

Statically Analyzing Information Flows

An Abstract Interpretation-based Hyperanalysis for Non-Interference

Isabella Mastroeni

University of Verona - Dept. of Computer Science
Verona, Italy
isabella.mastroeni@univr.it

Michele Pasqua*

University of Verona - Dept. of Computer Science
Verona, Italy
michele.pasqua@univr.it

ABSTRACT

In the context of systems security, information flows play a central role. Unhandled information flows potentially leave the door open to very dangerous types of attacks, such as code injection or sensitive information leakage. Information flows verification is based on the definition of Non-Interference [8], which is known to be an hyperproperty [7], i.e., a property of sets of executions. The sound verification of hyperproperties is not trivial [3, 16]: It is not easy to adapt classic verification methods, used for trace properties, in order to deal with hyperproperties. In the present work, we design an abstract interpretation-based static analyzer soundly checking Non-Interference. In particular, we define an hyper abstract domain, able to approximate the information flows occurring in the analyzed programs.

CCS CONCEPTS

• **Software and its engineering** → **Software verification**; • **Security and privacy** → *Information flow control*;

KEYWORDS

Static analysis, Information flows, Abstract interpretation, Hyperproperties verification, Non-Interference

ACM Reference Format:

Isabella Mastroeni and Michele Pasqua. 2019. Statically Analyzing Information Flows: An Abstract Interpretation-based Hyperanalysis for Non-Interference. In *The 34th ACM/SIGAPP Symposium on Applied Computing (SAC '19)*, April 8–12, 2019, Limassol, Cyprus. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3297280.3297498>

1 INTRODUCTION

Programs security often relies on how information is propagated during programs executions. Historically, in the context of confidentiality, *access control* has been the main means of preventing information from being disseminated. As the name indicates, access control verifies the program rights at entry-point. However, it is inadequate in many situations, in fact the program may

leak/compromise information after the access check. Instead *information flow control* tracks how information propagates through the program during execution to make sure that it handles the information safely. A security mechanism, given a program and a specification stating which flows of information are allowed and which are not, checks whether the program is secure w.r.t. the given specification. In general, unhandled information flows potentially leave the door open to very dangerous types of attacks, such as code injection or sensitive information leakage.

Information flow control is based on the definition of Non-Interference [8], where variables are marked as public or private, and flows from the public ones to the private ones are forbidden. Indeed, a program is said to be *non-interferent* if the values of the public outputs (i.e., at the end of the execution) do not depend on the values of the private inputs (i.e., at the beginning of the execution), in other words if whenever we start from states agreeing on the public inputs, the public observable outputs are the same.

The problem of Non-Interference verification is far from solved. Indeed, Non-Interference is a hard problem to solve due to two main reasons: it is undecidable and it is not even a trace property, in the sense that it cannot be verified on single executions, since it requires the comparison of multiple computations. In other words it is an undecidable *hyperproperty* [7].

Both these problems are tackled by approximation, in particular by designing *trace properties* stronger than Non-Interference, but which can be verified by means of standard techniques, inherited from classic programs verification for trace properties. Unfortunately, this means that we introduce two degrees of approximation, one for approximating Non-Interference as a trace property and one for making it decidable.

Whilst for trace properties standard analysis techniques can be used with an acceptable degree of approximation, the verification, and hence the approximation, of hyperproperties is a hard problem to solve in a sufficiently precise way. Being a hyperproperty means to be modeled as a set of sets, and the extra level of sets introduces a lot of technical problems for approximation-based verification methods. Hence, in order to obtain significant results w.r.t. analysis precision, we need new verification methods which approximate sets of sets (instead of just sets) of executions [16].

In [17], the authors propose a general framework for designing abstract interpretations for hyperproperties verification. In the present work, our aim is to instantiate this framework to the verification of Non-Interference, allowing us to design a prototype analyzer of Non-Interference for a toy programming language. In particular, we first design the abstract domain denoting the hyperproperty that has to be satisfied by the semantics of non-interfering programs, i.e., the values of public variables must be always the same single

*Corresponding author.

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SAC '19, April 8–12, 2019, Limassol, Cyprus

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5933-7/19/04...\$15.00

<https://doi.org/10.1145/3297280.3297498>

value. Unfortunately, this ideal domain (observing the precise single public values of all the computations starting from the same public input) is infinite and does not guarantee decidability of the analysis, hence we need to introduce further abstractions. Once we fix the domain of the abstract observable hyperproperty we define the abstract hypersemantics computing a program's executions on the abstract domain, exactly as it happens in classic static analysis, but with the only difference that we are abstracting a hypersemantics into a hyperdomain. To the best of our knowledge, this is the first attempt to provide a sound analyzer of Non-Interference exploiting the expressiveness of hyperproperties.

Finally, in order to test the feasibility of the proposed approach we have implemented a prototype (called nonInterfer) of the designed analyzer. This prototype exhibits a good trade-off between verification speed and precision. It can be tested at the link <http://bit.do/noninterfer>, via a web interface.

2 BACKGROUND

In this section, we set the background concepts needed to introduce our prototype analyzer. First, we define a toy programming language for which we design the analyzer. Then we make a brief introduction to hyperproperties and, finally, we have two little paragraphs, one for abstract interpretation and one for Non-Interference.

2.1 Programs

Programs are written in a simple deterministic imperative language Imp, whose grammar is the following:

$\text{Aexp} \ni a ::= x \mid n \mid a \oplus a \mid (a)$
 $\text{Bexp} \ni b ::= \text{tt} \mid \text{ff} \mid b \wedge b \mid b \vee b \mid \neg b \mid a \bowtie a \mid (b)$
 $\text{Imp} \ni P ::= \text{skip} \mid x := a \mid P ; P \mid \text{if } b \{P\} \text{ else } \{P\} \mid \text{while } b \{P\}$
 with $\oplus \in \{+, -, *\}$ and $\bowtie \in \{=, \neq, <, \leq\}$

Variables $x \in \text{Var}$ range over integer values, hence arithmetic expressions evaluate to values in \mathbb{Z} . Boolean expressions evaluate to boolean values in $\mathbb{B} \triangleq \{\text{tt}, \text{ff}\}$. The semantics of the language is given on top of *memories*, namely maps from variables to values. Let $\text{Mem} \triangleq \text{Var} \rightarrow \mathbb{Z}$ be the set of programs memories. We define the semantics of a program inductively from its syntax. In particular, it is built on top of the small-step operational semantics (SOS) of Imp. This latter models the execution of programs step by step and it is specified by a set of inference rules modifying *configurations*. A configuration $\langle m, P \rangle \in \text{Mem} \times \text{Imp}$ represents the current memory m in which the program P has to be executed. The SOS inference rules in Fig. 1 describe how configurations evolve during time, until a final configuration of the form $\langle m, \text{skip} \rangle$ is reached (if ever). The SOS rules rely on the semantics for arithmetic expressions $\llbracket a \rrbracket m \in \mathbb{Z}$ and for boolean expressions $\llbracket b \rrbracket m \in \mathbb{B}$. These latter are big-step semantics, since we are not interested in the intermediate steps of computation for expressions. In Fig. 2 we have the definition of the (big-step) semantics for expressions. We denote with \rightarrow^k the application of \rightarrow a number $k > 1$ of times, hence $\langle m, P \rangle \rightarrow^k \langle m', P' \rangle$ if there exists a sequence σ of k configurations such that: $\sigma_0 = \langle m, P \rangle$, $\sigma_{k-1} = \langle m', P' \rangle$ and $\forall i \in [0, k-1]. \sigma_i \rightarrow \sigma_{i+1}$. We denote with $\langle m, P \rangle \rightarrow^* \langle m', P' \rangle$ the fact that there exists a $k > 1$ such that $\langle m, P \rangle \rightarrow^k \langle m', P' \rangle$.

A program P , starting in the memory m , terminates, yielding the memory m' , iff $\langle m, P \rangle \rightarrow^* \langle m', \text{skip} \rangle$. Conversely, P diverges (on m) iff it is possible to apply \rightarrow to $\langle m, P \rangle$ infinitely many times.

2.2 Hyperproperties

In the field of information security, *verification* is the general process of checking if a system complies with a specification, i.e., a formal description of what systems are allowed and/or are not allowed to do. The majority of works about verification deal with particular specifications, those expressible with the so called *trace properties*, often simply called “properties”. They are defined in terms of single executions and hence cannot express specifications which need to take into account relations between executions. In [7], *hyperproperties* were introduced in order to formalize those specifications which are not trace properties. When systems are programs in Imp then program executions can be modeled as *traces* of state denotations. Examples of traces could be the set of all possible sequences of memories generated by the SOS of programs or the set of reachable memories, etc. The behavior (semantics) \mathcal{S}^P of a system (program) P is modeled as a set of traces, one for each possible input. In this setting, trace properties are sets of traces and hyperproperties are sets of sets of traces. The satisfiability relation is the set inclusion for trace property and the set membership for hyperproperties. Namely, P satisfies the trace property \mathcal{P} , written $P \models \mathcal{P}$, iff $\mathcal{S}^P \subseteq \mathcal{P}$ and it satisfies the hyperproperty $\mathcal{H}p$, written $P \models \mathcal{H}p$, iff $\mathcal{S}^P \in \mathcal{H}p$. The satisfiability for hyperproperties can be restated using set inclusion, in fact $\mathcal{S}^P \in \mathcal{H}p$ iff $\{\mathcal{S}^P\} \subseteq \mathcal{H}p$. The strongest program trace property of P is precisely \mathcal{S}^P and its strongest hyperproperty is $\{\mathcal{S}^P\}$ [10, 16]. Strongest here means that every property[hyperproperty] satisfied by P is implied by (i.e., contains) $\mathcal{S}^P[\{\mathcal{S}^P\}]$. In order to prove that a program P satisfies

$\frac{\langle m, P_1 \rangle \rightarrow \langle m', P_3 \rangle}{\langle m, P_1 ; P_2 \rangle \rightarrow \langle m', P_3 ; P_2 \rangle}$	$\frac{-}{\langle m, \text{skip} ; P \rangle \rightarrow \langle m, P \rangle}$
$\frac{\llbracket b \rrbracket m = \text{tt}}{\langle m, \text{if } b \{P_1\} \text{ else } \{P_2\} \rangle \rightarrow \langle m, P_1 \rangle}$	$\frac{\llbracket b \rrbracket m = \text{ff}}{\langle m, \text{if } b \{P_1\} \text{ else } \{P_2\} \rangle \rightarrow \langle m, P_2 \rangle}$
$\frac{\llbracket b \rrbracket m = \text{tt}}{\langle m, \text{while } b \{P\} \rangle \rightarrow \langle m, P ; \text{while } b \{P\} \rangle}$	$\frac{\llbracket b \rrbracket m = \text{ff}}{\langle m, \text{while } b \{P\} \rangle \rightarrow \langle m, \text{skip} \rangle}$
$\frac{\llbracket a \rrbracket m = n}{\langle m, x := a \rangle \rightarrow \langle m[x \leftarrow n], \text{skip} \rangle}$	

Figure 1: Small-step operational semantics of Imp.

Arithmetic expressions: $\llbracket a \rrbracket m \in \text{Mem} \rightarrow \mathbb{Z}$	
$\llbracket n \rrbracket m = n$	$\llbracket x \rrbracket m = m(x)$
$\llbracket a_1 \oplus a_2 \rrbracket m = \llbracket a_1 \rrbracket m \oplus \llbracket a_2 \rrbracket m$	where $\oplus \in \{+, -, *\}$
Boolean expressions: $\llbracket b \rrbracket m \in \text{Mem} \rightarrow \mathbb{B}$	
$\llbracket \text{tt} \rrbracket m = \text{tt}$	$\llbracket \text{ff} \rrbracket m = \text{ff}$
$\llbracket (b) \rrbracket m = \llbracket b \rrbracket m$	$\llbracket \neg b \rrbracket m = \neg \llbracket b \rrbracket m$
$\llbracket b_1 \wedge b_2 \rrbracket m = \llbracket b_1 \rrbracket m \wedge \llbracket b_2 \rrbracket m$	$\llbracket b_1 \vee b_2 \rrbracket m = \llbracket b_1 \rrbracket m \vee \llbracket b_2 \rrbracket m$
$\llbracket a_1 \bowtie a_2 \rrbracket m = \llbracket a_1 \rrbracket m \bowtie \llbracket a_2 \rrbracket m$	where $\bowtie \in \{=, \neq, <, \leq\}$

Figure 2: Big-step semantics for expressions.

a property[hyperproperty] we use an over-approximation of the strongest P property[hyperproperty].

Some specifications can be expressed as trace properties, like Access Control, and others cannot, like Non-Interference. Among all hyperproperties there are some easier to verify. In [17], the authors define k -bounded subset-closed hyperproperties: these specifications can be refuted just exhibiting a counterexample set consisting in at most k traces. It turns out that several interesting hyperproperties are k -bounded (like Non-Interference). In Sec. 4 we will see how we can exploit this fact in order to verify Non-Interference.

2.3 Abstract Interpretation

The (concrete) semantic of a program, namely its behavior, is a representation of all its possible executions by means of a set of mathematical objects. This set is, in general, not computable: It is not possible for a program to represent and to compute all possible executions of any program starting from all its possible inputs. Clearly all non trivial properties of the concrete semantics of a program are undecidable: It is not possible for a program to answer any question about the executions of any program. Abstract interpretation is born, as a theory for soundly approximating the semantics of discrete dynamic systems. The approximation consists in the observation of the semantics at a specified level of abstraction, focusing only on some important aspects of computations. In this setting, abstract interpretation allows us to compute an abstract semantics of the program, depending on the properties of interest. The approximation is sound by design, in the sense that what holds in the abstract holds also in the concrete (no false negative).

One of the fundamental aspects of abstract interpretation is that the majority of the features of the approximation process are specified only by the (abstract) domain of mathematical objects chosen for representing the properties of interest. A theory of domains for abstract interpretation was defined in [9], based on the notion of *Galois insertion*. A Galois insertion (C, α, γ, A) consists of two partially ordered sets (C, \leq_C) , (A, \leq_A) and two monotone functions $\alpha \in C \rightarrow A$, $\gamma \in A \rightarrow C$ such that for all c in C and a in A it holds: $\alpha(c) \leq_A a \Leftrightarrow c \leq_C \gamma(a)$ and $\alpha \circ \gamma = id$ (the identity function $\lambda x. x$). C is the concrete domain, A is the abstract domain, α is the abstraction function and γ is the concretization function. Sometimes, abstract interpretations are given by means of Galois connections (instead of Galois insertions), relaxing the constraints $\alpha \circ \gamma = id$. This does not restrict the generality of the framework, since every connection can be transformed in an insertion, eliminating the redundant elements in the abstract domain.

Let $f \in C \rightarrow C$ be a function on the concrete domain and $f^h \in A \rightarrow A$ be a function on the abstract domain. f^h is a sound (or correct) approximation of f if $f \circ \gamma = \gamma \circ f^h$ or, equivalently, if $\alpha \circ f = f^h \circ \alpha$ [9]. The *best correct approximation* of f in A , defined as $f^{bca} \triangleq \alpha \circ f \circ \gamma$, is sound by construction [9]. Hence, in order to prove that a given abstract function f^h is sound, it is sufficient to prove that it approximates f^{bca} , i.e., if $\forall a \in A. f^{bca}(a) \leq_A f^h(a)$.

2.4 Non-Interference

Computing system security relies on how information is propagated during system execution. Information flow control tracks how information propagates through the system during execution

to make sure that the system handles the information securely. In the last forty years, a lot of specifications for information flow control have been proposed (see [14] for a survey). They differ on the formal definitions but share the following common informal understanding of confidentiality: the lack of dependencies on confidential information. This is precisely the absence of strong dependency of [8] and it is the base of confidentiality information flow specifications. The idea of [8] is that there exists an information flow from variable x to variable y in a program P whenever variety in x is conveyed to y by the execution of P . The original formulation of Non-Interference [8] takes in consideration only two security levels: *private* (H), i.e., information that has to be kept secret, and *public* (L), i.e., information that could be freely released. A program is said non-interferent if there are no information flows from private to public variables, and interferent otherwise.

3 SIMPLIFYING NON-INTERFERENCE VERIFICATION

In the following, we define Non-Interference for programs in Imp and we show how to derive a verification mechanism for it. The classic notion of Non-Interference [8] checks the input/output relation between executions, so we represent programs computations with just the initial and the final memories. Furthermore, it is not termination sensitive, hence we ignore divergent computations. With these premises, the set of executions denotations is $\text{Mem} \times \text{Mem}$. It is well known that information flows are hyperproperties [7], indeed Non-Interference $\text{NI} \in \wp(\wp(\text{Mem} \times \text{Mem}))$ is defined as:

$$\text{NI} \triangleq \left\{ X \subseteq \text{Mem} \times \text{Mem} \mid \forall \langle m_1, m'_1 \rangle, \langle m_2, m'_2 \rangle \in X. \begin{array}{l} m_1 =_L m_2 \Rightarrow m'_1 =_L m'_2 \end{array} \right\}$$

The relation $=_L$ says that memories agree on public variables, i.e., $m =_L m'$ iff $\forall x. \Gamma(x) = L \Rightarrow m(x) = m'(x)$. Here $\Gamma \in \text{Var} \rightarrow \{L, H\}$ is the typing environment, assigning variables to security levels.

Given a program $P \in \text{Imp}$, its semantics (i.e., its strongest property) on $\text{Mem} \times \text{Mem}$ is $^*S^P \in \wp(\text{Mem}^P \times \text{Mem}^P)$, defined as $^*S^P \triangleq \{ \langle m, m' \rangle \mid \langle m, P \rangle \rightarrow^* \langle m', \text{skip} \rangle \}$. Its hypersemantics (i.e., its strongest hyperproperty) is $^*HS^P \triangleq \{ ^*S^P \}$. Hence we have that $P \models \text{NI}$ iff $^*S^P \in \text{NI}$ or, equivalently, iff $^*HS^P \subseteq \text{NI}$.

Now we will see how the verification process for Non-Interference can be made simpler. Since NI is a 2-bounded subset-closed hyperproperty [17], we have that $P \models \text{NI}$ iff $^*HS^{P\#2} \triangleq \{ X \subseteq ^*S^P \mid |X| = 2 \} \subseteq \text{NI}$. We can note that $^*HS^{P\#2}$ can be partitioned in

$$^*HS_{\neq}^{P\#2} \triangleq \{ \{ \langle m_1, m'_1 \rangle, \langle m_2, m'_2 \rangle \} \mid m_1 \neq_L m_2 \} \text{ and}$$

$$^*HS_{=}_L^{P\#2} \triangleq \{ \{ \langle m_1, m'_1 \rangle, \langle m_2, m'_2 \rangle \} \mid m_1 =_L m_2 \}$$

namely, $^*HS_{\neq}^{P\#2} \cup ^*HS_{=}_L^{P\#2} = ^*HS^{P\#2}$ and $^*HS_{\neq}^{P\#2} \cap ^*HS_{=}_L^{P\#2} = \emptyset$. At this point, as far as NI is concerned, we can observe that $^*HS_{\neq}^{P\#2} \subseteq \text{NI}$ always holds, hence we have the following result.

PROPOSITION 3.1. $P \models \text{NI}$ iff $^*HS_{=}_L^{P\#2} \subseteq \text{NI}$.

Example 3.2. Consider $P = x := y; y := 1$, with $\Gamma(x) = L$ and $\Gamma(y) = H$. Suppose that variables can only take values 0 and 1. In order to simplify, consider the following encoding of memories: $a \triangleq [x \mapsto 0, y \mapsto 0]$, $b \triangleq [x \mapsto 0, y \mapsto 1]$, $c \triangleq [x \mapsto 1, y \mapsto 0]$, $d \triangleq [x \mapsto 1, y \mapsto 1]$. The semantics of P is $^*S^P = \{ \langle a, b \rangle, \langle b, d \rangle, \langle c, b \rangle, \langle d, d \rangle \}$. Then $^*HS_{\neq}^{P\#2} = \{ \{ \langle a, b \rangle, \langle b, d \rangle \}, \{ \langle c, b \rangle, \langle d, d \rangle \} \}$ and $^*HS_{=}_L^{P\#2} = \{ \{ \langle a, b \rangle, \langle c, b \rangle \}, \{ \langle a, b \rangle, \langle d, d \rangle \}, \{ \langle b, d \rangle, \langle c, b \rangle \}, \{ \langle b, d \rangle, \langle d, d \rangle \} \}$. ■

The verification process for Non-Interference can be simplified even further. Take in consideration the following approximation $\alpha_r \in \wp(\wp(\text{Mem} \times \text{Mem})) \rightarrow \wp(\wp(\text{Mem}) \times \wp(\text{Mem}))$:

$$\alpha_r \triangleq \lambda X. \{ \langle \{a \mid \langle a, b \rangle \in X\}, \{b \mid \langle a, b \rangle \in X\} \rangle \mid X \in \mathcal{X} \}$$

which abstracts the relation between the input and the output memories of traces (i.e., we only keep the relation between the whole set of input memories and the whole set of output memories). Then consider a further approximation $\alpha_n \in \wp(\wp(\text{Mem}) \times \wp(\text{Mem})) \rightarrow \wp(\wp(\text{Mem})) \times \wp(\wp(\text{Mem}))$:

$$\alpha_n \triangleq \lambda X. \langle \{A \mid \langle A, B \rangle \in X\}, \{B \mid \langle A, B \rangle \in X\} \rangle$$

which abstracts the relation between sets of traces (i.e., we only keep the relation between the set of all the possible input sets of memories and the set of all the output sets of memories). Finally, setting $\alpha_{nr} \triangleq \alpha_n \circ \alpha_r$ we have that α_{nr} and its left adjoint α_{nr}^- form the Galois connection¹:

$$\langle \wp(\wp(\text{Mem} \times \text{Mem})), \subseteq \rangle \xleftrightarrow[\alpha_{nr}]{\alpha_{nr}^-} \langle \wp(\wp(\text{Mem})) \times \wp(\wp(\text{Mem})), \subseteq \rangle$$

Now we are able to state a simplified form of verification check for Non-Interference:

$$P \models \text{NI} \Leftrightarrow \alpha_{nr}(\times \mathcal{HS}_{=}^{P\#2}) \subseteq \alpha_{nr}(\text{NI})$$

We can note that $\alpha_{nr}(\times \mathcal{HS}_{=}^{P\#2}) \subseteq \alpha_{nr}(\text{NI})$ iff the set $\alpha_{nr}(\times \mathcal{HS}_{=}^{P\#2})_{\downarrow}$ contains only sets of memories $\{m, m'\}$ agreeing on the L variables, i.e., such that $m =_L m'$. Here, given a pair $\langle x, y \rangle$, we denote with $\langle x, y \rangle_{\downarrow} \triangleq y$ its projection on the second element (analogously, we denote with $\langle x, y \rangle_{\uparrow} \triangleq x$ its projection on the first element).

Example 3.3. Continuing Ex. 3.2, we have that $\alpha_{nr}(\times \mathcal{HS}_{=}^{P\#2})$ is equal to $\alpha_n(\{ \langle \{a, b\}, \{b, d\} \rangle, \langle \{c, d\}, \{b, d\} \rangle \})$, which is the pair $\langle \{ \{a, b\}, \{c, d\} \}, \{ \{b, d\} \} \rangle$. In this case, program P does not satisfy Non-Interference since $\alpha_{nr}(\times \mathcal{HS}_{=}^{P\#2})_{\downarrow} = \{ \{b, d\} \}$ and $b \neq_L d$. ■

Basically, we can verify Non-Interference just checking whether $\alpha_{nr}(\times \mathcal{HS}_{=}^{P\#2})_{\downarrow}$ satisfies a hyperproperty on the set of execution denotations Mem, as state by the following proposition.

PROPOSITION 3.4. $P \models \text{NI}$ iff $\alpha_{nr}(\times \mathcal{HS}_{=}^{P\#2})_{\downarrow} \subseteq \text{equiv}^L$, where $\text{equiv}^L \triangleq \{ X \subseteq \text{Mem} \mid \forall m, m' \in X. m =_L m' \}$.

This simplifies a lot the verification process for Non-Interference: We can build a hypersemantics computing on $\wp(\wp(\text{Mem}))$ instead of $\wp(\wp(\text{Mem} \times \text{Mem}))$.

4 ABSTRACT HYPERSEMANTICS

Classic methods for the verification of trace properties rely on the so called collecting semantics, namely a representation of a program's behaviors starting from every possible input. This semantics essentially computes the strongest program property. With a computable approximation of this latter the verification process is made feasible. Unfortunately, for hyperproperties this is not sufficient anymore: a collecting hypersemantics, computing the strongest program hyperproperty, is needed [3, 16]. Then the verification is performed with an approximation of this latter. In this section we give a collecting hypersemantics and its abstract version for the verification

¹We abuse notation denoting with \subseteq both set-inclusion and its component-wise extension to pairs, i.e., $\langle \mathcal{X}, \mathcal{Y} \rangle \subseteq \langle \mathcal{T}, \mathcal{Z} \rangle \triangleq (\mathcal{X} \subseteq \mathcal{T} \wedge \mathcal{Y} \subseteq \mathcal{Z})$.

of Non-Interference. In particular, we follow the construction introduced in [17] for k -bounded subset-closed hyperproperties (indeed Non-Interference is 2-bounded). First, we give the definition of the collecting hypersemantics, computing at the level of sets of sets. Then we instantiate the hyperlevel constants domain of [17] to Non-Interference. This latter, in its original formulation, is not machine-representable [11], namely it has an uncountable set of elements. Moreover, it contains infinite ascending chains (i.e., it is not ACC), inducing potential computation divergence. Hence, we approximate it in order to make its implementation feasible. Finally, we give the definition of the abstract semantics.

4.1 Collecting Hypersemantics

As we have seen in Sec. 3 (Prop. 3.4), we can verify Non-Interference in a simpler domain, namely we can collect memories instead of input/output traces of memories. We move from $\wp(\text{Mem} \times \text{Mem})$ to $\wp(\text{Mem})$, with the abstraction $\alpha_r \triangleq \lambda X. \{ m' \mid \langle m, m' \rangle \in X \}$ collecting all final memories of input/output traces. Hence, we have that the strongest trace property of P induced by its semantics on the domain Mem is $\mathcal{R}S^P \triangleq \alpha_r(\times \mathcal{S}^P) = \{ m' \mid \exists m. \langle m, P \rangle \rightarrow^* \langle m', \text{skip} \rangle \}$. Analogously, the strongest hyperproperty of P on Mem is $\mathcal{R}HS^P \triangleq \{ \alpha_r(\times \mathcal{S}^P) \} = \{ \mathcal{R}S^P \}$.

In order to compute this semantics, we follow the construction given in [17] (Sect. 3.1). We start with a classic collecting semantics $\llbracket P \rrbracket \in \wp(\text{Mem}) \rightarrow \wp(\text{Mem})$, defined inductively on programs statements, computing post-conditions:

$$\begin{aligned} \llbracket P \rrbracket \emptyset &\triangleq \emptyset & \llbracket P_1 ; P_2 \rrbracket X &\triangleq \llbracket P_2 \rrbracket \llbracket P_1 \rrbracket X & \llbracket \text{skip} \rrbracket X &\triangleq X \\ \llbracket x := a \rrbracket X &\triangleq \{ m \mid x \leftarrow \langle a \rangle m \mid m \in X \} \\ \llbracket \text{if } b \{ P_1 \} \text{ else } \{ P_2 \} \rrbracket X &\triangleq \llbracket P_1 \rrbracket \llbracket b \rrbracket X \cup \llbracket P_2 \rrbracket \llbracket \neg b \rrbracket X \\ \llbracket \text{while } b \{ P \} \rrbracket X &\triangleq \llbracket \neg b \rrbracket (lfp_{\emptyset}^{\subseteq} \lambda T. X \cup \llbracket P \rrbracket \llbracket b \rrbracket T) \end{aligned}$$

Here $\llbracket b \rrbracket \in \wp(\text{Mem}) \rightarrow \wp(\text{Mem})$ is the classic filtering function which selects the memories making b true, namely $\llbracket b \rrbracket X \triangleq \{ m \in X \mid \langle b \rangle m = \text{tt} \}$. Note that $\llbracket P \rrbracket$ computes exactly the strongest trace property of P, namely $\llbracket P \rrbracket \text{Mem} = \mathcal{R}S^P$. Then we have to lift this semantics on sets of sets, obtaining a collecting hypersemantics $\llbracket P \rrbracket^h \in \wp(\wp(\text{Mem})) \rightarrow \wp(\wp(\text{Mem}))$:

$$\begin{aligned} \llbracket P \rrbracket^h \emptyset &\triangleq \emptyset & \llbracket P_1 ; P_2 \rrbracket^h X &\triangleq \llbracket P_2 \rrbracket^h \llbracket P_1 \rrbracket^h X & \llbracket \text{skip} \rrbracket^h X &\triangleq X \\ \llbracket x := a \rrbracket^h X &\triangleq \{ \{ x := a \} X \mid X \in \mathcal{X} \} \\ \llbracket \text{if } b \{ P_1 \} \text{ else } \{ P_2 \} \rrbracket^h X &\triangleq \{ \llbracket P_1 \rrbracket \llbracket b \rrbracket X \cup \llbracket P_2 \rrbracket \llbracket \neg b \rrbracket X \mid X \in \mathcal{X} \} \\ \llbracket \text{while } b \{ P \} \rrbracket^h X &\triangleq \llbracket \neg b \rrbracket^h (lfp_{\emptyset}^{\subseteq} F) \text{ where} \\ F &\triangleq \lambda \mathcal{T}. X \cup \{ \llbracket P \rrbracket \llbracket b \rrbracket T \cup \llbracket \neg b \rrbracket T \mid T \in \mathcal{T} \} \end{aligned}$$

This latter semantics is very similar to the *hypercollecting semantics* of [3], instantiated on sets of sets of memories instead of sets of sets of input/output traces. Again, $\llbracket b \rrbracket^h \in \wp(\wp(\text{Mem})) \rightarrow \wp(\wp(\text{Mem}))$ is a filtering function, namely $\llbracket b \rrbracket^h X \triangleq \{ \{ m \in X \mid \langle b \rangle m = \text{tt} \} \mid X \in \mathcal{X} \} \setminus \{ \emptyset \}$. In this case $\llbracket P \rrbracket^h$ is not exactly the strongest hyperproperty of P, but it is a correct approximation, i.e., $\mathcal{R}HS^P \subseteq \llbracket P \rrbracket^h \{ \text{Mem} \}$ [17]. Luckily, for subset-closed hyperproperties, as equiv^L , this is not a problem because $\llbracket P \rrbracket^h$ is complete for such hyperproperties, namely $P \models \text{equiv}^L$ iff $\llbracket P \rrbracket^h \{ \text{Mem} \} \subseteq \text{equiv}^L$ [17]. Nevertheless, we are not interested in the strongest hyperproperty of P, since

it is not necessary for proving Non-Interference. Instead, due to Prop. 3.4, it is sufficient to compute $\alpha_{nr}(\times \mathcal{HS}^{\text{P}\#2})_{\perp}$ and this boils down to computing $\llbracket P \rrbracket^h$ starting from $\mathcal{I}^{2\perp} \triangleq \{\{m, m'\} \mid m =_{\perp} m'\}$. In fact, it is easy to see that $\alpha_{nr}(\times \mathcal{HS}^{\text{P}\#2})_{\perp}$ and $\llbracket P \rrbracket^h \mathcal{I}^{2\perp}$ coincide with the set:

$$\left\{ \{m'_1, m'_2\} \subseteq \text{Mem} \mid \begin{array}{l} \exists m_1, m_2. \langle m_1, P \rangle \rightarrow^* \langle m'_1, \text{skip} \rangle \wedge \\ \langle m_2, P \rangle \rightarrow^* \langle m'_2, \text{skip} \rangle \wedge \\ m_1 =_{\perp} m_2 \end{array} \right\}$$

This observation induces the following proposition.

PROPOSITION 4.1. $P \models \text{NI}$ iff $\llbracket P \rrbracket^h \mathcal{I}^{2\perp} \subseteq \text{equiv}^{\perp}$.

4.2 Towards Abstraction

Unfortunately, $\llbracket P \rrbracket^h \mathcal{I}^{2\perp}$ and equiv^{\perp} are not computable in general, hence we need approximations. In order to compute a sound approximation of $\llbracket P \rrbracket^h \mathcal{I}^{2\perp}$, we rely on abstract interpretation. The first step is to define the abstract domain used to verify Non-Interference.

4.2.1 The Abstract Domain for Non-Interference. The domain is an instance of the *hyperlevel (abstract) constants* of [17], and it checks whether a set of sets of values contains constant sets. In the following, all abstraction functions α are additive (i.e., they preserve the least upper bound of chains), hence their left adjoint α^- always exists². From now on, we will often use the operator $(\text{cond} ? \text{doTrue} : \text{doFalse})$ in order to denote in-line conditional definitions, namely its semantics is: if the condition cond is true then substitute the operator with doTrue otherwise substitute it with doFalse . Mathematically:

$$(\text{cond} ? \text{doTrue} : \text{doFalse}) = \begin{cases} \text{doTrue} & \text{if cond is true} \\ \text{doFalse} & \text{otherwise} \end{cases}$$

Furthermore, we will sometimes use the pointwise extension of relations and operators. Given a relation $\preceq \subseteq Y \times Y$, an operator $\gamma \in Y \times Y \rightarrow Y$ and a set X , the pointwise extension $\preceq \subseteq (X \rightarrow Y) \times (X \rightarrow Y)$, of \preceq , is defined as $f \preceq f' \triangleq (\forall x \in X. f(x) \preceq f'(x))$ and the pointwise extension $\dot{\gamma} \in (X \rightarrow Y) \times (X \rightarrow Y) \rightarrow (X \rightarrow Y)$, of γ , is defined as $f \dot{\gamma} f' \triangleq \lambda x \in X. f(x) \gamma f'(x)$.

The domain is an abstraction of $\wp(\wp(\mathbb{Z}))$. In [17], the authors claim that a hyperdomain, i.e., a domain suitable for the verification of hyperproperties, can be decomposed in an *inner* abstraction, approximating traces, and in an *outer* abstraction, approximating properties of traces. Following this idea, we use as inner abstraction the one for the classic constant propagation domain: it represents precisely the singletons $\{n\}$, for every $n \in \mathbb{Z}$, and it abstracts to \mathbb{Z} everything else. Instead, the outer abstraction checks whether the inner abstraction always returns constant values, not necessarily the same. For instance, $\{\{1\}, \{2\}\}$ contains only constant sets, instead $\{\{1\}, \{2, 3\}\}$ contains also a not-constant set (in this case the inner abstraction maps $\{2, 3\}$ to \mathbb{Z}).

We follow [17] defining $C_{\text{hc}} \triangleq \wp(\{\{n\} \mid n \in \mathbb{Z}\}) \cup \{\wp(\mathbb{Z})\}$, $\alpha_{\text{hc}} \triangleq \lambda X. (X \subseteq \{\{n\} \mid n \in \mathbb{Z}\} ? X : \wp(\mathbb{Z}))$ and $\gamma_{\text{hc}} \triangleq \text{id}$. Then we have the Galois insertion:

$$\langle \wp(\wp(\mathbb{Z})), \subseteq \rangle \xleftrightarrow[\alpha_{\text{hc}}]{\gamma_{\text{hc}}} \langle C_{\text{hc}}, \subseteq \rangle$$

This domain, as proved in [17], is sufficient for Non-Interference verification, but, as observed at the beginning of this section, it is

not machine-representable. For this reason we need to perform a further approximation.

The domain C_{hc} has an uncountable set of elements, so we define a simpler domain, which is machine-representable but still able to verify Non-Interference. Let $\bar{\mathbb{Z}}$ be a set isomorphic to \mathbb{Z} , aiming at representing sets containing only one singleton, i.e., of the form $\{\{n\}\}$, which is the information we want to observe precisely, and let $\bar{\cdot} \in \mathbb{Z} \rightarrow \bar{\mathbb{Z}}$ a bijection. Furthermore, we denote by κ the abstract element representing the set of all singletons. Then we define $C^h \triangleq \{\bar{n} \mid n \in \mathbb{Z}\} \cup \{\perp, \top, \kappa\}$, with the partial order $\leq \subseteq C^h \times C^h$ defined as $c_1 \leq c_2 \triangleq (c_1 = \perp \vee c_1 = c_2 \vee (c_1 = \bar{n} \wedge c_2 = \kappa) \vee c_2 = \top)$. Now consider the abstraction $\alpha_h \in C_{\text{hc}} \rightarrow C^h$:

$$\alpha_h(X) \triangleq \begin{cases} \perp & \text{if } X = \emptyset \\ \bar{n} & \text{if } X = \{\{n\}\} \\ \kappa & \text{if } X \subseteq \{\{n\} \mid n \in \mathbb{Z}\} \wedge |X| > 1 \\ \top & \text{otherwise} \end{cases}$$

Its corresponding concretization $\gamma_h \in C^h \rightarrow C_{\text{hc}}$ is such that $\gamma_h(\perp) = \emptyset$, $\gamma_h(\bar{n}) = \{\{n\}\}$, $\gamma_h(\kappa) = \{\{n\} \mid n \in \mathbb{Z}\}$ and $\gamma_h(\top) = \wp(\mathbb{Z})$. Then we have the following Galois insertion:

$$\langle C_{\text{hc}}, \subseteq \rangle \xleftrightarrow[\alpha_h]{\gamma_h} \langle C^h, \leq \rangle$$

The domain $\langle C^h, \leq, \nabla, \Delta, \perp, \top \rangle$ is a complete lattice where $\nabla, \Delta \in C^h \times C^h \rightarrow C^h$ are:

$$c_1 \nabla c_2 \triangleq \begin{cases} \top & \text{if } c_1 = \top \vee c_2 = \top \\ \perp & \text{if } c_1 = \perp \wedge c_2 = \perp \\ \bar{n} & \text{if } (c_1 = \perp \wedge c_2 = \bar{n}) \vee (c_1 = \bar{n} \wedge c_2 = \perp) \\ & \vee (c_1 = \bar{n} \wedge c_2 = \bar{n}) \\ \kappa & \text{otherwise} \end{cases}$$

$$c_1 \Delta c_2 \triangleq \begin{cases} \top & \text{if } c_1 = \top \wedge c_2 = \top \\ \perp & \text{if } c_1 = \perp \vee c_2 = \perp \vee (c_1 = \bar{n} \neq \bar{m} = c_2) \\ \bar{n} & \text{if } (c_1 = \bar{n} \wedge c_2 \in \{\kappa, \top\}) \vee (c_2 = \bar{n} \wedge c_1 \in \{\kappa, \top\}) \\ \kappa & \text{otherwise} \end{cases}$$

By composition, $\alpha_h \triangleq \alpha_{\text{hc}} \circ \alpha_{\text{hc}}$ and $\gamma_h \triangleq \gamma_{\text{hc}} \circ \gamma_{\text{hc}}$ form the insertion:

$$\langle \wp(\wp(\mathbb{Z})), \subseteq \rangle \xleftrightarrow[\alpha_h]{\gamma_h} \langle C^h, \leq \rangle$$

This domain approximates the set of sets of values a variable may have. Finally, in order to track information flows we need to work on memories instead of on values.

Consider a “double” non-relational abstraction for sets of sets of memories. Let $\text{Mem} \triangleq \text{Var} \rightarrow \wp(\wp(\mathbb{Z}))$ and $\alpha_{\text{nnr}} \in \wp(\wp(\text{Mem})) \rightarrow \text{Mem}$ be $\alpha_{\text{nnr}}(X) \triangleq \lambda x. \{\{m(x) \mid m \in X\} \mid X \in \wp(\text{Mem})\}$. We have the following Galois connection, with $\gamma_{\text{nnr}} \triangleq \alpha_{\text{nnr}}^-$:

$$\langle \wp(\wp(\text{Mem})), \subseteq \rangle \xleftrightarrow[\alpha_{\text{nnr}}]{\gamma_{\text{nnr}}} \langle \text{Mem}, \subseteq \rangle$$

Now we can compose point-wise (the machine-representable version of) the hyperlevel constants abstraction with α_{nnr} , obtaining $\alpha_m \triangleq \lambda X. \lambda x. \alpha_h(\alpha_{\text{nnr}}(X)(x))$. This latter forms, paired with $\gamma_m \triangleq \alpha_m^-$, the Galois connection:

$$\langle \wp(\wp(\text{Mem})), \subseteq \rangle \xleftrightarrow[\alpha_m]{\gamma_m} \langle \text{Var} \rightarrow C^h, \leq \rangle$$

²This is a sufficient condition in order to form a Galois connection.

We denote with Mem^h the set $\text{Var} \rightarrow C^h$ and we call its elements *m abstract memories*. In order to simplify the notation, we let $\sqsubseteq \triangleq \leq, \sqcup \triangleq \vee, \sqcap \triangleq \wedge, \perp \triangleq \lambda x. \perp$ and $\top \triangleq \lambda x. \top$. We have that $(\text{Mem}^h, \sqsubseteq, \sqcup, \sqcap, \perp, \top)$ is a complete lattice. This latter is indeed our abstract hyperdomain of computation.

Example 4.2. Let us show in a simple example how this abstraction works. Continuing Ex. 3.3, we have: $\mathcal{I}^{2L} = \{\{a, b\}, \{c, d\}\}$ and $\alpha_m(\{\{a, b\}, \{c, d\}\}) = [x \mapsto \alpha_h(\{\{0\}, \{1\}\}) y \mapsto \alpha_h(\{\{0, 1\}\})] = [x \mapsto \kappa y \mapsto \top]$. Indeed a, b both provide 0 to x and c, d both provide 1 to x , while a, b provide both 0 and 1 to y and c, d provide both 0 and 1 to y . Similarly, $\llbracket P \rrbracket^h \mathcal{I}^{2L} = \{\{b, d\}\}$ and $\alpha_m(\{\{b, d\}\}) = [x \mapsto \alpha_h(\{\{0, 1\}\}) y \mapsto \alpha_h(\{\{1\}\})] = [x \mapsto \top y \mapsto \bar{1}]$. ■

It is worth noting that, given $X \subseteq \text{equiv}^L$, then every set in X contains L-equivalent memories. This implies $\alpha_m(X)(x) \neq \top$, for each L variable x . So, the Non-Interference check $\llbracket P \rrbracket^h \mathcal{I}^{2L} \subseteq \text{equiv}^L$ becomes equivalent to checking whether $\alpha_m(\llbracket P \rrbracket^h \mathcal{I}^{2L})(x) \neq \top$, for each L variable x . Indeed the program of Ex. 4.2 is interferent, since computing $\alpha_m(\llbracket P \rrbracket^h \mathcal{I}^{2L})$ results in L variable x having value \top .

Finally, we can show how the abstract domain Mem^h can be used for Non-Interference verification, with the following theorem.

THEOREM 4.3. *Let $m^{\text{NI}} \in \text{Mem}^h$ be the abstract memory defined as $m^{\text{NI}}(x) \triangleq (\Gamma(x) = L ? \kappa : \top)$. Then $P \models \text{NI}$ iff $\alpha_m(\llbracket P \rrbracket^h \mathcal{I}^{2L}) \sqsubseteq m^{\text{NI}}$.*

PROOF. By definition of Galois connection: $\alpha_m(\llbracket P \rrbracket^h \mathcal{I}^{2L}) \sqsubseteq m^{\text{NI}}$ iff $\llbracket P \rrbracket^h \mathcal{I}^{2L} \subseteq \gamma_m(m^{\text{NI}})$. It is clear that $\gamma_m(m^{\text{NI}}) \subseteq \text{equiv}^L$, since the concretization of m^{NI} contains L-equivalent memories, by construction. Hence, due to Prop. 4.1, the theorem is proved. □

Hence, in order to verify Non-Interference it is sufficient an abstract semantics, computing on Mem^h , which approximates $\llbracket P \rrbracket^h$. Indeed we can prove $P \models \text{NI}$ by computing an over-approximation of $\llbracket P \rrbracket^h \mathcal{I}^{2L}$ in Mem^h .

4.2.2 The Abstract Semantics for Non-Interference. Finally, we have to show how to compute a program's collecting hypersemantics on the proposed abstract domain. The abstract semantics for programs relies on the abstract semantics for boolean and arithmetic expressions, given in Fig. 3, 4. The abstract semantics for arithmetic expressions $(\langle a \rangle)^h \in \text{Mem}^h \rightarrow C^h$ evaluates to an abstract value and it relies on the abstract (mathematical) operations given in Fig. 3. This semantics is such that abstract assignments are sound approximations of the concrete ones. We obtain this defining abstract operations \oplus^h such that they are sound w.r.t. the concrete ones \oplus , with $\oplus \in \{+, -, *\}$. This technically means that they satisfy the following constraint³:

$$\{X \ni n \oplus m \in Y \mid X \in \gamma_h(c_1), Y \in \gamma_h(c_2)\} \subseteq \gamma_h(c_1 \oplus^h c_2)$$

The constraint basically requires that every possible result obtained applying the concrete operation is contained in the concretization of the application of the abstract operator.

The abstract semantics for booleans $(\langle b \rangle)^h \in \text{Mem}^h \rightarrow \text{Mem}^h$ is an abstract filtering function. To simplify, we assume that all negations \neg have been removed using DeMorgan's laws and usual arithmetic laws: $\neg(b_1 \vee b_2) \equiv (\neg b_1) \wedge (\neg b_2)$, $\neg(a_1 < a_2) \equiv (a_2 \leq a_1)$,

³Recall that our abstract semantics is built after the "double" non-relational abstraction α_{nnr} , hence it abstracts $\text{Mem} = \text{Var} \rightarrow \wp(\wp(\mathbb{Z}))$.

Arithmetic expressions: $\langle a \rangle^h \in \text{Mem}^h \rightarrow C^h$

$\langle n \rangle^h m \triangleq \bar{n}$
 $\langle x \rangle^h m \triangleq m(x)$
 $\langle a \rangle^h m \triangleq \langle a \rangle^h m$

 $\langle a_1 \oplus a_2 \rangle^h m \triangleq \langle a_1 \rangle^h m \oplus^h \langle a_2 \rangle^h m$

with $\oplus \in \{+, -, *\}$

Abstract operations: $\oplus^h \in C^h \times C^h \rightarrow C^h$

\oplus^h	\perp	\bar{n}	κ	\top
\perp	\perp	\perp	\perp	\perp
\bar{m}	\perp	$\overline{m \oplus n}$	κ	\top
κ	\perp	κ	κ	\top
\top	\perp	\top	\top	\top

Figure 3: Abstract semantics for arithmetic expressions.

Boolean expressions: $(\langle b \rangle)^h \in \text{Mem}^h \rightarrow \text{Mem}^h$		$\bowtie \in \{=, \neq, <, \leq\}$
$(\langle \text{tt} \rangle)^h m \triangleq m$	$(\langle \text{ff} \rangle)^h m \triangleq \perp$	$(\langle b \rangle)^h m \triangleq (\langle b \rangle)^h m$
$(\langle b_1 \vee b_2 \rangle)^h m \triangleq \text{let } n = (\langle b_1 \rangle)^h m \sqcup (\langle b_2 \rangle)^h m \text{ in } \lambda x. (m(x) = \top \wedge x \in \text{vars}(b_1) \cap \text{vars}(b_2) ? \top : n(x))$		
$(\langle b_1 \wedge b_2 \rangle)^h m \triangleq (\langle b_1 \rangle)^h m \sqcap (\langle b_2 \rangle)^h m$		
$(\langle a_1 \bowtie a_2 \rangle)^h m \triangleq \text{let } \langle c_1, c_2 \rangle = (\langle a_1 \rangle)^h m \bowtie^h (\langle a_2 \rangle)^h m \text{ in } \sqcup \{n \sqsubseteq m \mid (\langle a_1 \rangle)^h n \leq c_1\} \sqcap \sqcup \{n \sqsubseteq m \mid (\langle a_2 \rangle)^h n \leq c_2\}$		
Abstract comparators $\bowtie^h \in C^h \times C^h \rightarrow C^h \times C^h$:		
$c_1 \bowtie^h c_2 \triangleq \begin{cases} \langle \perp, \perp \rangle & \text{if } c_1 = \perp \text{ or } c_2 = \perp \\ \langle \bar{n}, \bar{n} \rangle & \text{if } c_1 = c_2 = \bar{n}, \bowtie \in \{=, \leq\} \\ \langle \bar{n}, \bar{m} \rangle & \text{if } c_1 = \bar{n}, c_2 = \bar{m}, n < m, \bowtie \in \{<, \leq, \neq\} \\ \langle \bar{n}, \bar{m} \rangle & \text{if } c_1 = \bar{n}, c_2 = \bar{m}, n > m, \bowtie \in \{\neq\} \\ \langle \top, \top \rangle & \text{otherwise} \end{cases}$		

Figure 4: Abstract semantics for boolean expressions.

etc. This semantics must be sound w.r.t. the collecting hypersemantics for booleans, namely $\llbracket b \rrbracket^h \gamma_m(m) \subseteq \gamma_m(\llbracket b \rrbracket)$. In order to obtain this, we have defined the abstract comparators \bowtie^h such that they are sound w.r.t. the concrete ones \bowtie , with $\bowtie \in \{=, \neq, <, \leq\}$. This technically means that they satisfy the following constraint:

$$\begin{aligned} \text{let } X = \{ \{ \langle n \in X, m \in Y \mid n \bowtie m \rangle \mid X \in \gamma_h(c_1), Y \in \gamma_h(c_2) \} \text{ in } \\ \{ \{ n \mid \langle n, m \rangle \in X \} \mid X \in \mathcal{X} \} \subseteq \gamma_h((c_1 \bowtie^h c_2)_\top) \\ \{ \{ m \mid \langle n, m \rangle \in X \} \mid X \in \mathcal{X} \} \subseteq \gamma_h((c_1 \bowtie^h c_2)_\top) \end{aligned}$$

The constraint basically requires that every possible pair of values making true the concrete comparator is contained in the concretization of the application of the abstract comparator. In Fig. 4, the term $\sqcup \{n \sqsubseteq m \mid (\langle a \rangle)^h n \leq c\}$ can be computed with a *backward* abstract semantics for expressions (as we actually do in the implementation) but we omit here its definition for space reasons. The function $\text{vars}(b)$ returns the set of variables occurring in b .

Example 4.4. Let us see how to compute $(\langle x < 2 \rangle)^h m$ where $m = [x \mapsto \bar{1} y \mapsto \kappa]$. We have $(\langle x \rangle)^h m = \bar{1}$, $(\langle 2 \rangle)^h m = \bar{2}$ and $\bar{1} <^h \bar{2} = \langle \bar{1}, \bar{2} \rangle$. Then $\sqcup \{n \sqsubseteq m \mid (\langle x \rangle)^h n \leq \bar{1}\} = [x \mapsto \bar{1} y \mapsto \kappa]$ and $\sqcup \{n \sqsubseteq m \mid (\langle 2 \rangle)^h n \leq \bar{2}\} = [x \mapsto \bar{1} y \mapsto \kappa]$. The intersection is, indeed, equal to m . Suppose now to compute the negation of $x < 2$, namely we want to compute $(\langle 2 \leq x \rangle)^h m$. In this case we have $\bar{2} \leq^h \bar{1} = \langle \top, \top \rangle$ and $\sqcup \{n \sqsubseteq m \mid (\langle 2 \rangle)^h n \leq \top\} = \sqcup \{n \sqsubseteq m \mid (\langle x \rangle)^h n \leq \top\} = [x \mapsto \bar{1} y \mapsto \kappa]$. Hence we obtain $[x \mapsto \bar{1} y \mapsto \kappa]$ as result. ■

Finally, we need two auxiliary functions $\text{vars}_m^\top(b)$ and $\text{vars}^=(P)$, returning the set of variables occurring in b having value \top when evaluated in m and the set of *modified variables* in P , respectively. The first is easy to compute: $\text{vars}_m^\top(b) \triangleq \{x \in \text{vars}(b) \mid m(x) = \top\}$ and $\text{vars}(b)$ is just a syntactic check. The second involves semantic information, hence it is not trivial to compute. At the moment, we adopt a simple syntactic approach for approximating the set of variables which may be modified during P executions. Indeed, the function $\text{vars}^=(P)$ returns the set of variables occurring in P on the left-hand side of an assignment, which is easy implementable as a syntactic check. We plan to enhance our abstract semantics with a semantic check for modified variables, in order to increase precision, as a future work.

Now we have all the ingredients needed to define the abstract collecting hypersemantics $\llbracket P \rrbracket_h^h \in \text{Mem}^h \rightarrow \text{Mem}^h$, which is:

$$\begin{aligned} \llbracket P \rrbracket_h^h \perp_h &\triangleq \perp_h & \llbracket P_1 ; P_2 \rrbracket_h^h m &\triangleq \llbracket P_2 \rrbracket_h^h \llbracket P_1 \rrbracket_h^h m \\ \llbracket \text{skip} \rrbracket_h^h m &\triangleq m & \llbracket x := a \rrbracket_h^h m &\triangleq m[x \leftarrow \langle a \rangle_h m] \\ \llbracket \text{if } b \{ P_1 \} \text{ else } \{ P_2 \} \rrbracket_h^h m &\triangleq w \text{ where} \\ \text{let } n &= \llbracket P_1 \rrbracket_h^h \langle b \rangle_h m \sqcup \llbracket P_2 \rrbracket_h^h \langle \neg b \rangle_h m \text{ in} \\ w &= \lambda x. \begin{cases} n(x) & \text{if } x \notin \text{vars}^=(P_1) \cup \text{vars}^=(P_2) \vee \\ & (n(x) \leq \kappa \wedge n(x) \neq \kappa) \vee \text{vars}_m^\top(b) = \emptyset \\ \top & \text{otherwise} \end{cases} \\ \llbracket \text{while } b \{ P \} \rrbracket_h^h m &\triangleq \langle \neg b \rangle_h \llbracket \text{if } b \{ P \} \text{ else } \{ \text{skip} \} \rrbracket_h^h n \end{aligned}$$

The abstract semantics is quite standard for all statements, except for conditionals. We will explain here only this latter, which exploits the following idea. For every variable, we make the join between its value resulting after the execution of the true branch and its value resulting after the execution of the false branch. This is done in order to track the forbidden flows (implicit or explicit) generated inside the two branches. In fact, a L variable has value \top after the join if in at least one of the branches it has value \top (meaning that there is a forbidden flow). After this check we need to take in consideration the implicit flows generated by the conditional statement itself. Indeed, first we suppose that if there is at least one variable with value \top *before* the boolean guard is evaluated, then all variables modified in the conditional branches have a forbidden flow (a variable has value \top only if it is a H variable or if it has been “influence” by a H variable). This is done setting to \top all modified variables. Note that if, for some reasons, a H variable is not \top during this check, the flow is correctly not set. This procedure is sound but not so precise. In order to enhance precision, we exploit our abstract domain. In particular, we do not set to \top the variables which have the same constant value \bar{n} in both branches (this is the condition $(n(x) \leq \kappa \wedge n(x) \neq \kappa)$) because this means that at the end of the conditional statement the variable has always a constant value. If there are no \top -valued variables into b (this is the condition $\text{vars}_m^\top(b) = \emptyset$), then no variables are set to \top : the resulting flows are those generated into the two branches of the conditional.

We proved that our abstract semantics is sound w.r.t. the collecting hypersemantics, and that it can be used for Non-Interference verification. This is stated in the following two theorems.

THEOREM 4.5 (SOUNDNESS). *The abstract collecting hypersemantics is a sound approximation of the collecting hypersemantics:*

$$\forall X \in \wp(\wp(\text{Mem})) . \alpha_m(\llbracket P \rrbracket^h X) \subseteq \llbracket P \rrbracket_h^h \alpha_m(X)$$

PROOF. (Sketch) We just have to prove that the abstract collecting hypersemantics approximates the best correct approximation of the collecting hypersemantics in Mem^h , namely $\alpha_m \llbracket P \rrbracket^h \gamma_m \subseteq \llbracket P \rrbracket_h^h$. The proof relies on the soundness of the abstract semantics for arithmetic and boolean expressions and it is for structural induction. We omit here the full proof, we just show as example the case for programs sequences:

$$\begin{aligned} \alpha_m \llbracket P_1 ; P_2 \rrbracket^h \gamma_m(m) &= \llbracket \text{definition of } \llbracket \cdot \rrbracket^h \\ \alpha_m \llbracket P_2 \rrbracket^h \llbracket P_1 \rrbracket^h \gamma_m(m) &\subseteq \llbracket \text{extensivity of } \gamma_m \alpha_m \\ \alpha_m \llbracket P_2 \rrbracket^h \gamma_m \alpha_m \llbracket P_1 \rrbracket^h \gamma_m(m) &\subseteq \llbracket \text{inductive hypothesis: } \alpha_m \llbracket P_i \rrbracket^h \gamma_m \subseteq \llbracket P_i \rrbracket_h^h \text{ with } i \in \{1, 2\} \\ \llbracket P_2 \rrbracket_h^h \llbracket P_1 \rrbracket_h^h \gamma_m(m) &= \llbracket \text{definition of } \llbracket \cdot \rrbracket_h^h \\ \llbracket P_1 ; P_2 \rrbracket_h^h \gamma_m(m) & \end{aligned}$$

□

. With the abstract semantics just introduced we can define an effective verification method for Non-Interference.

THEOREM 4.6 (NON-INTERFERENCE VERIFICATION). *We have that a program satisfies Non-interference $P \models \text{NI}$ if $\llbracket P \rrbracket_h^h m^{\text{NI}} \subseteq m^{\text{NI}}$.*

PROOF. Note that $\alpha_m(\mathcal{I}^{2L}) \subseteq m^{\text{NI}}$ since \mathcal{I}^{2L} contains only sets of memories agreeing on L variables. This means that $\alpha_m(\mathcal{I}^{2L})(x) \leq \kappa = m^{\text{NI}}(x)$, for each L variable x . For H variables y , $\alpha_m(\mathcal{I}^{2L})(y) \leq \top = m^{\text{NI}}(y)$ trivially holds. Then the proof is given by the following implications:

$$\begin{aligned} \llbracket P \rrbracket_h^h m^{\text{NI}} &\subseteq m^{\text{NI}} \\ \Downarrow \llbracket \text{monotonicity of } \llbracket P \rrbracket_h^h \text{ and } \alpha_m(\mathcal{I}^{2L}) \subseteq m^{\text{NI}} \\ \llbracket P \rrbracket_h^h \alpha_m(\mathcal{I}^{2L}) &\subseteq \llbracket P \rrbracket_h^h m^{\text{NI}} \subseteq m^{\text{NI}} \\ \Downarrow \llbracket \text{soundness of } \llbracket P \rrbracket_h^h \text{ (Thm. 4.5)} \\ \alpha_m(\llbracket P \rrbracket^h \mathcal{I}^{2L}) &\subseteq \llbracket P \rrbracket_h^h \alpha_m(\mathcal{I}^{2L}) \subseteq \llbracket P \rrbracket_h^h m^{\text{NI}} \subseteq m^{\text{NI}} \\ \Downarrow \llbracket \text{Thm. 4.3} \\ P &\models \text{NI} \end{aligned}$$

□

This means that we can check Non-Interference simply by checking that each set of computations, starting from L -equivalent memories, provides only singletons as results.

5 THE PROTOTYPE ANALYZER

With the only aim of proving the feasibility of the proposed approach, and in particular of the abstract hypersemantics, we have written a prototype analyzer in Java SE 10 for Imp programs, which implements the abstract collecting semantics of Sec. 4. The tool,

Program P ₁ : <code>a := 0 ;</code> <code>if (b < x) { b := a * 3 }</code> <code>else { b := a - ((2 * a) - a) }</code>	Program P ₃ : <code>a := x - x</code>
Program P ₂ : <code>a := x ;</code> <code>if (b < a) { b := 2 }</code> <code>else { b := 3 }</code>	Program P ₄ : <code>a := 0 ;</code> <code>while (x < y) {</code> <code> x := x + 1 ;</code> <code> while (a < x) { a := a + 2 ;</code> <code> a := 0</code> <code>}</code>

Figure 5: Example programs.

called nonInterfer, can be tested on-line, through a web interface, at the link <http://bit.do/noninterfer>.

5.1 Validation

Since our analyzer is built for a toy language, there are no benchmark test sets. So we have measured speed and precision of the tool building our own test set. We have written 25 non-interferent programs and 25 interferent programs, with different levels of complexity. As expected, the prototype does not output false negatives, i.e., all interferent programs are discovered. For what concerns precision, the analyzer marks 3 programs as interferent even if they actually satisfy Non-Interference. In Fig. 5 we have four example programs, where variables a, b are public and variables x, y are private. The initial abstract memory m^{NI} is $[a \mapsto \kappa \ b \mapsto \kappa \ x \mapsto \top \ y \mapsto \top]$. We have that $\llbracket P_1 \rrbracket_h^m m^{NI} = [a \mapsto \bar{0} \ b \mapsto \bar{0} \ x \mapsto \top \ y \mapsto \top] \sqsubseteq m^{NI}$, meaning that the analyzer correctly marks P_1 as non-interferent. Analyzing program P_2 , the verifier is able to catch an implicit indirect flow, in fact $\llbracket P_2 \rrbracket_h^m m^{NI} = [a \mapsto \top \ b \mapsto \top \ x \mapsto \top \ y \mapsto \top] \not\sqsubseteq m^{NI}$ (i.e., P_2 is correctly marked as interferent). Unfortunately, our analyzer signals a false alarm in program P_3 , indeed $\llbracket P_3 \rrbracket_h^m m^{NI} = [a \mapsto \top \ b \mapsto \kappa \ x \mapsto \top \ y \mapsto \top] \not\sqsubseteq m^{NI}$, even if the program is non-interferent. Finally, we have a precise result on the more complex program P_4 , indeed $\llbracket P_4 \rrbracket_h^m m^{NI} = [a \mapsto \bar{0} \ b \mapsto \kappa \ x \mapsto \top \ y \mapsto \top] \sqsubseteq m^{NI}$.

The finite height of the hyper abstract domain C^h guarantees the termination of the analysis. Furthermore, the structure of the domain allows us to compute loops fixpoints quickly, hence there is no need for a widening operator in order to speed-up the analysis. On our test set, which comprises quite small hand-made programs, the analyzer is very fast. The analysis time is around 120 milliseconds in average, running on a commodity hardware⁴. We also tested the prototype on bigger programs, generated automatically with the tool Grammarinator[15]. The analyzer is able to handle programs with hundreds of lines of code, basically with the same speed time. As expected, the analyzer shows some slowdowns when programs use lots of variables. Nevertheless, its running time is lower than 600 milliseconds even on programs with more than 500 variables.

Indeed the analyzer exhibits a good trade-off between verification speed and precision (in general). Unfortunately, our analyzer is not precise in some trivial situations, like in P_3 , hence in the next subsection we show how it is possible to get better results.

⁴Laptop with Arch Linux 64-bit (kernel 4.17.5-1), Intel Core i7-7700HQ CPU, 8GiB RAM and SSD storage.

5.2 Improving Precision

The analyzer has a good precision overall but it signals false alarms in some, sometimes very trivial, cases. We mentioned in the previous section that our current approach for the approximation of modified variables is a very simple syntactic check. With a more semantic analysis we can gain precision and do not arise false alarms for programs like P_3 of Fig. 5. Apart from this detail, the sources of imprecision of our semantics are basically two: the approximation added making the hyperlevel constants domain machine-representable and the lack of relational information between variables. In this section we deal with these two issues.

5.2.1 Tuning the Hyperlevel Constants Domain. The original hyperlevel constants domain of [17] contains all the elements of the powerset of $\{\{n\} \mid n \in \mathbb{Z}\}$, meaning that every possible combination of constant sets is taken into account. This makes the domain very precise but not machine-representable, as already observed. In our implementation we have chosen to represent precisely only the singletons $\{\{n\}\}$, abstracted to \bar{n} , and the set of all singleton sets $\{\{n\} \mid n \in \mathbb{Z}\}$, abstracted to κ . In order to enhance precision we could extend our domain C^h with pairs of constant sets, namely we can represent sets of the form $\{\{n\}, \{m\}\}$, with $n, m \in \mathbb{Z}$. But we can gain more precision taking into account triples, quadruples and so on. Hence we can infinitely tune the precision of the analyzer. Clearly the more elements we add to the domain and the more space is consumed by the analyzer and the more abstract operations are complex. So, the trade-off between precision and performance of analysis depends on the analyzer's context of application.

5.2.2 Add Relational Information. Our analysis is not-relational, meaning that we do not explicitly track relations between different variables. We can increase the precision pairing Mem^h with a relational abstraction of $\wp(\wp(\text{Mem}))$. For instance we can define an abstract domain tracking equalities between variables. This latter, combined with a numerical domain such as the one for intervals will improve the precision w.r.t. implicit flows.

Take, as example, the program here on the left. Our analyzer signals a false alarm, since the program is non-interferent but our analysis outputs an abstract memory assigning \top to all variables. With an interval analysis we are able to find that variable a is equal to $[1, 1]$ at the end of the while and with a domain tracking equalities we can deduce the same for variable b . Hence, at the end of the program we can better our analysis obtaining the abstract memory $[a \mapsto \bar{1} \ b \mapsto \bar{1} \ x \mapsto \top]$, allowing us to prove that the program is non-interferent.

6 RELATED WORK

The closest related works are [3] and [20], which both deal with hyperproperties by means of abstract interpretation. In the second, the authors define a hyperproperty called Input Data Usage, claiming that it generalizes a lot of notion of information flows, comprising Non-Interference. They propose an ad-hoc hypersemantics useful to verify that hyperproperty. Then they show how it is possible

to obtain, by abstraction of their semantics, some known verification methods for information flows. The work of [3] is focused on Non-Interference, indeed they propose two abstract semantics: one for (qualitative) Non-Interference and one for quantitative Non-Interference. Quantitative information flows go beyond the scope of the present work, hence we do not take in consideration that case. For the other case, we have that our abstract domain is not directly comparable with the “dependences” abstract domain of [3]. Nevertheless our abstract semantics is able to state correctly the non-interference of the program in Listing 5 of [3] without any tuning for precision. The dependences abstract semantics of [3] needs to add the Intervals domain in order to reach this level of precision. Finally, to the best of our knowledge, none of the previous works have an implementation, even for a toy language, supporting their theoretical results.

Another related work is [2], where authors propose a new methodology for proving the absence of timing channels. This work is based on the idea of “decomposition instead of self-composition”. The idea is to partition the program semantics and to analyze each partition with standard methods. Their approach is similar to our method for simplifying the verification of Non-Interference. Nevertheless, for each partition they verify a classic trace property, instead we verify a hyperproperty. This lead us to better results w.r.t. precision.

Non-Interference is a k -hypersafety (with $k = 2$) hence it can be verified with a classic mechanism for safety trace properties on the k times self-composed system. The self-composition can be sequential, parallel or in an interleaving manner and a lot of works applied this methodology [5, 18, 19]. Unfortunately, this approach seems to be computationally to expensive to be used in practice [2]. Besides the reduction to safety, in [1] the authors introduce a runtime refutation method for k -hypersafety, based on a three-valued logic. Similarly, [6, 13] define hyperlogics (HyperLTL and HyperCTL/CTL*), i.e., extensions of temporal logic able to quantify over multiple traces. Some algorithms for model-checking in these extended temporal logics exist, but only for particular decidable fragments, since the model-checking problem for these logics is, in general, undecidable.

Classic methods for Non-Interference verification, which do not take in consideration its hyperproperty nature, comprise the type systems à la Volpano [21]. These latter perform just syntactic checks, with our approach we have more precision since we can exploit semantic information. The new logic-based approaches showed up recently seem very promising, like the epistemic temporal logic of [4] and SecLTL [12]. They extend classic temporal logics with modalities useful for the verification of Non-Interference. It is not so easy to compare our work with these latter, we let as a future work to deepen the link between these logics and our abstract semantics.

7 CONCLUSION AND FUTURE DIRECTIONS

In this work, we show step by step how to build a static analyzer for Non-Interference, based on abstract interpretation. The tool is sound, meaning that it will not signals false negative, i.e., if the analyzer returns that a program is non-interferent then it is guaranteed that it satisfies Non-Interference. We have soundness by design, exploiting the framework of abstract interpretation. We follow the theoretical work of [17] and we made their abstract

domain computer-representable, hence we show how to make their analysis feasible. Furthermore, we simplified the process of Non-Interference verification, moving from a semantics computing on input/output traces to a simpler semantics computing on memories.

We implemented our analyzer, called nonInterfer in order to validate our abstract semantics. The testing on our prototype has lead to very promising results, in particular w.r.t. analysis speed. Non-Interference verification is undecidable, hence we have obviously false negative, namely sometimes our analyzer marks as interferent a program which actually satisfies Non-interference. Despite its simplicity and speed, our analyzer is quite precise, at least as precise as classic type systems for Non-Interference.

As a future work, we will increase the precision of nonInterfer, adding the possibility to track relational information between different variables. Finally, since the results on Imp are promising, we planned the porting of nonInterfer to a real-world programming language, in order to see its performance on more challenging tests.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous referees for their valuable comments and helpful suggestions.

REFERENCES

- [1] S. Agrawal and B. Bonakdarpour. 2016. Runtime Verification of k -Safety Hyperproperties in HyperLTL. In *Proc. of CSF*. 239–252.
- [2] T. Antonopoulos, P. Gazzillo, M. Hicks, E. Koskinen, T. Terauchi, and S. Wei. 2017. Decomposition instead of self-composition for proving the absence of timing channels. In *Proc. of PLDI*. 362–375.
- [3] M. Assaf, D. A. Naumann, J. Signoles, E. Totel, and F. Tronel. 2017. Hypercollecting semantics and its application to static analysis of information flow. In *Proc. of POPL*. 874–887.
- [4] M. Balliu, M. Dam, and G. Le Guernic. 2011. Epistemic Temporal Logic for Information Flow Security. In *Proc. of PLAS*. 1–12.
- [5] G. Barthe, P. R. D’Argenio, and T. Rezk. 2004. Secure information flow by self-composition. In *Proc. of CSF*. 100–114.
- [6] M. R. Clarkson, B. Finkbeiner, M. Koleini, K. K. Micinski, M. N. Rabe, and C. Sánchez. 2014. Temporal Logics for Hyperproperties. In *Proc. of POST*.
- [7] M. R. Clarkson and F. B. Schneider. 2010. Hyperproperties. *J. of Comp. Security* 18 (2010), 1157–1210.
- [8] E. Cohen. 1977. Information Transmission in Computational Systems. *Oper. Syst. Rev.* 11 (1977), 133–139.
- [9] P. Cousot and R. Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Proc. of POPL*. 238–252.
- [10] P. Cousot and R. Cousot. 2012. An Abstract Interpretation Framework for Termination. In *Proc. of POPL*. 245–258.
- [11] P. Cousot, R. Giacobazzi, and F. Ranzato. 2018. Program Analysis Is Harder Than Verification: A Computability Perspective. In *Proc. of CAV*. 75–95.
- [12] R. Dimitrova, B. Finkbeiner, M. Kovács, M. N. Rabe, and H. Seidl. 2012. Model Checking Information Flow in Reactive Systems. In *Proc. of VMCAI*. 169–185.
- [13] B. Finkbeiner, M. N. Rabe, and C. Sánchez. 2015. Algorithms for Model Checking HyperLTL and HyperCTL. In *Proc. of CAV*. 30–48.
- [14] D. Hedin and A. Sabelfeld. 2012. A Perspective on Information-Flow Control. *Software Safety and Security* 33 (2012), 319–347.
- [15] R. Hodovan and A. Kiss. 2018. Grammarinator - ANTLRv4 grammar-based test generator. <https://github.com/renatahodovan/grammarinator>
- [16] I. Mastroeni and M. Pasqua. 2017. Hyperhierarchy of Semantics - A Formal Framework for Hyperproperties Verification. In *Proc. of SAS*. 232–252.
- [17] I. Mastroeni and M. Pasqua. 2018. Verifying Bounded Subset-Closed Hyperproperties. In *Proc. of SAS*. 263–283.
- [18] M. Sousa and I. Dillig. 2016. Cartesian Hoare Logic for Verifying K -safety Properties. In *Proc. of PLDI*. 57–69.
- [19] T. Terauchi and A. Aiken. 2005. Secure Information Flow As a Safety Problem. In *Proc. of SAS*. 352–367.
- [20] C. Urban and P. Müller. 2018. An Abstract Interpretation Framework for Input Data Usage. In *Proc. of ESOP*. 683–710.
- [21] D. Volpano, C. Irvine, and G. Smith. 1996. A Sound Type System for Secure Flow Analysis. *J. Comput. Secur.* 4 (1996), 167–187.