

A Typed Assembly Language for Confidentiality

Dachuan Yu Nayeem Islam

DoCoMo Communications Laboratories USA

{yu,nayeem}@docomolabs-usa.com

Abstract

Language-based information-flow analysis is promising in protecting data confidentiality. Although much work has been carried out in this area, relatively little has been done for assembly code. Techniques at a source level do not generalize straightforwardly to assembly code, because assembly code does not readily present certain abstraction about the program structure that is crucial to information-flow analysis. Nonetheless, low-level information-flow analysis is desirable, because it yields a small trusted computing base. Furthermore, many (untrusted) applications are distributed in native code; their verification should not be overlooked.

We present a simple yet effective solution for this problem. Our observation is that the missing abstraction in assembly code can be restored using annotations. Following the philosophy of certifying compilation, these annotations are generated by a compiler, used for static validation, and erased before execution. In particular, we propose a type system for low-level information-flow analysis. Our system is compatible with Typed Assembly Language, and models key features including a call stack, memory tuples and first-class code pointers. A noninterference theorem articulates that well-typed programs respect confidentiality. We also present a security-type preserving translation that targets our system, together with its soundness theorem. This illustrates the application of certifying compilation for noninterference.

1. Introduction

With the growing reliance on networked information systems, the protection of confidential data becomes increasingly important. The problem is especially subtle for a computing system which both manipulates sensitive data and requires access to public information channels. Simple policies that restrict the access to either the sensitive data or the public channels (or a combination therefrom) often prove too restrictive. A more desirable policy is that no information about the sensitive data can be inferred from observing the public channels, even though a computing system is granted access to both. Such a regulation of the flow of information is often referred to as *information flow*, and the policy that sensitive data should not affect public data is often called *noninterference*.

Whereas it is relatively easy to detect and prevent naive violations that directly give out sensitive data, it is much more difficult to prevent applications from sending out information that is sophisticatedly encoded. Conventional security mechanisms such as access control, firewalls, encryption and anti-virus fall short on enforcing the noninterference policy [23]. On the one hand, noninterference posts seemingly conflicting requirements: it allows the access to sensitive information, but restricts the flow of it. On the other hand, the violation of noninterference cannot be observed from monitoring a single execution of the program [25], yet such execution monitoring is the basis of many conventional mechanisms.

In recent years, much effort has been put on enforcing noninterference using techniques based on programming language theory

and implementation. These techniques are promising, because they directly inspect or instrument the program code, and hence have the potential of learning all possible run-time behavior of the program. Unfortunately, the vast amount of language-based research on information flow [23] does not address well the problem for assembly code. The challenge there, as we will elaborate later, largely lies in working with the lack of high-level abstractions and managing the extreme flexibility offered by assembly code.

Nonetheless, it is desirable to enforce noninterference directly at a low-level. On the one hand, high-level programs must be translated into low-level code before executed on a real machine; compilation or optimization bugs may invalidate the security guarantee established for the source programs. On the other hand, some applications are distributed (*e.g.*, native code for mobile computation) or even directly written (*e.g.*, core libraries for embedded systems) in assembly code; enforcement at a low-level is a must for them.

This paper presents some important steps of a project tackling information flow at the assembly level. The contributions are:

- We propose a Typed Assembly Language for Confidentiality (TAL_C) for information-flow analysis and present its proof of noninterference. Our abstract machine is generic and close to real architectures. To reuse existing results on low-level verification, our system is designed to be compatible with Typed Assembly Language (TAL) [18]. It thus approaches a unified framework for conventional type safety and security.
- Our system models key features of an assembly language, including heap, call stack and register file, memory tuples (aliasing), and first-class code pointers (higher-order functions). Because assembly code is often arduous to work with, we present our formal result with a core language supporting the above features for ease of understanding, but also informally discuss extensions such as polymorphic and existential types.
- Although desirable to directly verify at an assembly level, it is more practical to develop programs in high-level languages. We present a translation from a security-typed imperative source language with first-order procedures to TAL_C . This illustrates the application of certifying compilation for noninterference. We also present a type-preservation theorem for our translation.

This paper does not address covert channels (*e.g.*, termination [28, 1] and timing [30, 2]) or abstract-violation attacks (*e.g.*, cache [3]). Section 2 provides background on language-based approaches for information flow, places our work in the context of existing researches, and points out the extra difficulties for noninterference at an assembly level. An informal overview of our approach is given in Section 3. Sections 4 and 5 present the TAL_C system and a certifying compilation scheme, focusing on core features that illustrate ideas pertinent to information flow. Section 6 helps better understand TAL_C in comparison with previous work on linear continuations. Orthogonal issues and practical extensions are discussed in Section 7. Section 8 concludes.

2. Background

2.1 Information Flow

The problem of information flow can be abstracted as a program that operates on data of different security levels, *e.g.*, *low* and *high*. Low (low security) data are public data that may be observed by all principals; high (high security) data are secret data whose access is restricted. An information-flow policy requires that no information about high inputs can be inferred from observing low outputs. In general, the security levels can be generalized to a lattice [29].

Such a policy concerns tracking the flow of information inside a target system. Whereas it is easy to detect explicit flows (*e.g.*, through an assignment from a secret h to a public l with $l=h$), it is much harder to detect various forms of implicit flow. For example, the statement $l=0; \text{if } h \text{ then } l=1$ involves an implicit flow from h to l . At run-time, if the *then* branch is not taken, a conventional security mechanism based on execution monitoring will not detect any violation. However, information about h can indeed be inferred from the result of l .

Instead of observing a single execution, language-based techniques derive assurance about a program's behavior by examining, and possibly instrumenting, the program code. In the above example, the information essentially leaks through the program counter (often referred to as *pc*)—the fact that a branch is taken reflects information about the guard of the conditional. In response, a security-type system typically tags the program counter with a security label. If the guard of a conditional concerns high data, then the branches are verified under a program counter with a high security label. Furthermore, assignments to low variables are prohibited under such a high program counter.

2.2 Security Types

Figure 1 shows a two-level security-type system for a simple imperative language with first-order procedures. It is extended from a demonstrative system of Sabelfeld and Myers' [23], and will serve as the source of our translation. A program comprises a list of procedure declarations F_i and a main command C . A procedure declaration documents the security level of the program counter with pc , indicating that the procedure body will only update variables with security levels no less than pc . A procedure also declares a list of arguments x_i under call-by-reference semantics. Commands C consist of assignments, sequential compositions, conditional statements, while-loops, and procedure calls. Variables V cover both global variables v and procedure arguments x . Expressions E are formed by constants (i), variables, and their additions.

Rules [E1–4] relate expressions to security types (levels). Any expression may have type *high* (it is secure to treat any data as sensitive). Constants and low variables may have type *low*. An addition expression may have type *low* if both sub-expressions may have type *low*.

Rules [C1–7] track the security level of the program counter (pc) when verifying the commands. Assignments to high variables are always valid (Rule [C1]). However, an assignment to a low variable is valid only if both the expression and the pc are low (Rule [C2]). For a conditional (Rule [C3]), the security level of the sub-commands must match the security level of the guard expression; together with Rule [C2], this guarantees that low variables are not modified within a branch under a high guard. After a conditional, it is useful to reset the pc to low, avoiding a form of label creep [10] where monotonically increasing security labels are too restrictive to be generally useful. Such a context reset is achieved with a subsumption rule (Rule [C4]); intuitively, if it is secure to execute a command in a sensitive context, then it is also secure in an insensitive one. A sequential composition is verified so that both sub-commands are valid under the given pc (Rule [C5]). The

<i>(Typ)</i>	$t, pc ::= \text{low} \mid \text{high}$
<i>(Var)</i>	$V ::= v \mid x$
<i>(Env)</i>	$\Phi ::= \circ \mid V : t, \Phi \mid f : \langle pc \rangle (t_1, \dots, t_n) \rightarrow \text{void}, \Phi$
<i>(Exp)</i>	$E ::= i \mid V \mid E_1 + E_2$
<i>(Comm)</i>	$C ::= V := E \mid C_1; C_2 \mid \text{if } E \text{ then } C_1 \text{ else } C_2$ $\quad \mid \text{while } E \text{ do } C \mid f(V_1, \dots, V_n)$
<i>(Fun)</i>	$F ::= f \langle pc \rangle (x_1 : t_1, \dots, x_n : t_n) \{C\}$
<i>(Prog)</i>	$P ::= \{F_1; \dots; F_n; C\}$
[E1–2]	$\Phi \vdash E : \text{high} \quad \Phi \vdash i : t$
[E3–4]	$\frac{\Phi(V) = \text{low}}{\Phi \vdash V : \text{low}} \quad \frac{\vdash E_1 : \text{low} \vdash E_2 : \text{low}}{\vdash E_1 + E_2 : \text{low}}$
[C1–2]	$\frac{\Phi(V) = \text{high}}{\Phi; [pc] \vdash V := E} \quad \frac{\Phi(V) = \text{low} \quad \Phi \vdash E : \text{low}}{\Phi; [\text{low}] \vdash V := E}$
[C3]	$\frac{\Phi \vdash E : pc \quad \Phi; [pc] \vdash C_1 \quad \Phi; [pc] \vdash C_2}{\Phi; [pc] \vdash \text{if } E \text{ then } C_1 \text{ else } C_2}$
[C4–5]	$\frac{\Phi; [\text{high}] \vdash C}{\Phi; [\text{low}] \vdash C} \quad \frac{\Phi; [pc] \vdash C_1 \quad \Phi; [pc] \vdash C_2}{\Phi; [pc] \vdash C_1; C_2}$
[C6]	$\frac{\Phi \vdash E : pc \quad \Phi; [pc] \vdash C}{\Phi; [pc] \vdash \text{while } E \text{ do } C}$
[C7]	$\frac{\Phi(f) = \langle pc \rangle (t_1, \dots, t_n) \rightarrow \text{void} \quad \Phi \vdash V_i : t_i \quad \forall i \in \{1 \dots n\}}{\Phi; [pc] \vdash f(V_1, \dots, V_n)}$
[F1]	$\frac{x_1 : t_1, \dots, x_n : t_n, \Phi; [pc] \vdash C}{\Phi \vdash f \langle pc \rangle (x_1 : t_1, \dots, x_n : t_n) \{C\}}$
[P1]	$\frac{F_i = f_i \langle pc_i \rangle (x_{1i} : t_{1i}, \dots, x_{n_i i} : t_{n_i i}) \{C_i\} \quad \Phi(F_i) = \langle pc \rangle (t_{1i}, \dots, t_{n_i i}) \rightarrow \text{void} \quad \Phi \vdash F_i \quad \Phi; [\text{low}] \vdash C \quad \forall i \in \{1 \dots n\}}{\Phi \vdash \{F_1; \dots; F_n; C\}}$

Figure 1. A simple security-type system

handling of a while-loop is similar to that of a conditional statement (Rule [C6]). A procedure call is valid if pc matches the expected security level, and the arguments have the expected types (Rule [C7]); note that only variables (v or x) may server as the arguments, which are handled by reference (also know as “in-out” arguments in the previous work of Volpano and Smith [29]).

Finally, a procedure declaration is valid if the body can be verified under the expected pc and arguments (Rule [F1]). A program is valid if all procedure declarations and the main command are valid (Rule [P1]).

2.3 Related Work

Although there has been much work applying language-based techniques to information flow [23], most of it focused on high-level languages. Many high-level abstractions have been formally studied, including functions [11], exceptions [22], objects [6], and concurrency [27, 1, 12], and practical implementation is within reach [19]. Nonetheless, enforcing information flow at only a high level puts the compiler into the trusted computing base (TCB) [24]. Furthermore, we should not overlook the verification of software distributed (or written) directly in low-level code. For example, a user may wish to verify that a GPS software, downloaded from an untrusted party in the form of native code to a mobile device, does not send out the location information through the network.

Barthe *et al.* [7] presented a security-type system for a byte-code language and a translation that preserves security types. Their stack-based language is much different from the RISC architecture that we model. More importantly, their verification circumvents a main difficulty—the lack of program structures at a low-level—by introducing a (currently trusted) component that computes the dependence regions and postdominators [5] for conditionals. The separate checking of these information is under investigation.

Avvenuti *et al.* [4] applied abstract interpretation to enforce information flow for a stack-based bytecode language. Besides the difference in the machine models, their work also relied on the (trusted) computation of control flow graphs and postdominators.

Zdancewic and Myers [33] used linear continuations to enforce noninterference at a low-level. Their language is based on variables and still much different from assembly language. In particular, linear continuations, although useful in enforcing a stack discipline that helps information-flow analysis, are absent from conventional assembly code. Hence further (trusted) compilation to native code is required. Nonetheless, we borrowed some ideas from Zdancewic and Myers, including the handling of memory aliasing and code pointers, although these features are modeled as different constructs in our system. A more thorough discussion of the connection between linear continuations and our solution is given in Section 6, after presenting our system.

Bonelli *et al.* [8] explored the realization of linear continuations in an assembly language SIFTAL. Two new instructions are introduced in correspondence with the operations on linear continuations as proposed by Zdancewic and Myers [33]. These two instructions enforce structured control flows that are missing from normal assembly code with the help of a continuation stack (this stack is different from the one for function calls). One instruction pushes a linear continuation onto the stack, the other pops a linear continuation off the stack and transfers the control to it. Such a mechanism maintains structured control flow in assembly code, thus helps information-flow analysis. Unfortunately, conventional assembly programming and machine models do not contain such a special continuation stack and the instructions manipulating it.

Recently, Medel *et al.* [16] improved SIFTAL to SIF, using a stack of labels to simplify the above approximation of linear continuations. Unlike SIFTAL, SIF resorts to static type annotations to enforce noninterference, and no longer requires a stack of linear continuations during execution. This is in spirit similar to our solution of TAL_C . However, SIF supports only a minimal set of language features (arithmetic, memory update, branching and direct jumps), and does not address how the type annotations can be produced. In contrast, our language TAL_C further supports code pointers and a call stack. We also present a type-preserving translation to TAL_C from a security-typed source language, where the support for procedure calls introduces extra subtleties for noninterference. We will discuss this in more detail in Section 6.

This paper targets RISC-style assembly code. We introduce a type system to verify the unstructured control flow, which in turn helps information-flow analysis. Type annotations are used to recover information about high-level program structures, and no trusted component is required for computing postdominators. This contrasts with the above work on bytecode languages. Furthermore, we do not rely on extra constructs such as linear continuations or a continuation stack. An erasure semantics trivially reduces programs in our language to normal assembly code.

We also provide a formal model of a certifying compiler for noninterference. Certifying compilation [14] has mostly been studied for conventional type safety (*e.g.*, Typed Assembly Language [18], Proof-Carrying Code [21, 20] and Efficient Code Certification [13]). Walker [31] applied certifying compilation to security policies. However, being based on security automata,

Walker’s system cannot enforce noninterference. Besides the work on security-type preserving compilation by Barthe *et al.* [7] as discussed above, Honda and Yoshida [12] also studied related issues for π -calculus with security types.

2.4 Assembly Code

Whereas enforcing information flow for assembly code is important, it poses many new challenges.

First, high-level languages make use of a virtually infinite number of variables, each of which can be assigned a fixed security label. In assembly code, the use of memory cells is similar. However, a finite number of registers are reused for different source-level variables. As a result, one cannot assign a fixed security label to a register.

Second, the control flow of an assembly program is not as structured. The body of a conditional is often not obvious, and generally undecidable, from the program code. Hence the idea of using a security context to prevent implicit flow through conditionals cannot be easily carried out.

Third, assembly languages are very expressive. Aliasing between memory cells are difficult to reason about [26]. The support for first-class code pointers (the reflection of higher-order functions at the assembly level) is very subtle. A code pointer may direct a program to different execution paths, even though no branching instruction is present.

Fourth, since it is not practical to always directly program in an assembly language, a low-level type system must be designed so that the type annotations can be generated automatically, *e.g.*, through certifying compilation. The type system must be at least as expressive as a high-level type system, so that any well-typed source program can be translated into well-typed assembly code.

Finally, it is desirable to achieve an erasure semantics where type annotations have no effect at runtime. A security mechanism can not be generally applied in practice if it incurs too much overhead. Similarly, it is also undesirable to change the programming model for accommodating the verification needs. Such a model change indicates either a trusted compilation process or a different target machine.

3. Our Approach

3.1 Explicit Assignment

An obvious kind of information flow is through assignment. As discussed in Section 2.2, variables in a high-level language can be “tagged” with security labels; the security-type system prevents label mismatch for assignments. At an assembly level, memory cells can be tagged similarly. When storing into a memory cell, a typing rule ensures that the security label of the source matches that of the target.

Registers need to be regulated differently, because they can be reused for different variables with different security labels (registers cannot be aliased, which makes it safe to update their types). Since variable and liveness information is not available at an assembly level, one can not easily base the enforcement upon that.

In fact, a similar problem arises even for normal type safety. A register in TAL can have different types at different program points. These types are essentially inferred from the computation itself. For instance, in an addition instruction $\text{add } r_d, r_s, r_t$, the register r_d is given the type int , because only int can be valid here. Similarly, when loading from a memory cell, the target register is given the type of the source memory cell. Adapting such inference for security labels is straightforward. In the addition $\text{add } r_d, r_s, r_t$, the label of r_d is obtained by joining the labels of r_s and r_t , because the result in r_d reflects information from both r_s and r_t . Moving and memory reading instructions are handled similarly.

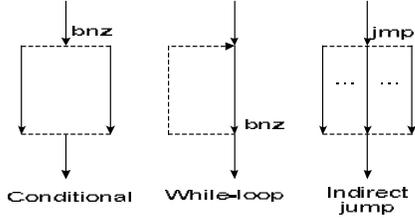


Figure 2. Flow through program structure

3.2 Program Structure

A conditional statement in a high-level program can be verified so that both sub-commands respect the security level of the guard expression. Such verification becomes difficult in assembly code, where the “flattened” control flow provides little help in identifying the program structure. A conditional is typically translated into a branching instruction (`bnz r, l`) and some code blocks, where the postdominator of the two branches are no longer apparent.

We use annotations to restore the program structure by pointing out the postdominators whenever they are needed. Note that high-level programs provide sufficient information for deciding the postdominators, and these postdominators can always be statically determined. For instance, the end of a conditional command is the postdominator of the two branches. Hence a compiler can generate the annotations automatically based on a securely typed source program. In our system, our postdominator annotation is essentially a static code label paired with a security label.

Since branching instructions (`bnz r, l`) are the only instructions that could directly result in different execution paths, it would appear that one should enhance branching instructions with postdominators. The typing rule then checks both branches under a proper security context that takes into account the guard expression. Such a security context terminates when the postdominator is reached.

Although plausible, this approach is awkward. Figure 2 demonstrates three scenarios. Besides the conditional scenario, branching instructions are also used to implement while-loops, where the postdominator is exactly the beginning of one of the branches. In this case, only the other branch should be checked under a new security context. If we directly annotate the branching instruction, the corresponding typing rule would be “overloaded.” More importantly, an assembly program may contain “implicit branches” where no branching instruction is present. The third scenario illustrates that an indirect jump may lead the program to different paths based on the value of its operand register. A concrete example will appear in Section 3.5.

Inspiration of a better solution lies in the simple system of Figure 1. Note that the subsumption rule [C4] is not tied to any particular commands. It essentially marks a region of computation where the security level is raised from low to high. The end of the region is exactly a postdominator. Following this, our approach is to mimic the high-level subsumption rule with two low-level *raising* and *lowering* operations that explicitly manipulate the security context and mark the beginning and the end of the secured region.

3.3 Memory Aliasing

Aliasing of memory cells present another channel for information transfer. In Figure 3, a low pointer `p_l` and a high pointer `p_h` are aliases of the same cell. This is useful if a high principal wishes to observe a low computation. The code in this figure may change the aliasing relation based on some high variable `h` by letting `p_h` point to another cell. Further modification through `p_h` may or may not change the value in the original cell. As a result, observing through the low pointer `p_l` reveals information about the high variable `h`.

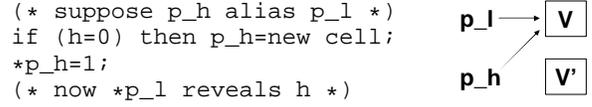


Figure 3. Flow through aliasing

```
fun f0 () = (l:=0; ())
fun f1 () = (l:=1; ())
let f = (if h then f1 else f0) in f()
```

Figure 4. Flow through code pointer

```
fun f0 () = (h' := 0; ())
fun f1 () = (h' := 1; ()) ...
if h then f := f1 else f := f0;
l := 1; !f(); l := l * 2; ...
```

Figure 5. Context coercion without branching

The problem lies in the assignment through the high pointer `p_h`, because it reveals information about the aliasing relation. The solution, following previous work [6, 33], is to tag pointers with two security labels. One is for the pointer itself, and the other is for the data being referenced. Assignments to low data through high pointers are not allowed. This is a conservative approach—all pointers are considered as potential aliases.

3.4 Code Pointers

Code pointers further complicate information flow. Figure 4 shows a piece of functional code where `f` represents different functions based on a high variable `h`. In its reflection at an assembly level, different code blocks will be executed based on the value of `h`. Naturally, `f` contains sensitive information and should be labeled high. However, the actual functions `f0` and `f1` can only be executed under a low context, because they modify a low variable `l`. In this case, the invocation to `f` should be prohibited.

In our system, similar to data pointers, code pointers are also given two security labels. The typing rules ensure that no low function is called through a high code pointer.

3.5 Security Context Coercion

Finally, Figure 5 shows a piece of code where a mutable code pointer complicates the flow analysis. Functions `f0` and `f1` only modify high data. A reference cell `f` is assigned different code pointers within a high conditional. Later, `f` is dereferenced and invoked in a low context.

This code is safe with respect to information flow. At a high level, a subsumption rule like Rule [C4] in Figure 1 allows calling the high function `!f()` in a low context. However, in its assembly counterparts, both the calling to `f` and the returning from `f` are implemented as indirect jumps. The calling sequence transfers the control from a low context to a high context, whereas the returning sequence does the opposite. Since the function invocation is no longer atomic at an assembly level, one cannot directly devise a subsumption rule. Furthermore, there is no explicit branching instruction present when `f` is dereferenced and invoked (the third scenario of Figure 2).

In our system, the raising and lowering operations explicitly mark the boundary of the subsumption rule. During certifying compilation, the source-level typing and program structure provide sufficient information for generating the target-level annotations. When a subsumption rule is applied in the source code, the corresponding target code is generated within a pair of raising and lowering operations.

(contexts)	$\kappa ::= \bullet \mid \theta \triangleright w$
(pre-type)	$\tau ::= \text{int} \mid \langle \sigma_1, \dots, \sigma_n \rangle \mid \forall[\Delta].\langle \kappa \rangle \Gamma$
(types)	$\sigma ::= \tau_\theta \mid ns$
(stack ty)	$\Sigma ::= \rho \mid nil \mid \sigma :: \Sigma$
(var env)	$\Delta ::= \circ \mid \rho \Delta \mid \alpha \Delta$
(type arg)	$\psi ::= \Sigma \mid w$
(heap ty)	$\Psi ::= \{l_1 : \sigma_1, \dots, l_n : \sigma_n\}$
(reg file ty)	$\Gamma ::= \{r_1 : \sigma_1, \dots, r_n : \sigma_n, \text{sp} : \Sigma\}$
(registers)	$r ::= r_1 \mid r_2 \mid \dots$
(word val)	$w ::= \alpha \mid l \mid i \mid ns \mid w[\psi]$
(small val)	$v ::= r \mid w \mid v[\psi]$
(heap val)	$h ::= \langle w_1, \dots, w_n \rangle \mid \text{code}[\Delta]\langle \kappa \rangle \Gamma.I$
(heaps)	$H ::= \{l_1 \mapsto h_1, \dots, l_n \mapsto h_n\}$
(reg files)	$R ::= \{r_1 \mapsto w_1, \dots, r_n \mapsto w_n, \text{sp} \mapsto S\}$
(stacks)	$S ::= nil \mid w :: S$
(instr)	$\iota ::= \text{add } r_d, r_s, v \mid \text{ld } r_d, r_s(i) \mid \text{st } r_d(i), r_s$ $\mid \text{mov } r_d, v \mid \text{bnz } r, v \mid \text{salloc } i \mid \text{sfree } i$ $\mid \text{sld } r_d, \text{sp}(i) \mid \text{sst } \text{sp}(i), r_s \mid \text{raise } \kappa$
(instr seq)	$I ::= \iota; I \mid \text{lower } w \mid \text{jmp } v \mid \text{halt } [\sigma]$
(prog)	$P ::= (H, R, I)_\kappa$

Figure 6. Syntax of TAL_C

4. TAL_C

4.1 Abstract Machine

Our language TAL_C is designed to resemble TAL [18] for ease of understanding. We introduce some new constructs for confidentiality, and accommodate a stack following STAL [17] for supporting procedure calls. In the interest of simplicity, we maintain just enough features for demonstrating the certifying compilation of security types from our source language in Figure 1, removing from TAL and STAL features that are orthogonal to it.

We assume that security labels form a lattice \mathcal{L} . We use \perp and \top as the bottom and top of the lattice, \cup and \cap as the lattice join and meet operations, and \subseteq as the lattice ordering. The syntactic constructs of TAL_C can be understood in three steps as follows.

Type constructs The top portion of Figure 6 presents the type constructs. Security contexts κ follow the idea of Section 3.2. An empty security context (\bullet) represents an program counter with the lowest security label. A concrete context ($\theta \triangleright w$) is made up of a security label θ (the current security level) and a postdominator w . The postdominator w has the syntax of a word value, but its use is restricted by the semantics to be eventually an instantiated code label, *i.e.*, the ending point of the current security level. The postdominator w could also be a variable α ; this is useful for compiling procedures, which can be called in different contexts with different postdominators.

Pre-types (τ) reflect the normal types as seen in TAL, including integer types, tuple types, and code types. In comparison with TAL, our code type requires an extra security context (κ) as part of the interface. A type (σ) is either a pre-type tagged with a security label or a nonsense type (ns) for uninitialized stack slots. A stack type (Σ) is either a variable (ρ), or a (possibly empty) sequence of types. The variable context (Δ) is used for typing polymorphic code; it documents stack type variables (ρ) and postdominator variables (α). Stack types and postdominators are also generally referred to as type arguments ψ . Finally, heap types (Ψ) or register file types (Γ) are mappings from heap labels or registers to types; the sp in the register file represents the stack.

Judgment	Meaning
$\Delta \vdash \kappa$	κ is a valid context
$\Delta \vdash \tau$	τ is a valid pre-type
$\Delta \vdash \sigma$	σ is a valid type
$\Delta \vdash \Sigma$	Σ is a valid stack type
$\vdash \Psi$	Ψ is a valid heap type
$\Delta \vdash \Gamma$	Γ is a valid register file type
$\Delta \vdash \Gamma_1 \subseteq \Gamma_2$	Register file type Γ_1 weakens Γ_2
$\vdash H : \Psi$	Heap H has type Ψ
$\Psi \vdash S : \Sigma$	Stack S has type Σ
$\Psi \vdash R : \Gamma$	Register file R has type Γ
$\Psi \vdash h : \sigma$	Heap value h has type σ
$\Psi; \Delta \vdash w : \sigma$	Word value w has type σ
$\Psi; \Delta; \Gamma \vdash v : \sigma$	Small value v has type σ
$\Psi; \Delta; \Gamma; \kappa \vdash I$	I is a valid sequence of instructions
$\Psi; \Gamma \vdash P$	P is a valid program

Figure 8. TAL_C typing judgments

Note that the type constructs provide two layers of security labels for a data pointer (*e.g.*, $\langle \text{int}_{\theta_2} \rangle_{\theta_1}$; see Section 3.3) or a code pointer (*e.g.*, $\langle \forall[\circ].\langle \theta_2 \triangleright l \rangle \Gamma \rangle_{\theta_1}$; see Section 3.4)—one (θ_1) for the pointer itself, the other (θ_2) for the data or code being referenced.

Value constructs The middle portion of Figure 6 shows the value constructs. A word value w is either a variable, a heap label l , an immediate integer i , a nonsense value for an uninitialized stack slot, or another word value instantiated with a type argument. Small values v serve as the operands of some instructions; they are either registers r , word values w , or instantiated small values. Heap values h are either tuples or typed code sequences; they are the building blocks of the heap H . Note that a value does not carry a security label. This is consistent with the philosophy that a value is never intrinsically sensitive—it is sensitive only if it comes from a sensitive location [29], which is documented in the corresponding types (Ψ and Γ). Finally, a register file R stores the contents of all registers and the stack, where the stack is a (possibly empty) sequence of word values.

Code constructs Code constructs are given in the bottom portion of Figure 6. We retain a minimal set of instructions from TAL and STAL, and introduce two new instructions ($\text{raise } \kappa$ and $\text{lower } l$) for manipulating the security context as discussed in Section 3. A program is the usual triple tagged with a security context. The security context facilitates the formal soundness proof, but does not affect the computation.

In the operational semantics (Figure 7), there are only two cases that modify the security context: $\text{raise } \kappa'$ updates the security context to κ' , and $\text{lower } w$ picks up a new security context from the interface of the target code w . In all other cases, the security context remains the same, and the semantics is standard. It is easy to see that this operational semantics mimics the behavior of a real machine, and does not prohibit bad flows. One can obtain a conventional machine by removing the security contexts and $\text{raise } \kappa$ instructions, and replacing $\text{lower } w$ with $\text{jmp } w$.

4.2 Typing Rules

The static semantics consists of judgment forms summarized in Figure 8. A security context appears in the judgment of a valid instruction sequence. Heap and register file types are made explicit in the judgment of a valid program for facilitating the noninterference theorem. All other judgment forms closely resemble those of TAL and STAL.

$(H, R, I)_\kappa \mapsto P$ where		
if $I =$	then $P =$	
add $r_d, r_s, v; I'$	$(H, R\{r_d \mapsto (i + i')\}, I')_\kappa$	where $R(r_s) = i$ and $\hat{R}(v) = i'$
ld $r_d, r_s(i); I'$	$(H, R\{r_d \mapsto w_i\}, I')_\kappa$	where $R(r_s) = l$ and $H(l) = \langle w_0, \dots, w_{n-1} \rangle$ with $0 \leq i < n$
mov $r_d, v; I'$	$(H, R\{r_d \mapsto \hat{R}(v)\}, I')_\kappa$	
bnz $r, v; I'$	$(H, R, I')_\kappa$ when $\hat{R}(r) = 0$	
bnz $r, v; I'$	$(H, R, I'[\vec{\psi}/\Delta])_\kappa$ when $R(r) = i$	where $i \neq 0$ and $\hat{R}(v) = l[\vec{\psi}]$ and $H(l) = \text{code}[\Delta]\langle \kappa' \rangle \Gamma. I''$
st $r_d(i), r_s; I'$	$(H\{l \mapsto \langle w_0, \dots, w_{i-1}, R(r_s), w_{i+1}, \dots, w_{n-1} \rangle\}, R, I')_\kappa$	where $R(r_d) = l$ and $H(l) = \langle w_0, \dots, w_{n-1} \rangle$ with $0 \leq i < n$
salloc $i; I'$	$(H, R\{\text{sp} \mapsto \underbrace{ns :: \dots :: ns}_{i} :: R(\text{sp})\}, I')_\kappa$	
sfree $i; I'$	$(H, R\{\text{sp} \mapsto S\}, I')_\kappa$	where $R(\text{sp}) = w_1 :: \dots :: w_i :: S$
sld $r_d, \text{sp}(i); I'$	$(H, R\{r_d \mapsto w_i\}, I')_\kappa$	where $R(\text{sp}) = w_0 :: \dots :: w_i :: S$ and $i \geq 0$
sst $\text{sp}(i), r_s; I'$	$(H, R\{\text{sp} \mapsto w_0 :: \dots :: w_{i-1} :: R(r_s) :: S\}, I')_\kappa$	where $R(\text{sp}) = w_0 :: \dots :: w_i :: S$ and $i \geq 0$
raise $\kappa; I'$	$(H, R, I')_{\kappa'}$	
lower w	$(H, R, I'[\vec{\psi}/\Delta])_{\kappa'}$	where $w = l[\vec{\psi}]$ and $H(l) = \text{code}[\Delta]\langle \kappa' \rangle \Gamma. I'$
jmp v	$(H, R, I'[\vec{\psi}/\Delta])_\kappa$	where $\hat{R}(v) = l[\vec{\psi}]$ and $H(l) = \text{code}[\Delta]\langle \kappa' \rangle \Gamma. I'$

$$\text{where } \hat{R}(v) = \begin{cases} R(r) & \text{when } v = r \\ w & \text{when } v = w \\ \hat{R}(v')[\psi] & \text{when } v = v'[\psi] \end{cases}$$

Figure 7. Operational semantics of TAL_C

The typing rules are given in Figures 9 and 10. A type construct is valid (top six judgment forms in Figure 8) if all free type variables are documented in the type environment. Heap values and integers may have any security label. The types of heap labels and registers are as described in the heap type and the register file type respectively. All other rules for non-instructions are straightforward extensions of those in TAL and STAL.

We use $SL(\kappa)$ to refer to the security label component of κ . $SL(\bullet)$ is defined to be \perp . The typing rules for **add**, **ld** and **mov** instructions infer the security labels for the destination registers; they take into account the security labels of the source and target operands and the current security context.

The rule for **bnz** first checks that the guard register r is an integer and the target value v is a code label. It then checks that the current security context is high enough to cover the security levels of the guard (preventing flows through program structures; Section 3.2) and the target code (preventing flows through code pointers; Section 3.4). Lastly, the checks on the register file and the remainder instruction sequence make sure that both branches are secure to execute.

The rule for **st** concerns four security labels. The label of the target cell must be higher than or equal to those of the context (Section 3.2), the containing tuple (Section 3.3), and the source value (Section 3.1).

The rules for the stack instructions follow similar ideas. In essence, the stack can be viewed as an infinite number of registers. Instruction **salloc** or **sfree** add new slots to or remove existing slots from the slot, so the rules check the remainder instruction sequence under an updated stack type. The rule for instruction **sld** or **sst** can be understood following that of the **mov** instruction.

The rule for **raise** checks that the new security context is higher than the current one. Moreover, it looks at the postdominator w' of the new context, and makes sure that the security context at w' matches the current one. The remainder instruction sequence is checked under the new context.

Since the rule for **raise** already checked the validity of the ending label of a secured region, the task for ending the region is relatively simple. The rule for **lower** checks that its operand label matches that dictated by the security context. This guarantees that

a secured region be enclosed within a **raise-lower** pair. The rule also makes sure that the code at w is safe to execute, which involves checking the security labels (Section 3.4) and the register file types.

The rule for **jmp** checks that the target code is safe to execute. Similar checks also appeared in the rule for **bnz**. In these two rules, the security context of the target code is always the same as the current one. This is because context changes are separated from conventional instructions in our system. For example, one may enclose high target code within **raise** and **lower** before calling it in a low context.

Finally, halting is valid only if the security context is empty, and the value in r_1 has the expected type σ .

Interested readers are referred to Appendix B for a simple example that demonstrates the use of security labels and contexts.

4.3 Soundness

TAL_C enjoys conventional type safety (memory and control-flow safety), which can be established following the preservation and progress lemmas. The proofs of these lemmas are similar to those of TAL and STAL and omitted.

Lemma 1 (Preservation) If $\Psi; \Gamma \vdash P$ and $P \mapsto P'$, then there exists Γ' such that $\Psi; \Gamma' \vdash P'$.

Lemma 2 (Progress) If $\Psi; \Gamma \vdash P$ then either:

- (1) there exists P' such that $P \mapsto P'$, or
- (2) P is of the form $(H, R\{r_1 \mapsto w\}, \text{halt } [\sigma])_\bullet$, where $\vdash H : \Psi$ and $\Psi; \circ \vdash w : \sigma$.

Before presenting the noninterference theorem, we define the equivalence of two programs with respect to a security level θ .

Definition 1 (Heap Equivalence) $\Psi \vdash H_1 \approx_\theta H_2 \iff$ for every $l \in \text{dom}(\Psi)$, if $\Psi(l) = \tau_{\theta'}$ and $\theta' \subseteq \theta$ then $H_1(l) = H_2(l)$.

Definition 2 (Stack Equivalence) $\Sigma \vdash S_1 \approx_\theta S_2 \iff$ for every slot $i \in \text{dom}(\Sigma)$, if $\Sigma(i) = \tau_{\theta'}$ and $\theta' \subseteq \theta$ then $S_1(i) = S_2(i)$.

Definition 3 (Register File Equivalence) $\Gamma \vdash R_1 \approx_\theta R_2 \iff$ (1) $\Gamma(\text{sp}) \vdash R_1(\text{sp}) \approx_\theta R_2(\text{sp})$, and (2) for every $r \in \text{dom}(\Gamma)$, if $\Gamma(r) = \tau_{\theta'}$ and $\theta' \subseteq \theta$, then $R_1(r) = R_2(r)$.

$$\begin{array}{c}
\Delta \vdash \bullet \quad \frac{\text{freevar}(w) \subseteq \Delta}{\Delta \vdash \theta \triangleright w} \quad \Delta \vdash \text{int} \\
\\
\frac{\Delta \vdash \sigma_i}{\Delta \vdash \langle \sigma_1, \dots, \sigma_n \rangle} \quad \frac{\Delta \vdash \kappa \quad \Delta \vdash \Gamma}{\Delta \vdash \forall[\Delta].\langle \kappa \rangle \Gamma} \\
\frac{\Delta \vdash \tau}{\Delta \vdash \tau_\theta} \quad \Delta \vdash ns \\
\\
\frac{\rho \in \Delta}{\Delta \vdash \rho} \quad \Delta \vdash nil \quad \frac{\Delta \vdash \sigma \quad \Delta \vdash \Sigma}{\Delta \vdash \sigma :: \Sigma} \\
\frac{\circ \vdash \sigma_i}{\vdash \{l_1 : \sigma_1 \dots l_n : \sigma_n\}} \quad \frac{\Delta \vdash \sigma_i \quad \Delta \vdash \Sigma}{\Delta \vdash \{r_1 : \sigma_1 \dots r_n : \sigma_n, \text{sp} : \Sigma\}} \\
\\
\frac{\Delta \vdash \sigma_i \quad \Delta \vdash \Sigma \quad m \leq n}{\Delta \vdash \{r_1 : \sigma_1 \dots r_m : \sigma_m, \text{sp} : \Sigma\} \subseteq \{r_1 : \sigma_1 \dots r_n : \sigma_n, \text{sp} : \Sigma\}} \\
\\
\frac{\vdash \Psi \quad \Psi = \{l_1 : \sigma_1, \dots, l_n : \sigma_n\} \quad \Psi \vdash h_i : \sigma_i}{\vdash \{l_1 \mapsto h_1, \dots, l_n \mapsto h_n\} : \Psi} \\
\\
\frac{\Gamma = \{r_1 : \sigma_1, \dots, r_n : \sigma_n, \text{sp} : \Sigma\} \quad \Psi \vdash w_i : \sigma_i \quad \Psi \vdash S : \Sigma}{\Psi \vdash \{r_1 \mapsto w_1, \dots, r_n \mapsto w_n, \text{sp} \mapsto S\} : \Gamma} \\
\\
\frac{\Psi \vdash w_i : \sigma_i}{\Psi \vdash \langle w_1, \dots, w_n \rangle : \langle \sigma_1, \dots, \sigma_n \rangle_\theta} \\
\\
\frac{\Delta \vdash \kappa \quad \Delta \vdash \Gamma \quad \Psi; \Delta; \Gamma; \kappa \vdash I}{\Psi \vdash \text{code}[\Delta]\langle \kappa \rangle \Gamma.I : (\forall[\Delta].\langle \kappa \rangle \Gamma)_\theta} \\
\\
\Psi; \Delta \vdash i : \text{int}_\theta \quad \Psi; \Delta \vdash ns : ns \quad \frac{\Psi(l) = \sigma}{\Psi; \Delta \vdash l : \sigma} \\
\\
\frac{\Delta \vdash \Sigma \quad \Psi; \Delta \vdash w : \forall[\rho\Delta'].\langle \kappa \rangle \Gamma}{\Psi; \Delta \vdash w[\Sigma] : \forall[\Delta'].\langle \kappa[\Sigma/\rho] \rangle \Gamma[\Sigma/\rho]} \\
\\
\frac{\Delta \vdash \Sigma \quad \Psi; \Delta \vdash w : \forall[\alpha\Delta'].\langle \kappa \rangle \Gamma}{\Psi; \Delta \vdash w[w'] : \forall[\Delta'].\langle \kappa[w'/\alpha] \rangle \Gamma[w'/\alpha]} \\
\\
\frac{\Gamma(r) = \sigma}{\Psi; \Delta; \Gamma \vdash r : \sigma} \quad \frac{\Psi; \Delta \vdash w : \sigma}{\Psi; \Delta; \Gamma \vdash w : \sigma} \\
\\
\frac{\Delta \vdash \Sigma \quad \Psi; \Delta; \Gamma \vdash v : \forall[\rho\Delta'].\langle \kappa \rangle \Gamma'}{\Psi; \Delta; \Gamma \vdash v[\Sigma] : \forall[\Delta'].\langle \kappa[\Sigma/\rho] \rangle \Gamma'[\Sigma/\rho]} \\
\\
\frac{\Delta \vdash \Sigma \quad \Psi; \Delta; \Gamma \vdash v : \forall[\alpha\Delta'].\langle \kappa \rangle \Gamma'}{\Psi; \Delta; \Gamma \vdash v[w] : \forall[\Delta'].\langle \kappa[w/\alpha] \rangle \Gamma'[w/\alpha]} \\
\\
\frac{\circ \vdash \kappa \quad \vdash H : \Psi \quad \Psi \vdash R : \Gamma \quad \Psi; \circ; \Gamma; \kappa \vdash I}{\Psi; \Gamma \vdash (H, R, I)_\kappa}
\end{array}$$

Figure 9. TAL_C typing rules (non-instructions)

Definition 4 (Program Equivalence) $\Psi; \Gamma \vdash P_1 \approx_\theta P_2 \iff P_1 = (H_1, R_1, I_1)_{\kappa_1}, P_2 = (H_2, R_2, I_2)_{\kappa_2}, \Psi \vdash H_1 \approx_\theta H_2, \Gamma \vdash R_1 \approx_\theta R_2,$ and either:

- (1) $\kappa_1 = \kappa_2, SL(\kappa_1) \subseteq \theta,$ and $I_1 = I_2,$ or
- (2) $SL(\kappa_1) \not\subseteq \theta, SL(\kappa_2) \not\subseteq \theta.$

It is easy to see that the above relations are all reflexive, symmetrical, and transitive. Our noninterference theorem relates the executions of two equivalent programs that both start in a low security context (relative to the security level of concern). If both executions terminate, then the result programs must also be equivalent.

The idea of the proof is intuitive. Given a security level of concern, the executions can be phased into “low steps” and “high steps.” It is easy to relate the two executions under a low step, because they involve the same instructions. Under a high step, the two executions are no longer in lock step. Recall that `raise` and

$$\begin{array}{c}
\frac{SL(\kappa) = \theta \quad \Gamma(r_s) = \text{int}_{\theta_1} \quad \Psi; \Delta; \Gamma \vdash v : \text{int}_{\theta_2} \quad \Psi; \Delta; \Gamma\{r_d : \text{int}_{\theta \cup \theta_1 \cup \theta_2}\}; \kappa \vdash I}{\Psi; \Delta; \Gamma; \kappa \vdash \text{add } r_d, r_s, v; I} \\
\\
\frac{SL(\kappa) = \theta \quad \Gamma(r_s) = \langle \sigma_1, \dots, \sigma_n \rangle_{\theta_1} \quad \sigma_i = \tau_{\theta_2} \quad \Psi; \Delta; \Gamma\{r_d : \tau_{\theta \cup \theta_1 \cup \theta_2}\}; \kappa \vdash I}{\Psi; \Delta; \Gamma; \kappa \vdash \text{ld } r_d, r_s(i); I} \\
\\
\frac{SL(\kappa) = \theta \quad \Psi; \Gamma \vdash v : \tau_{\theta'} \quad \Psi; \Delta; \Gamma\{r_d : \tau_{\theta \cup \theta'}\}; \kappa \vdash I}{\Psi; \Delta; \Gamma; \kappa \vdash \text{mov } r_d, v; I} \\
\\
\frac{SL(\kappa) = \theta \quad \Gamma(r) = \text{int}_{\theta_1} \quad \Psi; \Delta; \Gamma \vdash v : (\forall[\circ].\langle \kappa \rangle \Gamma')_{\theta_2} \quad \theta_1 \cup \theta_2 \subseteq \theta \quad \Delta \vdash \Gamma' \subseteq \Gamma \quad \Psi; \Delta; \Gamma; \kappa \vdash I}{\Psi; \Delta; \Gamma; \kappa \vdash \text{bnz } r, v; I} \\
\\
\frac{SL(\kappa) = \theta \quad \Gamma(r_d) = \langle \sigma_1, \dots, \sigma_n \rangle_{\theta_1} \quad \sigma_i = \tau_{\theta'} \quad \Gamma(r_s) = \tau_{\theta_2} \quad \theta \cup \theta_1 \cup \theta_2 \subseteq \theta' \quad \Psi; \Delta; \Gamma; \kappa \vdash I}{\Psi; \Delta; \Gamma; \kappa \vdash \text{st } r_d(i), r_s; I} \\
\\
\frac{\Gamma(\text{sp}) = \Sigma \quad \Psi; \Delta; \Gamma\{\text{sp} : \overbrace{ns :: \dots :: ns}^i :: \Sigma\}; \kappa \vdash I}{\Psi; \Delta; \Gamma; \kappa \vdash \text{salloc } i; I} \\
\\
\frac{\Gamma(\text{sp}) = \sigma_1 :: \dots :: \sigma_i :: \Sigma \quad \Psi; \Delta; \Gamma\{\text{sp} : \Sigma\}; \kappa \vdash I}{\Psi; \Delta; \Gamma; \kappa \vdash \text{sfree } i; I} \\
\\
\frac{SL(\kappa) = \theta \quad \Gamma(\text{sp}) = \sigma_0 :: \dots :: \sigma_i :: \Sigma \quad \sigma_i = \tau_{\theta'} \quad \Psi; \Delta; \Gamma\{r_d : \tau_{\theta \cup \theta'}\}; \kappa \vdash I}{\Psi; \Delta; \Gamma; \kappa \vdash \text{sld } r_d, \text{sp}(i); I} \\
\\
\frac{SL(\kappa) = \theta \quad \Gamma(\text{sp}) = \sigma_0 :: \dots :: \sigma_i :: \Sigma \quad \Gamma(r_s) = \tau_{\theta'} \quad \Psi; \Delta; \Gamma\{\text{sp} : \sigma_0 :: \dots :: \sigma_{i-1} :: \tau_{\theta \cup \theta'} :: \Sigma\}; \kappa \vdash I}{\Psi; \Delta; \Gamma; \kappa \vdash \text{sst } \text{sp}(i), r_s; I} \\
\\
\frac{\kappa = \theta \triangleright w \quad \kappa' = \theta' \triangleright w' \quad \theta \subseteq \theta' \quad \Psi; \Delta \vdash w' : (\forall[\circ].\langle \kappa \rangle \Gamma')_{\theta_1} \quad \Psi; \Delta; \Gamma; \kappa' \vdash I}{\Psi; \Delta; \Gamma; \kappa \vdash \text{raise } \kappa'; I} \\
\\
\frac{\kappa = \theta \triangleright w \quad \Psi; \Delta \vdash w : (\forall[\circ].\langle \kappa' \rangle \Gamma')_{\theta_1} \quad \theta_1 \subseteq SL(\kappa') \quad \Delta \vdash \Gamma' \subseteq \Gamma}{\Psi; \Delta; \Gamma; \kappa \vdash \text{lower } w} \\
\\
\frac{SL(\kappa) = \theta \quad \Psi; \Delta; \Gamma \vdash v : (\forall[\circ].\langle \kappa \rangle \Gamma')_{\theta_1} \quad \theta_1 \subseteq \theta \quad \Delta \vdash \Gamma' \subseteq \Gamma}{\Psi; \Delta; \Gamma; \kappa \vdash \text{jmp } v} \\
\\
\frac{\kappa = \bullet \quad \Delta \vdash \sigma \quad \Gamma(r_1) = \sigma}{\Psi; \Delta; \Gamma; \kappa \vdash \text{halt } [\sigma]}
\end{array}$$

Figure 10. Typing rules of TAL_C instructions

`lower` mark the beginning and the end of a secured region. We relate the program states before the `raise` and after the `lower`, circumventing directly relating two executions under high steps.

We give the formal details in three lemmas and a noninterference theorem. Lemma 3 indicates that a security context in a high step can be changed only with `raise` or `lower`. Lemma 4 says that a terminating program must reduce to a step that discharges the current security context with a `lower`. Lemma 5 articulates the lock step relation between two equivalent programs in a low step. Theorem 1 of noninterference then follows. In the following, \mapsto^* represents the reflexive and transitive closure of \mapsto . $\Sigma \succeq_\theta \Sigma'$ means that $\Sigma(i) = \Sigma'(i)$ for every i such that $\Sigma'(i) = \tau_{\theta'}$ and $\theta' \subseteq \theta$. $\Gamma \succeq_\theta \Gamma'$ means that $\Gamma(\text{sp}) \succeq_\theta \Gamma'(\text{sp})$ and $\Gamma(r) = \Gamma'(r)$ for every r such that $\Gamma'(r) = \tau_{\theta'}$ and $\theta' \subseteq \theta$. We use Q in addition to P to denote programs when comparing two executions. The proofs are given in Appendix A.

Lemma 3 (High Step) If $P = (H, R, I)_\kappa, SL(\kappa) \not\subseteq \theta, \Psi; \Gamma \vdash P$, then either: (1) there exists Γ_1 and $P_1 = (H_1, R_1, I_1)_\kappa$ such that $P \mapsto P_1, \Psi; \Gamma_1 \vdash P_1, \Gamma \succeq_\theta \Gamma_1$, and $\Psi; \Gamma_1 \vdash P \approx_\theta P_1$, or (2) I is of the form $(\text{raise } \kappa'; I')$ or $(\text{lower } w)$.

Lemma 4 (Context Discharge) If $P = (H, R, I)_{\theta \triangleright w}, \theta \not\subseteq \theta', \Psi; \Gamma \vdash P, P \mapsto^* (H_0, R_0, \text{halt } [\sigma])_\bullet$, then there exists Γ' and $P' = (H', R', \text{lower } w)_{\theta \triangleright w}$, such that $\Psi; \Gamma' \vdash P', P \mapsto^* P', \Gamma \succeq_{\theta'} \Gamma'$, and $\Psi; \Gamma' \vdash P \approx_{\theta'} P'$.

Lemma 5 (Low Step) If $P = (H, R, I)_\kappa, SL(\kappa) \subseteq \theta, \Psi; \Gamma \vdash P, \Psi; \Gamma \vdash Q, \Psi; \Gamma \vdash P \approx_\theta Q, P \mapsto P_1, Q \mapsto Q_1$, then exists Γ_1 such that $\Psi; \Gamma_1 \vdash P_1, \Psi; \Gamma_1 \vdash Q_1$ and $\Psi; \Gamma_1 \vdash P_1 \approx_\theta Q_1$.

Theorem 1 (Noninterference) If $P = (H, R, I)_\kappa, SL(\kappa) \subseteq \theta, \Psi; \Gamma \vdash P, \Psi; \Gamma \vdash Q, \Psi; \Gamma \vdash P \approx_\theta Q, P \mapsto^* (H_p, R_p, \text{halt } [\sigma_p])_\bullet$, and $Q \mapsto^* (H_q, R_q, \text{halt } [\sigma_q])_\bullet$, then exists Γ' such that $\Psi; \Gamma' \vdash (H_p, R_p, \text{halt } [\sigma_p])_\bullet \approx_\theta (H_q, R_q, \text{halt } [\sigma_q])_\bullet$.

5. Certifying Compilation

Certifying compilation for a realistic language typically involves a complex sequence of transformations, including CPS and closure conversion, heap allocation, and code generation [18, 15]. In this paper, we choose the simple security-type system of Figure 1 as our source language. This allows a concise presentation, yet suffices in demonstrating a main contribution: the separation of security-context operations (raise and lower) from conventional instructions and mechanisms (e.g., stack convention for procedure calls).

The low-high security hierarchy of Figure 1 defines a simple lattice consisting of two elements: \perp and \top . We use $|\mathfrak{t}|$ to denote the translation of source type \mathfrak{t} in TAL_C : $|\text{low}| \equiv \text{int}_\perp$ and $|\text{high}| \equiv \text{int}_\top$. We also translate the procedure types from the source language into TAL_C as follows:

$$|\langle \text{pc} \rangle (\mathfrak{t}_1, \dots, \mathfrak{t}_n) \rightarrow \text{void}| = (\forall [\Delta]. \langle \kappa \rangle \{ \text{sp} : \Sigma \})_\perp$$

where $(\Delta, \kappa) = \begin{cases} (\rho_\circ, \bullet) & \text{if } \text{pc} = \text{low} \\ (\alpha \rho_\circ, \top \triangleright \alpha) & \text{if } \text{pc} = \text{high} \end{cases}$

and $\Sigma = (\forall [\Delta]. \langle \kappa \rangle \{ \text{sp} : \rho \})_\perp :: \langle |\mathfrak{t}_1| \rangle_\perp :: \dots :: \langle |\mathfrak{t}_n| \rangle_\perp :: \rho$

This procedure type translation assumes a calling convention where the caller pushes a return pointer and the location of the arguments (implementing the call-by-reference semantics of the source language) onto the stack, and the callee deallocates the current stack frame upon return. The stack type Σ refers to a variable ρ because the procedure may be called under different stacks, as long as the current stack frame is as expected. The security context κ is empty if pc is low , or $\top \triangleright \alpha$ if pc is high . Postdominator variable α is used because the procedure may be called in security contexts with different postdominators. The type environment Δ simply collects all the needed type variables.

We assume that the program translation starts in a heap H_0 and a heap type Ψ_0 which satisfy $\vdash H_0 : \Psi_0$ and contain entries for all the variables and procedures of the source program. For any source variable v that $\mathfrak{f}(v) = \mathfrak{t}$, there exists a location l_v in the heap such that $\Psi(l_v) = \langle |\mathfrak{t}| \rangle_\perp$. For any source procedure f that $\mathfrak{f}(f) = \langle \text{pc} \rangle (\mathfrak{t}_1, \dots, \mathfrak{t}_n) \rightarrow \text{void}$, there exists a location l_f in the heap such that $\Psi(l_f) = |\langle \text{pc} \rangle (\mathfrak{t}_1, \dots, \mathfrak{t}_n) \rightarrow \text{void}|$. We use $\mathfrak{f} \sim \Psi$ to refer to this correspondence.

The above heap H_0 can be constructed with dummy slots for the procedures—the code in there simply jumps to itself. This suffices for typing the initial heap, thus facilitating the type-preservation proof. It creates locations for all source procedures and allows the translation of the actual code to refer to them.

[TREconst]	$ i = \text{mov } r, i \parallel r$
[TREvar]	$ v = \text{mov } r, l_v; \text{ld } r', r(0) \parallel r'$
[TREarg]	$ x_i = \text{sld } r, \text{sp}(i); \text{ld } r', r(0) \parallel r'$
[TREadd]	$\frac{ E = \vec{v} \parallel r \quad E' = \vec{v}' \parallel r' \quad \vec{v}' \text{ does not use } r}{ E + E' = \vec{v}; \vec{v}'; \text{add } r'', r, r' \parallel r''}$

Figure 11. Expression translation

The translation details are given in Figures 11, 12 and 13, based on the structure of the typing derivation of the source program. Which translation rule to apply is determined by the last typing rule used to check the source construct (program, procedure, or command). We use TD to denote (possibly multiple) typing derivations.

We define expression translation of the form $|E| = \vec{v} \parallel r$ in Figure 11. The instruction vector \vec{v} computes the value of E and the result is put in the register r . For a global variable, the value is loaded from the heap using its corresponding heap label. For a procedure argument, the location of the actual entity is loaded from the stack, and the value is then loaded from the heap.

In Figure 12, when translating a program [Rule [TRP1]], we translate all the procedure declarations, add halting code as the ending point of the program, and proceed to translate the main command. The result triple contains the updated heap type and heap, and a starting label l which leads to the starting point of the program. Procedure translation (Rule [TRF1]) takes care of part of the calling convention. It adds epilogue code that loads the return pointer, deallocates the current stack frame and transfers the control to the return pointer. It then resorts to command translation to translate the procedure body, providing the label to the epilogue code as the ending point of the procedure body.

In Figure 13, we define command translation of the form

$$\left| \frac{\text{TD}}{[\text{pc}] \vdash C} \right| \left[\begin{array}{c} \Psi \\ H \\ l_{start}; l_{end}; \Delta; \kappa; \Sigma \end{array} \right] = \left[\begin{array}{c} \Psi' \\ H' \end{array} \right].$$

This command translation takes 7 arguments: a code heap type (Ψ), a code heap (H), starting and ending labels (l_{start} and l_{end}) for the computation of C , a type environment (Δ), a security context (κ), and a stack type (Σ). It generates the extended code heap type (Ψ') and code heap (H'). Unsurprisingly, this translation appears complex, because it provides a formal model of a certifying compiler. Nonetheless, it is easy to follow if we remember some invariants maintained by the translation:

- H is well-typed under Ψ and contains entries for all source variables and procedures;
- Ψ and H already contain the continuation code labeled l_{end} ;
- The new code labeled l_{start} will be put into Ψ' and H' ;
- The security context κ must match pc ;
- The stack type Σ contains entries for all procedure arguments, if the command being compiled is in the body of a procedure;
- The environment Δ contains all free type variables in κ and Σ .

Most of the command translation rules simply put Δ, κ and Σ in place for the generated code types, and further propagate them to the translation of sub-components. The only rule that non-trivially manipulates the security context is Rule [TRC4]—when a subsumption rule is used for typing a source command, the translation generates code that is enclosed in a raise-lower pair. The translation of the sub-component is carried out in an updated heap with a new ending label l_1 . The code at l_1 restores the security context and transfers the control to the given ending label l' . After

$$\begin{array}{c}
\text{[TRP1]} \quad \frac{\left| \frac{\text{TD}_i}{\Phi \vdash F_i} \right| \left[\frac{\Psi_{i-1}}{H_{i-1}} \right] = \left[\frac{\Psi_i}{H_i} \right] \quad \forall i \in \{1 \dots n\} \quad l_{halt} \text{ is a fresh label}}{\left| \frac{\text{TD}}{\Phi; [\text{low}] \vdash C} \right| \left[\frac{\Psi_n \{l_{halt} : (\forall[o]. \langle \bullet \rangle \{sp : nil\})_{\perp}\}}{H_n \{l_{halt} \mapsto \text{code}[o] \langle \bullet \rangle \{sp : nil\}. \text{mov } r_1, l; \text{halt } [\text{int}_{\perp}]\}} \right]} = \left[\frac{\Psi}{H} \right]} \\
\frac{\left| \frac{\text{TD}_i}{\Phi \vdash F_i} \quad \forall i \in \{1 \dots n\} \quad \left| \frac{\text{TD}}{\Phi; [\text{low}] \vdash C} \right| \dots \right| \left[\frac{\Psi_0}{H_0} \right] = \left[\frac{\Psi}{H; l} \right]}{\Phi \vdash \{F_1; \dots; F_n; C\}} \\
\text{[TRF1]} \quad \frac{\left| \frac{\text{TD}}{\mathbf{x}_1 : \mathbf{t}_1, \dots, \mathbf{x}_n : \mathbf{t}_n, \Phi; [\text{pc}] \vdash C} \right| \left[\frac{\Psi \{l : (\forall[\Delta]. \langle \kappa \rangle \{sp : \Sigma\})_{\perp}\}}{H \{l \mapsto \text{code}[\Delta] \langle \kappa \rangle \{sp : \Sigma\}. \text{sld } r, sp(0); \text{sfree } (n+1); \text{jmp } r\}} \right]} = \left[\frac{\Psi'}{H'} \right]}{\text{where } (\Delta, \kappa) = \begin{cases} (\rho \circ \bullet) & \text{if pc = low} \\ (\alpha \rho \circ \top \triangleright \alpha) & \text{if pc = high} \end{cases} \quad l \text{ is a fresh label}} \\
\frac{\Sigma = (\forall[\Delta]. \langle \kappa \rangle \{sp : \rho\})_{\perp} :: \langle \mathbf{t}_1 \rangle_{\perp} \dots \dots \langle \mathbf{t}_n \rangle_{\perp} :: \rho}{\left| \frac{\text{TD}}{\Phi \vdash f(\text{pc})(\mathbf{x}_1 : \mathbf{t}_1, \dots, \mathbf{x}_n : \mathbf{t}_n)\{C\}} \right| \left[\frac{\Psi}{H} \right] = \left[\frac{\Psi'}{H'} \right]}
\end{array}$$

Figure 12. Program and procedure declaration translation

the translation of the sub-component, code is added at the starting label l to raise the security context to the expected level.

Procedure call translation is given as Rule [TRC7]. It creates “prologue” code that allocates a stack frame, pushes the return pointer and the arguments onto the stack, and jumps to the procedure label. Note that the corresponding epilogue code is generated by the procedure declaration translation in Rule [TRF1].

The translation of while-loops is also interesting (Rule [TRC6]). When translating the loop body, we need to prepare the continuation block, which happens to be the code for the loop test. We make use of a dummy block labeled l to serve as the continuation block when translating the body C . This block is introduced for maintaining the above invariants. It facilitates the type-preservation proof of the translation. After the translation of the loop body, this dummy block is replaced with the actual code that implements the loop test, as shown on the bottom right side of Rule [TRC6].

Lemma 6 (Expression Translation) If $\Phi \sim \Psi$, $\Phi \vdash E : \mathbf{t}$, $|E| = \vec{v} || r$, and $\Psi; \Delta; \{r : |\mathbf{t}|, sp : \Sigma\}; \kappa \vdash I$, then $\Psi; \Delta; \{sp : \Sigma\}; \kappa \vdash \vec{v}; I$.

Lemma 7 (Command Translation) If $\Phi \sim \Psi$, $\Phi; [\text{pc}] \vdash C$,

$$\left| \frac{\text{TD}}{\Phi; [\text{pc}] \vdash C} \right| \left[\frac{\Psi}{H} \right]_{l_{start}; l_{end}; \Delta; \kappa; \Sigma} = \left[\frac{\Psi'}{H'} \right],$$

$\Psi(l_{end}) = (\forall[\Delta]. \langle \kappa \rangle \{sp : \Sigma\})_{\perp}$, $SL(\kappa) = [\text{pc}]$, $\vdash H : \Psi$, then $\Phi \sim \Psi'$, $\vdash H' : \Psi'$ and $\Psi'; \Delta \vdash l_{start} : (\forall[\Delta]. \langle \kappa \rangle \{sp : \Sigma\})_{\perp}$.

The proofs for the above two lemmas are straightforward by structural induction on the derivation of the translation. Type preservation of procedure translation can be derived from Lemma 7 based on Rule [TRF1]. Type preservation of program translation then follows based on Rule [TRP1].

Lemma 8 (Procedure Translation) If $\Phi \sim \Psi$, $\Phi \vdash F, \vdash H : \Psi$,

$$\left| \frac{\text{TD}}{\Phi \vdash F} \right| \left[\frac{\Psi}{H} \right] = \left[\frac{\Psi'}{H'} \right], \text{ then } \Phi \sim \Psi' \text{ and } \vdash H' : \Psi'.$$

Theorem 2 (Program Translation) If $\Phi \sim \Psi_0$, $\Phi \vdash P$,

$$\vdash H_0 : \Psi_0, \left| \frac{\text{TD}}{\Phi \vdash P} \right| \left[\frac{\Psi_0}{H_0} \right] = \left[\frac{\Psi}{H; l} \right], \text{ then } \Phi \sim \Psi \text{ and } \Psi; \{sp : nil\} \vdash (H, \{sp : nil\}, \text{jmp } l)_{\bullet}.$$

6. Discussions

Linear Continuations Zdancewic and Myers [33] introduced a notion of ordered linear continuations to facilitate the information-flow analysis at a low level (we use ZM to refer to their system). An important requirement of such analysis is that one needs to allow a high-security conditional to be surrounded by low-security computation. In ZM, before the conditional statement, a linear continuation is created to capture the computation after the conditional. Such a linear continuation must be called exactly once at the end of either branch of the conditional. Furthermore, the linear continuation records the security context in which it is created, allowing the security context to be reset properly when the branches meet.

As a higher-order analog to postdominators in a control-flow graph, ordered linear continuations enforce a stack discipline that allows security contexts to be reset at the join points of program branches. The static semantics ensures that the linear continuations are properly nested, and at any time only the top continuation on the (virtual) continuation stack is available. The linearity is enforced because the continuation is essentially popped off the stack when used. In particular, every value in ZM is tagged with a security label. The operational semantics keeps track of the security context during the execution, and ensures that security labels of the values are propagated correctly.

It may help to view our solution as an adaptation of linear continuations for the RISC architecture (we emphasize that there is not a loss of expressiveness; interested readers are referred to Appendix C for some details). A postdominator of program branches is essentially expressed as a static code label. The security operations **raise** and **lower** correspond to the creation and elimination of linear continuations. At any program point, our static semantics keeps track of only the top element of the (virtual) continuation stack. The typing rule for **raise** ensures that the security context at the postdominator matches the current one, thus enforcing the stack discipline.

$$\begin{array}{c}
\text{[TRC1]} \quad \frac{|E| = \vec{l} \parallel r}{\left| \frac{\Phi(V) = \text{high}}{\Phi; [\text{pc}] \vdash V := E} \right| \left[\frac{\Psi}{H} \right]_{l; l'; \Delta; \kappa; \Sigma} = \left[\frac{\Psi\{l : (\forall[\Delta]. \langle \kappa \rangle \{\text{sp} : \Sigma\})_{\perp}\}}{H\{l \mapsto \text{code}[\Delta]\langle \kappa \rangle \{\text{sp} : \Sigma\}. \vec{l}; \text{mov } r', l_v; \text{st } r'(0), r; \text{jmp } l'[\Delta]\}} \right]}
\\
\text{[TRC2]} \quad \frac{|E| = \vec{l} \parallel r}{\left| \frac{\Phi(V) = \text{low} \quad \Phi \vdash E : \text{low}}{\Phi; [\text{low}] \vdash V := E} \right| \left[\frac{\Psi}{H} \right]_{l; l'; \Delta; \bullet; \Sigma} = \left[\frac{\Psi\{l : (\forall[\Delta]. \langle \bullet \rangle \{\text{sp} : \Sigma\})_{\perp}\}}{H\{l \mapsto \text{code}[\Delta]\langle \bullet \rangle \{\text{sp} : \Sigma\}. \vec{l}; \text{mov } r', l_1; \text{st } r'(0), r; \text{jmp } l'[\Delta]\}} \right]}
\\
\text{[TRC3]} \quad \frac{|E| = \vec{l} \parallel r \quad l_1, l_2 \text{ are fresh labels}}{\left| \frac{\text{TD}_1}{\Phi; [\text{pc}] \vdash C_1} \right| \left[\frac{\Psi}{H} \right]_{l_1; l'; \Delta; \kappa; \Sigma} = \left[\frac{\Psi_1}{H_1} \right] \quad \left| \frac{\text{TD}_2}{\Phi; [\text{pc}] \vdash C_2} \right| \left[\frac{\Psi_1}{H_1} \right]_{l_2; l'; \Delta; \kappa; \Sigma} = \left[\frac{\Psi_2}{H_2} \right]}
\\
\text{[TRC3]} \quad \frac{\Phi \vdash E : \text{pc} \quad \frac{\text{TD}_1}{\Phi; [\text{pc}] \vdash C_1} \quad \frac{\text{TD}_2}{\Phi; [\text{pc}] \vdash C_2}}{\Phi; [\text{pc}] \vdash \text{if } E \text{ then } C_1 \text{ else } C_2} \left[\frac{\Psi}{H} \right]_{l; l'; \Delta; \kappa; \Sigma} = \left[\frac{\Psi_2\{l : (\forall[\Delta]. \langle \kappa \rangle \{\text{sp} : \Sigma\})_{\perp}\}}{H_2\{l \mapsto \text{code}[\Delta]\langle \kappa \rangle \{\text{sp} : \Sigma\}. \vec{l}; \text{bnz } r, l_1[\Delta]; \text{jmp } l_2[\Delta]\}} \right]}
\\
\text{[TRC4]} \quad \left| \frac{\text{TD}}{\Phi; [\text{high}] \vdash C} \right| \left[\frac{\Psi\{l_1 : (\forall[\Delta]. \langle \top \triangleright l'[\Delta] \rangle \{\text{sp} : \Sigma\})_{\perp}\}}{H\{l_1 \mapsto \text{code}[\Delta]\langle \top \triangleright l'[\Delta] \rangle \{\text{sp} : \Sigma\}. \text{lower } l'[\Delta]\}} \right]_{l_0; l_1; \Delta; \top \triangleright l'[\Delta]; \Sigma} = \left[\frac{\Psi'}{H'} \right]_{l_0, l_1 \text{ are fresh labels}}
\\
\text{[TRC4]} \quad \left| \frac{\text{TD}}{\Phi; [\text{high}] \vdash C} \right| \left[\frac{\Psi}{H} \right]_{l; l'; \Delta; \bullet; \Sigma} = \left[\frac{\Psi'\{l : (\forall[\Delta]. \langle \bullet \rangle \{\text{sp} : \Sigma\})_{\perp}\}}{H'\{l \mapsto \text{code}[\Delta]\langle \bullet \rangle \{\text{sp} : \Sigma\}. \text{raise } \top \triangleright l'[\Delta]; \text{jmp } l_0[\Delta]\}} \right]}
\\
\text{[TRC5]} \quad \frac{l_1 \text{ is a fresh label}}{\left| \frac{\text{TD}_2}{\Phi; [\text{pc}] \vdash C_2} \right| \left[\frac{\Psi}{H} \right]_{l_1; l'; \Delta; \kappa; \Sigma} = \left[\frac{\Psi_1}{H_1} \right] \quad \left| \frac{\text{TD}_1}{\Phi; [\text{pc}] \vdash C_1} \right| \left[\frac{\Psi_1}{H_1} \right]_{l; l_1; \Delta; \kappa; \Sigma} = \left[\frac{\Psi_2}{H_2} \right]}
\\
\text{[TRC5]} \quad \left| \frac{\text{TD}_1}{\Phi; [\text{pc}] \vdash C_1} \quad \frac{\text{TD}_2}{\Phi; [\text{pc}] \vdash C_2} \right| \left[\frac{\Psi}{H} \right]_{l; l'; \Delta; \kappa; \Sigma} = \left[\frac{\Psi_2}{H_2} \right]}
\\
\text{[TRC6]} \quad \frac{|E| = \vec{l} \parallel r \quad \left| \frac{\text{TD}}{\Phi; [\text{pc}] \vdash C} \right| \left[\frac{\Psi\{l : (\forall[\Delta]. \langle \kappa \rangle \{\text{sp} : \Sigma\})_{\perp}\}}{H\{l \mapsto \text{code}[\Delta]\langle \kappa \rangle \{\text{sp} : \Sigma\}. \text{jmp } l[\Delta]\}} \right]_{l_1; l; \Delta; \kappa; \Sigma} = \left[\frac{\Psi'}{H'} \right]}{\left| \frac{\Phi \vdash E : \text{pc} \quad \frac{\text{TD}}{\Phi; [\text{pc}] \vdash C}}{\Phi; [\text{pc}] \vdash \text{while } E \text{ do } C} \right| \left[\frac{\Psi}{H} \right]_{l; l'; \Delta; \kappa; \Sigma} = \left[\frac{\Psi'\{l : (\forall[\Delta]. \langle \kappa \rangle \{\text{sp} : \Sigma\})_{\perp}\}}{H'\{l \mapsto \text{code}[\Delta]\langle \kappa \rangle \{\text{sp} : \Sigma\}. \vec{l}; \text{bnz } r, l_1[\Delta]; \text{jmp } l'[\Delta]\}} \right]}
\\
\text{[TRC7]} \quad \frac{\begin{array}{l} |\mathbf{V}_i| = \vec{l}_i \parallel r_i \text{ where } \vec{l}_i \text{ does not use } r_j \text{ if } j < i \quad \forall i \in \{1 \dots n\} \\ \Psi' = \Psi\{l : (\forall[\Delta]. \langle \kappa \rangle \{\text{sp} : \Sigma\})_{\perp}\} \\ H' = \Psi\{l \mapsto \text{code}[\Delta]\langle \kappa \rangle \{\text{sp} : \Sigma\}. \vec{l}_1; \dots \vec{l}_n; \text{salloc } (n+1); \text{sst } \text{sp}(0), l'[\vec{\psi}]; \\ \text{sst } \text{sp}(1), r_1; \dots \text{sst } \text{sp}(n), r_n; \text{jmp } l_f[\vec{\psi}]\} \end{array}}{\left| \frac{\Phi(f) = \langle \text{pc} \rangle (t_1, \dots, t_n) \rightarrow \text{void} \quad \frac{\text{TD}_i}{\Phi \vdash \mathbf{V}_i : t_i} \quad \forall i \in \{1 \dots n\}}{\Phi; [\text{pc}] \vdash f(\mathbf{V}_1, \dots, \mathbf{V}_n)} \right| \left[\frac{\Psi}{H} \right]_{l; l'; \Delta; \kappa; \Sigma} = \left[\frac{\Psi'}{H'} \right]} \quad \text{where } \vec{\psi} = \begin{cases} \Sigma & \text{if } \kappa = \bullet \\ w\Sigma & \text{if } \kappa = \top \triangleright w \end{cases}
\end{array}$$

Figure 13. Command translation

```

int double(x:int) { x=x*2; } ...
if h<10 then double(h);
double(1); ...

```

Figure 14. Security-polymorphic function

We wish to point out, nonetheless, that such an adaptation yields a simple, practical and well-grounded solution to the identified problem of information-flow analysis for assembly code. In particular, it bridges the gap between the functional abstraction of linear continuations and the raw assembly code running on actual machines. In comparison with ZM, our system TAL_C models the use of registers and assembly instructions, and hence is closer to the actual RISC architecture. We do not attach security labels to values; this makes it trivial to see that security annotations do not affect computation. In fact, the enforcement of noninterference in TAL_C is cleanly separated from normal program execution.¹ It is also obvious that security operations in TAL_C are orthogonal from conventional instructions (*e.g.*, branching and jumping) and mechanisms (*e.g.*, call stack), which allows our approach to be carried further with other language extensions. Consequently, we consider TAL_C as a good first step toward a scaled-up typed assembly language for noninterference.

SIF SIF [16] is developed independently from TAL_C . These two systems are similar in spirit—both use static types for information-flow analysis. However, SIF is based on a minimal language where relatively simple annotations, namely a stack of static code labels, suffice. In a more realistic language, a single function (even if monomorphic with respect to security levels) can be called at different program points. The security contexts of these program points may be different with respect to (1) the postdominator of the current context (SIF tracks this with the top stack element), and (2) the “enclosing contexts” (SIF tracks these with the stack tail). Since the label stack of SIF is made up of static code labels, one cannot reuse the same code at different program points with different contexts.

TAL_C only maintains the current security context at any program point, and we show that it suffices for establishing noninterference. With such a treatment, the code types are naturally polymorphic with respect to enclosing contexts. We also allow post-dominators to be polymorphic. The certifying compilation scheme further demonstrates that TAL_C is expressive enough for supporting the source language.

7. Extensions and Future Work

Orthogonal features For ease of understanding, TAL_C focuses on a minimal set of language features. Nonetheless, polymorphic and existential types, as seen in TAL, are orthogonal and can be introduced with little difficulty. Furthermore, since TAL_C is compatible with TAL, it is also possible to accommodate other features of the TAL family. For instance, alias types [26] may provide a more accurate alias analysis, improving the current conservative approach that considers every pointer as a potential alias. In the following, we will also discuss the use of singleton types [32].

Security polymorphism TAL_C relies on a security context $\theta \triangleright w$ to identify the current security level θ and its ending point w . It is monomorphic with respect to security, because the security level of a code block is fixed. In practice, security-polymorphic code can also be useful.

Figure 14 gives an example. The function `double` can be invoked with either low or high input. It is safe to invoke `double`

in a context if only the security level of the input matches that of the context. In a security polymorphic TAL_C -like type system, `double` can be given the type

$$(\forall[\theta, \alpha].(\theta \triangleright \alpha)\{r_1 : \text{int}_\theta, r_0 : (\forall[].(\theta \triangleright \alpha)\{r_1 : \text{int}_\theta\})_\perp\})_\perp.$$

Here r_1 is the argument register, r_0 stores the return pointer, and the meta-variable θ is reused as a variable.

It is straightforward to support this kind of polymorphism. In fact, most of the required constructs are already present in TAL_C . We omitted such polymorphism simply because it complicates the presentation without providing additional insights. Nonetheless, the expressiveness of such polymorphism is still limited. Since the label α is not known until instantiated, the code of `double` has no knowledge about α . Hence the security context $\theta \triangleright \alpha$ cannot be discharged within the body of `double`.

It is not obvious why one would wish to discharge the security context within a polymorphic function. Indeed, it is always possible to wrap a function call inside a secured region by symmetric `raise` and `lower` operations from the caller’s side. However, the asymmetric discharging of security context may be desirable for *certifying optimization*. For instance, in Figure 14, `double` is called as the last statement of the body of a high conditional. In this case, directly discharging the security context when `double` returns would remove a superfluous `lower` from the caller’s side. Such a discharging requires `lower` to operate on small values—since the return label is not statically fixed, it must be passed in through a register.

It may require singleton and intersection types to support such a `lower` operation. For example, a `double` function that discharges its security context can have type

$$\left(\forall[\theta, \alpha].(\theta \triangleright \alpha) \left\{ \begin{array}{l} r_1 : \text{int}_\theta, \\ r_0 : \text{ sint}(\alpha)_\perp \wedge (\forall[].(\bullet)\{r_1 : \text{int}_\theta\})_\perp \end{array} \right\} \right)_\perp.$$

At the end of the function, `lower` r_0 discharges the security context and transfers the control to the return code. For type checking, the singleton integer type `sint`(α) matches the register r_0 with the label in the security context, and the code type ensures that the control flow to the return point is safe.

Full erasure With the powerful type constructs above, one can achieve a full erasure for the `lower` operation. Instead of treating `lower` as an instruction, one can treat it as a transformation on small values. This is in spirit similar to the *pack* operation of existential types in TAL. Such a `lower` transformation bridges the gap between the current security context and the security level of the target label. The actual control flow transfer is then completed with a conventional jump instruction (*e.g.*, `jmp (lower r_0)`).

One can also achieve a full erasure for `lower` even without singleton types. The idea is to separate the jump instruction into direct jump and indirect jump. This is also consistent with real machine architectures. The `lower` operation transforms word values (eventually, direct labels). Lowered labels, similar to packed values, may serve as the operand of direct jump. Indirect jump, on the other hand, takes normal small values. This is expressive enough for certifying compilation, yet may not be sufficient for certifying optimization as discussed above.

Other future work It is a challenging task to study how the features above yield a system expressive enough for certifying optimization. In TAL_C , `lower` erases to a direct jump, hence consecutive `lower` operations result in superfluous jumps. Ideally, these should be combined into a single jump whose operand is a nested lowered value. Similarly, it is also desirable to combine `lower` with an adjacent jump instruction. In practice, certifying optimization are sometimes considered for conventional type safety [9].

¹The extra subscript of security context in a program P is only for facilitating the noninterference proof; it can be completely ignored for computation.

A security-type preserving translation for a full-fledged source language is another challenging task. The formal translation and type-preservation theorem in this paper are based on a concise source language for demonstrative purposes. Practical embodiment requires much further work. Existing work on type-preserving translation [18, 15] and high-level information-flow analysis [23] may shed light on the support of more advanced language features.

8. Conclusion

We have presented a language TAL_C for enforcing data confidentiality in assembly code. The main idea is to use type annotations to restore high-level abstractions that are crucial to information-flow analysis. In TAL_C , operations related to security are kept orthogonal from other language features. As a result, it is possible to accommodate existing results on low-level verification, such as the TAL family. We have also presented a translation from a high-level security language with first-order procedures to TAL_C . A soundness theorem shows that the translation preserves security types. We consider this as a useful step toward a certifying compiler for noninterference.

Acknowledgments

We wish to thank Eduardo Bonelli, Adriana Compagnoni, Ricardo Medel, Greg Morrisett, Steve Zdancewic and the anonymous referees for helpful comments on previous drafts of this paper.

References

- [1] M. Abadi, A. Banerjee, H. Heintze, and J. G. Riecke. A core calculus of dependency. In *Proc. 26th ACM Symposium on Principles of Programming Languages*, pages 147–160, San Antonio, TX, Jan. 1999.
- [2] J. Agat. Transforming out timing leaks. In *Proc. 27th ACM Symposium on Principles of Programming Languages*, pages 40–53, Boston, MA, Jan. 2000.
- [3] J. Agat. *Type Based Techniques for Covert Channel Elimination and Register Allocation*. PhD thesis, Chalmers University of Technology and Gothenburg University, Gothenburg, Sweden, Dec. 2000.
- [4] M. Avvenuti, C. Bernardeschi, and F. Francesco. Java bytecode verification for secure information flow. *ACM SIGPLAN Notices*, 38(12):20–27, Dec. 2003.
- [5] T. Ball. What’s in a region? Or computing control dependence regions in near-linear time for reducible control flow. *ACM Letters on Programming Languages and Systems*, 2(1–4):1–16, Mar.–Dec. 1993.
- [6] A. Banerjee and D. A. Naumann. Secure information flow and pointer confinement in a Java-like language. In *Proc. 15th IEEE Computer Security Foundations Workshop*, pages 253–267, June 2002.
- [7] G. Barthe, A. Basu, and T. Rezk. Security types preserving compilation. In *Proc. 5th International Conference on Verification, Model Checking and Abstract Interpretation*, volume 2937 of LNCS, pages 2–15, Venice, Italy, Jan. 2004.
- [8] E. Bonelli, A. Compagnoni, and R. Medel. SIFTAL: A typed assembly language for secure information flow analysis. Technical report, Stevens Institute of Technology, Hoboken, NJ, July 2004.
- [9] J. Chen, D. Wu, A. W. Appel, and H. Fang. A provably sound TAL for back-end optimization. In *Proc. 2003 ACM Conference on Programming Language Design and Implementation*, pages 208–219, San Diego, CA, June 2003.
- [10] D. E. R. Denning. *Cryptography and Data Security*. Addison-Wesley, Boston, MA, 1982.
- [11] N. Heintze and J. G. Riecke. The SLam calculus: Programming with security and integrity. In *Proc. 25th ACM Symposium on Principles of Programming Languages*, pages 365–377, San Diego, CA, Jan. 1998.
- [12] K. Honda and N. Yoshida. A uniform type structure for secure information flow. In *Proc. 29th ACM Symposium on Principles of Programming Languages*, pages 81–92, Portland, OR, Jan. 2002.
- [13] D. Kozen. Efficient code certification. Technical Report TR98-1661, Cornell University, Jan. 1998.
- [14] D. Kozen. Language-based security. In *Proc. 24th International Symposium on Mathematical Foundations of Computer Science*, volume 1672 of LNCS, pages 284–298, Szklarska Poreba, Poland, Sept. 1999.
- [15] C. League, Z. Shao, and V. Trifonov. Precision in practice: A type-preserving Java compiler. In *Proc. 12th International Conference on Compiler Construction*, volume 2622 of LNCS, pages 106–120, Warsaw, Poland, Apr. 2003.
- [16] R. Medel, A. Compagnoni, and E. Bonelli. Non-interference for a typed assembly language. In *Proc. 2005 Workshop on Foundations of Computer Security*, Chicago, IL, June 2005.
- [17] G. Morrisett, K. Crary, N. Glew, and D. Walker. Stack-based typed assembly language. *Journal of Functional Programming*, 12(1):43–88, Jan. 2002.
- [18] G. Morrisett, D. Walker, K. Crary, and N. Glew. From system F to typed assembly language. *ACM Transactions on Programming Languages and Systems*, 21(3):527–568, Nov. 1999.
- [19] A. C. Myers. JFlow: Practical mostly-static information flow control. In *Proc. 26th ACM Symposium on Principles of Programming Languages*, pages 228–241, San Antonio, TX, 1999.
- [20] G. C. Necula. *Compiling with Proofs*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, Sept. 1998.
- [21] G. C. Necula and P. Lee. The design and implementation of a certifying compiler. In *Proc. 98 ACM Conference on Programming Language Design and Implementation*, pages 333–344, Montreal, Canada, June 1998.
- [22] F. Pottier and V. Simonet. Information flow inference for ML. *ACM Transactions on Programming Languages and Systems*, 25(1):117–158, Jan. 2003.
- [23] A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, Jan. 2003.
- [24] J. H. Saltzer and M. D. Schroeder. The protection of information in computer systems. *Proc. of the IEEE*, 63(9), Sept. 1975.
- [25] F. B. Schneider, G. Morrisett, and R. Harper. A language-based approach to security. In *Informatics: 10 Years Back, 10 Years Ahead*, volume 2000 of LNCS, pages 86–101, 2001.
- [26] F. Smith, D. Walker, and G. Morrisett. Alias types. In *Proc. 9th European Symposium on Programming*, volume 1782 of LNCS, pages 366–381, Berlin, Germany, Apr. 2000.
- [27] G. Smith and D. Volpano. Secure information flow in a multi-threaded imperative language. In *Proc. 25th ACM Symposium on Principles of Programming Languages*, pages 355–364, San Diego, CA, Jan. 1998.
- [28] D. Volpano and G. Smith. Eliminating covert flows with minimum typings. In *10th IEEE Computer Security Foundations Workshop*, pages 156–169, Washington, DC, June 1997.
- [29] D. Volpano and G. Smith. A type-based approach to program security. In *Proc. 7th International Joint Conference CAAP/FASE on Theory and Practice of Software Development*, LNCS, pages 607–621, Lille, France, Apr. 1997.
- [30] D. Volpano and G. Smith. Probabilistic noninterference in a concurrent language. In *Proc. 11th IEEE Computer Security Foundations Workshop*, pages 34–43, Washington, DC, June 1998.
- [31] D. Walker. A type system for expressive security policies. In *Proc. 27th ACM Symposium on Principles of Programming Languages*, pages 254–267, Boston, MA, Jan. 2000.
- [32] H. Xi and R. Harper. A dependently typed assembly language. In *Proc. 6th ACM International Conference on Functional Programming*, pages 169–180, Florence, Italy, Sept. 2001.
- [33] S. Zdancewic and A. C. Myers. Secure information flow via linear continuations. *Higher-Order and Symbolic Computation*, 15(2–3):209–234, Sept. 2002.

A. Noninterference Proof of TAL_C

Lemma 3 (High Step) If $P = (H, R, I)_\kappa, SL(\kappa) \not\subseteq \theta, \Psi; \Gamma \vdash P$, then either: (1) there exists Γ_1 and $P_1 = (H_1, R_1, I_1)_\kappa$ such that $P \mapsto P_1, \Psi; \Gamma_1 \vdash P_1, \Gamma \succeq_\theta \Gamma_1$, and $\Psi; \Gamma_1 \vdash P \approx_\theta P_1$, or (2) I is of the form $(\text{raise } \kappa'; I')$ or $(\text{lower } w)$.

Proof sketch: By case analysis on the first instruction of I . I cannot be `halt`, because the typing rule for `halt` requires the context to be \bullet . If I is not `halt`, `raise` or `lower`, by the operational semantics and inversion on the typing rules, one can get Γ_1 and P_1 for the next step. The typing rules prohibit writing into a low heap cell, hence low heap cells remain the same after the step. When a register or stack slot is updated, Γ_1 gives it a type whose security label takes $SL(\kappa)$ into account, hence that register or stack slot has a high type in Γ_1 . As a result, $\Gamma \succeq_\theta \Gamma_1$ and $\Psi; \Gamma_1 \vdash P \approx_\theta P_1$. \square

Lemma 4 (Context Discharge) If $P = (H, R, I)_{\theta \triangleright w}, \theta \not\subseteq \theta', \Psi; \Gamma \vdash P, P \mapsto^* (H_0, R_0, \text{halt } [\sigma])_\bullet$, then there exists Γ' and $P' = (H', R', \text{lower } w)_{\theta \triangleright w}$, such that $\Psi; \Gamma' \vdash P', P \mapsto^* P', \Gamma \succeq_{\theta'} \Gamma'$, and $\Psi; \Gamma' \vdash P \approx_{\theta'} P'$.

Proof sketch: By generalized induction on the number of steps of the derivation $P \mapsto^* (H_0, R_0, \text{halt } [\sigma])_\bullet$.

The base case of zero step is not possible, because the security contexts do not match. In the inductive case, suppose the execution consists of n steps and the proposition holds for any step number less than n . There are two cases to consider, following Lemma 3.

In the case where the first instruction of I is not `raise` or `lower`, by Lemma 3, there exists Γ_1 and P_1 such that $P \mapsto P_1, \Psi; \Gamma_1 \vdash P_1, \Gamma \succeq_{\theta'} \Gamma_1, \Psi; \Gamma_1 \vdash P \approx_{\theta'} P_1$, and the security context of P_1 is the same as that of P . P_1 must be a step in between P and $(H_0, R_0, \text{halt } [\sigma])_\bullet$ because the operational semantics is deterministic. Hence by induction hypothesis on P_1 , there exists Γ' and P' such that $\Psi; \Gamma' \vdash P', P_1 \mapsto^* P', \Gamma_1 \succeq_{\theta'} \Gamma'$ and $\Psi; \Gamma' \vdash P_1 \approx_{\theta'} P'$. Putting the above together, $P \mapsto^* P', \Gamma \succeq_{\theta'} \Gamma'$ because $\succeq_{\theta'}$ is transitive by definition, and $\Psi; \Gamma' \vdash P \approx_{\theta'} P'$ by definition and the fact that $\Gamma_1 \succeq_{\theta'} \Gamma'$.

Case $I = \text{raise } \theta_1 \triangleright w_1; I_1$. By definition of the operational semantics, $P \mapsto P_1$ where $P_1 = (H, R, I_1)_{\theta_1 \triangleright w_1}$. By inversion on $\Psi; \Gamma \vdash P$ and the typing rule of `raise`, $\theta \subseteq \theta_1$ and $\Psi; \Gamma; \theta_1 \triangleright w_1 \vdash I_1$. By definition of well-typed programs, $\Psi; \Gamma \vdash P_1$. By induction hypothesis on P_1 , there exists Γ_2 and $P_2 = (H_2, R_2, \text{lower } w_1)_{\theta_1 \triangleright w_1}$ such that $\Psi; \Gamma_2 \vdash P_2, P_1 \mapsto^* P_2, \Gamma \succeq_{\theta'} \Gamma_2, \Psi; \Gamma_2 \vdash P_1 \approx_{\theta'} P_2$. $\Psi; \Gamma_2 \vdash P \approx_{\theta'} P_2$ then follows because the heaps and register files in P and P_1 are the same.

Further by the operational semantics, $P_2 \mapsto P_3$ where $P_3 = (H_2, R_2, I_3)_\kappa$ and I_3 is the instantiated code of w_1 whose security context is κ . By inversion on the well-typedness of I (i.e., $\text{raise } \theta_1 \triangleright w_1; I_1$), $\kappa = \theta \triangleright w$. By induction hypothesis on P_3 , there exists Γ' and $P' = (H', R', \text{lower } w)_{\theta \triangleright w}$ such that $\Psi; \Gamma' \vdash P', P_3 \mapsto^* P', \Gamma_2 \succeq_{\theta'} \Gamma'$, and $\Psi; \Gamma' \vdash P_3 \approx_{\theta'} P'$. Putting the above together, the original proposition holds for case $I = \text{raise } \theta_1 \triangleright w_1; I_1$.

Case $I = \text{lower } w_1$. By inversion on the typing rule of `lower`, $w = w_1$. Let $P' = P$, the proposition holds. \square

Lemma 5 (Low Step) If $P = (H, R, I)_\kappa, SL(\kappa) \subseteq \theta, \Psi; \Gamma \vdash P, \Psi; \Gamma \vdash Q, \Psi; \Gamma \vdash P \approx_\theta Q, P \mapsto P_1, Q \mapsto Q_1$, then exists Γ_1 such that $\Psi; \Gamma_1 \vdash P_1, \Psi; \Gamma_1 \vdash Q_1$ and $\Psi; \Gamma_1 \vdash P_1 \approx_\theta Q_1$.

Proof sketch: By case analysis on the first instruction of I . By $SL(\kappa) \subseteq \theta$ and the definition of \approx_θ , P and Q contain the same instruction sequence. The case of raising does not change the state, hence trivially maintains the equivalence. All other cases

maintain that the security context is lower than θ . Inspection on the typing rules shows that low locations in the heap can only be assigned low values. Once a register or stack slot is given a high value, its type in Γ_1 will change to high. In the case of branching, the guard must be low, so both P and Q branch to the same code. Hence the two programs remain equivalent after one step. \square

Theorem 1 (Noninterference) If $P = (H, R, I)_\kappa, SL(\kappa) \subseteq \theta, \Psi; \Gamma \vdash P, \Psi; \Gamma \vdash Q, \Psi; \Gamma \vdash P \approx_\theta Q, P \mapsto^* (H_p, R_p, \text{halt } [\sigma_p])_\bullet$, and $Q \mapsto^* (H_q, R_q, \text{halt } [\sigma_q])_\bullet$, then exists Γ' such that $\Psi; \Gamma' \vdash (H_p, R_p, \text{halt } [\sigma_p])_\bullet \approx_\theta (H_q, R_q, \text{halt } [\sigma_q])_\bullet$.

Proof sketch: By generalized induction on the number of steps of the derivation $P \mapsto^* (H_p, R_p, \text{halt } [\sigma_p])_\bullet$. The base case of zero step is trivial. The inductive case is done by case analysis on the first instruction of I .

Consider the case where $I = \text{raise } \theta_1 \triangleright w_1; I_1$ and $\theta_1 \not\subseteq \theta$. By definition of the operational semantics and the typing rules, $P \mapsto P_1$ where $P_1 = (H, R, I_1)_{\theta_1 \triangleright w_1}$ and $\Psi; \Gamma \vdash P_1$. By Lemma 4, there exists Γ_2 and $P_2 = (H_2, R_2, \text{lower } w_1)_{\theta_1 \triangleright w_1}$ such that $\Psi; \Gamma_2 \vdash P_2, P_1 \mapsto^* P_2, \Gamma \succeq_\theta \Gamma_2$, and $\Psi; \Gamma_2 \vdash P_1 \approx_\theta P_2$. Hence $\Psi \vdash H \approx_\theta H_2$ and $\Gamma_2 \vdash R \approx_\theta R_2$.

By the operational semantics, $P_2 \mapsto P_3$ where $w_1 = l_1[\vec{\psi}]$, $P_3 = (H_2, R_2, I_3[\vec{\psi}/\Delta])_{\kappa_3}$, and $H(l_1) = \text{code}[\Delta]_{\langle \kappa_3 \rangle} \Gamma_3. I_3$. By inversion on the derivation of $\Psi; \Gamma_2 \vdash P_2, \Gamma_3 \subseteq \Gamma_2$ and $\Psi; \Gamma_3 \vdash P_3$. It follows that $\Gamma_3 \vdash R \approx_\theta R_2$. By inversion on the derivation of $\Psi; \Gamma \vdash P$ where the instruction sequence of P is $\text{raise } \theta_1 \triangleright w_1; I_1, \kappa_3 = \kappa$.

By similarly reasoning, $Q \mapsto^* Q_3$ where $Q_3 = (H'_2, R'_2, I_3)_{\kappa_3}$, $\Psi \vdash H \approx_\theta H'_2, \Gamma_3 \vdash R \approx_\theta R'_2$ and $\Psi; \Gamma_3 \vdash Q_3$. By transitivity of the equivalence relations, $\Psi \vdash H_2 \approx_\theta H'_2$ and $\Gamma_3 \vdash R_2 \approx_\theta R'_2$. Hence $\Psi; \Gamma \vdash P_3 \approx_\theta Q_3$. The case then follows by induction hypothesis.

All other cases remain low after a step. By Lemma 5, the two programs in the next step are equivalent and well-typed. The proof then follows by induction hypothesis. \square

B. Example

Figure 15 gives a simple example to demonstrate the use of security labels and contexts. The high-level pseudo-code program involves a low variable a and two high variables b and c . In a corresponding TAL_C program, we use heap cells labeled l_a, l_b and l_c to represent these variables. The TAL_C program starts from the code labeled l_0 in a low security context. After the initial setup, it raises the security context to $\top \triangleright l_3$. The control is then transferred to the code labeled l_1 , which contains a test on the high variable b and directs the execution to two separate branches. In either branch of the conditional, the high variable c is updated, and the security context is restored with `lower` l_3 . The code at l_3 is then free to update the low variable a again.

A closer look at the code labeled l_1 reveals several interesting issues. When checking the first load instruction (`ld` $r_4, r_2(0)$), the security level for r_4 is inferred to be high (\top). The following branching instruction (`bnz` r_4, l_2) type-checks because the current security context ($\top \triangleright l_3$) is high enough to cover the security level of r_4 . The next store instruction (`st` $r_3(0), r_0$) is also valid, because it is ok to update a high variable in a high context. In comparison, the store instruction would fail to type-check if c was a low variable. Finally, the high security context is ended with a lower instruction (`lower` l_3) that directs the control flow to the postdominator of the conditional.

```

A pseudo-code program:      a = 0;
                             if (b <> 0) then c = 1 else c = 0;
                             a = 1

A corresponding TALC program:  (H, {sp : nil}, jmp l0)•  where H = {
la ↦ ...
lb ↦ ...
lc ↦ ...
l0 ↦ code[o](•){sp : nil}.
      mov r0, 0;           % r0 ← 0
      mov r1, la;        % r1 ← la
      mov r2, lb;        % r2 ← lb
      mov r3, lc;        % r3 ← lc
      st r1(0), r0;      % la ← 0
      raise ⊤ ▷ l3;      % raise security context
      jmp l1

l1 ↦ code[o](⊤ ▷ l3){r0 : <int⊥>⊥, r1 : <int⊥>⊥, r2 : <int⊥>⊥, r3 : <int⊥>⊥, sp : nil}.
      ld r4, r2(0);
      bnz r4, l2;        % go to l2 if content of lb is not zero
      st r3(0), r0;      % the else branch: lc ← 0
      lower l3          % restore security context and go to l3

l2 ↦ code[o](⊤ ▷ l3){r0 : <int⊥>⊥, r1 : <int⊥>⊥, r2 : <int⊥>⊥, r3 : <int⊥>⊥, sp : nil}.
      mov r0, 1;
      st r3(0), r0;      % the then branch: lc ← 1
      lower l3          % restore security context and go to l3

l3 ↦ code[o](•){r1 : int⊥, sp : nil}.
      mov r0, 1;
      st r1(0), r0;      % la ← 1
      halt [int⊥]
}

```

Figure 15. TAL_C example

C. Translating Linear Continuations

It may appear that TAL_C is not as expressive as the language of Zdancewic and Myers’ [33] (ZM), because the security context of TAL_C uses static labels. Nonetheless, these static labels are only used to refer to code (e.g., that of linear continuations in ZM) whose locations can be statically determined. Indeed, their source level counterparts are the ending points of conditional structures, which are always statically known. Therefore, there is not a loss of expressiveness. We demonstrate this by speculating a translation from ZM to TAL_C.

In ZM, there are two expressions manipulating linear continuations: creation and elimination. The creation of a linear continuation essentially has the form `letlin y = λ(pc)(x : σ).e in e’`. A corresponding elimination has the form `lgoto y v`.

The translation can be carried out following Morrisett *et al.* [18]. The step of CPS conversion is not needed because ZM is already in CPS. During closure conversion, the abstraction `λ(pc)(x : σ).e` (which corresponds to the code at a postdominator) will be assigned a static code label. This code label is exactly the static postdominator needed for raising the security context in TAL_C. In a formal translation, this label can be used to generate a `raise` instruction when a corresponding branching point is reached. The typing (in particular, the security labels) of a ZM program is sufficient for identifying the branching point.

The elimination of linear continuation (`lgoto y v`) is relatively straightforward. Suppose the code of `y` (the lambda abstraction)

declared using `letlin` is assigned the heap label l_y during closure conversion, the elimination expression `lgoto y v` can be translated as a `lower ly` preceded with appropriate code computing the argument `v`.

For better understanding the relationship between linear continuations in ZM and security contexts in TAL_C, we further look into an example of nested `letlin` declarations:

```
letlin y1 = lv1 in letlin y2 = lv2 in e.
```

Once the second `letlin` is declared, the first linear continuation y_1 should be accessible only from inside lv_2 . Therefore, ZM requires that e type checks under y_2 , and lv_2 type checks under y_1 . This essentially enforces a stack discipline.

TAL_C has a similar mechanism. Suppose the current security context is $\theta_1 \triangleright l_1$ and the current instruction sequence is `raise θ2 ▷ l2; I`. The type system of TAL_C checks I under $\theta_2 \triangleright l_2$, and checks that the code type at l_2 respects $\theta_1 \triangleright l_1$. This enforces a similar stack discipline as in ZM; note that only the top stack element is apparent at any given time.

In summary, it is possible to conduct a formal translation from ZM to TAL_C. Nonetheless, we believe that the simple security-type system in Figure 1 of Section 2.2 is more accessible and better facilitates understanding. We leave a formal translation from ZM to TAL_C as future work.