

Transforming out Timing Leaks

Johan Agat

Department of Computing Science
Chalmers University of Technology and Göteborg University
agat@cs.chalmers.se

Abstract

One aspect of security in mobile code is privacy: private (or secret) data should not be leaked to unauthorised agents. Most of the work on secure information flow has until recently only been concerned with detecting direct and indirect flows. Secret information can however be leaked to the attacker also through covert channels. It is very reasonable to assume that the attacker, even as an external observer, can monitor the timing (including termination) behaviour of the program. Thus to claim a program secure, the security analysis must take also these into account.

In this work we present a surprisingly simple solution to the problem of detecting timing leakages to external observers. Our system consists of a type system in which well-typed programs do not leak secret information directly, indirectly or through timing, and a transformation for removing timing leakages. For any program that is well typed according to Volpano and Smith [VS97a], our transformation generates a program that is also free of timing leaks.

1 Introduction

As the use of Internet and mobile code increases, prevention of security leakages in multilevel secure systems becomes a concern for the everyday user. In a multilevel secure systems, data with different security levels are processed and computed. The security levels are partially ordered and normally form a lattice. To maintain confidentiality in such a system, secret data must not be leaked to unauthorised users or flow downwards in the security lattice. This work is motivated by the need for privacy of secret data when code from an untrusted source is downloaded and run on, typically, a naive users computer. Since virtually everybody stores data they consider private (or secret) on their computers, some kind of guarantee for the privacy of that data is needed when untrusted code is executed.

The scenario we imagine is the following: A user downloads a program from an untrusted web-site. To compute some information valuable to the user, the program needs

access to the users private data. What is problematic is that the program might also need to access some databases over the Internet in order to function. Is it safe for the user to run the program or could the private data be leaked?

A concrete example of such a program could be an automated financial advisor that needs the users private financial information as input and will fetch stock market rates, bank loan interest rates etc. over the Internet while computing financial advice for the user.

In this paper, we will deal with sequential programs and security leakage to an attacker external to the system running the untrusted code. Secret or private information may be leaked to the external attacker in a multitude of different ways:

- *Direct leakage*, the simplest and most blunt way of leaking information, is when the secret data is just passed to the attacker as is.
- *Indirect leakage* (sometimes called leakage through a *covert storage channel*), is when the secret data is encoded in the observable behaviour of the program. The program might for instance perform different kinds of database accesses depending on the value of some secret data.
- *Timing leakage* (or leakage through a *covert timing channel*) occurs when the program encodes the secret data in its timing behaviour or manipulates some shared resource in such a way that the attacker can observe the manipulation by measuring the availability of the resource during a particular time interval.
- Programs may also leak information through their termination behaviour. That is given that the secret data satisfies some condition the program may terminate abnormally or go into a infinite loop. We consider leaking through nontermination to be a special case of timing leakage.

Use of covert timing channels is maybe the most cunning way to leak data and perhaps also the most difficult to detect and prevent. Luckily, network latencies etc., introduce a lot of noise that must be dealt with by an implementation of a timing channel over the Internet. This means that the capacity of such a channel is probably quite low. However, timing channels cannot be discarded when reasoning about the security of a system dealing with sensitive information. Even if the capacity of the timing channel is as low as 2 bits per minute, leaking a 16 digit VISA card number can

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

POPL 2000 Boston MA USA

Copyright ACM 2000 1-58113-125-9/00/1...\$5.00

```

volvoValue := 0;
i := 1;
while (i <= DBsize) {
  let share := sharesDB[i].name in
  let value := lookupVal(share) * sharesDB[i].no in
  if (isVolvoShare(share))
    volvoValue := volvoValue + value;
  i := i + 1
}

```

Figure 1: A program with a timing leak.

```

volvoValue := 0;
i := 1;
while (i <= DBsize) {
  let share := sharesDB[i].name in
  let value := lookupVal(share) * sharesDB[i].no in
  if (isVolvoShare(share))
    volvoValue := volvoValue + value
  else
    skipAsn volvoValue (volvoValue + value);
  i := i + 1
}

```

Figure 2: A padded, secure version of the program.

be done in less than 30 minutes¹. Also, combined with the ability to make repeated attacks, timing can disclose disastrously sensitive information to the attacker. For example, Kocher showed [Koc96], that some implementations of the RSA encryption algorithm leak information about the encryption key through their timing behaviour. By making a series of encryptions and measuring their times, the attacker could figure out the entire key.

1.1 Example

To illustrate how timing leaks can arise and also how they can be closed, we consider two program fragments. Figure 1 presents a program that computes the total value of the user's Volvo shares into the secret variable `volvoValue`. The program loops through a database, `sharesDB`, implemented as an array of records with two secret components: a string `name` and an integer `no`. The length of the database is public. No secret data will be leaked to the public variable `i`, but since the amount of computation performed in the loop is dependent on whether `sharesDB[i].name` is a Volvo share, the program will leak this information through its timing behaviour.

This timing leak can be closed simply by padding the program with dummy computation, as shown in Figure 2. The command `skipAsn x e` takes the same time to execute as `x := e` but does not do any assignment. Assuming that the functions `lookupVal` and `isVolvoShare` both execute in constant time, the program is secure since the contents of `sharesDB` will not influence the execution time of the `while`-loop body.

¹One decimal digit requires roughly 3.3 bits to encode ($2^{3.3} \approx 10$).

1.2 Contribution

We present a type-system and a type-directed transformation that removes timing leaks from programs to make them secure with respect to a semantic security condition based on bisimulation. The security condition is very strong and the programs that are considered secure do not leak any secret information directly, indirectly, by termination behaviour or through covert timing channels to external observers. Programs that are well typed in our type-system satisfy this security condition. Our type-system improves over existing ones [VS97a], in that it allows more secure programs according to our condition. The generality of our type-system makes type checking undecidable, but we provide a simple, decidable, type-directed transformation that gives well-typed (secure) programs. The transformation removes timing leaks from programs without direct and indirect leaks, by padding with dummy instructions where needed. Together, our transformation and type-system provide the first realistic solution for closing timing leaks to external observers.

2 A Semantic Security Condition

Essential to analysing the security properties of a program is to have a semantic notion of security. The condition of security used must be strong enough to capture all important leakages possible in a realistic implementation of the system. In this section, we first point to problems with the realism and strength of the security condition used in some related work. We then discuss what a realistic security condition should capture, present the language we use and finally define our condition of security.

To simplify the presentation of our system we consider only two security classes: *L* for low security (public) and *H* for high security (secret). We thus have a two-point security lattice: $L \leq H$.

2.1 Related Work

The most commonly used semantic security condition is that of *noninterference* [GM82], which has been adopted in many recent papers on secure flow analysis [VS96, VS97b, HR98]. A program satisfies the noninterference property if its low security outputs do not depend on the high security inputs. This can be formulated as:

$$\forall E_1, E_2. E_1 =_{\text{low}} E_2 \Rightarrow P(E_1) =_{\text{low}} P(E_2)$$

where $E_1 =_{\text{low}} E_2$ means that if the environments E_1 and E_2 are defined, all their low-security components are equal. Here the program is seen as a function on environments.

With noninterference as the underlying notion of security, type-based analyses capable of detecting direct and indirect security leakages have been proposed by Volpano and Smith [VS97b] (a reformulation of Dennings work [Den76, DD77]), and Heintze and Riecke [HR98]. A big drawback with noninterference, however, is the extensional view of programs as functions from input to output. Thus a program that is secure in this sense can still leak information through its timing behaviour.

To close leakage through nontermination, Volpano and Smith enforce a condition of termination agreement, where both the termination behaviour and low output of a secure program are independent of the high inputs [VS97a]. Practically, this was done by disallowing looping conditions to

depend on high security data. The same restriction was enforced in [SV98] to prevent nontermination to be used as a method of leaking information between parallel processes in a multi-threaded language. By forcing the arguments to partial primitive operators, like division, to be of lowest security, leakages through abnormal termination can be closed [VS97a].

When it comes to detecting leakages through covert timing channels, Volpano and Smith have taken two different approaches [VS97a, VS98]. Neither of these deal with timing leakage to external observers in a feasible way. In [VS97a], a theorem on *timing agreement* is formulated. This theorem states that sequential programs with both looping- and branching conditions independent of high data will execute in lock-step with the low security part of the environment independent of the high security data. The consequences of disallowing also branching on high data will be discussed in Section 3. Although not a security condition in itself, the timing agreement theorem contains some of the necessary ingredients for a timing-aware security condition.

Internal timing leakages between concurrent threads arising from the probabilistic behaviour of the scheduler is studied in [VS98]. A type system and a semantics based on Markov chains of probabilistic states and stochastic transition matrices are presented, but no explicit security criterion is given. The system's security is based on the use of a 'protect *C*' statement that ensures atomic execution of the command *C*. Properly implemented and used on all if-commands branching on high data, the protect-statement can eliminate internal timing leakage. However, (quote) "*if external observation of the running program is allowed, then of course covert channels of the kind discussed [...] remain possible*".

The language JFlow, presented by Myers in [Mye99], is an extension of Java with annotations for secure information flow. The annotations can be checked mostly-statically but some tests need to be performed at runtime. Leakages through timing or nontermination are not detected in JFlow. As a way of escaping the restrictiveness of strict information flow, the system allows declassification of secret data, which makes formal reasoning about the security of the system hard. No semantic security condition or soundness theorem is given.

Banâtre, Bryce and Le Métayer [BBL94], present an information flow logic in which proofs of (potential) flows between variables are derived. The correctness criterion given is essentially noninterference on a variable by variable basis. The logic does not deal with external observers or flows resulting from timing or termination behaviour.

Methods of detecting potential timing channels when both the sending and receiving processes are available for analysis have been proposed by He and Gligor in [HG92]. This work is rather informal and does not contain any formal semantic condition of security.

A method described for example in [Heh84, Nie84], is that of using a *programmed counter* to measure execution time. The idea is to transform the program so that a special program variable introduced to record execution time is incremented after each statement. Rustan, Leino and Joshi [RLJ98], suggest using this approach to reason about covert flows involving timing behaviour. Introducing a low-security counter variable that is incremented after each command can certainly give a timing-aware semantic security condition but this approach works poorly in combination with

existing type system for security. The problem is that incrementing a low-security counter variable after each command conflicts with the requirement that low-security variables are not assigned in the branches of high if-commands. This means that unless the incrementation of the counter variable is given some special treatment, all high if-commands will seem to have indirect leaks. We have thus chosen to build a notion of execution time into the semantics, as presented in Section 2.3.

2.2 Requirements on Secure Programs

To arrive at a realistic and adequately strong security condition, we must analyse which kinds of information leakages should be prevented. In our setting, code is down-loaded from an untrusted source, given secret data as input and also allowed to access resources over Internet. The attacker in this setting is the author of the code and possibly in control of the site from which it is down-loaded. The attacker is also external to the system running the code. When the program runs, it can communicate with the attacker in two different ways:

Immediate communication, where data is sent directly to the attacker, that is the site from which the program was loaded. This is the only communication allowed by Java applets (see e.g. [DFWB97]).

Indirect communication, by use of a third party. The communication to the attacker is made by sending data to or manipulating some another site on the net in a way that the attacker can monitor.

Since the attacker is external to the system running the program, and thus cannot be controlled by that system, we must assume that the attacker can observe both *what* is communicated by the program and *when* this communication is made. From the attackers view, the execution of the program can be seen as a (possibly infinite) sequence of time, value pairs that correspond to observable actions and when they are made. Our semantic security condition must thus force this sequence to be independent of the high security data that the program manipulates.

Forcing the observable actions of a program to be independent of high data poses no new problems. This is just the noninterference property. To make sure also that the time at which these actions are performed is independent of high data, we must in some way reason about the execution time of programs. It is hard to determine how fine-grained this reasoning should be. Obviously, the number of instructions and which operations that are performed affect the execution time and so does the hardware (and interpreter) on which the program is run. The kind of instructions executed must probably also be considered, but even this might not be enough. A subtle way of leaking would still be available: the program can exhibit different cache, and thereby time, behaviour depending on some high data. Consider the following piece of C code:

```
if (h)
  for(n=0; n<N; n++) xs[n]++; /* loop 1 */
else
  for(n=0; n<N; n++) ys[n]++; /* loop 2 */
for(n=N; n>=0; n--) xs[n]++; /* loop 3 */
```

The same number and the same kind of instructions will be executed by both branches but if loop 1 is executed the

Operators	$op ::= + \mid * \mid - \mid = \mid != \mid < \mid <=$
Expressions	$e ::= l \mid e \ op \ e \mid le$
Initialisers	$ie ::= e \mid \text{mkarray}(e) \ ie \mid \{x_1 = ie_1, \dots, x_n = ie_n\}$
Commands	$C, D ::= le := e \mid \text{skipAsn } le \ e \mid \text{if } (e) \ C \ \text{else } D \mid \text{skipIf } e \ C \mid \text{let } x := ie \text{ in } C \mid \text{while } (e) \ C \mid C ; D \mid \text{output } x$
Left-expressions	$le ::= x \mid le.x \mid le[e]$
Left-values	$lv ::= x \mid lv.x \mid lv[n]$
Values	$v ::= l \mid \{x_1 = v_1, \dots, x_n = v_n\} \mid [v_0, \dots, v_n]$
Basic Values	$l ::= n \mid \text{true} \mid \text{false}$

Figure 3: Syntax of expressions, commands and values.

array xs will probably be cached when loop 3 is reached. It is thus likely that the entire execution will be longer if h is false². The variance in timing here might not be big but it is probably big enough to implement a timing channel! The code shown above is not secure with our criterion given in Definition 2, nor is it well-typed in our type-system. Our type system will only allow a high if-command when the two branches have exactly the same pattern of allocations and variable references. Thus the example above will be considered insecure even without loop 3.

To avoid building a limitation into the system, we have implicitly parameterised our semantics and security condition on the interpretation of time. This is discussed somewhat more in Sections 2.3 and 2.4.

2.3 The Language

The language we use, with syntax given in Figure 3, has assignments, sequencing, conditionals, local bindings, while loops, output commands and two kinds of skip-commands: $\text{skipAsn } le \ e$ and $\text{skipIf } e \ C$. The transformation described in Section 4 inserts these skip-commands to remove timing leakages. Values consists of records, arrays, integers and booleans. An expression cannot construct a record or array but this can be done by an initialising expression when a new binding is introduced. The output construct implements externally observable actions by outputting the value of an integer variable to the attacker.

To give the semantics of programs we use two partial functions: a big-step natural semantics for expressions and a labelled small-step transition semantics for commands. Semantic rules for expressions, initialisers and left-expressions are standard and given in Figure 4. Evaluation judgements for commands are of the forms

$$\begin{aligned} \langle E \mid C \rangle &\xrightarrow{as} \langle E \mid D \rangle \\ \langle E \mid C \rangle &\xrightarrow{ts, \checkmark} E \end{aligned}$$

where E is the environment, associating variables with values. The annotations on the transition arrow is a possibly empty sequence of time-expressions, ts , and for as also output actions. These describe the time and observable actions

of making that transition. Output actions, ranged over by o , are integer values, n and the symbol \checkmark which is used to signal termination. There is at most one output action on any given transition. We write $as_1 \cdot as_2$ for the concatenation of two sequences and we consider a single action or time-expression as a singleton sequence. We use ts to range over sequences consisting only of time-expressions, which are:

$t_e \ (t_{le})$	The time it takes to evaluate e (or le).
t_{asn}	Time for making an assignment.
t_{br}	Time to inspect and branch on a value.
$t_{pu \ ie}$	Time to evaluate ie and push a new binding of the resulting value on the evaluation environment.
t_{po}	Time to pop the innermost binding from the evaluation environment.

To capture the time of evaluating expressions, the time-expressions t and t_{pu} are indexed on expressions and initialisers. Since only basic values can be assigned (assignment of arrays- and record-values is not allowed) it is enough to use a constant t_{asn} for the time of an assignment.

Note that time-expressions do not specify time directly. They are only descriptions of time and have to be given some kind of interpretation if we need to know the real time they denote. Such an interpretation can be arbitrarily complex. By looking at the entire history of the execution, an interpretation could even model data-cache behaviour to some extent. The semantics is parameterised on the interpretation of time. Any interpretation can be used as long as primitive operations are given constant times. This is discussed in Section 2.4.

Rules for the evaluation of commands are given in Figure 5. The environment E is a sequence of bindings but for convenience we also treat it as a mapping, writing $E(x)$ to access the value in the rightmost (innermost) binding of x . We write $E[lv = l]$ for updating E at the binding and component indexed by lv to the basic value l and we define $\text{dom}(E)$ as $\{lv \mid \exists l. E(lv) = l\}$.

The two skip-constructs, skipAsn and skipIf , are designed to have the same timing behaviour as an assignment and an if-branching respectively. In the type-system described in Section 3, skipAsn and skipIf are used in describing timing behaviour. Programs can be padded with appropriate skip-commands as done by the transformation presented in Section 4, thereby removing timing leakages.

All other rules for evaluation of commands are fairly standard. Evaluation of a let introduces a new local binding. We have chosen to represent local bindings syntactically with a let-like construct called local , and they are moved to and from the environment in the rule (Local). We use the notation $E, x = v$ for adding a new binding to E and also to match out the innermost binding of an environment. The time for pushing a binding on the evaluation environment is paid in the (Let)-rule when the binding is first introduced and the time for popping the binding is paid when the body of the let terminates. There are no time-annotations for the cost of shuffling bindings in and out of the environment made in the (Local) rule. This is of course due to the fact that the local construct is nothing but a notational trick to keep track of local and global variables and it has no corresponding implementation overhead. Another technical trick we use is that the output command evaluates to a skip that then simply terminates, thus we avoid the problem of more than one output action on any given transition.

² A simple test based on the code above showed that a difference of 2-3 seconds in execution time could be measured from a program that ran in 1 minute. The test was made on a 296 MHz Sun Ultra 4.

$$\begin{array}{c}
\text{(Lit)} \frac{}{E \vdash l \Downarrow l} \quad \text{(Var)} \frac{E(x) = v}{E \vdash x \Downarrow v} \quad \text{(Op)} \frac{E \vdash e_1 \Downarrow l_1 \quad E \vdash e_2 \Downarrow l_2}{E \vdash e_1 \text{ op } e_2 \Downarrow l_1[l_2]} \quad \text{(Ix)} \frac{E \vdash e_1 \Downarrow [v_0, \dots, v_n, \dots, v_m] \quad E \vdash e_2 \Downarrow n}{E \vdash e_1[e_2] \Downarrow v_n} \\
\\
\text{(Sel)} \frac{E \vdash e \Downarrow \{\dots, x = v, \dots\}}{E \vdash e.x \Downarrow v} \quad \text{(MkA)} \frac{E \vdash e \Downarrow n \quad E \vdash ie \Downarrow v}{E \vdash \text{mkarray}(e) \text{ ie } \Downarrow [v, \dots, v] \text{ (} n \text{ occurrences)}} \quad \text{(L-Ix)} \frac{E \vdash le \Downarrow^L lw \quad E \vdash c \Downarrow n}{E \vdash le[e] \Downarrow^L lw[n]} \\
\\
\text{(MkRec)} \frac{E \vdash ie_1 \Downarrow v_1 \dots E \vdash ie_n \Downarrow v_n}{E \vdash \{x_1 = ie_1, \dots, x_n = ie_n\} \Downarrow \{x_1 = v_1, \dots, x_n = v_n\}} \quad \text{(L-Sel)} \frac{E \vdash le \Downarrow^L lw}{E \vdash le.x \Downarrow^L lw.x} \quad \text{(L-Var)} \frac{}{E \vdash x \Downarrow^L x}
\end{array}$$

Figure 4: Big-step semantics for expressions and initialisers (\Downarrow) and left-expressions (\Downarrow^L).

$$\begin{array}{c}
\text{(Assign)} \frac{E \vdash e \Downarrow l \quad E \vdash le \Downarrow^L lw}{\langle E | le := e \rangle \xrightarrow{t_e \cdot t_{le} \cdot t_{asn} \cdot \checkmark} E[lw = l]} \quad lw \in \text{dom}(E) \quad \text{(Let)} \frac{E \vdash ie \Downarrow v}{\langle E | \text{let } x := ie \text{ in } C \rangle \xrightarrow{t_{pu} \cdot t_e} \langle E | \text{local } x := v \text{ in } C \rangle} \\
\\
\text{(Seq)} \frac{\langle E | C \rangle \xrightarrow{ts \cdot \checkmark} E' \quad \langle E | C; D \rangle \xrightarrow{ts} \langle E' | D \rangle}{\langle E | C \rangle \xrightarrow{as} \langle E' | C' \rangle} \quad \text{(Local)} \frac{\langle E, x = v | C \rangle \xrightarrow{as} \langle E', x = v' | D \rangle}{\langle E | \text{local } x := v \text{ in } C \rangle \xrightarrow{as} \langle E' | \text{local } x := v' \text{ in } D \rangle} \\
\frac{\langle E | C \rangle \xrightarrow{as} \langle E' | C' \rangle}{\langle E | C; D \rangle \xrightarrow{as} \langle E' | C'; D \rangle} \quad \frac{\langle E, x = v | C \rangle \xrightarrow{ts \cdot \checkmark} E', x = v'}{\langle E | \text{local } x := v \text{ in } C \rangle \xrightarrow{ts \cdot t_{po} \cdot \checkmark} E'} \\
\\
\text{(If)} \frac{E \vdash e \Downarrow \text{true}}{\langle E | \text{if } (e) \text{ } C \text{ else } D \rangle \xrightarrow{t_e \cdot t_{br}} \langle E | C \rangle} \quad \text{(While)} \frac{E \vdash e \Downarrow \text{false}}{\langle E | \text{while } (e) \text{ } C \rangle \xrightarrow{t_e \cdot t_{br} \cdot \checkmark} E} \\
\frac{E \vdash e \Downarrow \text{false}}{\langle E | \text{if } (e) \text{ } C \text{ else } D \rangle \xrightarrow{t_e \cdot t_{br}} \langle E | D \rangle} \quad \frac{E \vdash e \Downarrow \text{true}}{\langle E | \text{while } (e) \text{ } C \rangle \xrightarrow{t_e \cdot t_{br}} \langle E | C; \text{while } (e) \text{ } C \rangle} \\
\\
\text{(SkipAsn)} \frac{E \vdash e \Downarrow l \quad E \vdash le \Downarrow^L lw}{\langle E | \text{skipAsn } le \text{ } e \rangle \xrightarrow{t_e \cdot t_{le} \cdot t_{asn} \cdot \checkmark} E} \quad lw \in \text{dom}(E) \quad \text{(SkipIf)} \frac{E \vdash e \Downarrow l}{\langle E | \text{skipIf } e \text{ } C \rangle \xrightarrow{t_e \cdot t_{br}} \langle E | C \rangle} \quad l \in \{\text{true}, \text{false}\} \\
\\
\text{(Skip)} \frac{}{\langle E | \text{skip} \rangle \xrightarrow{\checkmark} E} \quad \text{(Output)} \frac{E(x) = n}{\langle E | \text{output } x \rangle \xrightarrow{t_x \cdot n} \langle \text{skip} | E \rangle}
\end{array}$$

Figure 5: Small-step semantics for commands.

2.4 A Bisimulation-based Security Condition

As a foundation for our semantic security condition, we use a partial bisimulation on commands, \sim_Γ . We index this relation on a typing environment, Γ , to distinguish low- and high-security left-values. Informally, two commands are Γ -bisimilar if they behave stepwise identically with respect to execution time, outputs and the manipulation of the low variables (according to Γ). Also, this behaviour must be independent of the high data in the environment.

Definition 1 (Γ -bisimulation)

\sim_Γ is the largest symmetric relation on commands that satisfies:

$C_1 \sim_\Gamma C_2$ if $\forall E_1, E_2$ such that $E_1 =_\Gamma E_2$ we have that

$$\begin{aligned}
\langle E_1 | C_1 \rangle \xrightarrow{as} \langle E'_1 | D_1 \rangle &\Rightarrow \langle E_2 | C_2 \rangle \xrightarrow{as} \langle E'_2 | D_2 \rangle \wedge \\
&\quad E'_1 =_\Gamma E'_2 \wedge D_1 \sim_\Gamma D_2 \\
\langle E_1 | C_1 \rangle \xrightarrow{ts \cdot \checkmark} E'_1 &\Rightarrow \langle E_2 | C_2 \rangle \xrightarrow{ts \cdot \checkmark} E'_2 \wedge E'_1 =_\Gamma E'_2
\end{aligned}$$

where $=_\Gamma$ is defined as

$$E_1 =_\Gamma E_2 \text{ iff } \text{dom}(E_1) = \text{dom}(E_2) \wedge \text{dom}(E_1) \subseteq \text{dom}(\Gamma) \wedge \forall lv \in \text{dom}(E_1). \Gamma(lv) = \bar{\tau}_L \Rightarrow E_1(lv) = E_2(lv)$$

Here, Γ is a typing environment and $\bar{\tau}_L$ is a base type of low security level. Both will be introduced in Section 3.1

By relating commands to themselves by Γ -bisimulation we get a sort of timing-aware, stepwise noninterference property, which suits us well as the definition of security:

Definition 2 (Γ -security)

C is Γ -secure if $C \sim_\Gamma C$

This security condition is semantically a lot stronger than actually necessary, due to the definition of Γ -bisimulation. For example, due to the statelessness of \sim_Γ the program

$h := 0; 1 := h$ is not secure, although it would not leak any secret information in our sequential setting. Weaker and more elaborate versions of the bisimulation could be defined. The relation could be made stateful by dropping the local quantification of the environments and relate configurations instead of commands. Moreover, as it stands, \sim_Γ , requires the evaluation of $\langle E_1 | C_1 \rangle$ and $\langle E_2 | C_2 \rangle$ to both make a step with syntactically equal time-expression sequences, which essentially forces them to make the same computations (but not the same assignments!). This condition could be relaxed by making the interpretation of time-expression sequences explicit and by forcing the two commands to evaluate in times that have the same interpretation, possibly involving the entire history of the two executions. Any weaker version of the bisimulation would clearly be implied by \sim_Γ . To avoid any unnecessarily complicated definition and since the type-system and transformation presented in Sections 3 and 4 will be sound with respect to \sim_Γ , we have chosen this, stronger definition. Also, the statelessness of \sim_Γ gives a composable security criterion and extends nicely to satisfy the Hook-up property in parallel programs [McC87], as shown by Sands and Sabelfeld [SS99].

It is of course very hard to statically predict cache behaviour. However by forcing two Γ -bisimilar commands to have syntactically equal time-expression sequences, the two commands must obey the same pattern of allocations and variable references. This severely reduces the possibilities for timing leaks based on cache behaviour.

The only assumption on the interpretation of time inherent in the definition of \sim_Γ , is that a time-expression sequence ts must denote the same time regardless of the values bound to the free high-security variables in ts . This means that primitive operations must be performed in constant time, independent of the argument values of the operator, and that the length of an array cannot be secret.

It might be worth pointing out that \sim_Γ is only a partial bisimulation in that it is *not* reflexive. Insecure commands, like $1 := h$, are not related to themselves, which is also the reason that \sim_Γ is not a congruence. Γ -bisimulation thus differs from ordinary bisimulation by not being an equivalence relation, but rather a partial such:

Lemma 1 (\sim_Γ is a PER)

The relation \sim_Γ is a Partial Equivalence Relation, i.e. it is symmetric, transitive but not necessarily reflexive.

Proof: \sim_Γ is symmetric by definition, and transitivity is easy to show since $=_\Gamma$ is transitive.

Due to the symmetry and transitivity of \sim_Γ , if a command is Γ -bisimilar to anything, it is also Γ -bisimilar to itself, and hence Γ -secure.

2.5 Example of a Secure Program

Even though Γ -bisimulation is a very strong security criterion, it does not rule out useful programs. If we informally extend the language with functions and strings, the padded version of the Volvo shares example, presented in Figure 2, is Γ -secure. Looping through the database does not leak since the length of the database is public. Branching on the (secret) value of `isVolvoShare(share)` is also secure since the two branches are Γ -bisimilar. Of course, the example is only Γ -secure provided that the two functions `lookupVal` and `isVolvoShare` are both Γ -secure.

$$\frac{\frac{s_1 \leq s_2}{\bar{\tau}_{s_1} \leq \bar{\tau}_{s_2}} \quad \frac{\tau_1 \leq \tau_2}{A \tau_1 \leq A \tau_2}}{\frac{\tau_1 \leq \tau'_1 \dots \tau_n \leq \tau'_n}{\{x_1 : \tau_1, \dots, x_n : \tau_n\} \leq \{x_1 : \tau'_1, \dots, x_n : \tau'_n\}}}$$

Figure 6: The subtyping relation

3 Typing Secure Programs

In [VS97a], Volpano and Smith present a type-system in which well-typed programs are secure with respect to timing leaks to external observers. By a simple inductive reasoning, it can be shown that programs typed according to section 5 of [VS97a] are also Γ -secure. Our critique to this type-system is that it is too restrictive to be of any practical use: it requires the condition of both while-loops and if-commands to be of lowest security. With this restriction, high data can only be copied around and passed to (total) primitive operators, which means that only programs that are essentially parametric in their high inputs are considered secure by the type-system. Our observation to relax this restrictiveness is the following:

A secure program may safely branch on high data as long as the external observer cannot determine which branch was taken.

Our type-system closes timing and termination channels in well-typed programs by requiring looping conditions to be of low security and branches of high if-commands to have the same externally observable behaviour.

3.1 The Type System

We use the following language of types:

Security levels $s ::= L \mid H$ (with $L \leq H$ and $s \leq s$)
 Base types $\bar{\tau} ::= \text{Int} \mid \text{Bool}$
 Security types $\tau ::= \bar{\tau}_s \mid A \tau \mid \{x_1 : \tau_1, \dots, x_n : \tau_n\}$

Security types, ranged over by τ , are base types annotated with a security level, array-, or record types. The components of records can all have different security types and thereby security levels. The type for arrays contains the type for the elements in the array, but there is no security level for the length of the array, since this must always be low. Allowing high values to specify the length of an array will open up for timing leakages of the same kind as loops on high data, since the array is initialised upon creation. Also, array update and indexing are non-total operations, so we cannot allow high-security values to be used in specifying the index in these operations. This problem, and our treatment of it, is analogous to that with the division operator in [VS97a]. In Figure 6, we extend the ordering on security levels to a subtyping relation.

The typing of expressions calculate an upper bound on the security levels of the variables in the expression. This typing has no time component since expressions evaluate atomically and are side-effect free in our semantics. Typing rules for expressions, initialisers and left-expressions are

$$\begin{array}{c}
\text{(LitInt)} \frac{}{\Gamma \vdash_{\triangleleft} n : \text{Int}_s} \quad \text{(LitBool)} \frac{}{\Gamma \vdash_{\triangleleft} \{\text{true}, \text{false}\} : \text{Bool}_s} \quad \text{(Var)} \frac{\Gamma(x) \triangleleft \tau}{\Gamma \vdash_{\triangleleft} x : \tau} \quad \text{(Op-Arithm)} \frac{\Gamma \vdash_{\leq} e_1 : \text{Int}_s \quad \Gamma \vdash_{\leq} e_2 : \text{Int}_s}{\Gamma \vdash_{\triangleleft} e_1 \{+, -, *\} e_2 : \text{Int}_s} \\
\\
\text{(Ix)} \frac{\Gamma \vdash_{\triangleleft} e_1 : A \tau \quad \Gamma \vdash_{\leq} e_2 : \text{Int}_L}{\Gamma \vdash_{\triangleleft} e_1[e_2] : \tau} \quad \text{(Sel)} \frac{\Gamma \vdash_{\triangleleft} e : \{\dots, x : \tau, \dots\}}{\Gamma \vdash_{\triangleleft} e.x : \tau} \quad \text{(Op-Cmp)} \frac{\Gamma \vdash_{\leq} e_1 : \text{Int}_s \quad \Gamma \vdash_{\leq} e_2 : \text{Int}_s}{\Gamma \vdash_{\triangleleft} e_1 \{=, <=, <\} e_2 : \text{Bool}_s} \\
\\
\text{(MkA)} \frac{\Gamma \vdash_{\leq} e : \text{Int}_L \quad \Gamma \vdash_{\leq} ie : \tau}{\Gamma \vdash_{\triangleleft} \text{mkarray}(e) \text{ } ie : A \tau} \quad \text{(MkRec)} \frac{\Gamma \vdash_{\leq} ie_1 : \tau_1 \dots \Gamma \vdash_{\leq} ie_n : \tau_n}{\Gamma \vdash_{\triangleleft} \{x_1 = ie_1, \dots, x_n = ie_n\} : \{x_1 : \tau_1, \dots, x_n : \tau_n\}}
\end{array}$$

Figure 7: Typing rules for expressions, initialisers and left-expressions. (\triangleleft ranges over $=$ and \leq)

$$\begin{array}{c}
\text{(Assign}_H\text{)} \frac{\Gamma \vdash_{\leq} e : \bar{\tau}_s \quad \Gamma \vdash_{= } le : \bar{\tau}_H \quad s \leq H}{\Gamma \vdash le := e : \text{skipAsn } le \text{ } e} \quad \text{(SkipAsn)} \frac{}{\Gamma \vdash \text{skipAsn } le \text{ } e : \text{skipAsn } le \text{ } e} \\
\\
\text{(Assign}_L\text{)} \frac{\Gamma \vdash_{\leq} e : \bar{\tau}_L \quad \Gamma \vdash_{= } le : \bar{\tau}_L}{\Gamma \vdash le := e : le := e} \quad \text{(SkipIf)} \frac{\Gamma \vdash C : C_L}{\Gamma \vdash \text{skipIf } e \text{ } C : \text{skipIf } e \text{ } C_L} \\
\\
\text{(Seq)} \frac{\Gamma \vdash C : C_L \quad \Gamma \vdash D : D_L}{\Gamma \vdash C; D : C_L; D_L} \quad \text{(Let)} \frac{\Gamma \vdash_{\leq} ie : \tau \quad \Gamma, x : \tau \vdash C : C_L}{\Gamma \vdash \text{let } x := ie \text{ in } C : \text{let } x := ie \text{ in } C_L} \\
\\
\text{(If}_H\text{)} \frac{\Gamma \vdash_{\leq} e : \text{Bool}_H \quad \Gamma \vdash C : C_L \quad \Gamma \vdash D : D_L}{\Gamma \vdash \text{if } (e) \text{ } C \text{ else } D : \text{skipIf } e \text{ } C_L} \quad C_L \sim_{\Gamma} D_L \quad \text{(While)} \frac{\Gamma \vdash_{\leq} e : \text{Bool}_L \quad \Gamma \vdash C : C_L}{\Gamma \vdash \text{while } (e) \text{ } C : \text{while } (e) \text{ } C_L} \\
\\
\text{(If}_L\text{)} \frac{\Gamma \vdash_{\leq} e : \text{Bool}_L \quad \Gamma \vdash C : C_L \quad \Gamma \vdash D : D_L}{\Gamma \vdash \text{if } (e) \text{ } C \text{ else } D : \text{if } (e) \text{ } C_L \text{ else } D_L} \quad \text{(Output)} \frac{\Gamma(x) = \text{Int}_L}{\Gamma \vdash \text{output } x : \text{output } x}
\end{array}$$

Figure 8: Typing rules for commands

given in Figure 7. These rules are parameterised on whether sub typing is allowed for variables or not, using \triangleleft to range over $=$ and \leq . The Assign rules in Figure 8 use the parameterisation to prevent sub typing from being used in typing left-expressions. Apart from using a richer type language, our way of typing expressions is essentially identical to that of Volpano and Smith [VS97b].

Typing environments, mapping variables to security types, are ranged over by Γ and we write $\Gamma, x : \tau$ for the environment that maps x to τ and otherwise behaves like Γ . We write $\Gamma(x) = \tau$ if x is mapped to τ by Γ , and $\Gamma(x) \leq \tau$ if there exists some τ' such that $\Gamma(x) = \tau'$ and $\tau' \leq \tau$. We write $\Gamma(lv) = \tau$ as a shorter notation for $\Gamma \vdash_{= } lv : \tau$ and we define $\text{dom}(\Gamma)$ as $\{lv \mid \exists \tau. \Gamma(lv) = \tau\}$.

Typing rules for commands are given in Figure 8. The typing judgements are of the form:

$$\Gamma \vdash C : C_L$$

The “type” of a command C is its *low-slice*, C_L . The low-slice is syntactically identical to C but only contains assignments to low security left-values. All assignments to high security left-values and branching on high security data are replaced with the appropriate skips. For example, with the environment $\Gamma = \{h : \text{Int}_H, l : \text{Int}_L\}$, we can derive:

$$\Gamma \vdash (h := h+1; l := l*4) : (\text{skipAsn } h \text{ } (h+1); l := l*4)$$

As stated in Lemma 2, the low-slice has the same observable

behaviour as the original command with respect to low left-values. Note that although there are similarities, the low-slice is not a program slice in the conventional sense (see e.g. [RT96]). The low-slice and the type-systems construction of it has more in common with the extraction of static program parts as described by Mogensen in [Mog89].

The typing rules are all rather straightforward. Direct leakage is prevented by the Assign_L-rule. By forcing the looping condition to be of low security, the While-rule prevents leakage through nontermination and also blocks a simple kind of timing leaks. The most interesting rule is If_H, which stops both indirect- and timing leaks. It allows branching on high data provided that the branches have Γ -bisimilar low-slices and thus have the same externally observable behaviour. The low slice is $\text{skipIf } e \text{ } C_L$ but could just as well have been $\text{skipIf } e \text{ } D_L$ or $\text{if } (e) \text{ } C_L \text{ else } D_L$ instead, since they are all Γ -bisimilar. By using $\text{skipIf } e \text{ } C_L$ instead of $\text{if } (e) \text{ } C_L \text{ else } D_L$ in the transformation described in Section 4, we avoid getting an exponential blow-up in code size.

3.2 Usefulness in Practice

Our system guarantees that well-typed programs do not leak information to an external observer even through timing behaviour. Unlike the system described in [VS97a], ours is not too restrictive to be used in practice since we do not disallow well-typed programs to compute with, and branch on secret

data. For example, the program in Figure 2 is well-typed in our system but not in [VS97a].

The side condition $C_L \sim_\Gamma D_L$ in the If_H -rule makes the type system undecidable since \sim_Γ is undecidable. This does not prevent sound but incomplete and conservative type checking algorithms to be made though. Given an $O(n)$ method of computing $C_L \sim_\Gamma D_L$, type checking will not be harder than normal [VS97b].

It can be argued that no programmers write programs where all if-commands branching on high data have Γ -bisimilar branches. As much as this is true, it is also necessary to produce such programs if we want to avoid timing leaks. To reduce the burden for the programmer, the next section presents a transformation that removes timing leaks by making branches of high if-commands Γ -bisimilar.

3.3 Soundness of the Type System

The type system described in Figure 8 is sound with respect to Γ -bisimulation. The soundness theorem is a corollary of the following lemma, stating that the low-slice of a command has the same observable behaviour as the command itself.

Lemma 2 (Typing and Γ -bisimulation)

If $\Gamma \vdash C : C_L$ then $C \sim_\Gamma C_L$

A detailed proof of this lemma is given in Appendix A.

Theorem 1 (Well-typed programs are secure)

If $\Gamma \vdash C : C_L$ then C is Γ -secure.

Proof: Follows from Lemma 2 and the symmetry and transitivity of \sim_Γ according to Lemma 1.

4 Transforming out Timing Leaks

Paying with performance, we can transform out the timing leakages of programs by padding the branches of high if-commands with dummy instructions so that they get the same timing and otherwise externally observable behaviour. In Figure 9 we give an inductively defined algorithm that transforms out timing leakages. The algorithm subsumes the type-system and works by recursively performing a crosswise padding of each branch of a high if-commands with the low-slice of the other branch. Thereby producing a program where the low-slices of the branches of high if-commands are Γ -bisimilar. Transformation judgements are of the form:

$$\Gamma \vdash C \hookrightarrow D \mid D_L$$

In a given type environment, the command C is transformed to an almost semantically equivalent Γ -secure command, D , and to the low-slice, D_L , of this command. The commands D and C are semantically equivalent in the sense that they will perform the same sequence of outputs and assignments to global variables given identical environments to start the evaluation in. They differ in that D might need more evaluation steps than C to do a particular assignment. Also, D might go into a nonterminating, non-productive loop even if C does not. We argue that this is acceptable since the extra nontermination in D will be due to potential information leaks by nontermination in C . We state the semantic soundness of the transformation formally in Theorem 3.

The crosswise padding with low-slices is done in the If_H -rule. The function $\text{ge}(C)$, inductively defined in Figure 10, is used to ensure that the low-slices of the transformed branches are free from outputs and assignments to other than variables local to the branches. Thus, the low-slices used for padding will add only time, not observable actions, to the computation – a requirement which is essential to the semantic soundness of the transformation. The premises $\text{ge}(D_{1L}) = \emptyset$ and $\text{ge}(D_{2L}) = \emptyset$ essentially requires both branches to be without output actions and assignments to low security variables bound outside the branches. This requirement stops programs with indirect leaks and is but a variant of that made in [VS97b], where only assignments to high variables are allowed in the branches of a high if-command. Since assignments to low variables locally let-bound in the branches is not restricted, they may contain arbitrary loop-structures (which then also will be present in the low-slices of the branches). The extra nontermination that may be introduced by the transformation is due to non-terminating loops in one branch being copied over into the other.

4.1 The Cost in Performance and Code Size

Performing a crosswise copying of the low-slices is a simple but not very refined way of making the branches Γ -bisimilar. Ideally, we would like to pad each branch with the difference to the maximum execution time, rather than taking the sum of the execution times of the two branches. Computing the difference is undecidable in general but a more refined method of padding than cross copying can certainly be defined.

To try to give some kind of formal argument about the slow-down of the transformed program relative the original one is pointless. The best we can say is that the transformed program can take arbitrarily longer time to execute, since even code that was semantically dead in the original program may be executed in the transformed one. For a tree-structure of nested high if-commands, the original program would execute only one path in the tree whereas the transformed program will execute *all* paths. As discussed in the previous section, this can lead to nontermination of the transformed program.

One can easily be misled to think³ that a combination of nested high if-commands and the cross copying made in the If_H rule of the transformation would lead to exponential blow-up in code size. However, the worst case blow-up in code size from cross copying is actually only by a factor linear in the nesting depth of high if-commands. What saves us from exponential blow-up in code size is that the `skipIf` introduced as low-slice of a transformed high if has roughly half the size of a corresponding if-command with Γ -bisimilar branches. To reason formally about the code size blow-up, we begin by defining a simple measurement of code size.

Definition 3 (Code size)

Define $\#C$ as:

$$\begin{aligned} \#(\text{if } (e) C \text{ else } D) &= 1 + \#C + \#D & \#(le := e) &= 1 \\ \#(C; D) &= \#C + \#D & \#(\text{while } (e) C) &= 1 + \#C \\ \#(\text{skipIf } e C) &= 1 + \#C & \#(\text{let } x := ie \text{ in } C) &= 1 + \#C \\ \#(\text{skipAsn } le e) &= 1 & \#(\text{output } x) &= 1 \end{aligned}$$

³As the author was in an earlier draft.

(Assign _H)	$\frac{\Gamma \vdash_{\leq} e : \bar{\tau}_s \quad \Gamma \vdash_{=} le : \bar{\tau}_H \quad s \leq H}{\Gamma \vdash le := e \hookrightarrow le := e \mid \text{skipAsn } le \ e}$
(Assign _L)	$\frac{\Gamma \vdash_{\leq} e : \bar{\tau}_L \quad \Gamma \vdash_{=} le : \bar{\tau}_L}{\Gamma \vdash le := e \hookrightarrow le := e \mid le := e}$
(Seq)	$\frac{\Gamma \vdash C_1 \hookrightarrow D_1 \mid D_{1L} \quad \Gamma \vdash C_2 \hookrightarrow D_2 \mid D_{2L}}{\Gamma \vdash C_1; C_2 \hookrightarrow D_1; D_2 \mid D_{1L}; D_{2L}}$
(If _H)	$\frac{\Gamma \vdash_{\leq} e : \text{Bool}_H \quad \Gamma \vdash C_1 \hookrightarrow D_1 \mid D_{1L} \quad \Gamma \vdash C_2 \hookrightarrow D_2 \mid D_{2L} \quad \text{ge}(D_{1L}) = \emptyset \quad \text{ge}(D_{2L}) = \emptyset}{\Gamma \vdash \text{if } (e) \ C_1 \ \text{else } C_2 \hookrightarrow \text{if } (e) \ D_1; D_{2L} \ \text{else } D_{1L}; D_2 \mid \text{skipIf } e \ (D_{1L}; D_{2L})}$
(If _L)	$\frac{\Gamma \vdash_{\leq} e : \text{Bool}_L \quad \Gamma \vdash C_1 \hookrightarrow D_1 \mid D_{1L} \quad \Gamma \vdash C_2 \hookrightarrow D_2 \mid D_{2L}}{\Gamma \vdash \text{if } (e) \ C_1 \ \text{else } C_2 \hookrightarrow \text{if } (e) \ D_1 \ \text{else } D_2 \mid \text{if } (e) \ D_{1L} \ \text{else } D_{2L}}$
(SkipAsn)	$\frac{}{\Gamma \vdash \text{skipAsn } le \ e \hookrightarrow \text{skipAsn } le \ e \mid \text{skipAsn } le \ e}$
(SkipIf)	$\frac{\Gamma \vdash C \hookrightarrow D \mid D_L}{\Gamma \vdash \text{skipIf } e \ C \hookrightarrow \text{skipIf } e \ D \mid \text{skipIf } e \ D_L}$
(Let)	$\frac{\Gamma \vdash_{\leq} ie : \tau \quad \Gamma, x : \tau \vdash C \hookrightarrow D \mid D_L}{\Gamma \vdash \text{let } x := ie \text{ in } C \hookrightarrow \text{let } x := ie \text{ in } D \mid \text{let } x := ie \text{ in } D_L}$
(While)	$\frac{\Gamma \vdash_{\leq} e : \text{Bool}_L \quad \Gamma \vdash C \hookrightarrow D \mid D_L}{\Gamma \vdash \text{while } (e) \ C \hookrightarrow \text{while } (e) \ D \mid \text{while } (e) \ D_L}$
(Output)	$\frac{\Gamma \vdash_{\leq} x : \text{Int}_L}{\Gamma \vdash \text{output } x \hookrightarrow \text{output } x \mid \text{output } x}$

Figure 9: An algorithm for transforming out timing leakages

$\frac{}{\text{ge}(x) = \{x\}}$	$\frac{\text{ge}(le) = \sigma}{\text{ge}(le.x) = \sigma}$	$\frac{\text{ge}(le) = \sigma}{\text{ge}(le[e]) = \sigma}$	$\frac{\text{ge}(le) = \sigma}{\text{ge}(le := e) = \sigma}$
$\frac{\text{ge}(C_1) = \sigma_1 \quad \text{ge}(C_2) = \sigma_2}{\text{ge}(C_1; C_2) = \sigma_1 \cup \sigma_2}$	$\frac{\text{ge}(C_1) = \sigma_1 \quad \text{ge}(C_2) = \sigma_2}{\text{ge}(\text{if } (e) \ C_1 \ \text{else } C_2) = \sigma_1 \cup \sigma_2}$	$\frac{\text{ge}(C) = \sigma}{\text{ge}(\text{let } x := ie \text{ in } C) = \sigma \setminus \{x\}}$	
$\frac{\text{ge}(C) = \sigma}{\text{ge}(\text{while } (e) \ C) = \sigma}$	$\frac{}{\text{ge}(\text{skipAsn } le \ e) = \emptyset}$	$\frac{\text{ge}(C) = \sigma}{\text{ge}(\text{skipIf } e \ C) = \sigma}$	$\frac{}{\text{ge}(\text{output } x) = \{\bullet\}}$

Figure 10: $\text{ge}(C)$ – the global effects of C – free assigned variables and indication of output actions.

Our first observation is that the low slice computed by the transformation will always have the same size as the original command.

Lemma 3 (Size of computed low-slice)

If $\Gamma \vdash C \hookrightarrow D \mid D_L$ then $\#C = \#D_L$.

Proof: Simple by induction on the height of the transformation derivation.

We will write $\text{nd}_\Gamma(C)$ for the maximal nesting depth of high (according to Γ) if-commands in C .

Theorem 2 (Code size blow-up wrt. nesting depth)

If $\Gamma \vdash C \hookrightarrow D \mid D_L$ and $\text{nd}_\Gamma(C) \leq n$ then $\#D \leq (n+1)\#C$.

Proof: By induction on the height of the transformation derivation. Cases Assign_H , Assign_L , SkipAsn and Output are all trivial. Cases Seq , If_L , SkipIf , Let and While all follow immediately from the induction hypothesis and some simple arithmetic. In case If_H we have that C is of the form $\text{if}(e) C_1 \text{ else } C_2$ and $D = \text{if}(e) D_1; D_{2L} \text{ else } D_{1L}; D_2$ where $\Gamma \vdash C_1 \hookrightarrow D_1 \mid D_{1L}$ and $\Gamma \vdash C_2 \hookrightarrow D_2 \mid D_{2L}$. From the assumption that $\text{nd}_\Gamma(C) \leq n$ we know that $\text{nd}_\Gamma(C_1) \leq n-1$ and $\text{nd}_\Gamma(C_2) \leq n-1$. Now from Definition 3 we have

$$\begin{aligned} \#D &= 1 + \#D_1 + \#D_{2L} + \#D_{1L} + \#D_2 \\ &\leq 1 + \#D_{1L} + \#D_{2L} + n\#C_1 + n\#C_2 \quad (\text{by ind. hyp.}) \\ &\leq 1 + \#C_1 + \#C_2 + n(\#C_1 + \#C_2) \quad (\text{by Lemma 3}) \\ &\leq (n+1)\#C \quad (\text{by Definition 3}) \end{aligned}$$

In the worst case scenario, the high if-commands in a command C are nested to a depth which is linear in the size of C . Such nesting occurs for example in a multiple choice implemented by nesting if-commands in the else-branch and the transformation will in that case blow up C to a command with size in the order of $O(\#C^2)$. To avoid such quadratic blow-up, deeply nested if-commands can often be flattened. Consider the following program:

```
if (e1) C1
else if (e2) C2
...
else Cn
```

Transforming this program will give a quadratic blow-up in code size, given that all C_i 's are relatively small. By introducing a fresh high security variable x we can flatten the program as follows:

```
let x := true
in if (e1 && x) (C1; x := false);
   if (e2 && x) (C2; x := false);
   ...
   if (x) Cn;
```

The size of this program is linear in the size of the original one and the transformation will only blow it up to roughly the double of its size. The amount of computation made in the transformed versions of these two programs is essentially the same. In both cases, one transformed C_i and the low-slices of all transformed C_j , where $i \neq j$, will be executed.

4.2 Semantic Soundness of the Transformation Algorithm

To reason about the semantic soundness of the transformation algorithm, we abstract from the evaluation time. We first introduce a multiple step evaluation \rightarrow .

Definition 4 (Multiple step evaluation)

We define \rightarrow inductively as:

$$\begin{array}{c} \frac{\langle E \mid C \rangle \xrightarrow{ts, \checkmark} E'}{\langle E \mid C \rangle \xrightarrow{ts, \checkmark} E'} \quad \frac{\langle E \mid C \rangle \xrightarrow{as} \langle E' \mid D \rangle}{\langle E \mid C \rangle \xrightarrow{as} \langle E' \mid D \rangle} \\ \frac{\langle E \mid C \rangle \xrightarrow{as} \langle E' \mid D \rangle \quad \langle E' \mid D \rangle \xrightarrow{as', \checkmark} E''}{\langle E \mid C \rangle \xrightarrow{as \cdot as', \checkmark} E''} \\ \frac{\langle E \mid C \rangle \xrightarrow{as} \langle E' \mid C' \rangle \quad \langle E' \mid C' \rangle \xrightarrow{as'} \langle E'' \mid C'' \rangle}{\langle E \mid C \rangle \xrightarrow{as \cdot as'} \langle E'' \mid C'' \rangle} \end{array}$$

To reason about observable evaluation steps we define \xRightarrow{o} to be a multiple step evaluation with exactly one output action, namely o .

Definition 5 (Observable evaluation)

Define \xRightarrow{o} as:

$$\begin{aligned} \langle E \mid C \rangle &\xRightarrow{o} \langle E' \mid D \rangle \quad \text{iff } \exists ts. \langle E \mid C \rangle \xrightarrow{ts, o} \langle E' \mid D \rangle \\ \langle E \mid C \rangle &\xRightarrow{\checkmark} E' \quad \text{iff } \exists ts. \langle E \mid C \rangle \xrightarrow{ts, \checkmark} E' \end{aligned}$$

Definition 6 (Weak simulation)

Define \sqsubseteq as the largest relation that satisfies

$$\begin{aligned} C &\sqsubseteq D \text{ if } \forall E. \\ \langle E \mid C \rangle &\xRightarrow{o} \langle E' \mid C' \rangle \Rightarrow \langle E \mid D \rangle \xRightarrow{o} \langle E' \mid D' \rangle \wedge C' \sqsubseteq D' \\ \langle E \mid C \rangle &\xRightarrow{\checkmark} E' \Rightarrow \langle E \mid D \rangle \xRightarrow{\checkmark} E' \end{aligned}$$

It is easy to show that \sqsubseteq is reflexive, transitive and that it is preserved by contexts (i.e. that it is a precongruence).

Our statement of semantic soundness of the transformation is that the transformed program, D , will be simulated by the original program: $D \sqsubseteq C$. The transformation is not complete so $C \sqsubseteq D$ does not hold in general due to the extra nontermination that might be introduced into D . Completeness could be achieved if the transformation was assisted by some kind of termination analysis, so that the algorithm would fail rather than introduce extra nontermination. In Section 4.4, we prove the transformation complete for a certain class of programs.

To prove semantic soundness, we need a few lemmas expressing that only time can be observed from computations that do not perform any outputs or assignments to global variables.

Lemma 4 (Assigned variables)

Write $E \setminus \sigma$ for E with all bindings of variables in σ removed. If $\text{ge}(C) = \sigma$, $\bullet \notin \sigma$ and either $\langle E \mid C \rangle \xrightarrow{as} \langle E' \mid D \rangle$ or $\langle E \mid C \rangle \xrightarrow{as, \checkmark} E'$, then $E \setminus \sigma = E' \setminus \sigma$ and $\exists ts. as = ts$.

Proof: The proof is simple by induction on the height on the derivation of $\text{ge}(C) = \sigma$.

Lemma 5

As a corollary to Lemma 4 we have that:

If $\text{ge}(C) = \emptyset$ and either $\langle E | C \rangle \xrightarrow{\text{as}} \langle E' | D \rangle$ or $\langle E | C \rangle \xrightarrow{\text{as}, \checkmark} E'$, then $E = E'$ and $\exists ts. as = ts$

Lemma 6 (Dummy computation)

If $\text{ge}(C) = \emptyset$ then $D; C \sqsubseteq D$ and $C; D \sqsubseteq D$.

Proof: Follows from Lemma 5. We use the same technique as in the proof of Lemma 2 to show that $\{(D; C, D)\} \subseteq \sqsubseteq$ and $\{(C; D, D)\} \subseteq \sqsubseteq$.

Theorem 3 (Transformation & semantic soundness)

If $\Gamma \vdash C \hookrightarrow D \mid D_L$ then $D \sqsubseteq C$.

Proof: We use induction on the height of the transformation derivation. The cases Assign_L , Assign_H , SkipAsn , SkipIf and Output are all immediate from the reflexivity of \sqsubseteq . Cases Seq , If_L , Let and While follow from use of the induction hypothesis and the precongruence property of \sqsubseteq .

In the case of If_H we know that C is of the form $\text{if } (e) C_1 \text{ else } C_2$ and $D = \text{if } (e) D_1; D_{2L} \text{ else } D_{1L}; D_2$. We also have that $\Gamma \vdash C_1 \hookrightarrow D_1 \mid D_{1L}$, $\Gamma \vdash C_2 \hookrightarrow D_2 \mid D_{2L}$, $\text{ge}(D_{1L}) = \emptyset$ and $\text{ge}(D_{2L}) = \emptyset$. By induction, $D_1 \sqsubseteq C_1$ and $D_2 \sqsubseteq C_2$. So by transitivity and Lemma 6 we have $D_1; D_{2L} \sqsubseteq C_1$ and $D_{1L}; D_2 \sqsubseteq C_2$. Now the precongruence property gives us our goal: $D \sqsubseteq C$.

4.3 Correctness of the Transformation Algorithm

The transformation algorithm is correct in the sense that transformed programs are secure. To prove this we need the idempotence property of the type system.

Lemma 7 (Low-slice idempotence)

If $\Gamma \vdash C : C_L$ then $\Gamma \vdash C_L : C_L$

Proof: Trivial by induction on the height of the typing derivation.

Lemma 8 (Transformation and well-typing)

If $\Gamma \vdash C \hookrightarrow D \mid D_L$ then $\Gamma \vdash D : D_L$.

Proof: A simple inductive reasoning on the depth of the transformation derivation proves this statement. Virtually all cases in the proof are immediate from use of the induction hypothesis and the typing rule corresponding to the transformation rule used.

The only case where some more elaborate reasoning has to be made is for the If_H rule. We then have:

$$\Gamma \vdash e : \text{Bool}_H \quad \text{ge}(D_{1L}) = \emptyset \quad \text{ge}(D_{2L}) = \emptyset$$

$$\Gamma \vdash C_1 \hookrightarrow D_1 \mid D_{1L} \quad \Gamma \vdash C_2 \hookrightarrow D_2 \mid D_{2L}$$

$$\Gamma \vdash \text{if } (e) C_1 \text{ else } C_2 \hookrightarrow \text{if } (e) D_1; D_{2L} \text{ else } D_{1L}; D_2 \mid \text{skipIf } e (D_{1L}; D_{2L})$$

Using the induction hypothesis we get $\Gamma \vdash D_1 : D_{1L}$ and $\Gamma \vdash D_2 : D_{2L}$. From Lemma 7 and typing rule Seq we get $\Gamma \vdash D_1; D_{2L} : D_{1L}; D_{2L}$ and $\Gamma \vdash D_{1L}; D_2 : D_{1L}; D_{2L}$. By Lemma 2 and the symmetry and transitivity of \sim_Γ according to Lemma 1, we thus have $D_{1L}; D_{2L} \sim_\Gamma D_{1L}; D_{2L}$. We can then use typing rule If_H to get our goal.

Theorem 4 (Transformation gives secure programs)

If $\Gamma \vdash C \hookrightarrow D \mid D_L$ then D is Γ -secure.

Proof: Follows from Theorem 1 and Lemma 8.

4.4 Completeness w.r.t. Volpano and Smith's system

In section 2 of [VS97a], Volpano and Smith present a type-system in which well typed programs are noninterfering, free of nontermination leakages but may contain timing leaks. Our transformation is complete with respect to this type-system in the sense that well typed programs can be transformed without inserting any nontermination:

If C is well-typed according to [VS97a], the transformation will succeed and yield a semantically equivalent program D that is free of timing leaks.

To state this formally we first define the notion of semantic equivalence we will use.

Definition 7 (Weak bisimulation)

$C \sim D$ iff $C \sqsubseteq D$ and $D \sqsubseteq C$.

Typing environments in [VS97a] are split into two components, $\lambda; \gamma$, mapping *locations* and *identifiers* to security levels. We will not distinguish locations and identifiers, and refer to both as variables.

The language used by Volpano and Smith differs from ours in that integers are the only values used and that observable actions are modelled by assignments to global variables of low security. Moreover, our language has no construct corresponding to the `try` command that guard uses of the division operator. Let P range over programs in [VS97a] not containing `try` commands. Define $T(P)$ to be the translation to our syntax that takes each `if`- and `while` condition, e , to $e != 0$ and each assignment to a global low variable, $x := e$, to the sequence $x := e; \text{output } x$.

Theorem 5

If $\lambda; \gamma \vdash P : s \text{ cmd}$ in the system of [VS97a], $C = T(P)$ and $\Gamma = \{x : \text{Int}_s \mid \lambda(x) = s \vee \gamma(x) = s\}$, then $\Gamma \vdash C \hookrightarrow D \mid D_L$ and $C \sim D$.

Proof: It suffices to show $\Gamma \vdash C \hookrightarrow D \mid D_L$ and $C \sqsubseteq D$. The well typedness of P ensures that all `while` loops have a condition of lowest security and that no loops exists in the branches of `if`-commands with a high condition. We thus use induction on the height of the typing derivation.

4.5 Example: RSA encryption

RSA encryption is based on computing $A = M^e \bmod n$, where M is the clear text message and e is the encryption key. To decrypt, $M = A^d \bmod n$ is computed, where d is the decryption key. To efficiently compute $X^k \bmod n$, the modular exponentiation algorithm can be used, but as shown by Kocher [Koc96], a careless implementation will leak k through timing. In Figure 11, we give an example of such an implementation. We represent k as a `w` elements long array of secret booleans, with the most significant bit of k at `k[0]`. The program shown is not Γ -secure since the two branches of the `if` have different timing behaviours, neither is the program type correct in our system. The transformation algorithm presented in Figure 9 will close these timing

```

s := 1;
i := 0;
while (i < w) {
  if (k[i])
    r := (s*x) mod n
  else
    r := s;
    s := r*r;
    i := i+1
}

```

(The result is now in r)

Figure 11: An implementation of the modular exponentiation algorithm that leaks through timing.

```

s := 1;
i := 0;
while (i < w) {
  if (k[i])
    r := (s*x) mod n;
    skipAsn r s
  else {
    skipAsn r ((s*x) mod n);
    r := s
  };
  s := r*r;
  i := i+1
}

```

Figure 12: The output of our transformation: a secure implementation of the modular exponentiation algorithm.

leaks by transforming the program into the one given in Figure 12. The security of this, padded program relies on constant time execution of the operators $*$ and mod . Since r , s and x are bignums in a realistic example, multiplication and modulo will be implemented by subroutines. To get secure implementations of these, the transformation can be used again.

5 Conclusions and Future Work

We have presented a simple and realistic solution to removing timing leaks to external observers, by using a combination of typing and transformation. Our solution is formalised and proved correct. It improves over previous work in that it relaxes the requirement that branching conditions must be of lowest security, thereby allowing a much larger class of programs to be typed.

By parameterising over the interpretation of time, our system becomes very precise. Even some timing leaks implemented by varying data-cache behaviour are detected and deemed insecure by our security criterion and type system. A key assumption built into our system is that primitive operators execute in constant time, regardless of the values given as arguments.

A minor oversimplification in our semantics is that it neglects the unconditional jumps made in the machine code that a compiler produces for if- and while-commands. The

code generated for the command `if (e) C_1 else C_2 ; D` will typically have the following form:

$\llbracket e \rrbracket$	Code for conditional expression
<code>brt L1</code>	Branch to L1 if the condition was true
$\llbracket C_2 \rrbracket$	Code for else-branch
<code>jmp L2</code>	Unconditional jump to L2
L1: $\llbracket C_1 \rrbracket$	Code for then-branch
L2: $\llbracket D \rrbracket$	Code for the rest of the program

Thus, even if C_1 and C_2 execute with the same timing behaviour, the value of e can probably be leaked given that the time for executing `jmp L2` can be noticed. Although performing an unconditional jump often does not take any time at all for modern, heavily pipelined processors, the jump would certainly be noticeable for an interpreted language like Java byte-code. The simplest way to close this possibility to leak secret information is to generate a (redundant) `jmp L2` instruction also after the code for C_1 .

As every semantic model abstracts from the real world, so does ours. This means that in an implementation of our system, there will always be well-typed programs that leak information by utilising some aspect of reality not covered by the semantic model. One weakness of our system is that the semantics has no concept of program counter or where in memory a piece of code is physically stored. It is thus impossible for the security condition to tell whether two syntactically identical pieces of code originate from the same location in memory. This sadly opens up for timing leaks implemented by utilising the behaviour of the instruction cache. If and how such leaks can be closed without completely turning off the cache remains to be investigated.

6 Acknowledgements

Thanks to David Sands, Andrei Sabelfeld, Jörgen Gustavsson and the Multi-group at Chalmers in general for encouragement and helpful comments during the development of this work. Also thanks to Pablo Giambiagi, Walid Taha and the anonymous reviewers for suggesting improvements and for spotting several typos and misplaced squiggles in preliminary versions of this paper.

References

- [BBL94] J.-P. Banatre, C. Bryce, and D. Le Metayer. Compile-time detection of information flow in sequential programs. *Lecture Notes in Computer Science*, 875:55–73, 1994.
- [DD77] D. E. Denning and P. J. Denning. Certification of programs for secure information flow. *Communications of the ACM*, 20(7):504–513, July 1977.
- [Den76] D. E. Denning. A lattice model of secure information flow. *Communications of the ACM*, 19(5):236–243, May 1976.
- [DFWB97] D. Dean, E. W. Felten, D. . Wallach, and D. Balanz. Java security: Web browsers and beyond. Technical Report TR-566-97, Princeton University, Computer Science Department, February 1997.

- [GM82] J. Goguen and J. Meseguer. Security policies and security models. In *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, April 1982.
- [Gor94] A. Gordon. A tutorial on co-induction and functional programming. In *Glasgow functional programming workshop*, pages 78–95. Springer Workshops in Computing, 1994.
- [Heh84] E. C. R. Hehner. Predicative programming part 1. *Communications of the ACM*, 27(2):134–143, February 1984.
- [HG92] J. He and V. D. Gligor. Formal methods and automated tool for timing-channel identification in tcb source code. In *Proceedings 2nd European Symposium on Research in Computer Security*, LNCS 648, pages 57–75, November 1992.
- [HR98] N. Heintze and J. G. Riecke. The SLam calculus: programming with secrecy and integrity. In *Conference Record of the Twenty-Fifth Annual ACM Symposium on Principles of Programming Languages*, pages 365–377. ACM, 1998.
- [Koc96] P. C. Kocher. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In Neal Koblitz, editor, *Advances in Cryptology – CRYPTO’96*, volume 1109 of LNCS, pages 104–113. Springer-Verlag, 1996.
- [McC87] D. McCullough. Specifications for multi-level security and hook-up property. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 161–166. IEEE Computer Society Press, 1987.
- [Mog89] T.Æ. Mogensen. Separating binding times in language specifications. In *Proceedings of the Conference on Functional Programming Languages and Computer Architecture ’89, Imperial College, London*, pages 12–25, New York, NY, 1989. ACM.
- [Mye99] A.C. Myers. JFlow: practical mostly-static information flow control. In ACM, editor, *POPL ’99. Proceedings of the 26th ACM SIGPLAN-SIGACT on Principles of programming languages, January 20–22, 1999, San Antonio, TX*, pages 228–241, New York, NY 10036, USA, 1999. ACM Press.
- [Nie84] H.R. Nielson. *Hoare Logic’s for Run-time Analysis of Programs*. Ph.D. thesis, CST-30-84, Edinburgh University, 1984.
- [RLJ98] K. Rustan, M. Leino, and R. Joshi. A semantic approach to secure information flow. *Lecture Notes in Computer Science*, 1422:254–271, 1998.
- [RT96] T. Reps and T. Turnidge. Program specialization via program slicing. In O. Danvy, R. Glueck, and P. Thiemann, editors, *Dagstuhl Seminar on Partial Evaluation*, volume 1110 of *Lecture Notes in Computer Science*, pages 409–429. Springer-Verlag, Feb 1996.
- [SS99] A. Sabelfeld and D. Sands. Probabilistic noninterference for multi-threaded programs. Unpublished (<http://www.cs.chalmers.se/~dave/papers/prob-sabelfeld-sands.ps>), June; Revised October 1999.
- [SV98] G. Smith and D. Volpano. Secure information flow in a multi-threaded imperative language. In *Conference Record of POPL ’98: The 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 355–364, San Diego, California, 19–21 January 1998.
- [VS97a] D. Volpano and G. Smith. Eliminating covert flows with minimum typings. *Proc. 10th IEEE Computer Security Foundations Workshop*, pages 156–168, June 1997.
- [VS97b] D. Volpano and G. Smith. A type-based approach to program security. In *TAPSOFT’97*, volume 1214 of LNCS, pages 607–621. Springer-Verlag, April 1997.
- [VS98] D. Volpano and G. Smith. Probabilistic noninterference in a concurrent language. *Proc. 11th IEEE Computer Security Foundations Workshop*, pages 34–43, June 1998.
- [VSI96] D. Volpano, G. Smith, and C. Irvine. A sound type system for secure flow analysis. *J. Computer Security*, 4(3):1–21, 1996.

A Proof of Lemma 2

To prove Lemma 2 we will need the following small additional lemmas.

Lemma 9 (High left-expressions)

If $\Gamma \vdash_{\leq} le : \bar{\tau}_H$, $E_1 =_{\Gamma} E_2$ and $E_1 \vdash le \Downarrow^L lv$ then $E_2 \vdash le \Downarrow^L lv$.

Proof: By induction over the height of the evaluation derivation. Uses the fact that only expressions of type Int_L may be used for indexing.

Lemma 10 (Well-typed expressions)

If $\Gamma \vdash_{\leq} ie : \tau$, $E_1 \vdash ie \Downarrow v$ and $E_1 =_{\Gamma} E_2$ then $\Gamma \vdash_{\leq} v : \tau$, $E_2 \vdash ie \Downarrow v'$ and $\Gamma \vdash_{\leq} v' : \tau$,

Proof: By induction over the height of the evaluation derivation.

Lemma 2 (Typing and Γ -bisimulation)

If $\Gamma \vdash C : C_L$ then $C \sim_{\Gamma} C_L$

Proof: We prove this lemma by induction on the height of the typing derivation and by cases on its structure. In each case we provide a set X , such that $(C, C_L) \in X$. To prove $X \subseteq \sim_{\Gamma}$ and thereby $C \sim_{\Gamma} C_L$, we use the standard technique of proving that $X \subseteq F_{\Gamma}(X \cup \sim_{\Gamma})$, where F_{Γ} is the dense bisimulation operator that has \sim_{Γ} as its largest fix-point (see e.g. [Gor94]). F_{Γ} is defined according to:

$$\begin{aligned} F_{\Gamma}(X) = & \{(C_1, C_2) \mid \langle E_1 \mid C_1 \rangle \xrightarrow{ts, \checkmark} E'_1 \Rightarrow \\ & \langle E_2 \mid C_2 \rangle \xrightarrow{ts, \checkmark} E'_2 \wedge E'_1 =_{\Gamma} E'_2\} \cup \\ & \{(C_1, C_2) \mid \langle E_2 \mid C_2 \rangle \xrightarrow{ts, \checkmark} E'_2 \Rightarrow \\ & \langle E_1 \mid C_1 \rangle \xrightarrow{ts, \checkmark} E'_1 \wedge E'_1 =_{\Gamma} E'_2\} \cup \\ & \{(C_1, C_2) \mid \langle E_1 \mid C_1 \rangle \xrightarrow{as} \langle E'_1 \mid D_1 \rangle \Rightarrow \\ & \langle E_2 \mid C_2 \rangle \xrightarrow{as} \langle E'_2 \mid D_2 \rangle \wedge \\ & E'_1 =_{\Gamma} E'_2 \wedge (D_1, D_2) \in X\} \cup \\ & \{(C_1, C_2) \mid \langle E_2 \mid C_2 \rangle \xrightarrow{as} \langle E'_2 \mid D_2 \rangle \Rightarrow \\ & \langle E_1 \mid C_1 \rangle \xrightarrow{as} \langle E'_1 \mid D_1 \rangle \wedge \\ & E'_1 =_{\Gamma} E'_2 \wedge (D_1, D_2) \in X\} \end{aligned}$$

The proof is now by induction on the height of the typing derivation, with a case analysis on the outermost typing rule used.

Assign_L: C is of the form $le := e$ and $C_L = C$. Choose $X = \{(C, C_L)\}$.

Assign_H: C is of the form $le := e$ and $C_L = \text{skipAsn } le \ e$. We also have that $\Gamma \vdash_{\leq} e : \bar{\tau}_s$ and $\Gamma \vdash_{\leq} le : \bar{\tau}_H$, where $s \leq H$. Given E_1 and E_2 such that $E_1 =_{\Gamma} E_2$, we assume

$$\langle E_1 \mid le := e \rangle \xrightarrow{te, t_{le}, t_{asn}, \checkmark} E'_1$$

From semantic rule Assign we then have that $E_1 \vdash e \Downarrow l$, $E_1 \vdash le \Downarrow^L lv$, $lv \in \text{dom}(E_1)$ and $E_1[lv = l]$. The definition of $=_{\Gamma}$ gives us $lv \in \text{dom}(E_2)$. By Lemma 9 we get $E_2 \vdash le \Downarrow^L lv$. From Lemma 10 we get $E_2 \vdash e \Downarrow l'$, and semantic rule SkipAsn gives us

$$\langle E_2 \mid \text{skipAsn } le \ e \rangle \xrightarrow{ta, t_{le}, t_{asn}, \checkmark} E_2$$

Since lv is a high left-value w.r.t. Γ , we have that $E'_1 =_{\Gamma} E_2$. Thus it is enough to choose $X = \{(C, C_L)\}$.

Seq: C is of the form $C_1; C_2$ and $C_L = D_1; D_2$, where $\Gamma \vdash C_1 : D_1$ and $\Gamma \vdash C_2 : D_2$. By induction, $C_1 \sim_{\Gamma} D_1$ and $C_2 \sim_{\Gamma} D_2$. We can thus take:

$$X = \{(C_1; C_2, D_1; D_2) \mid C_1 \sim_{\Gamma} D_1, C_2 \sim_{\Gamma} D_2\}$$

If_L: C is of the form **if** (e) C_1 **else** C_2 and $C_L = \text{if } (e) \ D_1 \text{ else } D_2$, where $\Gamma \vdash_{\leq} e : \text{Bool}_H$, $\Gamma \vdash C_1 : D_1$ and $\Gamma \vdash C_2 : D_2$. By induction we have $C_1 \sim_{\Gamma} D_1$ and $C_2 \sim_{\Gamma} D_2$. Since the same branch will be taken by both C and C_L given that $E_1 =_{\Gamma} E_2$. We can thus choose $X = \{(C, C_L)\}$

If_H: C is of the form **if** (e) C_1 **else** C_2 and $C_L = \text{skipIf } e \ D_1$, where $\Gamma \vdash_{\leq} e : \text{Bool}_H$, $\Gamma \vdash C_1 : D_1$, $\Gamma \vdash C_2 : D_2$ and $D_1 \sim_{\Gamma} D_2$. Given E_1 and E_2 , $E_1 =_{\Gamma} E_2$, we assume $\langle E_2 \mid C_L \rangle \xrightarrow{te, t_{br}} \langle E_2 \mid D_1 \rangle$. We then have $E_2 \vdash e \Downarrow l$ where $l \in \{\text{true}, \text{false}\}$ and from Lemma 10 that $E_1 \vdash e \Downarrow l'$, where also $l' \in \{\text{true}, \text{false}\}$. By semantic rule **If** we now have two cases:

If $\langle E_1 \mid C \rangle \xrightarrow{te, t_{br}} \langle E_1 \mid C_1 \rangle$ we can choose $X = \{(C, C_L)\}$ since we have $C_1 \sim_{\Gamma} D_1$ by induction.

If $\langle E_1 \mid C \rangle \xrightarrow{te, t_{br}} \langle E_1 \mid C_2 \rangle$ we can also choose $X = \{(C, C_L)\}$ since since we have $C_2 \sim_{\Gamma} D_2$ by induction, and $C_2 \sim_{\Gamma} D_1$ from Lemma 1.

Thus, we take $X = \{(C, C_L)\}$.

SkipAsn: C is of the form **skipAsn** $le \ e$ and $C_L = C$. Choose $X = \{(C, C_L)\}$.

SkipIf: C is of the form **skipIf** $e \ D$ and $C_L = \text{skipIf } e \ D_L$ where $\Gamma \vdash D : D_L$. By induction we have $D \sim_{\Gamma} D_L$ and thus we can choose $X = \{(C, C_L)\}$.

Let: C is of the form **let** $x := ie$ **in** D and $C_L = \text{let } x := ie \text{ in } D_L$, where $\Gamma \vdash_{\leq} ie : \tau$ and $\Gamma, x : \tau \vdash D : D_L$. Assuming $E_1 =_{\Gamma} E_2$ and $\langle E_1 \mid C \rangle \xrightarrow{tpu, ie} \langle E_1 \mid \text{local } x := v \text{ in } D \rangle$, we have by Lemma 10 that $\langle E_2 \mid C_L \rangle \xrightarrow{tpu, ie} \langle E_2 \mid \text{local } x := v' \text{ in } D_L \rangle$, where v' differs from v only in the high parts according to τ . By induction, we have $D \sim_{\Gamma, x: \tau} D_L$, so we also have $\text{local } x := v \text{ in } D \sim_{\Gamma} \text{local } x := v' \text{ in } D_L$ and we can choose $X = \{(C, C_L)\}$.

While: C is of the form **while** (e) D , and $C_L = \text{while } (e) \ D_L$, where $\Gamma \vdash_{\leq} e : \text{Bool}_L$ and $\Gamma \vdash D : D_L$. By induction we have that $D \sim_{\Gamma} D_L$. With E_1 and E_2 such that $E_1 =_{\Gamma} E_2$, both C and C_L will either terminate or continue to loop, since e only depends on low variables. We have two cases:

If $\langle E_1 \mid C \rangle \xrightarrow{te, t_{br}, \checkmark} E_1$ then $\langle E_2 \mid C_L \rangle \xrightarrow{te, t_{br}, \checkmark} E_2$.

If $\langle E_1 \mid C \rangle \xrightarrow{te, t_{br}} \langle E_1 \mid D; C \rangle$ then $\langle E_2 \mid C_L \rangle \xrightarrow{te, t_{br}} \langle E_2 \mid D_L; C_L \rangle$

Choosing $X = \{(C, C_L)\} \cup \{(D; C, D_L; C_L) \mid D \sim_{\Gamma} D_L\}$ we have that $X \subseteq F_{\Gamma}(X \cup \sim_{\Gamma})$.

Output: C is of the form **output** x , and $C_L = C$. We can choose $X = \{(C, C_L), (\text{skip}, \text{skip})\}$, since $\Gamma(x) = \text{Int}_L$.