

Technical Report no. 2009-01

All Secrets Great and Small¹

Delphine Demange² David Sands

CHALMERS | GÖTEBORG UNIVERSITY



Department of Computing Science and Engineering
Chalmers University of Technology and Göteborg University
S-412 96 Göteborg, Sweden

Göteborg, March 2009

1. An abbreviated version of this article appears in *Programming Languages and Systems, the 18th European Symposium on Programming, ESOP 2009*.
2. Author's address: University of Rennes 1, France



Technical Report in Computing Science at
Chalmers University of Technology and Göteborg University

Technical Report no. 2004-06
ISSN: 1650-3023

Department of Computing Science
Chalmers University of Technology and Göteborg University
SE-412 96 Göteborg, Sweden

Göteborg, Sweden, 2004

Table of Contents

1	Introduction	4
2	A Refined Multilevel Lattice	5
3	Secret-Sensitive Noninterference	6
3.1	Characterising SSNI	10
3.2	Computational Security	12
4	Secret-Sensitive Noninterference by Typing	13
4.1	Type system	13
4.2	Type Soundness	16
5	Correlation Leaks	22
6	Conclusions	23

Abstract

Tools for analysing secure information flow are almost exclusively based on ideas going back to Denning’s work from the 70’s. This approach embodies an imperfect notion of security which turns a blind eye to information flows which are encoded in the termination behaviour of a program. In exchange for this weakness many more programs are deemed “secure”, using conditions which are easy to check. Previously it was thought that such leaks are limited to at most one bit per run. Recent work by Askarov et al (ESORICS’08) offers some bad news and some good news: the bad news is that for programs which perform output, the amount of information leaked by a Denning style analysis is not bounded; the good news is that if secrets are chosen to be sufficiently large and sufficiently random then they cannot be effectively leaked at all. The problem addressed in this paper is that secrets cannot always be made sufficiently large or sufficiently random. Contrast, for example, an encryption key with an “hasHIV”-field of a patient record. In recognition of this we develop a notion of *secret-sensitive non-interference* in which “small” secrets are handled more carefully than “big” ones. We illustrate the idea with a type system which combines a liberal Denning-style analysis with a more restrictive system according to the nature of the secrets at hand.

1 Introduction

Most tools for analysing information flow in programs such as Jif [MZZ⁺08] and Flow-Caml [Sim03] build upon ideas going back to Denning’s work from the 70’s [DD77]. These systems enforce an imperfect notion of information flow which has become known as *termination-insensitive noninterference* (TINI). Under this version of noninterference, information leaks are permitted if they are transmitted purely by the program’s termination behaviour (i.e., whether it terminates or not). This imperfection is the price to pay for having a security condition which is relatively liberal (e.g. allowing while-loops whose termination may depend on the value of a secret) and easy to check.

How bad is termination-insensitive noninterference? Previously there have been informal arguments that termination-insensitive noninterference leaks at most one bit: either a program terminates or it does not, so at most one bit of information can be encoded in the termination state. However, recent work by Askarov *et al* [AHSS08] shows that for programs which perform output, an arbitrary amount of information can be leaked. The following program outputs an ascending sequence of natural numbers on a public channel until the secret has been output, at which point it goes into a silent loop:

```
for i = 0 to maxNat (  
  output i on public_channel  
  if (i = secret) then (while true do skip)  
)
```

At the very least we can say that at each output step, the observer is able to narrow down the possible values of the secret. This program (in suitable variants) is accepted as secure by state-of-the-art information flow analysis tools such as Jif [MZZ⁺08], FlowCaml [Sim03], and the SPARK Examiner [BB03,CH04].

Askarov *et al* formalise the notion of termination-insensitive noninterference and show that although termination-insensitive noninterference can leak an arbitrary amount of information, it cannot do so any more efficiently than the above example. The revised intuition for programs performing public output is that the number of possible “termination states” that can be used to encode information is of the order of the number of public outputs performed by the program – since the program can diverge after 0 outputs, after 1 output, after 2 outputs, etc. This means that to leak n bits of information the program needs to perform 2^n outputs.

For Denning-style analyses this means that if secrets are sufficiently large and sufficiently random then programs are *computationally secure* in the sense that the probability of the attacker guessing the secret after observing a polynomial number of outputs (again, in the size of the secret) gives only a negligible advantage over guessing the secret without running the program.

What does this mean for information flow analysis in practice? Whereas previously the imperfections of a Denning-style analysis were viewed as a reasonable tradeoff between ease of analysis versus degree of security, we believe that in the light of [AHSS08] we need a different perspective. The leak caused by termination-insensitivity is only acceptable for sufficiently large and random secrets. But secrets, in general, are not always parametric: one cannot always freely choose to make a secret larger and

more random. For example, an application cannot decide that a credit card CCV number should be made larger. An encryption key, on the other hand, might be something that the application can control, and decide to scale up.

In this paper we consider the information flow problem in an arbitrary multi-level security lattice. We present a way (Section 2) of refining each security level in an information-flow lattice into two levels: *big secrets*, that are sufficiently large and randomized to abide some leakage, and *small secrets*, for which even slow leakage is unacceptable. Then, we define a two-level noninterference (Section 3), following Askarov et al’s recent work, which combines the demands of termination-insensitive noninterference (for big secrets) with the stricter requirements of termination-sensitive noninterference (for small secrets). A type system is provided (Section 4) that ensures this notion of noninterference. Additional novelties of the system are a somewhat more liberal treatment of small secrets than found in previous termination-sensitive type systems. Section 5 describes a strengthening of the definition of security to eliminate leakage correlations between big and small secrets.

2 A Refined Multilevel Lattice

In [AHSS08] a definition of termination-insensitive noninterference (**TINI**) was introduced which is suitable for programs with outputs, assuming only two security levels *low* and *high*. They proved that, even if programs verifying this condition can leak more than a bit of information, the attacker cannot reliably (i.e in a single run) learn a secret in polynomial time in the size of the secret. They also proved that, for programs satisfying TINI, if secrets are uniformly distributed, then a particular observation of a computation represents only a negligible hint for the attacker (Theorem 3).

The basic idea in this work is to refine the notion of *high* into two points *bhigh* and *shigh*. These will correspond to “big” secrets and “small” secrets respectively. We will define a notion of secret-sensitive noninterference which allows a low user to learn a little about big secrets, and nothing at all about small secrets (relative to the notion of observation that we model).

How are big and small secrets related? A key point here is that data labelled *bhigh* will depend only on *bhigh* or *low* data sources, whereas data labelled *shigh* might also depend on *shigh* data sources. Thus the label *bhigh* does not mean that the data *is* a large secret – it just means that it does not depend on (contain any information about) a small secret. We can then see that the resulting refined security lattice is as given in Figure 1.

Now we generalise this refinement to the case of an arbitrary multi-level lattice of information levels [Den76]. Denning’s lattice model of information considers an arbitrary complete lattice $\langle \mathcal{L}, \sqsubseteq_{\mathcal{L}}, \sqcup_{\mathcal{L}}, \sqcap_{\mathcal{L}}, \perp_{\mathcal{L}} \rangle$ where \mathcal{L} is the set of security *clearance levels* (henceforth just *levels*, ranged over by i, j), and $\sqsubseteq_{\mathcal{L}}$ is the ordering relation which determines when one level is higher than another. The idea is that a principal with a clearance level i is permitted to see data which is classified at level i or below according to the partial ordering. Information from any levels may be combined, in which case the classification for the resulting data is given deterministically by the least-upper-bound operation $\sqcup_{\mathcal{L}}$.

To refine this general case we note that we must split each level $i \in \mathcal{L}$, with the exception of the bottom level $\perp_{\mathcal{L}}$ (which can always be thought of as public data) into two points, corresponding to the big secrets (labelled b) and the small (labelled s). Thus any non-bottom element i will be refined to (i, b) and (i, s) . To define the appropriate order between lattice elements we first note that $(i, b) \sqsubseteq (i, s)$ – with the same motivation as given for the refined two-point lattice. Similarly, when comparing secrets of the same kind we have $(i, a) \sqsubseteq (j, a)$ only when $i \sqsubseteq_{\mathcal{L}} j$.



Fig. 1. The refined 2-point lattice

What about the relationship between two points (i, b) and (j, s) – when can information flow between these points? The idea is that information at level b is potentially leaked via a covert channel, so that it may be leaked to *any* level. Because of this we can only permit flow from (i, b) to (j, s) , and then only when $i \sqsubseteq_{\mathcal{L}} j$. If we permitted a small secret (i, s) to flow to any (j, b) for $(j \neq i)$ then we would be able to launder small secrets by first allowing them to flow to a big secret and then leaking via the covert channel

from there. In summary, we define the refinement of a given security lattice:

Definition 1. Let \mathcal{S} denote the 2-point lattice formed from b and s under the ordering $b \sqsubseteq s$. We define the refinement of a security lattice \mathcal{L} as the partial product of \mathcal{L} and \mathcal{S} , which is the standard product lattice $\mathcal{L} \times \mathcal{S}$, quotiented by the equivalence $(\perp_{\mathcal{L}}, b) \equiv (\perp_{\mathcal{L}}, s)$ – and this bottom element will be simply denoted by $\perp_{\mathcal{L}}$.

Example Consider the example where $\mathcal{L} = \{secret, financial, medical, public\}$ is the set of the four security levels a program has to deal with, ordered according to the Hasse diagram in Figure 2. Motivating a refinement of the lattice, there could be medical data that is encrypted – or simply very large (e.g. high resolution image data) that could be safely allowed to leak slowly, and other medical data that are to be handled with more care, such as an “hasHIV” boolean flag in a patient record. The partial product of lattices \mathcal{L} and \mathcal{S} is presented in Figure 3.




Fig. 2. Example \mathcal{L}




Fig. 3. The refinement of \mathcal{L}

3 Secret-Sensitive Noninterference

In this section we define the security goal for programs computing over data labelled with a refined lattice. This variant of the notion of noninterference, *secret-sensitive*

noninterference, combines the demands of termination-insensitive noninterference for *b*-data, and the stronger termination-sensitive noninterference for *s*-data. Further, we develop a bisimulation-style characterisation of *secret-sensitive noninterference* which provides a convenient proof method.

Operational Semantics We keep our presentation language independent, but we assume some basic structure for an operational semantics. We will consider simple imperative computation modelled by a standard small-step operational semantics defined over configurations of the form $\langle M, C \rangle$ where M is a memory (store) – a finite mapping from variables to values – and C (C' , D etc.) is a command. Each variable x is assumed to have a fixed policy denoted $\Gamma(x)$, which we take to be a member of the refinement of some lattice \mathcal{L} .

We assume an operational semantics consisting of deterministic labelled transitions between configurations, where a label u is either (i) an observable output $i(v)$, meaning that a value v is output on a channel observable at level $i \in \mathcal{L}$ or above, or (ii) a silent action labelled τ . We write e.g. $\langle M, C \rangle \xrightarrow{i(v)} \langle M', C' \rangle$.

On top of the basic labelled transitions we define a family of transition systems labelled by a particular level:

Definition 2 (*i*-observable transitions). We can define the transition relations \xrightarrow{u}_i , $i \in \mathcal{L}$ as:

$$\frac{\langle M, C \rangle \xrightarrow{j(v)} \langle M', C' \rangle \quad j \sqsubseteq_{\mathcal{L}} i}{\langle M, C \rangle \xrightarrow{v}_i \langle M', C' \rangle} \quad \frac{\langle M, C \rangle \xrightarrow{u} \langle M', C' \rangle \quad u = \tau \text{ or } u = j(n) \text{ where } j \not\sqsubseteq_{\mathcal{L}} i}{\langle M, C \rangle \xrightarrow{\tau}_i \langle M', C' \rangle}$$

Thus the *i*-observable transitions are obtained from the raw transitions by filtering out (replacing by τ) all output actions that are not visible at level *i*. Note that the non- τ transitions are just the value which is observed and not the channel on which it is observed.

Now we define the “big step” transitions $\langle M, C \rangle \xRightarrow{u}_i \langle M', C' \rangle$ as follows

$$\begin{aligned} \langle M, C \rangle \xrightarrow{\tau}_i \langle M', C' \rangle &\triangleq \langle M, C \rangle \xrightarrow{\tau}_i^* \langle M', C' \rangle \\ \langle M, C \rangle \xrightarrow{v}_i \langle M', C' \rangle &\triangleq \langle M, C \rangle \xrightarrow{\tau}_i^* \xrightarrow{v}_i \langle M', C' \rangle \end{aligned}$$

We also define the multi-step observations $\langle M, C \rangle \xRightarrow{\vec{v}}_i \langle M', C' \rangle$ with $\vec{v} = v_1 v_2 \cdots v_n$ as follows:

$$\langle M, C \rangle \xRightarrow{v_1}_i \langle M_1, C_1 \rangle \xRightarrow{v_2}_i \langle M_2, C_2 \rangle \xRightarrow{v_3}_i \cdots \xRightarrow{v_{n-1}}_i \langle M_{n-1}, C_{n-1} \rangle \xRightarrow{v_n}_i \langle M', C' \rangle$$

for some sequence of intermediate configurations $\langle M_i, C_i \rangle$. We define the multi-step reduction for the empty vector to be synonymous with $\xrightarrow{\tau}_i$.

Attacker’s knowledge Our presentation follows the style of Askarov *et al* [AHSS08] closely. The definition of noninterference developed here builds on the concept of *attacker knowledge* which is what an attacker (an observer of a given clearance level i) can deduce about the initial values of variables based on a particular observation of a program run.

The attacker i knows the initial low part of the memory. The low part of the memory from the perspective of a given level i is all variables with policy (i, s) or lower - and observes some output trace \vec{v} that is not necessarily maximal, knows the program and is able to make perfect deductions about the semantics of the program. For a memory M we let M^i denote the low part of the memory from the perspective of an observer at level i , i.e. the part of the memory that he can see.

Definition 3 (Observations). *Given a program C and a low memory M^i , the i -observations is the set of all possible sequences of observable outputs that could arise from a run of C with a memory compatible with M^i . It is defined:*

$$Obs_i(C, M^i) = \{\vec{v} | \langle N, C \rangle \xrightarrow{\vec{v}}_i \langle N', C' \rangle, \quad N^i = M^i\}$$

Definition 4 (Attacker’s knowledge). *Given a program C , an initial choice M^i of the low part of the memory (for level i) and a trace of i -observable outputs \vec{v} , the attacker’s knowledge gained from this observation is the set of all possible memories that could have lead to this observation.*

$$k_i(C, M^i, \vec{v}) = \{N | \langle N, C \rangle \xrightarrow{\vec{v}}_i \langle N', C' \rangle, \quad N^i = M^i\}$$

Note that increase in knowledge corresponds to a decrease in the size of the knowledge set. Knowledge increases with outputs: the more outputs the attacker observes, the more precise is his knowledge [AS07]:

$$\forall C, M^i, \vec{v}, v. \quad k_i(C, M^i, \vec{v}v) \subseteq k_i(C, M^i, \vec{v})$$

In order to distinguish between what is learnt about the “big” secrets (variables at levels (i, b)) from what is learnt about the “small secrets” (variables at levels (i, s)) we define the projections of knowledge sets to the b - and s -parts.

Definition 5 (b - and s -restricted memories). *Given a memory M , and a security size $a \in \mathcal{S}$, we define $M|_a^i$ to be the restriction of M to those variables x such that $\Gamma(x) = (j, a)$, $j \not\sqsubseteq i$ – i.e. the “ a -secrets” from i ’s perspective. We extend the definition pointwise to sets of memories.*

Definition 6 (b - and s -restricted knowledge). *Given a program C , a security size $a \in \mathcal{S}$ and an initial choice M^i of the low part of the memory and a trace of outputs \vec{v} , the a -restricted knowledge of the attacker i , written $k_i^a(C, M^i, \vec{v})$ is defined $(k_i(C, M^i, \vec{v}))|_a^i$.*

Informally, the restricted knowledge $k_i^a(C, M^i, \vec{v})$ is i ’s knowledge about the a -secrets (from i ’s perspective) after having observed \vec{v} from initial memory M^i .

The idea of “ i ’s secrets” can be illustrated using the lattice presented in Figure 3. For example, the projection $M|_s^{finance}$ restricts M to just those variables with classifications $(medical, s)$ or $(secret, s)$. The *finance*-perspective on the lattice is illustrated




Fig. 4. The *finance*-perspective on the example refined lattice

in Figure 4, where the *b*-secrets and *s*-secrets are marked. The low part of the lattice, from the *finance* perspective, is also marked.

The *s*-restricted knowledge for an attacker at level *finance* is thus the knowledge that can be deduced about the *s*-secret part of the memory.

Noninterference Several kinds of noninterference can be defined from the notion of knowledge. Here we adapt the definition of termination-(in)sensitive noninterference that was proposed in [AHSS08] and then propose a definition of a two-levelled noninterference.

Definition 7 (Termination-Sensitive Noninterference (TSNI)). A program C satisfies TSNI if for all i , whenever $\vec{v}v \in \text{Obs}_i(C, M^i)$ then

$$k_i(C, M^i, \vec{v}) = k_i(C, M^i, \vec{v}v).$$

TSNI means that at each step of output, nothing new about the high memory is learnt by the attacker.

Definition 8 (Termination-Insensitive Noninterference (TINI)). A program C satisfies TINI if for all i , whenever $\vec{v}v \in \text{Obs}_i(C, M^i)$ then

$$k_i(C, M^i, \vec{v}v) = \bigcup_{v'} k_i(C, M^i, \vec{v}v').$$

TINI allows leakage at each low output step, but only through the fact that there is *some* output step. The knowledge leaked by one output is the same as for any other.

In order to deal with our two different kinds of secret (*b* and *s*), the idea is here to combine both TSNI and TINI: although we only accept TSNI for *s*-data which must be handled with more care, we allow TINI for *b*-data, that abide some leakage since they are randomized and large enough.

Definition 9 (Secret-Sensitive Noninterference (SSNI)). A program C satisfies SSNI if for all i , whenever $\vec{v}v \in \text{Obs}_i(C, M^i)$ then the following two properties hold:

$$\begin{aligned} k_i^s(C, M^i, \vec{v}v) &= k_i^s(C, M^i, \vec{v}) && (s\text{-TSNI}) \\ k_i^b(C, M^i, \vec{v}v) &= \bigcup_{v'} k_i^b(C, M^i, \vec{v}v') && (b\text{-TINI}) \end{aligned}$$

3.1 Characterising SSNI

The knowledge based definitions are (in our opinion) lucid because they give a clear attacker perspective on the problem. However, for reasoning about secret-sensitive non-interference we find it convenient to work with a more conventional characterisation in terms of bisimulation relations. Here we develop this alternative characterisation, which we will employ in Section 4 in order to prove that the type system there guarantees secret-sensitive noninterference.

The basic idea is to establish the two components of SSNI via two forms of bisimulation relations between configurations.

Definition 10 (Termination-sensitive i -bisimulation (i -TSB)). *A symmetric relation \mathcal{R} on configurations is a termination-sensitive i -bisimulation, if $\langle M, C \rangle \mathcal{R} \langle N, D \rangle$ implies:*

- (i) $M^i = N^i$ and $M|_b^i = N|_b^i$, and
- (ii) whenever $\langle M, C \rangle \xrightarrow{u}_i \langle M', C' \rangle$ then $\langle N, D \rangle \xrightarrow{u}_i \langle N', D' \rangle$ with $\langle M', C' \rangle \mathcal{R} \langle N', D' \rangle$.

Two configurations are said to be i -TSB equivalent (denoted by \cong_i) if there exists a i -TSB relating them.

Here, the termination-sensitivity comes from the ability to produce the next output together with the symmetry of the relation.

Definition 11 (Termination-insensitive i -bisimulation (i -TIB)). *We say that a configuration $\langle M, C \rangle$ diverges for i , written $\langle M, C \rangle \uparrow_i$, if $\langle M, C \rangle$ cannot perform any i -observable output transition \xrightarrow{v}_i .*

A symmetric relation \mathcal{R} on configurations is defined to be a termination-insensitive i -bisimulation if whenever $\langle M, C \rangle \mathcal{R} \langle N, D \rangle$ we have

- (i) $M^i = N^i$ and
- (ii) if $\langle M, C \rangle \xrightarrow{u}_i \langle M', C' \rangle$ then either $\langle N, D \rangle \xrightarrow{u}_i \langle N', D' \rangle$ with $\langle M', C' \rangle \mathcal{R} \langle N', D' \rangle$, or $\langle N, D \rangle \uparrow_i$.

Two configurations are said to be i -TIB equivalent (denoted by \simeq_i) if there exists a i -TIB relating them.

Note that the notion of “divergence” used here is purely from the perspective of a remote observer who sees only the outputs on channels. We could make this more conventional if we made program termination an observable event for all levels. We have chosen not to do so, but the technical development in this paper does not depend in a crucial way on this fact.

Before we show these relations in Definitions 10 and 11 are sufficient to characterise SSNI, we need the following lemmas about i -TSB and i -TIB.

Lemma 1.

If $\langle M, C \rangle \cong_i \langle N, D \rangle$ and $\langle M, C \rangle \xrightarrow{\vec{u}}_i \langle M', C' \rangle$ then $\langle N, D \rangle \xrightarrow{\vec{u}}_i \langle N', D' \rangle$ with $\langle M', C' \rangle \cong_i \langle N', D' \rangle$.

Proof. We proceed by induction on the number of outputs (length of \vec{v}), and in the base case by induction on the length of the raw transition sequence.

- Base case: \vec{v} has length 0 so every transition is silent. We proceed by induction on the number n of silent steps:
 - Base case: trivial for $n = 0$.
 - Induction: here, we have the following $n + 1$ -step computation $\langle M, C \rangle \xrightarrow{\tau}_i \langle M_2, C_2 \rangle \xrightarrow{\tau}_i \dots \langle M_n, C_n \rangle \xrightarrow{\tau}_i \langle M_{n+1}, C_{n+1} \rangle = \langle M', C' \rangle$. Applying the induction hypothesis on the n first steps, we have that $\langle N, D \rangle \xrightarrow{\tau}_i \langle N_m, D_m \rangle$ with $\langle M_n, C_n \rangle \cong_i \langle N_m, D_m \rangle$. But $\langle M_n, C_n \rangle \xrightarrow{\tau}_i \langle M_{n+1}, C_{n+1} \rangle$, so by definition of an i -TSB, we have that $\langle N_m, D_m \rangle \xrightarrow{\tau}_i \langle N', D' \rangle$ with $\langle M_{n+1}, C_{n+1} \rangle \cong_i \langle N', D' \rangle$. We conclude concatenating computations: $\langle N, D \rangle \xrightarrow{\tau}_i \langle N_m, D_m \rangle \xrightarrow{\tau}_i \langle N', D' \rangle$.
- Induction: suppose \vec{v} has length $k + 1$. This can be written $\langle M, C \rangle \xrightarrow{v_1 \dots v_k}_i \langle M_{n_k}, C_{n_k} \rangle \xrightarrow{v_{k+1}}_i \langle M_{n_{k+1}}, C_{n_{k+1}} \rangle$. Applying the induction hypothesis on the k first output multisteps, we obtain $\langle N, D \rangle \xrightarrow{v_1 \dots v_k}_i \langle N_{m_k}, D_{m_k} \rangle$ with $\langle M_{n_k}, C_{n_k} \rangle \cong_i \langle N_{m_k}, D_{m_k} \rangle$. But $\langle M_{n_k}, C_{n_k} \rangle$ makes an output multistep $\xrightarrow{v_{k+1}}_i$. We conclude applying the definition of an i -TSB as many times as there are silent steps before v_{k+1} is output. \square

Lemma 2.

If $\langle M, C \rangle \simeq_i \langle N, D \rangle$ and $\langle M, C \rangle \xrightarrow{\vec{v}}_i \langle M', C' \rangle$ then $\langle N, D \rangle \xrightarrow{\vec{v}'}_i \langle N', D' \rangle$ for some \vec{v}' such that either $\vec{v} = \vec{v}'$ and $\langle M', C' \rangle \simeq_i \langle N', D' \rangle$, or \vec{v}' is a prefix of \vec{v} and $\langle N', D' \rangle \uparrow_i$.

Proof. Similar to Lemma 1. \square

The following propositions state that SSNI can be couched in the bisimulation setting. To simplify notations, we will denote by $L, L' \dots$ the low projections M^i of a memory M from the perspective of i , when i is clear from the context. Similarly, we will respectively denote by $B, B' \dots$ and $S, S' \dots$ the restrictions $M|_b^i$ and $M|_s^i$. The triple LBS will denote the necessarily unique memory M such that $M^i = L$, $M|_b^i = B$ and $M|_s^i = S$.

Proposition 1.

Suppose that for all levels i and all memories M and N such that $M^i = N^i$ and $M|_b^i = N|_b^i$ we have $\langle M, C \rangle \cong_i \langle N, C \rangle$. Then for all i , whenever $\vec{v}v \in \text{Obs}_i(C, M^i)$ then $k_i^s(C, M^i, \vec{v}v) = k_i^s(C, M^i, \vec{v})$.

Proof. Here we just focus on the inclusion \supseteq , since \subseteq comes from the monotonicity of knowledge.

Suppose $\vec{v}v \in \text{Obs}_i(C, L)$ and for all L, B, S and S' we have $\langle LBS, C \rangle \cong_i \langle LBS', C \rangle$. Let S_0 be an element of $k_i^s(C, L, \vec{v})$. We have to show that $S_0 \in k_i^s(C, L, \vec{v}v)$ (nothing is learned from the output v).

Since $\vec{v}v \in \text{Obs}_i(C, L)$, there exist B, S such that $\langle LBS, C \rangle \xrightarrow{\vec{v}v}_i \langle L'B'S', C' \rangle$. But we also have $\langle LBS, C \rangle \cong_i \langle LBS_0, C \rangle$. Thanks to Lemma 1 we then have that $\langle LBS_0, C \rangle \xrightarrow{\vec{v}v}_i \langle L''B''S'_0, C'' \rangle$, which means that $S_0 \in k_i^s(C, L, \vec{v}v)$. \square

Proposition 2.

Suppose that for all levels i and all M, N such that $M^i = N^i$ we have that $\langle M, C \rangle \simeq_i \langle N, C \rangle$. Then $\vec{v}v \in \text{Obs}_i(C, M_i)$ implies $k_i^b(C, M^i, \vec{v}v) = \bigcup_{v'} k_i^b(C, M^i, \vec{v}v')$.

Proof. Let $i \in \mathcal{L}$. Suppose $\vec{v}v \in \text{Obs}_i(C, L)$. Then, $\langle LBS, C \rangle \xrightarrow{\vec{v}v}_i \langle L_2B_2S_2, C_2 \rangle$. We decompose this computation as follows: $\langle LBS, C \rangle \xrightarrow{\vec{v}}_i \langle L_1B_1S_1, C_1 \rangle \xrightarrow{v}_i \langle L_2B_2S_2, C_2 \rangle$.

Let v' and B_0 such that $B_0 \in k_i^b(C, L, \vec{v}v')$. We have to show that $B_0 \in k_i^b(C, L, \vec{v}v)$ (what is learned from v is the same as from the output v').

Since B_0 is in $k_i^b(C, L, \vec{v}v')$, we have $\langle LB_0S_0, C \rangle \xrightarrow{\vec{v}}_i \langle L'B'_0S'_0, C' \rangle \xrightarrow{v'}_i \langle L''B''_0S''_0, C'' \rangle$ for some S_0 .

By hypothesis, we have that $\langle LBS, C \rangle \simeq_i \langle LB_0S_0, C \rangle$.

Thanks to Lemma 2, we have that $\langle L_1B_1S_1, C_1 \rangle \simeq_i \langle L'B'_0S'_0, C' \rangle$. By definition of a i -TIB, because C' does not diverge starting from $L'B'_0S'_0$ (since $\langle L'B'_0S'_0, C' \rangle \xrightarrow{v'}_i \langle L''B''_0S''_0, C'' \rangle$), then the output value is the same : $v = v'$.

So, $\langle L_0B_0S_0, C \rangle \xrightarrow{\vec{v}}_i \langle L'B'_0S'_0, C' \rangle \xrightarrow{v}_i \langle L''B''_0S''_0, C'' \rangle$, which means that $B_0 \in k_i^b(C, L, \vec{v}v)$. □

Clearly, then, putting the propositions together we get a proof technique for SSNI:

Corollary 1. C satisfies SSNI if, for all levels i and all M , and N we have

- $M^i = N^i$ implies $\langle M, C \rangle \simeq_i \langle N, C \rangle$, and
- $M^i = N^i$ and $M|_b^i = N|_b^i$ implies $\langle M, C \rangle \cong_i \langle N, C \rangle$.

3.2 Computational Security

Definition 9 clearly enforces termination-sensitive noninterference for s -data. Regarding b -data, we can provide the computational security guarantees of [AHSS08] to show that b -secrets, if chosen uniformly, cannot be leaked in polynomial time in their size. To argue this we can first reclassify all secrets as b -data (or equivalently assume that there are no s -secrets). Then we are back in the standard security lattice, and we simply need to generalise the results of [AHSS08] from a two-point lattice to an arbitrary one. This is, as usual, unproblematic since from the perspective of each individual level i there are only two levels of interest: the levels which can be seen (i.e. the levels less than or equal to i) and those which cannot. The main result is that if b -data is randomly chosen, then an observer at level i learns a negligible amount of information (as a function of the size of the b -data) about the data which i cannot see. We will not further develop the details of this argument in the present article. The differences from the development in [AHSS08] would be minor.

4 Secret-Sensitive Noninterference by Typing

In this section, we describe a type system that enforces the noninterference Definition 9: well-typed programs are secret-sensitive noninterfering. We study a classical deterministic while programming language defined with expressions and commands.

$$\begin{aligned}
 e &::= n \mid x \mid e \text{ op } e \\
 c &::= \text{skip} \mid x := e \mid c ; c \mid \text{if } e \text{ then } c \text{ else } c \mid \\
 &\quad \text{while } e \text{ do } c \mid \text{for } e \text{ do } c \mid \text{output}_i(e)
 \end{aligned}$$

Here n stands for any integer constant, x for any variable and op for any of the classical binary arithmetical operators. Booleans are represented by integers the classical way (0 is *false*, and everything else is *true*). We also assume that there are no exceptions raised: all binary operators are totally defined.

Note that the language provides two types of loops: **for** loops are always terminating, that is the guard expression is evaluated just once, leading to a constant that is decreased each time the end of the loop body is reached, and **while** loops are potentially non terminating. The distinction will be used in the type system to good effect.

The language includes the **output** _{i} primitive method that writes the value of its argument to a channel with level i . The operational semantics is standard and is given in Figure 5. We denote by $M(x)$ the value of the variable x .

4.1 Type system

This type system is based on the combination of a standard Denning-style analysis (in type system form [VSI96]) for enforcing the termination-insensitive security for b -secrets, and a more restrictive type system for handling the s -secrets. One such termination-sensitive type system is that described in [VS97], but that system is extremely restrictive: loops are only allowed if the guard does not refer to anything except data at the lowest lattice level, and if there is a branch on secret data at any level then no loops are allowed inside the branches. Instead we adapt an idea common to the type systems from [BC01] and [Smi01] for the termination-sensitive part. The idea is here to allow high while loops (i.e. loops with high guards or arbitrary while loops occurring in a high context) so long as no assignment or output to levels below the loop guards follows them.

The form of the typing judgements follows the style of [BC01] in that it handles indirect information flows by recording the write effect of a command (the lowest level to which it writes data). This gives the same power as Denning’s popular approach which uses a “program counter” level.

Consider both lattices \mathcal{L} and \mathcal{S} , and let \mathcal{P} be their partial product as previously defined. A type is either an expression type denoted $e : \tau$, or a command type written $(\tau, \sigma, \delta)cmd$, where both τ and σ are in \mathcal{P} , the set of security levels, and δ , the *termination flag* is a member of the set $\{\downarrow, \uparrow\}$, where we order the elements $\downarrow \leq \uparrow$.

Type judgments are of the form

$$\Gamma \vdash C : (\tau, \sigma, \delta)cmd$$

$$\begin{array}{c}
\frac{}{\langle M, n \rangle \Downarrow n} \text{E-CONST} \quad \frac{M(x) = n}{\langle M, x \rangle \Downarrow n} \text{E-VAR} \\
\frac{\langle M, e_1 \rangle \Downarrow n_1 \quad \langle M, e_2 \rangle \Downarrow n_2 \quad n_1 \text{ op } n_2 = n}{\langle M, e_1 \text{ op } e_2 \rangle \Downarrow n} \text{E-BINOP} \\
\frac{\langle M, e \rangle \Downarrow n}{\langle M, \mathbf{output}_i(e) \rangle \xrightarrow{i(n)} \langle M, \mathbf{skip} \rangle} \text{C-OUT} \\
\frac{\langle M, e \rangle \Downarrow n}{\langle M, x := e \rangle \xrightarrow{\tau} \langle M[x \leftarrow n], \mathbf{skip} \rangle} \text{C-ASSIGN} \\
\frac{\langle M, C_1 \rangle \xrightarrow{u} \langle M', C'_1 \rangle}{\langle M, C_1; C_2 \rangle \xrightarrow{u} \langle M', C'_1; C_2 \rangle} \text{C-SEQ1} \quad \frac{}{\langle M, \mathbf{skip}; C_2 \rangle \xrightarrow{\tau} \langle M, C_2 \rangle} \text{C-SEQ2} \\
\frac{\langle M, e \rangle \Downarrow n \quad n \neq 0}{\langle M, \mathbf{if } e \text{ then } C_1 \text{ else } C_2 \rangle \xrightarrow{\tau} \langle M, C_1 \rangle} \text{C-IF1} \quad \frac{\langle M, e \rangle \Downarrow 0}{\langle M, \mathbf{if } e \text{ then } C_1 \text{ else } C_2 \rangle \xrightarrow{\tau} \langle M, C_2 \rangle} \text{C-IF2} \\
\frac{}{\langle M, \mathbf{while } e \text{ do } C \rangle \xrightarrow{\tau} \langle M, \mathbf{if } e \text{ then } (C; \mathbf{while } e \text{ do } C) \text{ else skip} \rangle} \text{C-WHILE} \\
\frac{\langle M, e \rangle \Downarrow n \quad n_1 = |n|}{\langle M, \mathbf{for } e \text{ do } C \rangle \xrightarrow{\tau} \langle M, \mathbf{if } n_1 \text{ then } (C; \mathbf{for } n_1 - 1 \text{ do } C) \text{ else skip} \rangle} \text{C-FOR}
\end{array}$$

Fig. 5. The labelled transition system

where Γ is the typing environment i.e. a mapping from variables to variable types. In the following, Γ is kept implicit. The syntactic meaning of such a judgment is that

- τ is a lower bound on the security levels of variables that are assigned to in C .
- σ is the least upper bound on the levels of (for, if, while) guards occurring in C .
- δ is \downarrow if C contains no while loops, and is \uparrow otherwise.

The semantic implication of these typings is that

- τ is a lower bound on the the *write effect* of the command – i.e., the command only modifies variables of level τ or above, and
- σ is the *termination effect*: observing that C produces some output (i.e. “terminates”) give us knowledge about data at level at most σ .
- δ is a *termination flag*: if $\delta = \downarrow$ then the command always terminates.

With these intended meanings of τ , σ and δ , there is a natural partial order on types which is contravariant in the first component and covariant in the second and third:

$$(\tau, \sigma, \delta) \text{cmd} \leq (\tau', \sigma', \delta') \text{cmd} \text{ if } \tau' \sqsubseteq_{\mathcal{P}} \tau \text{ and } \sigma \sqsubseteq_{\mathcal{P}} \sigma' \text{ and } \delta \leq \delta'$$

This relation is not used in the type system, but is used in the statement of e.g. the subject reduction property below.

For elements of \mathcal{P} (the first two components of a command type in particular) we define the first and second projections in the obvious way: $fst(i, a) = i$ and $fst(\perp_{\mathcal{P}}) = \perp_{\mathcal{L}}$; $snd(i, a) = a$ and $snd(\perp_{\mathcal{P}}) = \perp_{\mathcal{S}} = b$.

Rules of the security type system are displayed in Figure 6, where we drop the subscript for the relation $\sqsubseteq_{\mathcal{P}}$.

Explicit flows are handled with rules for expressions, rules T-ASSIGN, and T-OUT, while implicit flows are treated in T-IF, T-WHILE and T-FOR which demand that their body is at least as high as their guard level.

Most of the action takes place in the sequential composition rules. The interesting case is T-SEQ2 where the termination effect σ_1 of C_1 is an s -secret, and C_1 is indeed potentially nonterminating. This means that we cannot allow arbitrary assignments in C_2 since these might leak information about the s -secrets which affected the termination of C_1 . Thus the write effect of C_2 is constrained so that it does not write below σ_1 , the termination effect of C_1 . For rule T-SEQ1 we are more liberal, since either the guards do not depend on s -secrets, or C_1 is always terminating.

The same reasoning is applied to while and for loops – their execution may be a sequential composition of the body of the loop and the loop itself.

$$\begin{array}{c}
\frac{}{\vdash n : \tau} \text{T-CONST} \quad \frac{\Gamma(x) = \tau \text{ var}}{\vdash x : \tau} \text{T-VAREXP} \\
\frac{\vdash e : \tau' \quad \tau' \sqsubseteq \tau}{\vdash e : \tau} \text{T-SUBEXP} \quad \frac{\vdash e_1 : \tau \quad \vdash e_2 : \tau}{\vdash e_1 \text{ op } e_2 : \tau} \text{T-BINOP} \\
\frac{}{\vdash \text{skip} : (\top_{\mathcal{P}}, \perp_{\mathcal{P}}, \downarrow) \text{cmd}} \text{T-SKIP} \quad \frac{\vdash e : \tau \quad \Gamma(x) = \tau \text{ var}}{\vdash x := e : (\tau, \perp_{\mathcal{P}}, \downarrow) \text{cmd}} \text{T-ASSIGN} \\
\frac{\vdash e : \tau \quad fst(\tau) \sqsubseteq_{\mathcal{L}} i}{\vdash \text{output}_i(e) : ((i, s), \perp_{\mathcal{P}}, \downarrow) \text{cmd}} \text{T-OUT} \\
\frac{\vdash C_i : (\tau_i, \sigma_i, \delta_i) \text{cmd} \quad snd(\sigma_1) = b \text{ or } \delta_1 = \downarrow}{\vdash C_1; C_2 : (\tau_1 \sqcap \tau_2, \sigma_1 \sqcup \sigma_2, \delta_1 \sqcup \delta_2) \text{cmd}} \text{T-SEQ1} \\
\frac{\vdash C_i : (\tau_i, \sigma_i, \delta_i) \text{cmd} \quad \sigma_1 \sqsubseteq \tau_2 \quad snd(\sigma_1) = s \quad \delta_1 = \uparrow}{\vdash C_1; C_2 : (\tau_1 \sqcap \tau_2, \sigma_1 \sqcup \sigma_2, \uparrow) \text{cmd}} \text{T-SEQ2} \\
\frac{\vdash e : \theta \quad \vdash C_i : (\tau_i, \sigma_i, \delta_i) \text{cmd} \quad \theta \sqsubseteq \tau_i}{\vdash \text{if e then } C_1 \text{ else } C_2 : (\tau_1 \sqcap \tau_2, \sigma_1 \sqcup \sigma_2 \sqcup \theta, \delta_1 \sqcup \delta_2) \text{cmd}} \text{T-IF} \\
\frac{\vdash e : \theta \quad \vdash C : (\tau, \sigma, \delta) \text{cmd} \quad \theta \sqsubseteq \tau \quad snd(\sigma) = s \Rightarrow \sigma \sqsubseteq \tau}{\vdash \text{while e do } C : (\tau, \sigma \sqcup \theta, \uparrow) \text{cmd}} \text{T-WHILE} \\
\frac{\vdash e : \theta \quad \vdash C : (\tau, \sigma, \delta) \text{cmd} \quad \theta \sqsubseteq \tau \quad snd(\sigma) = s \wedge \delta = \uparrow \Rightarrow \sigma \sqsubseteq \tau}{\vdash \text{for e do } C : (\tau, \sigma, \delta) \text{cmd}} \text{T-FOR}
\end{array}$$

Fig. 6. The security type system

4.2 Type Soundness

In this section we prove some results about well typed programs with regard to the type system in Figure 6. The main proposition establishes that the type system indeed enforces the secret-sensitive noninterference property we defined in Section 3.

The first property is the standard notion of *subject reduction* which guarantees that execution preserves types.

Theorem 1 (Subject reduction). *If $\vdash C : (\tau, \sigma, \delta)cmd$ and $\langle M, C \rangle \xrightarrow{u} \langle M', C' \rangle$, then $\vdash C' : (\tau', \sigma', \delta')cmd$ with $(\tau', \sigma', \delta')cmd \leq (\tau, \sigma, \delta)cmd$.*

Proof. We proceed by induction on the type derivation, and then by case analysis on the last rule of operational semantics.

- T-ASSIGN: Here, $\vdash x := e : (\tau, \perp_p, \downarrow)cmd$. On every memory M , $\langle M, x := e \rangle \xrightarrow{\tau} \langle M', \mathbf{skip} \rangle$. By T-SKIP, $\vdash \mathbf{skip} : (\top_P, \perp_P, \downarrow)cmd$. And $(\top_P, \perp_P, \downarrow)cmd \leq (\tau, \perp_p, \downarrow)cmd$, because $\tau \sqsubseteq_P \top_P$ for every τ .
- T-OUT: Here, $\vdash \mathbf{output}_i(e) : ((i, s), \perp_p, \downarrow)cmd$. On every memory M , we have $\langle M, \mathbf{output}_i(e) \rangle \xrightarrow{\tau} \langle M, \mathbf{skip} \rangle$. By T-SKIP, $\vdash \mathbf{skip} : (\top_P, \perp_P, \downarrow)cmd$. And $(\top_P, \perp_P, \downarrow)cmd \leq ((i, s), \perp_p, \downarrow)cmd$, because $(i, s) \sqsubseteq_P \top_P$ for every i .
- T-SEQ1: Here, $\vdash C_1; C_2 : (\tau_1 \sqcap \tau_2, \sigma_1 \sqcup \sigma_2, \delta_1 \sqcup \delta_2)cmd$ because $\vdash C_i : (\tau_i, \sigma_i, \delta_i)cmd$ and $snd(\sigma_1) \sqsubseteq_S b$ or $\delta_1 = \downarrow$.

We distinguish two cases according to whether C-SEQ1 or C-SEQ2 is used to make a step.

- C-SEQ1: Here, $\langle M, C_1; C_2 \rangle \xrightarrow{u} \langle M', C'_1; C_2 \rangle$ because $\langle M, C_1 \rangle \xrightarrow{u} \langle M', C'_1 \rangle$. Applying induction hypothesis on C_1 , we get $\vdash C'_1 : (\tau'_1, \sigma'_1, \delta'_1)cmd$ with $(\tau'_1, \sigma'_1, \delta'_1)cmd \leq (\tau_1, \sigma_1, \delta_1)cmd$. Moreover, $snd(\sigma'_1) \sqsubseteq b$ or $\delta'_1 = \downarrow$. Now, applying rule T-SEQ1 on C'_1 and C_2 , we get $\vdash C'_1; C_2 : (\tau'_1 \sqcap \tau_2, \sigma'_1 \sqcup \sigma_2, \delta'_1 \sqcup \delta_2)cmd$. Now we can conclude :
 - * We have $\tau_1 \sqcap \tau_2 \sqsubseteq \tau'_1 \sqcap \tau_2$. By case analysis, because $\tau_1 \sqsubseteq \tau'_1$.
 - * We have $\sigma'_1 \sqcap \sigma_2 \sqsubseteq \sigma_1 \sqcap \sigma_2$ because $\sigma'_1 \sqsubseteq \sigma_1$.
 - * Finally, $\delta'_1 \sqcup \delta_2 \sqsubseteq \delta_1 \sqcup \delta_2$ because $\delta'_1 \leq \delta_1$.
- C-SEQ2: Here, C_1 is **skip** and $\langle M, \mathbf{skip}; C_2 \rangle \xrightarrow{\tau} \langle M, C_2 \rangle$. We have $\vdash C_2 : (\tau_2, \sigma_2, \delta_2)cmd$, and we can conclude since :
 - * $\tau_1 \sqcap \tau_2 \sqsubseteq \tau_2$
 - * $\sigma_2 \sqsubseteq \sigma_1 \sqcup \sigma_2$
 - * $\delta_2 \sqsubseteq \delta_1 \sqcup \delta_2$
- T-SEQ2: Here, $\vdash C_1; C_2 : (\tau_1 \sqcap \tau_2, \sigma_1 \sqcup \sigma_2, \uparrow)cmd$ because $\vdash C_i : (\tau_i, \sigma_i, \delta_i)cmd$ and $\sigma_1 \sqsubseteq \tau_2$ and $snd(\sigma_1) = s \delta_1 = \uparrow$. Only rule C-SEQ1 can be used, since $\delta_1 = \uparrow$. Thus, $\langle M, C_1; C_2 \rangle \xrightarrow{u} \langle M', C'_1; C_2 \rangle$ because $\langle M, C_1 \rangle \xrightarrow{u} \langle M', C'_1 \rangle$. Applying induction hypothesis on C_1 , we get $\vdash C'_1 : (\tau'_1, \sigma'_1, \delta'_1)cmd$ with $(\tau'_1, \sigma'_1, \delta'_1)cmd \leq (\tau_1, \sigma_1, \delta_1)cmd$.

We distinguish two cases :

- If $\delta'_1 = \downarrow$, then we can apply T-SEQ1, obtaining $\vdash C'_1; C_2 : (\tau'_1 \sqcap \tau_2, \sigma'_1 \sqcup \sigma_2, \delta'_1 \sqcup \delta_2)cmd$. We again conclude by a case analysis as previously.

- If $\delta'_1 = \uparrow$, then we can apply T-SEQ2 because $\sigma'_1 \sqsubseteq \sigma_1 \sqsubseteq \tau_2$. We get $\vdash C'_1; C_2 : (\tau'_1 \sqcap \tau_2, \sigma'_1 \sqcup \sigma_2, \uparrow)cmd$ and conclude as previously.
- T-IF: Here, \vdash **if** e **then** C_1 **else** $C_2 : (\tau_1 \sqcap \tau_2, \sigma_1 \sqcup \sigma_2 \sqcup \theta, \delta_1 \sqcup \delta_2)cmd$ because $\vdash e : \theta$ and $\vdash C_i : (\tau_i, \sigma_i, \delta_i)cmd$ and $\theta \sqsubseteq \tau_i$.

Depending on whether C-IF1 or C-IF2 is used, there are two cases :

- If C-IF1 is used, C_1 will be the next instruction, then we have to show that $(\tau_1, \sigma_1, \delta_1)cmd \leq (\tau_1 \sqcap \tau_2, \sigma_1 \sqcup \sigma_2 \sqcup \theta, \delta_1 \sqcup \delta_2)cmd$. We proceed simply by definition of \sqcap and \sqcup .
- Similarly if C-IF2 is used.
- T-FOR: Here \vdash **for** e **do** $C : (\tau, \sigma, \delta)cmd$ because $\vdash e : \theta$, $\vdash C : (\tau, \sigma, \delta)cmd$, $\theta \sqsubseteq \tau$ and $snd(\sigma) = s \wedge \delta = \uparrow \Rightarrow \sigma \sqsubseteq \tau$.

According to C-FOR, we have to type **if** n_1 **then** C ; **for** $n_1 - 1$ **do** C **else skip** where $\langle M, e \rangle \Downarrow n$ and $n_1 = |n|$. But $\vdash C : (\tau, \sigma, \delta)cmd$ and \vdash **for** $n_1 - 1$ **do** $C : (\tau, \sigma, \delta)cmd$.

There are two cases to type the sequential composition :

- If $\delta = \downarrow$, then we use rule C-SEQ1 to get $\vdash C$; **for** $n_1 - 1$ **do** $C : (\tau, \sigma, \delta)cmd$. We conclude using T-IF with \vdash **if** n_1 **then** C ; **for** $n_1 - 1$ **do** C **else skip** : $(\tau, \sigma, \delta)cmd$ (because $\vdash n_1 : \perp_P$).
- If $\delta = \uparrow$, then we can use T-SEQ2 : in fact $\sigma \sqsubseteq \tau$ by the side condition of T-FOR because $snd(\sigma) = s$ (otherwise, we would use T-SEQ1) and $\delta = \uparrow$. We get $\vdash C$; **for** $n_1 - 1$ **do** $C : (\tau, \sigma, \uparrow)cmd$. We conclude using T-IF rule, obtaining \vdash **if** n_1 **then** C ; **for** $n_1 - 1$ **do** C **else skip** : $(\tau, \sigma, \delta)cmd$.
- T-WHILE: Here \vdash **while** e **do** $C : (\tau, \sigma \sqcup \theta, \uparrow)cmd$ because $\vdash e : \theta$, $\vdash C : (\tau, \sigma, \delta)cmd$, $\theta \sqsubseteq \tau$ and $snd(\sigma) = s \Rightarrow \sigma \sqsubseteq \tau$. According to C-WHILE, we have to type **if** e **then** $(C$; **while** e **do** $C)$ **else skip**. But $\vdash C : (\tau, \sigma, \delta)cmd$ and \vdash **while** e **do** $C : (\tau, \sigma \sqcup \theta, \uparrow)cmd$.

We have two cases for typing C ; **while** e **do** C .

- If $snd(\sigma) \sqsubseteq_P b$, we can use rule T-SEQ1. Thus, $\vdash C$; **while** e **do** $C : (\tau, \sigma \sqcup \theta, \uparrow)cmd$. Using rule T-IF, we get \vdash **if** e **then** $(C$; **while** e **do** $C)$ **else skip** : $(\tau, \sigma \sqcup \theta, \uparrow)cmd$.
- If $snd(\sigma) = s$ then, by the side condition of T-WHILE, we have $\sigma \sqsubseteq \tau$. Applying T-SEQ2, we get $\vdash C$; **while** e **do** $C : (\tau, \sigma \sqcup \theta, \uparrow)cmd$. By rule T-IF, we conclude with \vdash **if** e **then** $(C$; **while** e **do** $C)$ **else skip** : $(\tau, \sigma \sqcup \theta, \uparrow)cmd$.

□

The following lemmas (using the terminology from [VSI96]) confirm that the informal definitions we gave about both components of a command type in Section 4.1 are enforced by the type system.

Lemma 3 (Simple security). *If $\vdash e : \tau$ then every variable occurring in e has type τ' var where $\tau' \sqsubseteq \tau$.*

Lemma 4 (Confinement). *If $\vdash C : (\tau, \sigma, \delta)cmd$, then every variable assigned to in program C has type θ var with $\tau \sqsubseteq \theta$.*

Lemma 5 (Guard safety). *If $\vdash C : (\tau, \sigma, \delta)cmd$, then every while loop or conditional guard in program C has type θvar with $\theta \sqsubseteq \sigma$.*

Lemma 6 (Termination). *If $\vdash C : (\tau, \sigma, \downarrow)cmd$, then C terminates on all memories.*

The proofs of these four lemmas are straightforward inductions on the respective typing derivations; the details are omitted.

In the formal development that follows for simplicity's sake we only treat the case of the three point lattice in Figure 1. The following results can be extended to the general case: for a given clearance level i in \mathcal{L} , as was depicted in the example of *finance*'s perspective in Figure 4, the refinement of \mathcal{L} can be rethought of as a three point lattice - low level, *bhigh* and *shigh* secrets.

Recall that the three point lattice is the partial product of $\{low, high\}$ and S , which yields the three ordered points $low \sqsubseteq_{\mathcal{P}} (high, b) \sqsubseteq_{\mathcal{P}} (high, s)$. In the following we write *bhigh* and *shigh* for the latter two points. A command C is said to be *shigh* or *bhigh* if there exists τ and σ such that $\vdash C : (\tau, \sigma, \delta)cmd$ with respectively $\tau = shigh$ or $bhigh \sqsubseteq \tau$.

In order to further lighten the notation, we first introduce the notion of a -equivalence on memories, for each a in the three point-lattice, as follows:

Definition 12 (Memory a -equivalence).

Let a be in $\{low, bhigh, shigh\}$. Two memories M and N are said to be a -equivalent, denoted by $M =_a N$ if for all variables x , $\Gamma(x) \sqsubseteq a \Rightarrow M(x) = N(x)$.

The goal in the type soundness proof is to build for each i a i -TSB and i -TIB over configurations of the form $\langle M, C \rangle$ and $\langle N, C \rangle$, where C is typeable. The following two propositions identify classes of configurations that can be easily shown to be bisimilar.

Proposition 3 (l -TSBisimilarity of *shigh* commands).

The relation $\{(\langle M, C \rangle, \langle N, D \rangle) \mid M =_b N; C, D \text{ are shigh}\}$ is a l -TSB.

Proof. The relation is symmetric by definition. Let C and D be two *shigh* commands, and M, N be two memories such that $M =_b N$, so we have that $M =_l N$. Suppose we have $\langle M, C \rangle \xrightarrow{u}_l \langle M', C' \rangle$. Since C is *shigh*, then the transition is silent, hence we can match the step by $\langle N, D \rangle \xrightarrow{\tau}_l \langle N, D \rangle$. The two resulting configurations stay in the relation, because *shigh*-ness is preserved along the execution of C (Lemma 1), and $M' =_b N$ because $M' =_b M$ (by Lemma 4 and *shigh*-ness of C). □

Proposition 4 (l -TIBisimilarity of *bhigh* commands).

The relation defined $\{(\langle M, C \rangle, \langle N, D \rangle) \mid M =_l N; C, D \text{ bhigh}\}$ is a l -TIB.

Proof. Similar to Proposition 3, by matching with the empty transition. □

We also need the following lemma that states the l or b -equivalence of memories is preserved along the execution of a typeable command depending on its termination effect σ .

Lemma 7.

If $\vdash C : (\tau, low, \delta)cmd$ and $M =_l N$ then $\langle M, C \rangle \xrightarrow{u}_l \langle M', C' \rangle$ implies $\langle N, C \rangle \xrightarrow{u}_l \langle N', C' \rangle$ with $M' =_l N'$.

If $\vdash C : (\tau, bhigh, \delta)cmd$ and $M =_b N$ then $\langle M, C \rangle \xrightarrow{u}_l \langle M', C' \rangle$ implies $\langle N, C \rangle \xrightarrow{u}_l \langle N', C' \rangle$ with $M' =_b N'$.

Proof. By induction on the type derivation $\vdash C : (\tau, \sigma, \delta)cmd$. We only prove the first part, the other one is similar.

- $C = x := e$ and $\vdash e : \tau$ with $\Gamma(x) = \tau$. Then, $C' = \mathbf{skip}$. If $\tau = low$ then by Lemma 3 the variables in e must all be low, and hence evaluating e in M or N yields the same value so the resulting memories stay *low-equivalent*. Otherwise, if $\tau \neq low$ then x cannot be a low variable and thus low-equivalence of memories is trivially preserved by this command.
- $C = \mathbf{output}_i(e)$. If $i = low$ then, the same value is output during the computation step. Now if i is greater, then this is a silent transition on M and N . Here, memories are not modified and stay trivially *low-equivalent*.
- $C = C_1; C_2$. Only rule T-SEQ1 can be used, so $\vdash C_1 : (\tau_1, \sigma_1, \delta_1)cmd$ and $\vdash C_2 : (\tau_2, \sigma_2, \delta_2)cmd$ with $\sigma_1 = \sigma_2 = low$. If the transition is taken by mean of rule C-SEQ1, then $C' = C'_1; C_2$ with $\langle M, C_1 \rangle \xrightarrow{u}_l \langle M', C'_1 \rangle$. Induction hypothesis on C_1 gives that $\langle N, C_1 \rangle \xrightarrow{u}_l \langle N', C'_1 \rangle$ with $M' =_l N'$. We conclude by applying rule C-SEQ1 on this transition.
Now if rule C-SEQ2 is used, then $C_1 = \mathbf{skip}$ and there is no output. We again apply induction hypothesis on C_2 and rule C-SEQ2.
- $C = \mathbf{if } e \mathbf{ then } C_1 \mathbf{ else } C_2$. Here, we have that $\vdash e : l$. But $M =_l N$ so the same branch is executed starting either from M or N , in a silent transition. Memories are not modified and hence stay *low-equivalent*.
- $C = \mathbf{while } e \mathbf{ do } D$ and $C = \mathbf{for } e \mathbf{ do } D$. Same as conditional.

□

Now we can state the main result.

Proposition 5 (Noninterference of well typed commands).

If a command C is typable, i.e., $\vdash C : (\tau, \sigma, \delta)cmd$, then C satisfies SSNI.

Proof. We use the proof technique provided by Corollary 1. The first step is to show that $\vdash C : (\tau, \sigma, \delta)cmd$ implies $\langle C, M \rangle \cong_i \langle C, N \rangle$ for all levels i , to have the s-TSNI property of Definition 9.

The case $i = high$ is vacuous: memories and commands are in this case equal. The interesting case is $i = low$.

We show that $\vdash C : (\tau, \sigma, \delta)cmd$ implies $\langle C, M \rangle \cong_l \langle C, N \rangle$ for all M and N that are equal on their low and bhigh parts, defining a relation $\mathcal{R}_1: \langle M, C \rangle \mathcal{R}_1 \langle N, D \rangle$ if and only if C and D are typable, $M^l = N^l$ and $M|_b^l = N|_b^l$, and one of the following holds:

- (i) C and D are *shigh*

- (ii) $C = D$
- (iii) $C = C_1; C_2, D = D_1; C_2$ with $\langle M, C_1 \rangle \mathcal{R}_1 \langle N, D_1 \rangle$ and C_2 is *shigh*
- (iv) C is *shigh*, $D = D_1; D_2$ with $\langle M, \mathbf{skip} \rangle \mathcal{R}_1 \langle N, D_1 \rangle$ and D_2 is *shigh*

We have to show that \mathcal{R}_1 is a l -TSB. Suppose we have $\langle M, C \rangle \mathcal{R}_1 \langle N, D \rangle$. By definition of \mathcal{R}_1 , we have that $M =_l N$. Now we have to show that if $\langle M, C \rangle \xrightarrow{u}_l \langle M', C' \rangle$ then $\langle N, D \rangle \xrightarrow{u}_l \langle N', D' \rangle$ with $\langle M', C' \rangle \mathcal{R}_1 \langle N', D' \rangle$. We proceed by induction on the definition of \mathcal{R}_1 . The reflexivity of the relation \mathcal{R}_1 is clear, except for the clause (iv) in which it will be treated.

- (i) Here we simply use Proposition 3.
- (ii) Here, $C = D$. We proceed by induction on the structure of C .
 - $C = x := e$. If $\Gamma(x) \sqsubseteq b$, then we have $\langle M, e \rangle \Downarrow n$ and $\langle N, e \rangle \Downarrow n$ by the simple security lemma because $M =_b N$. So the transition $\langle M, C \rangle \xrightarrow{\tau}_l \langle M[x \leftarrow n], \mathbf{skip} \rangle$ is matched by $\langle N, C \rangle \xrightarrow{\tau}_l \langle N[x \leftarrow n], \mathbf{skip} \rangle$. We also have that $M[x \leftarrow n] =_b N[x \leftarrow n]$ and we stay in clause (ii). If $\Gamma(x) = \mathit{shigh}$ and evaluating e in M and N does not yield the same value (which is the interesting case), we match by the same step, since memories remain equal concerning their *bhigh* (or lower) parts.
 - $C = \mathbf{output}_i(e)$. There are two cases depending on whether $i \sqsubseteq l$ or not. In the former case, since $M =_b N$ $\langle M, e \rangle \Downarrow n$ and $\langle N, e \rangle \Downarrow n$ so the same value is output. We stay again in clause (ii) because the memories are not modified. If $i = \mathit{high}$, then the transition is silent, and computations stay together since the memories are not modified.
 - $C = \mathbf{if } e \mathbf{ then } C_1 \mathbf{ else } C_2$.
If $\vdash e : \tau$ with $\tau \sqsubseteq b$, then evaluating e on M or N yields the same value because $M =_b N$. The transition matching $\langle M, C \rangle \xrightarrow{\tau}_l \langle M, C_i \rangle$ is thus $\langle N, C \rangle \xrightarrow{\tau}_l \langle N, C_i \rangle$ and we stay in clause (ii).
Now if $\vdash e : s$, according to the typing rule, $\vdash C_i : (s, \sigma, \delta) \mathit{cmd}$. So, even if the branching is different depending on whether C is run on M or N , both configurations remain related, since memories are unchanged, and both commands are *shigh* (clause (i)).
 - $C = \mathbf{while } e \mathbf{ do } C'$ and $C = \mathbf{for } e \mathbf{ do } C'$. Here we proceed as in conditional case.
 - $C = C_1; C_2$. Depending on which transition rule is used, we have two cases :
 1. If C-SEQ1 is used, we hence have $\langle M, C_1; C_2 \rangle \xrightarrow{u}_l \langle M', C'_1; C_2 \rangle$ with $\langle M, C_1 \rangle \xrightarrow{u}_l \langle M', C'_1 \rangle$. We distinguish two cases whether C_2 is *shigh* or not.
 - (a) If $\vdash C_2 : (\tau_2, \sigma_2, \delta_2) \mathit{cmd}$ with $\tau_2 \sqsubseteq b$, only rule T-SEQ1 can be used. Then $\vdash C_1 : (\tau_1, \sigma_1, \delta_1) \mathit{cmd}$ with $\sigma_1 \sqsubseteq b$ so C_1 does not contain a *shigh* guard. Using Lemma 7, we have $\langle N, C_1 \rangle \xrightarrow{u}_l \langle N', C'_1 \rangle$ with $M' =_b N'$. Applying rule C-SEQ1 gives that $\langle N, C_1; C_2 \rangle \xrightarrow{u}_l \langle N', C'_1; C_2 \rangle$. We stay again in clause (ii).
 - (b) If C_2 is *shigh*. Applying the induction hypothesis on C_1 gives that $\langle N, C_1 \rangle \xrightarrow{u}_l \langle N', C''_1 \rangle$ with $M' =_b N'$ and $\langle M', C'_1 \rangle \mathcal{R}_1 \langle N', C''_1 \rangle$. We then obtain by rule C-SEQ1 that $\langle N, C_1; C_2 \rangle \xrightarrow{u}_l \langle N', C''_1; C_2 \rangle$. We finally have that $\langle M', C'_1; C_2 \rangle \mathcal{R}_1 \langle N', C''_1; C_2 \rangle$ by clause (iii).

2. If C-SEQ2 is used, then $\langle M, \mathbf{skip}; C_2 \rangle \xrightarrow{\tau}_l \langle M, C_2 \rangle$. The matching step is simply $\langle N, \mathbf{skip}; C_2 \rangle \xrightarrow{\tau}_l \langle N, C_2 \rangle$ and the configuration are related by clause (ii).
- (iii) $C = C_1; C_2$, $D = D_1; C_2$ with $\langle M, C_1 \rangle \mathcal{R}_1 \langle N, D_1 \rangle$ and C_2 is *shigh*.
1. If C-SEQ1 is used, then $\langle M, C_1; C_2 \rangle \xrightarrow{u}_l \langle M', C'_1; C_2 \rangle$ and $\langle M, C_1 \rangle \xrightarrow{u}_l \langle M', C'_1 \rangle$. Applying the induction hypothesis on D_1 leads to $\langle N, D_1 \rangle \xrightarrow{u}_l \langle N', D'_1 \rangle$ with $M' =_b N'$. We then apply rule C-SEQ1 and stay in clause (iii).
 2. If C-SEQ2 is used, then $C_1 = \mathbf{skip}$ and $\langle M, \mathbf{skip}; C_2 \rangle \xrightarrow{\tau}_l \langle M, C_2 \rangle$. The matching move is then $\langle N, D \rangle \xrightarrow{\tau}_l \langle N, D \rangle$. We have that $\langle M, C_2 \rangle \mathcal{R}_1 \langle N, D \rangle$ by clause (iv) because C_2 is *shigh* and $\langle M, \mathbf{skip} \rangle \mathcal{R}_1 \langle N, D_1 \rangle$.
- (iv) Here, C is *shigh*, $D = D_1; D_2$, with $\langle M, \mathbf{skip} \rangle \mathcal{R}_1 \langle N, D_1 \rangle$ and D_2 is *shigh*.
1. Suppose we have $\langle M, C \rangle \xrightarrow{u}_l \langle M', C' \rangle$. C is *shigh*, so $M =_b M'$ and $u = \tau$. The matching step is then $\langle N, D \rangle \xrightarrow{\tau}_l \langle N, D \rangle$ and we stay in clause (iv).
 2. Suppose now we have $\langle N, D \rangle \xrightarrow{u}_l \langle N', D' \rangle$. If rule C-SEQ1 is used, then $\langle N, D_1 \rangle \xrightarrow{u}_l \langle N', D'_1 \rangle$ and applying the induction hypothesis on D_1 gives us $\langle M, \mathbf{skip} \rangle \xrightarrow{u}_l \langle M', C'' \rangle$. The only transition from \mathbf{skip} is itself so $M' = M$ and $C'' = \mathbf{skip}$. It gives us also that $u = \tau$. We can thus match the transition with $\langle M, C \rangle \xrightarrow{\tau}_l \langle M, C \rangle$ and we stay in clause (iv).
If rule C-SEQ2 is used, then $D_1 = \mathbf{skip}$ and $\langle N, D \rangle \xrightarrow{u}_l \langle N, D_2 \rangle$. So $u = \tau$. The matching step is hence again $\langle M, C \rangle \xrightarrow{\tau}_l \langle M, C \rangle$ and we conclude using clause (i) because D_2 is *shigh*.

\mathcal{R}_1 is a l -TSB. By clause (ii) and Proposition 1, since $M =_b N$ implies $M^l = N^l$ and $M|_b^l = N|_b^l$, we have that in a well typed program, there is no flow from *shigh* data to *bhigh* and *low* data.

The next step is to prove that the type system ensures TINI concerning the *bhigh*-data.

We proceed a similar way, providing a l -TIB \mathcal{R}_2 over configurations. The relation \mathcal{R}_2 is defined: $\langle M, C \rangle \mathcal{R}_2 \langle N, D \rangle$ if and only if C and D are typable, $M =_l N$, and one of the following holds:

- (i) C and D are *bhigh*
- (ii) $C = D$
- (iii) $\langle M, C \rangle \mathcal{R}'_2 \langle N, D \rangle$, the relation \mathcal{R}'_2 being defined inductively as follows :

$$\frac{C, D \text{ bhigh}}{\langle M, C; C' \rangle \mathcal{R}'_2 \langle N, D; C' \rangle} \mathcal{R}^1 \quad \frac{\langle M, C \rangle \mathcal{R}'_2 \langle N, D \rangle}{\langle N, C; C' \rangle \mathcal{R}'_2 \langle N, D; C' \rangle} \mathcal{R}^2$$

We now have to prove that \mathcal{R}_2 is a l -TIB. We will use the property that \mathcal{R}_2 is preserved by sequential composition, *i.e.* $\langle M, C \rangle \mathcal{R}_2 \langle N, D \rangle$ implies $\langle M, C; C' \rangle \mathcal{R}_2 \langle N, D; C' \rangle$. This result can be easily proved by case analysis on the definition of the relation. The relation \mathcal{R}_2 is clearly symmetric. Suppose we have $\langle M, C \rangle \mathcal{R}_2 \langle N, D \rangle$. By definition of \mathcal{R}_2 , we have that $M =_l N$. Now we have to show that if $\langle M, C \rangle \xrightarrow{u}_l \langle M', C' \rangle$ then either $\langle N, D \rangle \xrightarrow{u}_l \langle N', D' \rangle$ with $\langle M', C' \rangle \mathcal{R}_2 \langle N', D' \rangle$ or $\langle N, D \rangle \uparrow_l$. We proceed by case analysis on the definition of \mathcal{R}_2 .

- (i) Here we simply use Proposition 4.
- (ii) $C = D$. We proceed by induction on the structure of C .
 - $C = C_1; C_2$. Depending on the semantic rule that is used, we have two cases:
 1. C-SEQ1: We have that $\langle M, C_1; C_2 \rangle \xrightarrow{u}_l \langle M', C'_1; C_2 \rangle$ because $\langle M, C_1 \rangle \xrightarrow{u}_l \langle M', C'_1 \rangle$. Applying the induction hypothesis on C_1 , $\langle N, C_1 \rangle \xrightarrow{u}_l \langle N', C''_1 \rangle$ with $\langle M, C'_1 \rangle \mathcal{R}_2 \langle N', C''_1 \rangle$. But \mathcal{R}_2 is preserved by sequential composition, so $\langle M, C'_1; C_2 \rangle \mathcal{R}_2 \langle N', C''_1; C_2 \rangle$.
 2. C-SEQ2: Here, $C_1 = \mathbf{skip}$ and $\langle M, C \rangle \xrightarrow{\tau}_l \langle M, C_2 \rangle$. We simply have that $\langle N, C \rangle \xrightarrow{\tau}_l \langle N, C_2 \rangle$ and the configurations stay together thanks to the clause (ii).
 - Other cases are similar to those of the relation \mathcal{R}_1 .
- (iii) We proceed by induction on the definition of \mathcal{R}'_2 .
 - R1: Here, $C = C_1; C_2$, $D = D_1; C_2$ and C_1, D_1 are *bhigh*. Here, only rule C-SEQ1 can be used, since **skip** cannot be typed as a *bhigh* command. We thus have that $\langle M, C_1; C_2 \rangle \xrightarrow{u}_l \langle M', C'_1; C_2 \rangle$ because $\langle M, C_1 \rangle \xrightarrow{u}_l \langle M', C'_1 \rangle$. But C_1 is *bhigh*, so $u = \tau$, $M =_l M'$ and C'_1 is *bhigh*. We can match this step by $\langle N, D_1; C_2 \rangle \xrightarrow{\tau}_l \langle N, D_1; C_2 \rangle$. Configurations are in the relation by rule R1.
 - R2: Note that only rule C-SEQ1 can be used: if $C_1 = \mathbf{skip}$, it cannot be related to D , since \mathcal{R}'_2 only relates sequential compositions. Then, we have that $\langle M, C_1; C_2 \rangle \xrightarrow{u}_l \langle M', C'_1; C_2 \rangle$ because $\langle M, C_1 \rangle \xrightarrow{u}_l \langle M', C'_1 \rangle$. But we have $\langle M, C_1 \rangle \mathcal{R}'_2 \langle N, D_1 \rangle$, so by induction hypothesis on D_1 :
 - * If $\langle N, D_1 \rangle \xrightarrow{u}_l \langle N', D'_1 \rangle$, we match using rule C-SEQ1 by $\langle N, D_1; C_2 \rangle \xrightarrow{u}_l \langle N', D'_1; C_2 \rangle$ and we can conclude by rule R2.
 - * If $\langle N, D_1 \rangle \uparrow_l$ then since D_1 cannot be **skip** (from the definition of \mathcal{R}'_2) we know that $\langle N, D_1 \rangle \xrightarrow{\tau}_l \langle N', D'_1 \rangle$ for some N', D'_1 . By the symmetry of relation \mathcal{R}'_2 and induction hypothesis, we get that $\langle M, C_1 \rangle \xrightarrow{\tau}_l \langle M'', C''_1 \rangle$ with $\langle N', D'_1 \rangle \mathcal{R}_2 \langle M'', C''_1 \rangle$ or $\langle M, C_1 \rangle \uparrow_l$. Thus, by rule C-SEQ1, the transition we had originally to match is a silent one, and we can match it by $\langle N, D_1; C_2 \rangle \xrightarrow{\tau}_l \langle N', D'_1; C_2 \rangle$. Configurations stay related by rule R2.

\mathcal{R}_2 is thus a *l*-TIB, and by clause (ii) and Proposition 2, we have the TINI property of well typed programs concerning their *bhigh*-data: there is no flow from *bhigh* data to *low* data except via the termination channel. \square

5 Correlation Leaks

In this section we mention a weakness in the definition of secret-sensitive noninterference which allows the attacker to observe *correlations* between big and small secrets. We show how the definition can be strengthened to remove such correlations, and conjecture that the type system guarantees correlation-freedom without need for modification.

Suppose that b is *bhigh* and s is *shigh* (in the lattice in Figure 1). Somewhat surprisingly the program $\mathbf{output}_{low}(b == s)$ is secret-sensitive noninterfering (note though that it is not typeable). This is because the low observer can say nothing about the value of e.g. s in isolation. The problem is that although the observer cannot deduce anything

about the individual kinds of secret, he can deduce information about their *correlation* (in this example whether they are equal or not).

To eliminate the possibility of learning something about the correlation of big and small secrets we need to demand that the knowledge learnt about big and small secrets together is the same as for the combined knowledge learnt about them independently. To express this precisely we need some additional notation.

In the definitions of secret-sensitive noninterference we have dealt with knowledge as sets of projections of memories. We say that a memory M is *full* if $\text{dom}(M)$ is the set of all variables. In order to easily compare and combine knowledge sets we need to work with full memories. Define M^* to be the set of full memories obtainable by completing M :

$$M^* = \{N \mid N|_{\text{dom}(M)} = M, N \text{ is full}\}.$$

Now lift \cdot^* to sets of memories K in the natural way by defining

$$K^* = \bigcup_{M \in K} M^*$$

Definition 13 (Correlation Freedom). *A program C is Correlation Free if for all $\vec{v} \in \text{Obs}_i(C, M^i)$, we have $k_i^{bs}(C, M^i, \vec{v})^* = k_i^b(C, M^i, \vec{v})^* \cap k_i^s(C, M^i, \vec{v})^*$, where $k_i^{bs}(C, M^i, \vec{v}) = \{M|_i \mid M \in k_i(C, M^i, \vec{v})\}$ and $M|_i$ is the complement of M^i – i.e., the projection of M to the variables not visible at level i .*

In the case that C is secret-sensitive noninterfering we can show that this condition is equivalent to $k_i^{bs}(C, M^i, \vec{v})^* = k_i^b(C, M^i, \vec{v})^*$, which says that nothing more is learnt about the big and small secrets together than can be deduced from the big secrets alone.

Conjecture 1 *Well-typed programs are correlation free.*

We leave the proof of this conjecture to further work; the intuition here is that any “correlation information” will always be typed as s -level data, and hence cannot be leaked at all.

6 Conclusions

In this article we provided a way to refine an arbitrary complex security lattice in order to distinguish two levels of secret, the big secrets b and the small ones s . Big secrets can be handled more liberally on the grounds that they can be made sufficiently large and random for slow leakage to be tolerable. We introduced an accompanying notion of secret-sensitive noninterference which combines the relative merits of termination-sensitive and termination-insensitive noninterferences. We illustrated the use of the definition in the soundness argument for a simple type system for verifying secret-sensitive noninterference.

Related Work As mentioned previously, the starting point of this work is [AHSS08]. Our interpretation of the results there is that we need to treat different kinds of secret in different ways, and to our knowledge this paper is the first to do so in a noninterference setting. It is, however, relatively common to give a special treatment to cryptographic keys as compared to other kinds of secret – e.g. [AHS06] – but usually the goal here is to deal with integrity (a key cannot be modified using a low value) or freshness (a key cannot be used more than once).

Our type system is essentially a fusion of a type-based version of Denning’s system [VSI96], and a stricter system based on [BC01]. The latter system is stricter than a Denning-style analysis for quite a different purpose: to deal with multi-threaded programs. Our system, in a sequential setting, improves on [BC01] by additionally tracking whether a program is terminating.

Further Work A natural and interesting next step would be to combine such a type system with cryptographic primitives (e.g. [Vol00][LV05][AHS06]). The notion of “big” and “small” secrets have a natural interpretation in the cryptographic setting, since “big” secrets correspond to e.g. cryptographic keys. In such a setting it might also be important to handle “size integrity”, so that one could know that a variable is not only independent of small secrets, but that it *is* a big secret.

Acknowledgements Thanks to Andrei Sabelfeld for pointing out the correlation problem discussed in Section 5, and to Niklas Broberg, Josef Svenningsson, David Pichardie, Thomas Jensen and the anonymous referees for very helpful comments on an earlier draft. This work was partly supported by grants from the Swedish funding agencies SSF, Vinnova (The Swedish Governmental Agency for Innovation Systems), VR, and by the European IST-2005-015905 MOBIUS project.

References

- AHS06. A. Askarov, D. Hedin, and A. Sabelfeld. Cryptographically-masked flows. In *Proc. Symp. on Static Analysis*, LNCS, pages 353–369. Springer-Verlag, August 2006.
- AHSS08. A. Askarov, S. Hunt, A. Sabelfeld, and D. Sands. Termination-insensitive noninterference leaks more than just a bit. In *ESORICS 2008, 13th European Symposium on Research in Computer Security*, volume 5283 of LNCS. Springer Verlag, 2008.
- AS07. A. Askarov and A. Sabelfeld. Gradual release: Unifying declassification, encryption and key release policies. In *Proc. IEEE Symp. on Security and Privacy*, pages 207–221, May 2007.
- BB03. J. Barnes and JG Barnes. *High Integrity Software: The SPARK Approach to Safety and Security*. Addison-Wesley Longman Publishing Co., Inc. Boston, MA, USA, 2003.
- BC01. G. Boudol and I. Castellani. Noninterference for concurrent programs. In *Proc. ICALP’01*, volume 2076 of LNCS, pages 382–395. Springer-Verlag, July 2001.
- CH04. R. Chapman and A. Hilton. Enforcing security and safety models with an information flow analysis tool. *ACM SIGAda Ada Letters*, 24(4):39–46, 2004.
- DD77. D. E. Denning and P. J. Denning. Certification of programs for secure information flow. *Comm. of the ACM*, 20(7):504–513, July 1977.
- Den76. D. E. Denning. A lattice model of secure information flow. *Comm. of the ACM*, 19(5):236–243, May 1976.
- LV05. P. Laud and V. Vene. A type system for computationally secure information flow. In *Proc. Fundamentals of Computation Theory*, volume 3623 of LNCS, pages 365–377, August 2005.
- MZZ⁺08. A. C. Myers, L. Zheng, S. Zdancewic, S. Chong, and N. Nystrom. Jif: Java information flow. Software release. Located at <http://www.cs.cornell.edu/jif>, July 2001–2008.
- Sim03. V. Simonet. The Flow Caml system. Software release. Located at <http://crystal.inria.fr/~simonet/soft/flowcaml/>, July 2003.
- Smi01. G. Smith. A new type system for secure information flow. In *Proc. IEEE Computer Security Foundations Workshop*, pages 115–125, June 2001.
- Vol00. D. Volpano. Secure introduction of one-way functions. In *CSFW ’00: Proceedings of the 13th IEEE workshop on Computer Security Foundations*, page 246, Washington, DC, USA, 2000. IEEE Computer Society.
- VS97. D. Volpano and G. Smith. Eliminating covert flows with minimum typings. *Proc. IEEE Computer Security Foundations Workshop*, pages 156–168, June 1997.
- VSI96. D. Volpano, G. Smith, and C. Irvine. A sound type system for secure flow analysis. *J. Computer Security*, 4(3):167–187, 1996.