \text{odd}(i) \Rightarrow x \times$, and by R'_0 the precondition $(\neg \text{odd}(i) \wedge \text{odd}(i) \vee \neg \text{odd}(i) \wedge \neg \text{odd}(i)) \Rightarrow \text{odd}(i) \times$ which is always true (two states satisfying $\neg \text{odd}(i)$ will agree on the value of $\text{odd}(i)$). Assuming VCgen is able to carry out such basic simplifications, it will return, for the body of the **while** loop, an R component given by

$$\begin{aligned} &(x \times, u, x \times), \\ &(\neg \text{odd}(i) \Rightarrow x \times, m, \text{odd}(i) \Rightarrow v \times) \end{aligned}$$

The noteworthy part is that even though the postcondition mentions $v \times$, and v is updated using h , VCgen generates a precondition which does not mention h , since it exploits the parity of i .

For a **while** loop (Fig. 5), VCgen checks whether the given postcondition Θ can indeed serve as a flow invariant. (As mentioned earlier this may fail in which case the user must strengthen the postcondition.) First we partition Θ into two sets, Θ_m and Θ_u ; a 2-assertion can be in the latter set if its consequent is not modified by the loop body. Now VC_2 serves a similar function as R'_0 did in the clause for conditionals: by demanding a precondition with the loop test B as consequent, it ensures that if one run stays in the loop and updates a variable on which the two runs must agree, then also the other run stays in the loop. When both runs stay in the loop, VC_1 ensures that the loop flow invariant is maintained.

The need for VC_3 , VC_4 and VC_5 is less obvious, but they are designed so as to establish an auxiliary result, stated below as Lemma 4.5. VC_3 demands that Θ_m contains an assertion θ_m with a “weakest” antecedent. (This is no serious restriction, since if $\Theta_m = \{\phi_i \Rightarrow E_i \times \mid i \in \{1 \dots n\}\}$ we can just add $(\phi_1 \vee \dots \vee \phi_n) \Rightarrow 0 \times$ to Θ_m .)

LEMMA 4.5. Assume $[VC]\{\Theta\} (R) \Leftarrow S \{\Theta'\}$ with $\models VC$. Given $\theta' \in \Theta'$, there exists $(\theta, \neg, \theta') \in R$ such that

- if $S = RS$: whenever $s, r \llbracket S \rrbracket s', r'$ and $s', r' \models \text{ant}(\theta')$ then $s, r \models \text{ant}(\theta)$;
- if $S = TS$: whenever $s, h \llbracket S \rrbracket s', h'$ and $s' \models \text{ant}(\theta')$ then $s \models \text{ant}(\theta)$.

For $S = \mathbf{while} B \mathbf{do} S_0$, if $\theta' \in \Theta_u$ we can use $\theta = \theta'$, otherwise we can use $\theta = \theta_m$.

We now address the clause for **open** x in RS **close**, where we first compute in Θ_0 a precondition for RS , given a postcondition that is augmented with \mathcal{I} (as the object invariant must be re-established at the end). Note that we must remove from Θ_0 any references to field names; for that purpose we assume that there is a function $ff^+ : \mathbf{1Assert} \rightarrow \mathbf{1Assert}$ such that if $\phi' = ff^+(\phi)$ then (i) ϕ' is field-free, and (ii) ϕ logically implies ϕ' . These demands are trivially fulfilled if $ff^+(\phi) = \text{true}$ for all ϕ , but a more precise solution is possible; then, e.g., ff^+ returns $x = 7$ given $x = 7 \wedge \neg(.f = 8)$. Thus, e.g., Θ will (by R_3) contain $x = 7 \Rightarrow y \times$ if Θ_0 contains $(x = 7 \wedge \neg(.f = 8)) \Rightarrow y \times$.

Equipped with ff^+ , we can explain the various clauses, first R_4 which “lifts out” assertions in Θ_0 that originate from a top-level assertion and whose consequents have not been modified. Now consider an assertion in Θ_0 whose consequent *has* been modified. If the resulting consequent is not field-free, we must demand that it follows from the object flow invariant, as expressed by VC_2 . Otherwise, it can be lifted out of the scope, as done by R_3 . A precondition, say $\text{true} \Rightarrow (.f + x) \times$ might need to be replaced by the two assertions $\text{true} \Rightarrow x \times$ and $\text{true} \Rightarrow .f \times$ which together are strictly stronger; the former can be lifted out, the latter must follow from \mathcal{I} . Also, assertions in \mathcal{I} whose consequents have not been modified (and therefore still contain field names) must follow from \mathcal{I} , as expressed by VC_1 . The role of R_1 and R_2 is to ensure that if a relevant variable (in Θ' or in \mathcal{I}) is modified, the two runs are indeed manipulating the same object.

Note that R_2 ensures that there are “ m ” tags going out from *all* 2-assertions in the postcondition of a command that modifies a consequent of a 2-assertion in \mathcal{I} . This property is required by the following Lemma:

LEMMA 4.6. Assume $[VC]\{\Theta\} (R) \Leftarrow TS \{\Theta'\}$ with $\models VC$, and that $\theta' \in \Theta'$ is such that if $(\neg, \gamma, \theta') \in$

R then $\gamma = u$. For $(\phi_0 \Rightarrow E_0 \times) \in \mathcal{I}$, if $s, h \llbracket TS \rrbracket s', h'$ then for all $l \in \text{dom}(h)$: (i) if $h'(l) \models \phi_0$ then $h(l) \models \phi_0$; (ii) $h(l)(f) = h'(l)(f)$ for all f in $\text{ff}(E_0)$.

To see why Lemma 4.6 is needed, recall that the correctness of **if** and **while** rests on Lemma 4.3 which ensures that if two runs follow different paths then they do not modify consequents of top-level assertions. Lemma 4.6 now further ensures that two such diverting runs do not invalidate object flow invariants.

The clause for **new** first computes in Θ_0 a precondition for RS , and then exploits that the semantics of **new** initializes all fields to a default value. So if Θ_0 contains say $.f = 1 \Rightarrow y \times$, we generate the (trivial) precondition $0 = 1 \Rightarrow y \times$; if Θ_0 contains say $\text{true} \Rightarrow (.f + y) \times$, we generate the precondition $\text{true} \Rightarrow (0 + y) \times$. We also want to eliminate x from the precondition; this is possible due to the freshness of the new location and the absence of pointer arithmetic: after object allocation, it can never hold that $x = A$, unless $A = x$. This is formalized by the function $rm_x : \mathbf{1Assert} \rightarrow \mathbf{1Assert}$ which is a homomorphism on the structure of ϕ , which maps $x = x$ to true , which maps $x = A$ and $A = x$ to false if $x \neq A$ and hence $x \notin \text{fv}(A)$, and which maps any B not containing x to itself. Concerning the consequents, we exploit that two runs will always agree on the value of x after allocation (as β can be extended to relate the fresh locations); this is formalized by the function $rm_x : \mathbf{Exp} \rightarrow \mathbf{Exp}$ which maps E into 0 if $x \in \text{fv}(E)$, and into E otherwise.

Strengthening and simplifying assertions. As can be seen by inspecting, e.g., the clause for conditionals, the preconditions generated by VCgen may contain a number of assertions which is exponential in the size of the program. Our implementation therefore needs to be able to simplify assertions, replacing a precondition with one which is equivalent. In particular, it is important (cf. Example 4.4) to recognize when a 2-assertion has an antecedent which is always false, or when it is of the form $\phi \Rightarrow B \times$ where ϕ implies B (or $\neg B$), since then it can be eliminated. Preliminary experiments with our prototype implementation indicate that a few such rules are sufficient to yield readable preconditions; this makes us hope for a running time which is close to linear though further experiments are needed.

Let us be a bit more formal about what must hold, apart from $\{\theta_1, \dots, \theta_n\} \triangleright^2 \theta$, when θ is replaced by

$\theta_1 \dots \theta_n$. Lemma 4.5 requires that for at least one $i \in \{1 \dots n\}$ we can verify $\text{ant}(\theta) \triangleright^1 \text{ant}(\theta_i)$. Moreover, we need to record in R that θ is related to each θ_i , and if we want to assign the tag u we must demand (due to Lemma 4.3) that $n = 1$ and $\text{con}(\theta) = \text{con}(\theta_1)$. These considerations suggest that rather than eliminating a 2-assertion which is always true, we replace it by a designated such assertion, e.g., $\text{true} \Rightarrow 0 \times$.

5. Worked Out Example

In this section we work out the example in Fig. 1(b). The other examples are worked out in the forthcoming technical report. We want to prove that the program satisfies the specification $\{\text{true} \Rightarrow x \times\} - \{\text{odd}(i) \Rightarrow \text{result} \times\}$. The object invariant, \mathcal{I} , is a conjunction of $\text{odd}(.idx) \Rightarrow .\text{val} \times$ and $\text{odd}(.idx) \Rightarrow .\text{src} \times$.

We first consider the last **open**, lines 9–13 of Fig. 1(b), where we must analyze the body (lines 10–12) with a postcondition which is $\text{odd}(i) \Rightarrow \text{result} \times$ conjoined with the object invariant. Using VCgen’s clauses for assignment, field update, and **assert**, this yields an empty set of VCs, and R_0 containing

$$\begin{aligned} &(\text{odd}(i) \wedge (.idx = i) \Rightarrow q \times, m, \text{odd}(i) \Rightarrow \text{result} \times) \\ &(\text{odd}(.idx) \wedge (.idx = i) \Rightarrow q \times, m, \text{odd}(.idx) \Rightarrow .\text{val} \times) \\ &(\text{odd}(.idx) \wedge (.idx = i) \Rightarrow .\text{src} \times, u, \text{odd}(.idx) \Rightarrow .\text{src} \times) \end{aligned}$$

Applying the clause in VCgen for **open** now generates the verification conditions: $VC_1 = \{\text{odd}(.idx) \wedge (.idx = i) \triangleright^1 \text{odd}(.idx)\}$ and $VC_2 = \{\}$. (To see why VC_2 is empty, note that the relevant assertions are of the form $_ \Rightarrow q \times$ but $q \times$ is field-free.) Also, it generates a set R which is the union of the sets R_1, R_2, R_3 below (since R_4 is empty).

$$\begin{aligned} R_1 &= \{(\text{odd}(i) \Rightarrow x \times, m, \text{odd}(i) \Rightarrow \text{result} \times)\} \\ R_2 &= \{\text{true} \Rightarrow x \times, m, \text{odd}(i) \Rightarrow \text{result} \times\} \\ R_3 &= \{(\text{odd}(i) \Rightarrow q \times, m, \text{odd}(i) \Rightarrow \text{result} \times)\} \end{aligned}$$

We have assumed that ff^+ maps $\text{odd}(.idx) \wedge (.idx = i)$ into $\text{odd}(i)$. Now the precondition of lines 9–13 can be read off from the above sets as

$$\{\text{odd}(i) \Rightarrow x \times, x \times, \text{odd}(i) \Rightarrow q \times\}$$

where the first assertion can be removed as it follows from the second. Lines 1–8 of Fig. 1(b) are next analyzed with the above as postcondition. For lines 5–8, it can be shown that, using the case for **open**, VCgen generates the verification conditions

$$\begin{aligned} &\text{odd}(.idx) \wedge (\text{odd}(i) \rightarrow \text{odd}(.idx)) \triangleright^1 \text{odd}(.idx), \\ &\mathcal{I} \triangleright^2 (\text{odd}(i) \wedge (\text{odd}(i) \rightarrow \text{odd}(.idx))) \Rightarrow .\text{val} \times \end{aligned}$$

and preconditions $\{\text{odd}(i) \Rightarrow y \times, \text{true} \Rightarrow x \times\}$. Finally, the analysis of lines 1–4 of Fig. 1(b) generates the VCs $\text{odd}(.idx) \triangleright^1 \text{odd}(.idx)$, $\mathcal{I} \triangleright^2 \text{odd}(.idx) \Rightarrow .\text{src} \times$ as well as the overall precondition of the program, $\text{true} \Rightarrow x \times$. The overall VCs generated for the program are the union of the VCs generated for each of the **open**. We note that the VCs are valid.

Fig. 7 in the Appendix shows the assertions that hold at each line in the program.

6. Discussion

A recently popular approach to information flow analysis is *self-composition*, first proposed by Barthe et al. [9] and later extended by, e.g., Terauchi and Aiken [24] and Naumann [22]. Self-composition works as follows: for a given program S , a copy S' is created with all variables renamed (primed); with the observable variables say x, y , then NI holds provided the sequential composition $S; S'$ when given precondition $x = x' \wedge y = y'$ also ensures postcondition $x = x' \wedge y = y'$.

Terauchi and Aiken [24] use self-composition to verify information flow automatically using the BLAST [20] tool. To obtain good experimental results, they introduce sound program transformations of self-composed programs; it is also often necessary to leverage the results of a standard information flow analyses, such as a security typing. In a sense, our approach is dual in that noninterference properties are explicit in our analysis but we can leverage standard assertions, inserted and/or checked by general verifiers. An interesting question is whether the 2-assertions generated by VCgen could be translated into assertions that would assist the self-composition approach.

Since [24] does not address heap-manipulating programs, the work most closely related to ours is the one by Naumann [22] whose goal was the verification of information flow using existing verifiers like Spec# [8] or ESC/Java2 [14], and whose contribution is to extend the theory of self-composition to account for manipulations of heap objects. In some cases, like for while loops, it is more practical (but not necessary) for the technique to perform program transformations. For heap-manipulating programs, the two copies of the programs involve different sets of objects and therefore the correspondence between the objects (“mates” in Naumann’s terminology) must be made explicit in the

$$\begin{aligned}
[VC]\{\Theta\} (R) &\Leftarrow \mathbf{skip} \{\Theta'\} \\
&\text{iff } R = \{(\theta, u, \theta) \mid \theta \in \Theta'\} \text{ and } \Theta = \Theta' \text{ and } VC = \emptyset \\
[VC]\{\Theta\} (R) &\Leftarrow \mathbf{assert}(\phi_0) \{\Theta'\} \\
&\text{iff } R = \{((\phi \wedge \phi_0) \Rightarrow E \times, u, \phi \Rightarrow E \times) \mid \phi \Rightarrow E \times \in \Theta'\} \\
&\text{and } \Theta = \text{dom}(R) \text{ and } VC = \emptyset \\
[VC]\{\Theta\} (R) &\Leftarrow x := A \{\Theta'\} \\
&\text{iff } R = \{(\phi[A/x] \Rightarrow E[A/x] \times, \gamma, \phi \Rightarrow E \times) \mid \phi \Rightarrow E \times \in \Theta'\} \\
&\quad \text{where } \gamma = m \text{ iff } x \in \text{fv}(E) \\
&\text{and } \Theta = \text{dom}(R) \text{ and } VC = \emptyset \\
[VC]\{\Theta\} (R) &\Leftarrow .f := A \{\Theta'\} \\
&\text{iff } R = \{(\phi[A/.f] \Rightarrow E[A/.f] \times, \gamma, \phi \Rightarrow E \times) \mid \phi \Rightarrow E \times \in \Theta'\} \\
&\quad \text{where } \gamma = m \text{ iff } f \in \text{ff}(E) \\
&\text{and } \Theta = \text{dom}(R) \text{ and } VC = \emptyset \\
[VC]\{\Theta\} (R) &\Leftarrow S_1 ; S_2 \{\Theta'\} \\
&\text{iff } [VC_2]\{\Theta''\} (R_2) \Leftarrow S_2 \{\Theta'\} \text{ and } [VC_1]\{\Theta\} (R_1) \Leftarrow S_1 \{\Theta''\} \\
&\text{and } R = \{(\theta, \gamma, \theta') \mid \exists \theta'', \gamma_1, \gamma_2 \bullet (\theta, \gamma_1, \theta'') \in R_1, (\theta'', \gamma_2, \theta') \in R_2\} \\
&\quad \text{where } \gamma = m \text{ iff } \gamma_1 = m \text{ or } \gamma_2 = m \\
&\text{and } VC = VC_1 \cup VC_2 \\
[VC]\{\Theta\} (R) &\Leftarrow \mathbf{if } B \mathbf{ then } S_1 \mathbf{ else } S_2 \{\Theta'\} \\
&\text{iff } [VC_1]\{\Theta_1\} (R_1) \Leftarrow S_1 \{\Theta'\} \text{ and } [VC_2]\{\Theta_2\} (R_2) \Leftarrow S_2 \{\Theta'\} \\
&\text{and } R = R'_1 \cup R'_2 \cup R'_0 \cup R_0 \\
&\quad \text{where } R'_1 = \{((\phi_1 \wedge B) \Rightarrow E_1 \times, m, \theta') \mid \theta' \in \Theta'_m, (\phi_1 \Rightarrow E_1 \times, -, \theta') \in R_1\} \\
&\quad \text{and } R'_2 = \{((\phi_2 \wedge \neg B) \Rightarrow E_2 \times, m, \theta') \mid \theta' \in \Theta'_m, (\phi_2 \Rightarrow E_2 \times, -, \theta') \in R_2\} \\
&\quad \text{and } R'_0 = \{(((\phi_1 \wedge B) \vee (\phi_2 \wedge \neg B)) \Rightarrow B \times, m, \theta') \\
&\quad \quad \mid \theta' \in \Theta'_m, (\phi_1 \Rightarrow E_1 \times, -, \theta') \in R_1, (\phi_2 \Rightarrow E_2 \times, -, \theta') \in R_2\} \\
&\quad \text{and } R_0 = \{(((\phi_1 \wedge B) \vee (\phi_2 \wedge \neg B)) \Rightarrow E \times, u, \theta') \\
&\quad \quad \mid \theta' \in \Theta'_u, (\phi_1 \Rightarrow E \times, u, \theta') \in R_1, (\phi_2 \Rightarrow E \times, u, \theta') \in R_2\} \\
&\quad \text{and } \Theta'_m = \{\theta' \in \Theta' \mid \exists (-, m, \theta') \in R_1 \cup R_2\} \\
&\quad \text{and } \Theta'_u = \Theta' \setminus \Theta'_m \\
&\text{and } \Theta = \text{dom}(R) \text{ and } VC = VC_1 \cup VC_2
\end{aligned}$$

Figure 4. The verification condition generator, part I

specification of the composed program. Our approach avoids program transformations, and our specifications do not need to specify mates: that is handled by the semantics of assertions. On the other hand, we cannot use an existing verifier like Spec# or ESC/Java2 directly; we must thus show how preconditions and VCs are actually generated.

Dufay et al. [16] use self-composition to check non-interference for data mining algorithms implemented in Java, using the Krakatoa tool, based on the Coq theorem prover and using JML [12]. However, they do not

provide details on how the heap is handled. Darvas et al. [15] use the KeY tool for interactive verification of noninterference. Information flow is modeled by a dynamic logic formula rather than by assertions as in self-composition.

Bergeretti and Carré [11] present a compositional method for inferring which variables are dependent on which variables; this technique forms the basis for the Spark Ada Examiner [13] which requires that each method is annotated with `derives` annotations like

`derives u from y,z, derives w from x,y`

$[VC]\{\Theta\} (R) \Leftarrow \text{while } B \text{ do } S_0 \{\Theta\}$
iff $[VC_0]\{\Theta_0\} (R_0) \Leftarrow S_0 \{\Theta\}$
and $R = \{(\theta, u, \theta) \mid \theta \in \Theta_u\} \cup \{(\theta_1, m, \theta_2) \mid \theta_1, \theta_2 \in \Theta_m\}$
and $VC = VC_0 \cup VC_1 \cup VC_2 \cup VC_3 \cup VC_4 \cup VC_5$
where $VC_1 = \{\Theta \triangleright^2 (\phi \wedge B) \Rightarrow E \times \mid (\phi \Rightarrow E \times, -, -) \in R_0\}$
and $VC_2 = \{\Theta \triangleright^2 \phi_m \Rightarrow B \times\}$
and $VC_3 = \{ant(\theta) \triangleright^1 \phi_m \mid \theta \in \Theta_m\}$
and $VC_4 = \{ant(\theta) \triangleright^1 \phi_m \mid (\theta, -, \theta_m) \in R_0\}$
and $VC_5 = \{ant(\theta_0) \triangleright^1 ant(\theta) \mid (\theta_0, -, \theta) \in R_0, \theta \in \Theta_u\}$
and $\Theta_m = \{\theta \in \Theta \mid \exists(-, m, \theta) \in R_0\}$
and $\Theta_u = \Theta \setminus \Theta_m$
and Θ_m contains a special element θ_m with $\phi_m = ant(\theta_m)$

$[VC]\{\Theta\} (R) \Leftarrow \text{open } x \text{ in } RS \text{ close } \{\Theta'\}$
iff $[VC_0]\{\Theta_0\} (R_0) \Leftarrow RS \{\Theta' \cup \mathcal{I}\}$
and $R = R_1 \cup R_2 \cup R_3 \cup R_4$
where $R_1 = \{(\mathit{ff}^+(\phi) \Rightarrow x \times, m, \theta') \mid \theta' \in \Theta', (\phi \Rightarrow - \times, m, \theta') \in R_0\}$
and $R_2 = \text{if exists } \theta \in \mathcal{I} \text{ with } (-, m, \theta) \in R_0 \text{ then } \{(true \Rightarrow x \times, m, \theta') \mid \theta' \in \Theta'\} \text{ else } \emptyset$
and $R_3 = \{(\mathit{ff}^+(\phi) \Rightarrow E \times, m, \theta') \mid \theta' \in \Theta', E \text{ field-free, } \exists \theta'_0 \in \mathcal{I} \cup \{\theta'\} \bullet (\phi \Rightarrow E \times, m, \theta'_0) \in R_0\}$
and $R_4 = \{(\mathit{ff}^+(\phi) \Rightarrow E \times, u, \theta') \mid \theta' \in \Theta', (\phi \Rightarrow E \times, u, \theta') \in R_0\}$
and $\Theta = dom(R)$ and $VC = VC_0 \cup VC_1 \cup VC_2$
where $VC_1 = \{ant(\theta) \triangleright^1 ant(\theta') \mid \theta' \in \mathcal{I}, (\theta, u, \theta') \in R_0\}$
and $VC_2 = \{\mathcal{I} \triangleright^2 \theta \mid (\theta, m, -) \in R_0, con(\theta) \text{ not field-free}\}$

$[VC]\{\Theta\} (R) \Leftarrow \text{new } x \text{ in } RS \text{ close } \{\Theta'\}$
iff $[VC_0]\{\Theta_0\} (R_0) \Leftarrow RS \{\Theta' \cup \mathcal{I}\}$
and $R = \{(rm_x(\phi[\overline{deflt}(f)/\bar{f}]) \Rightarrow rm_x(E[\overline{deflt}(f)/\bar{f}]) \times, \gamma, \theta') \mid (\phi \Rightarrow E \times, \gamma_0, \theta') \in R_0, \gamma = m \text{ iff } \gamma_0 = m \text{ or } x \in fv(E)\}$
and $\Theta = dom(R)$ and $VC = VC_0$

Figure 5. The verification condition generator, part II

It is interesting to observe that such “channels” of information flow is captured by our R component, as when

$$[VC]\{x \times, y \times, z \times\} (R) \Leftarrow S \{u \times, w \times\}$$

with R containing the elements $(y \times, -, u \times)$, $(z \times, -, u \times)$, $(x \times, -, w \times)$, $(y \times, -, w \times)$. Our approach is more general in that it also captures *conditional* channels; we plan to investigate how to extend the Spark Ada Examiner framework to express R elements like $(i > 5 \Rightarrow y \times, -, j > 7 \Rightarrow u \times)$. Also, we hope to investigate the relationship to the path conditions presented by Hammer et al. [19].

In the near future, we plan to experiment with the prototype implementation which is currently being developed by our undergraduate student Jonathan Hoag. Over the summer, we might try to integrate it with the Bogor tool [17] to generate and/or check standard as-

sertions that will increase precision. To ease expressiveness, we would like to allow multiple scopes to be simultaneously open.

An important long-term goal is to develop techniques for the automatic computation of flow (loop/object) invariants, thereby moving closer to an automatic information flow analysis, and to extend the framework to an interprocedural setting. We would also like a (sound and preferably complete) axiomatization of \triangleright^2 so as to automatically check whether the VCs generated are satisfiable; a trivial rule is that $\phi \Rightarrow x \times \wedge \phi \Rightarrow w \times \triangleright^2 \phi \Rightarrow (x + w) \times$ holds for all ϕ, x, w . Relatedly, we would like to investigate whether our analysis is in some sense “optimal”, with the preconditions being “weakest”.

References

- [1] Torben Amtoft, Sruthi Bandhakavi, and Anindya Banerjee. A logic for information flow in object-oriented programs. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 91–102, 2006. Extended version available as KSU CIS-TR-2005-1.
- [2] Torben Amtoft and Anindya Banerjee. Information flow analysis in logical form. In *SAS 2004 (11th Static Analysis Symposium)*, volume 3148 of *LNCS*, pages 100–115. Springer-Verlag, 2004.
- [3] Torben Amtoft and Anindya Banerjee. A logic for information flow analysis with an application to forward slicing of simple imperative programs. *Science of Computer Programming*, 64(1):3–28, 2007.
- [4] Torben Amtoft and Anindya Banerjee. Verification condition generation for conditional information flow. Technical Report CIS-TR-2007-2, Kansas State University, 2007.
- [5] Anindya Banerjee and David A. Naumann. Stack-based access control for secure information flow. *Journal of Functional Programming*, 15(2):131–177, 2005. Special issue on Language Based Security.
- [6] Anindya Banerjee, David A. Naumann, and Stan Rosenberg. Towards a logical account of declassification (short paper). In *PLAS*, 2007.
- [7] Michael Barnett, Robert DeLine, Manuel Fähndrich, K. Rustan M. Leino, and Wolfram Schulte. Verification of object-oriented programs with invariants. *Journal of Object Technology*, 3(6):27–56, 2004.
- [8] Michael Barnett, K. Rustan M. Leino, and Wolfram Schulte. The Spec# programming system: An overview. In *Proceedings of CASSIS*, volume 3362 of *Lecture Notes in Computer Science*, pages 49–69, 2004.
- [9] Gilles Barthe, Pedro R. D’Argenio, and Tamara Rezk. Secure information flow by self-composition. In *IEEE Computer Security Foundations Workshop (CSFW)*, 2004.
- [10] D.E. Bell and L.J. LaPadula. Secure computer systems: Mathematical foundations. Technical Report MTR-2547, MITRE Corp., 1973.
- [11] Jean-Francois Bergeretti and Bernard A. Carré. Information-flow and data-flow analysis of while-programs. *ACM Transactions on Programming Languages and Systems*, 7(1):37–61, January 1985.
- [12] Lilian Burdy, Yoonsik Cheon, David R. Cok, Michael D. Ernst, Joseph R. Kiniry, Gary T. Leavens, K. Rustan M. Leino, and Erik Poll. An overview of JML tools and applications. *STTT*, 7(3):212–232, 2005.
- [13] Roderick Chapman and Adrian Hilton. Enforcing security and safety models with an information flow analysis tool. In *SIGAda’04, Atlanta, Georgia*. ACM, November 2004.
- [14] David R. Cok and Joseph Kiniry. Esc/java2: Uniting ESC/Java and JML. In *Proceedings of CASSIS*, volume 3362 of *Lecture Notes in Computer Science*, pages 108–128, 2004.
- [15] Adam Darvas, Reiner Hähnle, and David Sands. A theorem proving approach to analysis of secure information flow. In *SPC*, volume 3362 of *Lecture Notes in Computer Science*, pages 151–171, 2005.
- [16] Guillaume Dufay, Amy Felty, and Stan Matwin. Privacy-sensitive information flow with JML. In *CADE*, 2005.
- [17] Matthew B. Dwyer, John Hatcliff, Matthew Hoosier, and Robby. Building your own software model checker using the Bogor extensible model checking framework. In *17th Conference on Computer-Aided Verification (CAV 2005)*, 2005.
- [18] Joseph Goguen and Jose Meseguer. Security policies and security models. In *Proc. IEEE Symp. on Security and Privacy*, pages 11–20, 1982.
- [19] Christian Hammer, Jens Krinke, and Gregor Snelting. Information flow control for java based on path conditions in dependence graphs. In *IEEE International Symposium on Secure Software Engineering (ISSSE 2006)*, pages 87–96, March 2006.
- [20] Thomas A. Henzinger, Ranjit Jhala, Rupak Majumdar, and Gregoire Sutre. Software verification with BLAST. In *10th SPIN Workshop on Model Checking Software (SPIN)*, volume 2648 of *Lecture Notes in Computer Science*, pages 235–239, 2003.
- [21] Andrew C. Myers. JFlow: Practical mostly-static information flow control. In *POPL*, 1999.
- [22] David A. Naumann. From coupling relations to mated invariants for secure information flow and data abstraction. In *ESORICS*, 2006.
- [23] Andrei Sabelfeld and Andrew C. Myers. A model for delimited information release. In *ISSS*, 2004.
- [24] Tachio Terauchi and Alex Aiken. Secure information flow as a safety problem. In *Static Analysis Symposium (SAS)*, volume 3672 of *Lecture Notes in Computer Science*, pages 352–367, 2005.
- [25] Dennis Volpano, Geoffrey Smith, and Cynthia Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3):167–187, 1996.

Appendix

$$\begin{aligned}
s, r \llbracket \text{assert}(\phi) \rrbracket s', r' & \text{ iff } s, r \models \phi \text{ and } s' = s \text{ and } r' = r \\
s, r \llbracket RS_1 ; RS_2 \rrbracket s', r' & \text{ iff } \exists s'', r'' \bullet s, r \llbracket RS_1 \rrbracket s'', r'' \text{ and } s'', r'' \llbracket RS_2 \rrbracket s', r' \\
s, h \llbracket \text{if } B \text{ then } TS_1 \text{ else } TS_2 \rrbracket s', h' & \text{ iff } (\llbracket B \rrbracket^s = \text{True} \text{ and } s, h \llbracket TS_1 \rrbracket s', h') \\
& \text{ or } (\llbracket B \rrbracket^s = \text{False} \text{ and } s, h \llbracket TS_2 \rrbracket s', h') \\
s, h \llbracket x := A \rrbracket s', h' & \text{ iff } \exists v \bullet v = \llbracket A \rrbracket^s \text{ and } s' = [s \mid x \mapsto v] \text{ and } h' = h \\
s, r \llbracket .f := A \rrbracket s', r' & \text{ iff } \exists v \bullet v = \llbracket A \rrbracket_r^s \text{ and } s' = s \text{ and } r' = [r \mid f \mapsto v] \\
s, h \llbracket \text{new } x \text{ in } RS \text{ close} \rrbracket s', h' & \text{ iff } \exists l, r, r' \bullet (l \notin \text{dom}(h) \cup \text{ran}(h) \cup \text{ran}(s) \text{ and } r = \text{deflt} \\
& \text{ and } [s \mid x \mapsto l], r \llbracket RS \rrbracket s', r' \text{ and } h' = [h \mid l \mapsto r']) \\
s, h \llbracket \text{open } x \text{ in } RS \text{ close} \rrbracket s', h' & \text{ iff } \exists l, r, r' \bullet (l = s(x) \text{ and } r = h(l) \\
& \text{ and } s, r \llbracket RS \rrbracket s', r' \text{ and } h' = [h \mid l \mapsto r']) \\
s, h \llbracket \text{while } B \text{ do } TS \rrbracket s', h' & \text{ iff } \exists i \geq 0 \bullet s, h f_i s', h' \text{ where } f_i \text{ is inductively defined by:} \\
& s, h f_0 s', h' \text{ iff } \llbracket B \rrbracket^s = \text{False} \text{ and } s' = s \text{ and } h' = h \\
& s, h f_{i+1} s', h' \text{ iff } \exists s'', h'' \bullet (\llbracket B \rrbracket^s = \text{True} \text{ and} \\
& \quad s, h \llbracket TS \rrbracket s'', h'' \text{ and } s'', h'' f_i s', h')
\end{aligned}$$

Figure 6. Command semantics, selected clauses

```

1.   {true ⇒ x×}
    open x in
      {odd(.idx) ⇒ .src×, true ⇒ x×, odd(.idx) ⇒ .val×}
2.     y := .src;
      //Case of field access: replace y by .src to obtain pre
      {odd(.idx) ⇒ y×, true ⇒ x×, odd(.idx) ⇒ .val×, odd(.idx) ⇒ .src×}
3.     i := .idx;
      //Case of field access: replace i by .idx to obtain pre
      {odd(i) ⇒ y×, true ⇒ x×, odd(.idx) ⇒ .val×, odd(.idx) ⇒ .src×}
      // Conjoin object invariant to post
4.   close;
      {odd(i) ⇒ y×, true ⇒ x×}
5.   open y in
      {(odd(i) → odd(.idx)) ⇒ x×,
       odd(i) ∧ (odd(i) → odd(.idx)) ⇒ .val×,
       odd(.idx) ∧ (odd(i) → odd(.idx)) ⇒ .val×,
       odd(.idx) ∧ (odd(i) → odd(.idx)) ⇒ .src×}
6.     assert (odd(i) → odd(.idx));
      //Conjoin assertion to obtain pre
      {true ⇒ x×, odd(i) ⇒ .val×, odd(.idx) ⇒ .val×, odd(.idx) ⇒ .src×}
7.     q := .val;
      //Case of field access: replace q by .val to obtain pre
      {true ⇒ x×, odd(i) ⇒ q×, odd(.idx) ⇒ .val×, odd(.idx) ⇒ .src×}
      // Conjoin object invariant to simplified post
8.   close;
      {odd(i) ⇒ x×, true ⇒ x×, odd(i) ⇒ q×}
9.   open x in
      {odd(i) ∧ (.idx = i) ⇒ q×, odd(.idx) ∧ (.idx = i) ⇒ q×,
       odd(.idx) ∧ (.idx = i) ⇒ .src×}
10.    assert (.idx = i);
      //Conjoin assertion to obtain pre
      {odd(i) ⇒ q×, odd(.idx) ⇒ q×, odd(.idx) ⇒ .src×}
11.    .val := q;
      //Case of field update: replace .val by q to obtain pre
      {odd(i) ⇒ .val×, odd(.idx) ⇒ .val×, odd(.idx) ⇒ .src×}
12.    result := .val;
      //Case of field access: replace result by .val to obtain pre
      {odd(i) ⇒ result×, odd(.idx) ⇒ .val×, odd(.idx) ⇒ .src×}
      // Conjoin object invariant to post
13.  close;
      {odd(i) ⇒ result×}

```

Figure 7. Applying VCgen to Fig. 1(b).