



 Latest updates: <https://dl.acm.org/doi/10.1145/3786347>

RESEARCH-ARTICLE

## **Abstract Interpretation-based Verification for Confidentiality: Information Hiding and Code Protection by Abstract Interpretation**

**ISABELLA MASTROENI**, University of Verona, Verona, VR, Italy

**MICHELE PASQUA**, University of Verona, Verona, VR, Italy

**Open Access Support** provided by:

**University of Verona**



PDF Download

3786347.pdf

04 April 2026

Total Citations: 0

Total Downloads: 254

**Published:** 11 February 2026

**Online AM:** 26 December 2025

**Accepted:** 01 December 2025

**Revised:** 26 September 2025

**Received:** 22 April 2025

[Citation in BibTeX format](#)

# Abstract Interpretation-based Verification for Confidentiality: Information Hiding and Code Protection by Abstract Interpretation

ISABELLA MASTROENI, Dept. of Computer Science, University of Verona, Verona, Italy  
MICHELE PASQUA, Dept. of Computer Science, University of Verona, Verona, Italy

In modern computing systems, preventing sensitive information leakage is a crucial issue. Indeed, to deploy secure computing systems, data protection is an aspect that cannot be ignored. Many security requirements are adopted in this respect, such as *opacity* and *non-interference*. The first assures that the truth value of a predicate is masked during computation, while the second prevents confidential information is leaked through uncontrolled system components. Unfortunately, despite their simple intended meaning, confidentiality notions are quite difficult requirements to enforce. In fact, they are actually *hyperproperties*, and thus require enforcing mechanisms that reason on multiple executions at a time.

To develop effective verification and validation mechanisms for confidentiality notions, it is crucial to precisely characterize the requirements of system executions they dictate. In this article, we investigate the relation between *abstract non-interference* (a weakening of non-interference observing properties of data instead of concrete values) and opacity through the lens of *abstract interpretation*. By adopting such a holistic, abstract approach, we show how to formally characterize the *structure* of confidentiality notions and to compare them in terms of the constraints on system executions they impose and verification complexity. In addition, we show how *code obfuscation* can be restated as a confidentiality problem by defining a corresponding confidentiality notion that can possibly be enforced.

Finally, by exploiting the recently proposed static analysis approach for verifying non-interference, based on *hypersemantics*, we show how to verify abstract non-interference, therefore opacity and other security requirements. Based on *abstract interpretation*, this yields an effective mechanism to enforce a broad range of confidentiality notions.

CCS Concepts: • **Security and privacy** → **Security requirements**; • **Software and its engineering** → *Software verification*;

Additional Key Words and Phrases: Information flow control, opacity, non-interference, abstract interpretation, hyperproperties verification, static analysis, code protection, obfuscation

## ACM Reference Format:

Isabella Mastroeni and Michele Pasqua. 2026. Abstract Interpretation-based Verification for Confidentiality: Information Hiding and Code Protection by Abstract Interpretation. *ACM Trans. Priv. Sec.* 29, 1, Article 12 (February 2026), 30 pages. <https://doi.org/10.1145/3786347>

This work was partially supported by the project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU; by PRIN2022PNRR “RAP-ARA” (PE6) - codice MUR: P2022HXNSC. Authors’ Contact Information: Isabella Mastroeni, Computer Science, University of Verona, Verona, Italy; e-mail: [isabella.mastroeni@univr.it](mailto:isabella.mastroeni@univr.it); Michele Pasqua (corresponding author), Computer Science, University of Verona, Verona, Italy; e-mail: [michele.pasqua@univr.it](mailto:michele.pasqua@univr.it).



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2026 Copyright held by the owner/author(s).

ACM 2471-2566/2026/02-ART12

<https://doi.org/10.1145/3786347>

## 1 Introduction

The increasing reliance on computing systems for storing, processing, and transmitting sensitive information has made confidentiality a paramount concern in modern society. From personal health records and financial transactions to national security data and intellectual property, the need to protect sensitive information from unauthorized access and disclosure is critical. Confidentiality, a fundamental principle of information security, aims to ensure that information is accessible only to authorized individuals or entities, preventing its exposure to unintended recipients. In this context, it is crucial to develop computing systems that are resilient to attacks aimed at breaking confidentiality. Either to build secure-by-design systems or to check compliance of already deployed systems, a fundamental issue is to precisely model the security requirements, e.g., confidentiality, of interest. This requires a precise, ideally formal, definition of the requirements the system should fulfill. In literature, two security requirements are well-established and adopted in many contexts: opacity and non-interference.

*Opacity* [6, 39] is a security requirement assuring the protection of sensitive aspects concerning the behavior of computing systems [41]. Indeed, it is a natural choice in several scenarios, such as privacy, anonymity [25], and geo-privacy [41]. As Bryans et al. [6] have observed, “the essential idea of opacity is that a predicate is made opaque if an observer of the system will never be able to determine the truth of that predicate”. In other words, a predicate is made opaque by a system (e.g., a computer program) when, by observing the system’s execution, it is not possible to deduce the predicate truth value (or both truth values in case of *symmetric opacity* [41]). Instead, *non-interference* [8] is a security requirement checking whether and/or how the system’s *sensitive* inputs interfere with the *observable* results of system execution. It is a natural choice when information leaks are concerned, as confidential inputs should not flow into non-confidential outputs.

As a common ground between opacity and non-interference, we find that both requirements model a variation of a system’s aspect as the effect of system execution, underlying a possible connection between the two notions. Indeed, Schoepe and Sabelfeld [41] intuitively observed that “a program satisfies non-interference if and only if it is opaque for almost all predicates”. Unfortunately, this intuition does not allow us to understand how such notions can be compared in the way they tackle confidentiality. Hence, a naturally arising question is: do these notions lay at different levels of formalization, as shown in Reference [41], or do they provide a different perspective on the same security issue? And, in turn: can we use (part of) the verification mechanisms for opacity to verify non-interference (or vice versa)?

To tackle these questions, we can consider *abstract non-interference* [21], a weakening of non-interference observing properties, namely abstractions, of data instead of concrete values (e.g., observing the parity of numbers instead of their actual values). Indeed, opacity already deals with observations/abstractions of system executions (on system outputs only in Reference [6], on system inputs and outputs in Reference [41]). In particular, when we aim at protecting a property of system inputs w.r.t. an observation of system executions (the so-called *initial opacity* [6]), avoiding interference between the variation of the input property and executions observation, then we fall into abstract non-interference. Similarly, when we aim at protecting a property of system outputs (the so-called *final opacity* [6]), then we fall into (semantic) code obfuscation [14] (code is transformed to mask output data properties). Not surprisingly, the latter is strongly related to abstract non-interference [15].

When verification comes into play, things become complex in the context of opacity and (abstract) non-interference. Here, with *verification* is meant the general process of checking whether a system complies with a requirement, i.e., a formal description of what systems are allowed and are not allowed to do. Indeed, the verification process is quite standard when dealing with *trace properties*, i.e., requirements that can be checked inspecting *single* system executions. Problems

arise when the system's aspect of interest requires comparing more than one execution for being verified, as happens in security requirements such as opacity and non-interference. In Reference [7], *hyperproperties* were introduced to formalize precisely those requirements which are not *trace properties*. Usually, a system's behavior (semantics) is modeled as a set of executions, one for each possible input. In this setting, trace properties are sets (properties) of executions, while hyperproperties are sets (properties) of collections of executions, making verification of the latter more challenging. In Reference [34], this challenge has been faced to verify non-interference statically. Specifically, [34] introduced a sound static analysis for checking non-interference based on abstract interpretation [12].

In this article, we will show that opacity and abstract non-interference represent a different perspective on the same security issue, and this strong relation provides us the perfect field where moving the verification approach presented in Reference [34] from non-interference to abstract non-interference and opacity. Indeed, later in the article, we will provide a static verification mechanism for confidentiality notions, such as abstract non-interference and opacity, based on abstract interpretation. More in detail, we will first introduce a formalism for describing many flavors of abstract non-interference and opacity in the same setting, allowing us to compare the two notions formally. Once the connection between opacity and abstract non-interference is clear, we will identify a version of the latter that implies the former. Then, by adapting the verification mechanism of Reference [34] to abstract non-interference, we can provide a static analysis technique for soundly verifying opacity. In addition, we will leverage the proposed formalism to define confidentiality notions modeling code obfuscation as a combination of opacity and abstract non-interference. The latter can be useful to check whether a system is already obfuscating.

*Structure of the article.* In Section 2, we provide the background material, introducing abstract interpretation, as well as presenting the classic notions of opacity and abstract non-interference. At the end of the section, we introduce the reader to hyperproperties, and we depict a high-level view of the verification process for hyperproperties using abstract interpretation. Section 3 provides a general formalism for defining and comparing confidentiality notions, where opacity and abstract non-interference are instantiated and compared. In Section 4, we use again the proposed formalism to define confidentiality notions in code protection, specifically code obfuscation. In Section 5, we deep dive into the verification process for abstract non-interference and opacity. In particular, we investigate how to apply abstract interpretation, adapting the approach followed by [34], to analyze generalized abstract non-interference and symmetric opacity. Finally, Section 6 concludes the article and discusses the related work.

## 2 Background

### 2.1 Equivalences and Closure Operators

The *equivalence relations* on a set  $C$  form a lattice  $\langle Eq(C), \sqsubseteq, \sqcap, \sqcup, Id_C, All_C \rangle$  where:  $Id_C \triangleq \{(c, c) \mid c \in C\}$  is the relation distinguishing all the elements in  $C$  (i.e., the bottom element);  $All_C \triangleq C \times C$  is the relation that cannot distinguish any elements in  $C$  (i.e., the top element); for all relations  $Q$  and  $R$  in  $Eq(C)$ ,  $Q \sqsubseteq R \triangleq Q \subseteq R$ ;  $Q \sqcap R \triangleq Q \cap R$ ;  $Q \sqcup R \triangleq (Q \cup R)^+$ , where  $(Q \cup R)^+$  is the transitive closure of  $Q \cup R$ . In the following,  $[c]_R \triangleq \{c' \mid c R c'\}$  denotes the equivalence class of  $c$  under the relation  $R$  in  $Eq(C)$ .

An *upper closure operator* on a poset  $\langle C, \leq_C \rangle$  is a function  $\rho : C \rightarrow C$  having the following properties.

- $\forall c, c' \in C$ , if  $c \leq_C c'$  then  $\rho(c) \leq_C \rho(c')$  (monotonicity);
- $\forall c \in C$ .  $\rho(c) = \rho(\rho(c))$  (idempotence);
- $\forall c \in C$ .  $c \leq_C \rho(c)$  (extensivity).

A closure<sup>1</sup>  $\rho : C \rightarrow C$  is uniquely identifiable with the set of its fixpoints  $\rho(C) \triangleq \{c \in C \mid \rho(c) = c\}$ . With a little abuse of notation, in the following, we will often use  $\rho$  to denote the closure as a function and to denote its set of fixpoints  $\rho(C)$ . If  $C$  is a complete lattice then  $\langle uco(C), \sqsubseteq, \sqcap, \sqcup, id, \top \rangle$  is the lattice of upper closure operators on  $C$ , where the top element is  $\top \triangleq \lambda c. \top_C$  (here  $\top_C$  is the top element of  $C$ ) and the bottom element is  $id \triangleq \lambda c. c$ . For every  $\rho$  and  $\eta$  in  $uco(C)$  and for any subset  $\{\rho_i\}_{i \in I}$  of  $uco(C)$  we have that:  $\rho \sqsubseteq \eta \triangleq \eta(C) \subseteq \rho(C)$ ;  $\bigsqcup_{i \in I} \rho_i \triangleq \bigcap_{i \in I} \rho_i$ ; and  $\bigsqcap_{i \in I} \rho_i \triangleq \mathcal{M}(\bigcup_{i \in I} \rho_i)$ . In particular,  $\mathcal{M}$  is the operation of closing a domain by a concrete greatest lower bound, e.g., intersection on power domains. Given a closure  $\rho$  in  $uco(C)$ ,  $X \subseteq C$  is the set of fixpoints of  $\rho$  iff  $X$  is a *Moore-family* of  $C$ , i.e.,  $X = \mathcal{M}(X) \triangleq \{\bigwedge_C Y \mid Y \subseteq X\}$ , where  $\bigwedge_C \emptyset = \top_C$  is in  $\mathcal{M}(X)$ .

When  $\rho : \wp(C) \rightarrow \wp(C)$ , meaning that  $\rho$  is an upper closure operator defined on a power set domain, then  $\rho$  is called *partitioning* [38] iff it is complemented<sup>2</sup>, i.e., for all subsets  $X \subseteq C$ , if  $X$  belongs to  $\rho(\wp(C))$  then also  $\bar{X}$  belongs to  $\rho(\wp(C))$ , where  $\bar{X} \triangleq C \setminus X$  is the complement of  $X$ . The *atoms* of a closure  $\rho$ , denoted  $Atom(\rho)$ , are the fixpoints covering the bottom of the corresponding Moore family.

We recall that there exists an isomorphism between equivalence relations on a domain  $C$  and a subset of upper closure operators [26, 38] on  $\wp(C)$ . For each equivalence relation  $R \subseteq C \times C$  on a set  $C$ , we can define a closure  $Clo^R$  in  $uco(\wp(C))$  and, vice versa, from each closure  $\rho$  in  $uco(\wp(C))$  we can define an equivalence relation  $Rel^\rho \subseteq C \times C$  on  $C$ . Consider a closure  $\rho$  in  $uco(\wp(C))$ , we define  $Rel^\rho \subseteq C \times C$  as: given  $c$  and  $c'$  in  $C$ ,  $c Rel^\rho c' \triangleq \rho(\{c\}) = \rho(\{c'\})$ . Proving that  $Rel^\rho$  is an equivalence relation is immediate and does not depend on the fact that  $\rho$  is a closure, but only on the fact that it is a function. Consider now an equivalence relation  $R \subseteq C \times C$ , we define  $Clo^R$  in  $uco(\wp(C))$  as: given  $X \subseteq C$ ,  $Clo^R(X) \triangleq \bigcup_{c \in X} [c]_R$ . Note that  $Clo^R$  is partitioning by construction, and  $[c]_R$ , given  $c$  in  $C$ , are its atoms. In [26], the authors observed that for each  $R$  in  $Eq(C)$  we have  $Rel^{Clo^R} = R$ . In the following, for the sake of readability, when a closure  $\rho$  in  $uco(\wp(C))$  is applied to a single element  $c$  in  $C$ , we abuse notation writing  $\rho(c)$  instead of  $\rho(\{c\})$ .

## 2.2 Abstract Interpretation

In the standard framework of *abstract interpretation* [12], abstract domains can be equivalently formulated either in terms of Galois connections or upper closure operators [13]. A pair of functions  $\alpha : C \rightarrow A$  and  $\gamma : A \rightarrow C$  on posets, denoted  $\langle C, \alpha, \gamma, A \rangle$ , forms an *adjunction* or a *Galois connection* (GC) if for any concrete element  $c$  in  $C$  and abstract element  $a$  in  $A$  we have  $\alpha(c) \leq_A a$  iff  $c \leq_C \gamma(a)$ . In this case,  $\alpha$  (resp.,  $\gamma$ ) is the *left-* (resp., *right-*) *adjoint* of  $\gamma$  (resp.,  $\alpha$ ), and it is additive (resp. co-additive), i.e., it preserves existing least upper bounds (resp., greatest lower bounds) of arbitrary non-empty sets. A *Galois insertion* (GI) is a Galois connection such that  $\alpha \circ \gamma$ , where  $\circ$  denotes function composition from left to right<sup>3</sup>, is the identity on  $A$ . Every Galois insertion  $\langle C, \alpha, \gamma, A \rangle$  is uniquely identifiable with a closure  $\gamma \circ \alpha$  in  $uco(C)$ . Since it is always possible to reduce a GC to a GI by eliminating redundant elements [13], we can identify any abstract domain of  $C$  with closure in  $uco(C)$ . In the following, given a concrete domain  $C$ , we will use the terms abstract domain, abstraction, and closure of  $C$  interchangeably.

Abstract interpretation is often used for abstracting computations, e.g., system/program semantics. Given a monotone function  $f : C \rightarrow C$  and an abstraction  $\rho$  in  $uco(C)$  of  $C$ , we call *best correct approximation* (bca for short) of  $f$  in  $\rho$  the function  $\rho \circ f \circ \rho$ . When the domains of input (say  $C$ ) and output (say  $D$ ) of  $f$  are different, that is,  $f : C \rightarrow D$ , then also the input abstraction  $\eta$  in  $uco(C)$

<sup>1</sup>We will often just say 'closure' in place of 'upper closure operator'.

<sup>2</sup>Here,  $\wp(\cdot)$  denotes the powerset operator, that is,  $\wp(C)$  is the st of all subsets of  $C$ .

<sup>3</sup>We will often omit the composition operator for brevity.

may be different from the output abstraction  $\rho$  in  $uco(D)$ . In this case, the bca is  $\rho \circ f \circ \eta$  [13]. Bca computation is always correct, i.e., for every  $c$  in  $C$ ,  $\rho \circ f(c) \leq \rho \circ f \circ \eta(c)$ , meaning that we may lose precision when computing on abstract inputs, but no erroneous information is added. The bca is *complete* when for every  $c$  in  $C$ ,  $\rho \circ f(c) = \rho \circ f \circ \eta(c)$  [13, 23], i.e., when there is no loss of information when computing on abstract inputs (no loss of precision in abstract computations). Completeness is termed *local* [5] w.r.t. a concrete element  $c$  in  $C$ , when it is verified on that fixed element, i.e., when  $\rho \circ f(c) = \rho \circ f \circ \eta(c)$ .

### 2.3 Opacity

Opacity is a security requirement that aims at protecting sensitive predicates on system executions [41], making them *opaque*. Intuitively, a predicate is opaque if, for any execution that satisfies the predicate, another execution does not satisfy the predicate and is indistinguishable by an attacker observing the system. Opacity is particularly useful in scenarios like privacy, anonymity [25] and *geo-privacy* [41]. For instance, opacity is the natural choice in the context of *geo-fences* for geo-tagging policies in sensitive areas [41].

In Reference [6] the authors provide an extensive study of the notion of opacity on labeled transition systems, but the same notion can be easily moved on any system semantics (e.g., program semantics). The original notion concerns the protection of a predicate that can be defined on initial states (*initial opacity*) or on final states (*final opacity*), or on the whole execution, namely on the execution trace (*total opacity*). When moving from transition systems to system semantics modeled by a function  $f$ , the last two notions can be collapsed into a single one (final opacity), since the results of the function can be either values or a trace of values (or any other denotation for modeling executions). We combine here the work of [6] with [41], where an input observation is added to a notion corresponding to initial opacity on program semantics and where *symmetric opacity* is introduced, i.e., opacity protecting both predicate truth values. We recall the original definitions of initial opacity [6], which we simply call Opacity as in [41]. When doing so, we use the notation we will formally introduce in Section 3. In particular, we assume program semantics to be modeled as functions  $f : \mathbb{I} \rightarrow \mathbb{O}$  from an input set  $\mathbb{I}$  to an output set  $\mathbb{O}$ . In this context, the restricted universal (resp., existential) quantification  $\forall_I x . \Pi$  (resp.,  $\exists_I x . \Pi$ ) limits the scope of the variable  $x$  in the formula  $\Pi$  to elements belonging to  $I \subseteq \mathbb{I}$ , that is,  $x$  is not quantified on the whole  $\mathbb{I}$  but on the elements of its subset  $I$  only (see Section 3 for more details).

*Definition 2.1 (Opacity).* Let  $\overset{\sim}{\sim}$  in  $Eq(\mathbb{I})$  and  $\overset{\sim}{\sim}$  in  $Eq(\mathbb{O})$  be equivalence relations on input  $\mathbb{I}$  and output  $\mathbb{O}$ , respectively. A predicate  $P \subseteq \mathbb{I}$  is *opaque* for  $f : \mathbb{I} \rightarrow \mathbb{O}$  w.r.t.  $\overset{\sim}{\sim}$  and  $\overset{\sim}{\sim}$ , denoted  $P \in Op(f, \overset{\sim}{\sim}, \overset{\sim}{\sim})$ , if:

$$\forall_P x_1 . \exists_{\bar{P}} x_2 . \left( x_1 \overset{\sim}{\sim} x_2 \wedge f(x_1) \overset{\sim}{\sim} f(x_2) \right), \quad (1)$$

where  $\bar{P} \triangleq \mathbb{I} \setminus P$ . When both  $P, \bar{P} \in Op(f, \overset{\sim}{\sim}, \overset{\sim}{\sim})$ ,  $P$  is said *symmetrically opaque*, written  $P \in SOp(f, \overset{\sim}{\sim}, \overset{\sim}{\sim})$ .

The writing in Equation (1) tantamount to say that for every input value  $v_1$  that satisfies  $P$  (i.e.,  $v_1$  is in  $P$ ) there exists another input value  $v_2$  that does not satisfy  $P$  (i.e.,  $v_2$  is not in  $P$  or, equivalently,  $v_2$  is in  $\bar{P}$ ) such that  $v_1 \overset{\sim}{\sim} v_2$  and  $f(v_1) \overset{\sim}{\sim} f(v_2)$ . This is exactly the definition of Opacity given in [41].

The classic Opacity [6, 39] without the input observation means that for any input whose execution satisfies the predicate, there is an execution that does not satisfy the predicate. The way the input observation is introduced in [41] consists in requiring that for each possible input observation, the predicate must be opaque. It is a strengthening process that requires Opacity inside each equivalence class of input observations. In the following, we instantiate Opacity to show its usefulness in geo-privacy scenarios.

*Example 2.2.* Consider the *Program 1* of [41], that we list below. Variables  $hX$  and  $hY$  store the user's (potentially sensitive) location coordinates. Suppose that a clinic represents a sensitive location that occupies a rectangular area with corners (200, 50) and (400, 150). The program outputs on a public channel, and it aims to prevent privacy leakage. In particular, when the user is inside the clinic, the program outputs a default location (100, 200) and otherwise the user's real location.

```

/* Program 1 */
clinicXmin := 200; clinicXmax := 400;
clinicYmin := 50; clinicYmax := 150;
if (hX >= clinicXmin and hX <= clinicXmax and hY >= clinicYmin and hY <= clinicYmax) {
  out (100, 200);
} else {
  out (hX, hY);
}

```

Given a pair of natural numbers  $(v, v')$ , we denote with  $(v, v')_+ = v$  and  $(v, v')_- = v'$  the projection on the first and the second element, respectively. The behavior of the program above can be modeled with the function  $f_1 : \mathbb{I} \rightarrow \mathbb{O}$ , where  $\mathbb{I} = \mathbb{O} = \mathbb{N} \times \mathbb{N}$  and  $\mathbb{N}$  are the natural numbers, as follows:

$$f_1(x) = \begin{cases} (100, 200) & \text{if } 200 \leq x_+ \leq 400 \text{ and } 50 \leq x_- \leq 150 \\ x & \text{otherwise} \end{cases}$$

The predicate to protect is the one indicating whether the user is inside the clinic, i.e., the predicate  $P \subseteq \mathbb{I}$ , defined as  $P \triangleq \{(v, v') \mid 200 \leq v \leq 400 \wedge 50 \leq v' \leq 150\} \subseteq \mathbb{I}$ .

Consider the equivalence relations  $\overset{\circ}{\sim}$  (on output) and  $\overset{\sim}{\sim}$  (on input):  $\overset{\circ}{\sim} \triangleq \{(v, v'), (v, v') \mid (v, v') \in \mathbb{O}\}$  and  $\overset{\sim}{\sim} \triangleq \mathbb{I} \times \mathbb{I}$ . In particular,  $\overset{\circ}{\sim}$  models the fact that the attacker can see everything about outputs, while  $\overset{\sim}{\sim}$  models the fact that the attacker cannot distinguish inputs. Then, we have that  $P$  is opaque (Definition 2.1) for  $f_1$ , i.e.,  $P \in \text{Op}(f_1, \overset{\sim}{\sim}, \overset{\circ}{\sim})$ , namely for every  $x_1$  in  $P$  there exists  $x_2$  not in  $P$  (or, equivalently, in  $\bar{P}$ ) such that  $x_1 \overset{\sim}{\sim} x_2$  and  $f_1(x_1) \overset{\circ}{\sim} f_1(x_2)$ . Indeed, for every  $x_1$  in  $P$  (for instance,  $x = (250, 100)$ ) take  $x_2 = (100, 200)$ , that is not in  $P$ . We have that  $x \overset{\sim}{\sim} (100, 200)$  and  $f_1(x) \overset{\circ}{\sim} f_1(100, 200)$ . Indeed,  $f_1(x) = (100, 200)$ , for any  $x$  belonging to  $P$ , and  $f_1(100, 200) = (100, 200)$ . Hence, Program 1 protects the predicate  $P$ .

Let now consider the predicate  $\bar{P} = \mathbb{I} \setminus P$ . In this case,  $\bar{P}$  is not opaque for  $f_1$ , i.e.,  $\bar{P} \notin \text{Op}(f_1, \overset{\sim}{\sim}, \overset{\circ}{\sim})$ . Indeed, consider  $x_1 = (50, 250)$ , which is in  $\bar{P}$ . For every  $x_2$  not in  $\bar{P}$ , i.e.,  $x_2$  is in  $P$ , we have  $f_1(x_2) = (100, 200)$ . In this case,  $f_1(x_1) = f_1(50, 250) = (50, 250) \not\overset{\circ}{\sim} (100, 200) = f_1(x_2)$ . Hence, Program 1 does not protect  $\bar{P}$ .

*End Example.*

*Example 2.3.* Consider the *Program 2* of [41], that we list below. It is similar to Program 1 except that a pseudo-random location outside the clinic is returned whenever the user is inside it. When observing the output, the attacker cannot deduce whether it is the user's real location outside the clinic or a (pseudo) randomly generated location while the user is inside.

```

/* Program 2 */
clinicXmin := 200; clinicXmax := 400;
clinicYmin := 50; clinicYmax := 150;
if (hX >= clinicXmin and hX <= clinicXmax and hY >= clinicYmin and hY <= clinicYmax) {
  randomize(x);
  randomize(y);
  while (x >= clinicXmin and x <= clinicXmax and hY >= clinicYmin and hY <= clinicYmax) {
    randomize(x);
    randomize(y);
  }
  out (x, y);
} else {
  out (hX, hY);
}

```

The behavior of the program can be modeled with the function  $f_2 : \mathbb{I} \rightarrow \mathbb{O}$ , where  $\mathbb{I} = \mathbb{O} = \mathbb{N} \times \mathbb{N}$ , as follows:

$$\text{let } Rnd \triangleq \text{rand}(\mathbb{N} \times \mathbb{N} \setminus [200..400] \times [50..150]) \text{ in}$$

$$f_2(x) = \begin{cases} Rnd & \text{if } 200 \leq x_1 \leq 400 \text{ and } 50 \leq x_2 \leq 150 \\ x & \text{otherwise} \end{cases}$$

Here,  $\text{rand}(N)$  returns a pseudo-random element of the set  $N$ , while  $[min..max]$  is the set of numbers in  $\mathbb{N}$  between  $min$  and  $max$  (included), following the usual less than or equal ordering between naturals. Note that, since  $\text{rand}()$  uses a pseudo-random number generator with fixed seed, hence deterministic,  $f_2$  is indeed a function.

Consider the equivalence relations  $\overset{\sim}{\sim}$  and  $\overset{\circ}{\sim}$  and the predicate  $P$  of Example 2.2. In this case, both  $P$  and  $\bar{P}$  are opaque (Definition 2.1) for  $f_2$ , since the attacker cannot realize whether a user is inside or outside the clinic. Hence,  $P$  is symmetrically opaque, i.e.,  $P \in \text{SOp}(f_2, \overset{\sim}{\sim}, \overset{\circ}{\sim})$ . In this case, Program 2 makes it more difficult for an attacker to learn anything about the user's location, since both the sensitive predicate (whether the user is inside the clinic) and its negation are protected.

*End Example.*

As far as verification is concerned, we have to remark that opacity is not a trace property since it requires comparing at least two system executions. Therefore, standard approaches are either too imprecise or not applicable, in the case of opacity verification. Indeed, opacity is a hyperproperty [7], that requires hyper-level verification approach [1, 34], e.g., a hyperanalysis, as introduced in Subsection 2.6.

## 2.4 Abstract Non-Interference

In Reference [21] the classic notion of *non-interference* [8], checking whether and/or how the sensitive input interferes with the observable result of a function, e.g., a program semantics, has been generalized, allowing to state *when* to check interference, *how* an attacker observes the result and *what* has to be protected. The central idea is to track interference between properties (or abstractions) of data instead of concrete values, where properties/abstractions are modeled by means of upper closure operators. In Reference [21], the most general (so far) version of non-interference, there called *abstract non-interference*, has been given in terms of the best correct approximation of the function to be checked. Similarly to Opacity, we use the notation we will formally introduce in Section 3 to recall abstract non-interference as defined in Reference [21].

*Definition 2.4 (Abstract Non-Interference).* Let  $\phi, \eta$  in  $\text{uco}(\wp(\mathbb{I}))$  be abstractions on input and  $\rho$  in  $\text{uco}(\wp(\mathbb{O}))$  be an abstraction on output. A function  $f : \mathbb{I} \rightarrow \mathbb{O}$  satisfies *Abstract Non-Interference* (ANI in short) w.r.t.  $\phi, \eta$  in input and  $\rho$  in output, written  $f \models \text{ANI}[\langle \phi, \eta \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ , if<sup>4</sup>:

$$\forall x_1 . \forall x_2 . (\phi(x_1) = \phi(x_2) \Rightarrow \rho f \eta(x_1) = \rho f \eta(x_2)). \quad (2)$$

The input observation  $\phi$  is a weakening fixing when (i.e., for which inputs) we ask for non-interference. It behaves as a selection function, namely, we check non-interference only for the inputs having the property  $\phi$ . Instead,  $\rho$  is the observation of the computation; again, it is a weakening, meaning that we can observe the property  $\rho$  of results and not concrete values. Finally,  $\eta$  models an input data abstraction whose variation has not to be observable, that is, variety  $\eta$  on inputs must not be conveyed to outputs. This is the input property protected when ANI holds [21]. Note that ANI, in how it is formalized, can only model the protection of *input* data properties, not output data property.

<sup>4</sup>Recall that, here we use the more concise notation  $\phi(x_1)$  in place of  $\phi(\{x_1\})$ .

*Example 2.5.* Consider a predicate  $P \subseteq \mathbb{I}$  on  $\mathbb{I} = \mathbb{N} \times \mathbb{N}$  (pairs of natural numbers) and let  $\mathbb{O} = \mathbb{I}$ . For instance,  $P$  could be the predicate protecting whether the user is inside the clinic of Example 2.2. Consider the following closures on  $\wp(\mathbb{I}) = \wp(\mathbb{O})$ , i.e., data abstractions<sup>5</sup>:  $\phi = \{\mathbb{I}\}$ ,  $\rho = \wp(\mathbb{I})$  and  $\eta = \{\emptyset, P, \bar{P}, \mathbb{I}\}$ . Here,  $\phi$  and  $\rho$  are the top and bottom abstractions, respectively:  $\phi$  cannot distinguish any element of  $\mathbb{I}$ , while  $\rho$  can distinguish every element of  $\mathbb{I}$ . In particular,  $\phi(\{(v_1, v_2)\}) = \mathbb{I}$  and  $\rho(\{(v_1, v_2)\}) = \{(v_1, v_2)\}$ , for every  $(v_1, v_2) \in \mathbb{I}$ . Furthermore,  $\eta(\{(v_1, v_2)\}) = P$  if  $(v_1, v_2)$  belongs to  $P$ , and  $\eta(\{(v_1, v_2)\}) = \bar{P}$  otherwise.

Consider ANI for the function  $f_2$  of the *Program 2* in Example 2.3. We have that  $f_2 \models \text{ANI}[\langle \phi, \eta \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ . Indeed, the premise of the implication in Equation (2) is always trivially true, since  $\phi(x_1) = \phi(x_2)$  holds for any  $x_1$  and  $x_2$  in  $\mathbb{I}$  when  $\phi = \{\mathbb{I}\}$ . Then, when  $\eta(x_1) = \eta(x_2)$  the whole implication is true, being  $f_2$  a function. Conversely, when  $\eta(x_1) \neq \eta(x_2)$ , we have that  $x_1$  is in  $P$  and  $x_2$  is in  $\bar{P}$ , or vice versa. But, by definition of  $f_2$ , we have that both<sup>6</sup>  $f_2(P)$  and  $f_2(\bar{P})$  are equal to  $\mathbb{N} \times \mathbb{N} \setminus [200..400] \times [50..150]$ . Hence, the whole implication is true, and *Program 2* satisfies ANI.

Note that,  $f_1 \models \text{ANI}[\langle \phi, \eta \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$  does not hold for  $f_1$  (of the *Program 1*) defined in Example 2.2. Consider the following counterexample. Take  $x_1 = (300, 100)$ , which is in  $P$ , and  $x_2 = (20, 20)$ , which is in  $\bar{P}$ . Then,  $\phi(x_1) = \phi(300, 100) = \mathbb{I} = \phi(20, 20) = \phi(x_2)$ ,  $\eta(x_1) = \eta(300, 100) = P$ ,  $\eta(x_2) = \eta(20, 20) = \bar{P}$ ,  $\rho f_1 \eta(x_1) = \rho f_1(P) = \{(100, 200)\}$  and  $\rho f_1 \eta(x_2) = \rho f_1(\bar{P}) = \mathbb{N} \times \mathbb{N} \setminus [200..400] \times [50..150]$ . Hence,  $\phi(x_1) = \phi(x_2)$  but  $\rho f \eta(x_1) \neq \rho f \eta(x_2)$ , and *Program 1* does not satisfy ANI.

*End Example.*

To tune non-interference for coping with different information-flow protection requirements, several flavors of Abstract Non-Interference have been defined in literature [20, 21, 30]. We report here the most relevant for our work.

A *narrow* version of ANI fixes under which input and output observations we have not to detect interference, resulting in a property of an abstract observation of the concrete executions (the abstraction  $\eta$  is not present).

*Definition 2.6 (Narrow Abstract Non-Interference).* Let  $\phi$  in  $\text{uco}(\wp(\mathbb{I}))$  be an abstraction on input and  $\rho$  in  $\text{uco}(\wp(\mathbb{O}))$  be an abstraction on output. A function  $f : \mathbb{I} \rightarrow \mathbb{O}$  satisfies *Narrow Abstract Non-Interference* (NANI in short) w.r.t.  $\phi$  in input and  $\rho$  in output, written  $f \models \text{NANI}[\phi, \rho]$ , if:

$$\forall x_1. \forall x_2. (\phi(x_1) = \phi(x_2) \Rightarrow \rho f(x_1) = \rho f(x_2))$$

It is easy to note that NANI is a special case of ANI, indeed  $f \models \text{NANI}[\phi, \rho]$  iff  $f \models \text{ANI}[\langle \phi, \text{id} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ . A more interesting notion is a *declassified* version of ANI, fixing which input data property must be protected, namely, which property must not affect the output observation. Note that, also in ANI, the property  $\eta$  is protected; what changes here is the way computation is observed.

*Definition 2.7 (Declassified Abstract Non-Interference).* Let  $\eta$  in  $\text{uco}(\wp(\mathbb{I}))$  be an abstraction on input and  $\rho$  in  $\text{uco}(\wp(\mathbb{O}))$  be an abstraction on output. A function  $f : \mathbb{I} \rightarrow \mathbb{O}$  satisfies *Declassified Abstract Non-Interference* (DANI in short) w.r.t.  $\eta$  in input and  $\rho$  in output, written  $f \models \text{DANI}[\eta, \rho]$ , if:

$$\forall x_1. \forall x_2. (\eta(x_1) \neq \eta(x_2) \Rightarrow \rho f(x_1) = \rho f(x_2))$$

In this case, DANI is not a special case of ANI for any *concrete* function  $f$ . Still, we can correlate the two notions when we consider *abstract* functions of the form  $f \circ \eta$ .

**PROPOSITION 2.8.** *Let  $\eta$  in  $\text{uco}(\wp(\mathbb{I}))$  be an abstraction on input and  $\rho$  in  $\text{uco}(\wp(\mathbb{O}))$  be an abstraction on output. Then we have that  $f \circ \eta \models \text{DANI}[\eta, \rho]$  iff  $f \circ \eta \models \text{ANI}[\langle \top, \eta \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ .*

<sup>5</sup>Recall that, a closure can be uniquely identified by the set of its fixpoints.

<sup>6</sup>In the following, we abuse notation denoting with  $f$  a function in  $\mathbb{I} \rightarrow \mathbb{O}$  and its additive lift to sets, defined as:  $f(X) \triangleq \{f(x) \mid x \in X\}$ , given  $X \subseteq \mathbb{I}$ .




Fig. 1. From left to right: Parity domain, Sign domain, and the reduced product between Parity and Sign domains.

## 2.5 Program Semantics and Static Analysis

With program verification, we mean the general process of checking whether a system complies with a requirement, i.e., (formal) description of what the system is allowed and not allowed to do. When considering programs as systems, the behavior (semantics) of a program  $P$  is usually modeled as *the observation* of a set of executions at a fixed degree of precision, one for each possible input for  $P$ . The semantics of a program is usually computed inductively on the program syntax by means of a language interpreter, computing the intended and observable behavior for each step of computation determined by the program instructions.

In this setting, the set-theoretic *semantics* of a program  $P$  manipulating data in  $\mathbb{D}$ , is a subset of  $\mathbb{I} \times \mathbb{O}$ , where  $\mathbb{I} = \mathbb{D}$  is usually the domain of program inputs, while  $\mathbb{O}$  corresponds to the denotation used to observe the program behavior. In standard cases, the semantics can be the small-step trace semantics, denoted  $\llbracket P \rrbracket^T$ , with  $\mathbb{O} = \mathbb{D}^*$  (i.e., each input is associated with a small-step finite execution trace) or the I/O semantics, denoted  $\llbracket P \rrbracket$ , with  $\mathbb{O} = \mathbb{D}$  (i.e., each input is associated with the computed output). In the following, we abuse notation by indicating with the same term a semantics in its set-theoretic form (e.g.,  $\llbracket P \rrbracket \subseteq \mathbb{D} \times \mathbb{D}$ ) and its equivalent functional form (e.g.,  $\llbracket P \rrbracket : \mathbb{D} \rightarrow \mathbb{D}$ ).

When dealing with program verification and analysis, a program semantics is usually lifted to sets of inputs, in order to analyze properties of executions.

For decidability reasons, such a (concrete) semantics can be approximated by building an *abstract* interpreter [12, 13]. In this case, we can fix an observation (i.e., an abstraction)  $\rho$  modeling what we aim to observe of program computations. Then, we can build the interpreter computing the program on properties of data instead of concrete values, namely on subsets of  $\mathbb{D}$ . Specifically, we consider computations on  $\rho(\wp(\mathbb{D}))$  yielding an abstract semantics  $\llbracket P \rrbracket^\rho : \rho(\wp(\mathbb{D})) \rightarrow \rho(\wp(\mathbb{D}))$ . Since we will use this construction only in a few examples, we avoid technical details and show the meaning with a simple example. Consider the following program  $P$ , and consider the classic parity Par and sign Sgn abstractions on numeric values reported in Figure 1 (on the left and center, respectively).

```

in num;
if (num % 2 == 0) {
  res := num * num;
} else {
  res := -(num * num);
}
out res;

```

The program manipulates integer numbers, thus  $\mathbb{D} = \mathbb{Z}$ . For every even integer  $x \in \mathbb{D}$ , we have that  $\llbracket P \rrbracket(x) = x^2$ , while for every odd integer  $x \in \mathbb{D}$ , we have that  $\llbracket P \rrbracket(x) = -x^2$ . That is, the I/O semantics of the program is  $\llbracket P \rrbracket = \{(x, x^2) \mid x \in \mathbb{Z}_{2k}\} \cup \{(x, -x^2) \mid x \in \mathbb{Z}_{2k+1}\}$ .

Let us compute the (abstract) semantics  $\llbracket P \rrbracket^{\text{Sgn}} : \text{Sgn} \rightarrow \text{Sgn}$  evaluating in the abstract the sign of expressions, only knowing the sign of input data. In this case, information may be lost, for instance, when dealing with sums or tests not implying any sign information, such as parity (the parity of a number does not provide any information on its sign). In particular, we have both even and odd numbers among positive and negative numbers, hence we cannot decide the test at the first line of program  $P$  when we only know the sign of the input. This means that we have to conservatively merge both branches: in the true branch  $\text{res}$  is always positive, while in the false one it is always negative, meaning that  $\llbracket P \rrbracket^{\text{Sgn}}(\mathbb{Z}^{+0}) = \mathbb{Z}$  and  $\llbracket P \rrbracket^{\text{Sgn}}(\mathbb{Z}^{-}) = \mathbb{Z}$ . That is, for any  $X \in \text{Sgn}$  in input, the result is always *any sign* (the sign is unknown), modeled by  $\mathbb{Z}$ . On the other hand, when computing  $\llbracket P \rrbracket^{\text{Par}} : \text{Par} \rightarrow \text{Par}$ , if the input is even ( $\mathbb{Z}_{2k}$ ), we can say that  $\text{res}$  is even, and the same for the odd input ( $\mathbb{Z}_{2k+1}$ ). Formally,  $\llbracket P \rrbracket^{\text{Par}}(\mathbb{Z}_{2k}) = \mathbb{Z}_{2k}$  and  $\llbracket P \rrbracket^{\text{Par}}(\mathbb{Z}_{2k+1}) = \mathbb{Z}_{2k+1}$ .

Finally, consider the reduced product of the parity and sign domains, depicted in Figure 1 (on the right), keeping track of both parity and sign of variables. When computing the (abstract) semantics  $\llbracket P \rrbracket^{\text{ParSgn}} : \text{ParSgn} \rightarrow \text{ParSgn}$ , then we know that if the input is even (either positive or negative), then the result is positive even, while if it is odd (either positive or negative), the result is negative odd. That is,  $\llbracket P \rrbracket^{\text{ParSgn}}(X) = \mathbb{Z}_{2k}^{+0}$  when  $\text{Par}(X) = \mathbb{Z}_{2k}$ , and  $\llbracket P \rrbracket^{\text{ParSgn}}(X) = \mathbb{Z}_{2k+1}^{-}$  when  $\text{Par}(X) = \mathbb{Z}_{2k+1}$ . For the sake of simplicity, we considered here as abstract semantics  $\llbracket P \rrbracket^{\rho}$  the abstract I/O computation (returning the property at the exit program point), but, in general, we can compute the data abstract property at each program point [13].

## 2.6 Hyper Static Analysis

In Reference [7], *hyperproperties* were introduced to formalize those requirements which are not *trace properties*, i.e., that cannot be checked observing *single* system executions. In this setting, hyperproperties are sets of sets of executions, and no more sets of executions (like trace properties). Hence, a program  $P$  satisfies a trace property  $P$  iff  $\llbracket P \rrbracket^{\mathcal{T}} \subseteq P$ , and it satisfies a hyperproperty  $HP$  iff  $\llbracket P \rrbracket^{\mathcal{T}} \in HP$  or, equivalently, iff  $\{\llbracket P \rrbracket^{\mathcal{T}}\} \subseteq HP$ . The same holds when we observe the I/O only instead of the whole execution trace. That is, the same lift can be done for any chosen execution denotation.

In Reference [33], the authors define  $k$ -bounded subset-closed hyperproperties: these requirements can be refuted just by exhibiting a counterexample set consisting in at most  $k$  executions. It turns out that many interesting hyperproperties are  $k$ -bounded, like Abstract Non-Interference (which is 2-bounded). Conversely, Opacity is not a subset-closed hyperproperty, hence its verification is particularly challenging. In Section 5, we will see how we can approximate the verification of Opacity, using a verification mechanism for Abstract Non-Interference. The latter is obtained by adapting the abstract interpretation-based verification mechanism for non-interference proposed in Reference [34]. We report here the background needed to understand the constructions in Section 5.

*Hypersemantics.* In the rest of the section, we briefly describe the principles of hyperproperties verification [33]. From now on, we consider as concrete semantics the I/O semantics  $\llbracket P \rrbracket$  of a program  $P$ . As outlined in Reference [33], in order to verify hyperproperties, we need to lift the concrete semantics to sets of sets, namely, we need a *hypersemantics*  $\llbracket P \rrbracket_h : \wp(\wp(\mathbb{D})) \rightarrow \wp(\wp(\mathbb{D}))$ . In Reference [33], the authors show how to define a *correct* hypersemantics, starting from a given concrete semantics. Correct here means that it contains the concrete semantics, namely  $\llbracket P \rrbracket(X) \triangleq \{\llbracket P \rrbracket(x) \mid x \in X\} \in \llbracket P \rrbracket_h(\{X\})$ , for any  $X \subseteq \mathbb{D}$ . With an over-approximation of a (correct) hypersemantics, we can soundly verify hyperproperties. The over-approximation is given by an *abstract* hypersemantics, computing on a suitable abstract domain, exploiting *abstract interpretation* [12]. Indeed, abstract interpretation is a framework for approximating programs semantics, which is sound (i.e., correct) by construction. An actual verification mechanism requires a (computable) abstract semantics computed on a (machine representable) abstract domain. In particular, we first

have to define an abstract domain of computation  $\langle \mathbb{D}^{\natural}, \subseteq^{\natural} \rangle$ , forming a Galois connection with the concrete domain of computation  $\langle \wp(\wp(\mathbb{D})), \subseteq \rangle$ , namely<sup>7</sup>:

$$\langle \wp(\wp(\mathbb{D})), \subseteq \rangle \xleftrightarrow[\alpha]{\gamma} \langle \mathbb{D}^{\natural}, \subseteq^{\natural} \rangle$$

Then, with an over-approximation  $\llbracket P \rrbracket_h^{\natural} : \mathbb{D}^{\natural} \rightarrow \mathbb{D}^{\natural}$  of the concrete hypersemantics, i.e., such that  $\alpha(\llbracket P \rrbracket_h(\{X\})) \subseteq^{\natural} \llbracket P \rrbracket_h^{\natural}(\alpha(\{X\}))$ , for any  $X \subseteq \mathbb{D}$ , and an under-approximation  $\text{HP}^{\natural}$  of the hyperproperty to check, i.e., such that  $\gamma(\text{HP}^{\natural}) \subseteq \text{HP}$ , we can verify HP totally in the abstract domain [34]:

$$\llbracket P \rrbracket_h^{\natural}(\alpha(\{\mathbb{D}\})) \subseteq^{\natural} \text{HP}^{\natural} \text{ implies } \{\llbracket P \rrbracket\} \subseteq \text{HP}$$

### 3 Generalizing Confidentiality Notions

In this section, we show how existing notions of confidentiality on systems, in particular Abstract Non-Interference and Opacity, can be generalized (abstracted) and how they correlate. Without loss of generality, from now on we consider programs as the systems of interest. The idea is to build up what we will call confidentiality properties in terms of simpler execution requirements, that is, constraints on program executions. This process has the two-fold objective of better understanding the *essence* of known confidentiality notions and, contingently, of highlighting their mutual correlation.

In the following, we provide a general process of defining semantic (program) properties, that we then instantiate to confidentiality. Let  $\mathcal{S} \triangleq \mathbb{I} \rightarrow \mathbb{O}$  be the set of semantic functions of the programs under consideration. An *execution requirement* of  $\mathcal{S}$  is a formula  $\pi$  defined on single functions (i.e., programs) in  $\mathcal{S}$ . Formally, when  $\pi$  contains occurrences of a free variable, say  $x$ , and the function symbol  $f$ , then  $\pi$  stipulates a constraint for the (semantic) function  $f$  in  $\mathcal{S}$  executed on  $x$  when  $x$  is substituted with an input value in  $\mathbb{I}$ . More generally,  $\pi$  may contain more than one free variable, stipulating a constraint on multiple program executions (dubbed *hyperproperties* [7] in literature). We assume an underlying first-order logic with equality where  $f$  is the only uninterpreted function symbol and variables are interpreted in  $\mathbb{I}$ . We denote with  $\text{free}(\pi) = \{x_1, \dots, x_n\}$  the set of variables appearing free in  $\pi$  and we denote by  $x \in \text{free}(\pi)$  the fact the  $x$  is among such variables. Given an execution requirement  $\pi$ , having  $x$  as a free variable, we denote with  $\pi[V/x]$  the fact that the variable  $x$  has been substituted with the value  $v$  belonging to  $\mathbb{I}$  inside  $\pi$ . We then define a *semantic formula*  $\Pi$  inductively on its free variables as follows.

- An execution requirement  $\pi$  is a semantic formula.
- If  $\Pi$  is a semantic formula and  $x \in \text{free}(\Pi)$ , then the following are semantic formulae:  $\Pi[V/x]$ , with  $v$  belonging to  $\mathbb{I}$ ;  $\exists x . \Pi$ ; and  $\forall x . \Pi$ .

A semantic formula  $\Pi$  is *closed* when it does not contain free variables, namely when  $\text{free}(\Pi) = \emptyset$ . In this setting, we define a *semantic property* as a closed semantic formula, and we denote with  $f \models \Pi$  the fact that the function  $f$  in  $\mathcal{S}$  satisfies (i.e., it is a model of) the semantic property  $\Pi$ . Intuitively:  $f$  satisfies a semantic property of the form  $\Pi[V/x]$  if  $f$  fulfills the constraints stipulated by  $\Pi$  when a given value  $v$  belonging to  $\mathbb{I}$  is substituted to  $x$  inside  $\Pi$ ;  $f$  satisfies a semantic property of the form  $\exists x . \Pi$  if there exists at least one value  $v$  belonging to  $\mathbb{I}$  that substituted to  $x$  inside  $\Pi$  makes  $f$  fulfilling the constraints stipulated by  $\Pi$ ; and  $f$  satisfies a semantic property of the form  $\forall x . \Pi$  if every value  $v$  belonging to  $\mathbb{I}$  substituted to  $x$  inside  $\Pi$  makes  $f$  fulfilling the constraints stipulated by  $\Pi$ .

We call *local semantic property* a closed semantic formula in which universal ( $\forall$ ) and existential ( $\exists$ ) quantifiers do not appear. Formally, given an execution requirement  $\pi$ , with  $\text{free}(\pi) = \{x_1, \dots, x_n\}$ ,

<sup>7</sup>The following is an alternative representation for Galois connections

we have that  $\Pi \triangleq \pi[V_1, \dots, V_n/x_1, \dots, x_n]$ , for some values  $v_1, \dots, v_n$  belonging to  $\mathbb{I}$ , is a local semantic property.

In addition, we define a notion of *restricted semantic formula*, where variables are interpreted in  $I \subseteq \mathbb{I}$  in place of  $\mathbb{I}$ . This intuitively means that the constraints imposed by  $\Pi$  should hold for a subset of input values only. Of course, the restriction makes sense only when  $\Pi$  contains free variables. Formally, given a semantic formula  $\Pi$  with  $x \in \text{free}(\Pi)$ , we have that  $\exists_1 x . \Pi$  and  $\forall_1 x . \Pi$  are restricted semantic formulae whose intended meaning is as follows:  $f$  satisfies  $\exists_1 x . \Pi$  if there exists at least one value  $v$  belonging to  $I \subseteq \mathbb{I}$  that substituted to  $x$  inside  $\Pi$  makes  $f$  fulfilling the constraints stipulated by  $\Pi$ ; and  $f$  satisfies  $\forall_1 x . \Pi$  if every value  $v$  belonging to  $I \subseteq \mathbb{I}$  substituted to  $x$  inside  $\Pi$  makes  $f$  fulfilling the constraints stipulated by  $\Pi$ .

In this work, we will focus on semantic properties modeling confidentiality, which we call *confidentiality properties*, but the schema above can also be applied to other contexts. For instance, precision in static analysis is formalized by a notion of *completeness* [23] w.r.t. an input observation  $\eta$  and an output observation  $\rho$ , and it can be defined in terms of the execution requirement  $\pi_c \triangleq \rho f \eta(x) = \rho f(x)$ , with  $\text{free}(\pi_c) = \{x\}$ . Then, the classic notion of Completeness [13] is the semantic property  $\Pi_{uc} \triangleq \forall x . \pi_c$ , while Local Completeness [5] on the input value  $v$  belonging to  $\mathbb{I}$  is the (local) semantic property<sup>8</sup>  $\Pi_{lc} = \pi_c[V/x]$ . Note that, the semantic property  $\Pi_{ec} = \exists x . \pi_c$  asks a different question from Local Completeness. The former checks completeness for a given, fixed input value, while the latter checks whether such an input value exists or not, and, to the best of our knowledge, it has not been investigated in the literature.

### 3.1 Generalizing Abstract Non-Interference

We now generalize Abstract Non-Interference, in that we define a confidentiality property capturing the Abstract Non-Interference notions introduced before as special cases. To do so, we introduce a new general notion of Abstract Non-Interference, modeling all different confidentiality aspects considered so far into a unique setting. Following the schema described above, we build a confidentiality property by identifying the execution requirement prerogative of the most general version of ANI and then quantifying the free variables.

In the general case, we can encapsulate into Abstract Non-Interference three confidentiality dimensions: (*what*), the property of input values to protect; (*where*), the input values for which non-interference should hold; and (*who*), the observation capabilities (e.g., of an attacker) on output values. Formally: let  $\eta$  in  $uco(\wp(\mathbb{I}))$  be an abstraction modeling the property of input values that we want to protect, namely whose variation has not to affect the output observation (*what*); let  $\phi$  in  $uco(\wp(\mathbb{I}))$  be an abstraction fixing for which input values we aim at checking (abstract) non-interference (*where*); and let  $\rho$  in  $uco(\wp(\mathbb{O}))$  be an abstraction characterizing what is observable about output values (*who*). The (*where*) abstraction can model declassification since we do not care whether the variation between values with the same  $\phi$  abstraction yields interference in the output observation.

Being Abstract Non-Interference a 2-bounded hyperproperty [33], the execution requirement must characterize the pairs of executions of  $f : \mathbb{I} \rightarrow \mathbb{O}$  that can be classified as non-interfering:

$$\pi_{ANI} \triangleq \phi(x_1) = \phi(x_2) \wedge \eta(x_1) \neq \eta(x_2) \Rightarrow \rho f(x_1) = \rho f(x_2)$$

The above requirement has two free variables hence we can define different confidentiality properties starting from  $\pi_{ANI}$  by adding a quantification on its free variables  $x_1$  and  $x_2$  or by instantiating them with a value.

<sup>8</sup>It can also be defined, equivalently, in terms of the restricted semantic property  $\forall_{\{v\}} x . \pi_c$ .

First, we can fix two specific executions, having input values  $v_1$  and  $v_2$  in  $\mathbb{I}$ , obtaining the local confidentiality property  $\pi_{\text{ANI}}[V_1, V_2/x_1, x_2]$ . This is not a very significant notion, simply requiring that two fixed computations are not exposing any (abstract) interference.

Then, we can fix one execution and quantify the other. An outer existential quantification would have no particular meaning. Indeed, given an input value  $v$  in  $\mathbb{I}$ , the confidentiality property  $\exists x_1. \pi_{\text{ANI}}[V/x_2]$ , unfolded as  $\exists x_1. (\phi(x_1) = \phi(v) \wedge \eta(x_1) \neq \eta(v) \Rightarrow \rho f(x_1) = \rho f(v))$ , is satisfied by any  $f$  when  $\phi \neq \top$  or  $\eta \neq \text{id}$ . Things are different when we adopt a universal outer quantification, i.e., by considering the confidentiality property  $\forall x_1. \pi_{\text{ANI}}[V/x_2]$ , which indeed corresponds to the abstract version of an existing non-interference requirement, that is, single non-interference [16].

*Definition 3.1 (Single Generalized Abstract Non-Interference).* Let  $\phi, \eta$  in  $\text{uco}(\wp(\mathbb{I}))$  be abstractions on input and  $\rho$  in  $\text{uco}(\wp(\mathbb{O}))$  be an abstraction on output. Given an input value  $v$  in  $\mathbb{I}$ , a function  $f : \mathbb{I} \rightarrow \mathbb{O}$  satisfies *Generalized Abstract Non-Interference on  $v$* , w.r.t.  $\phi, \eta$  in input and  $\rho$  in output, written  $f \models \text{GANI}(v) [\langle \phi^{\#}, \eta^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ , if:

$$\forall x. (\phi(x) = \phi(v) \wedge \eta(x) \neq \eta(v) \Rightarrow \rho f(x) = \rho f(v))$$

This notion can be seen as a “controlled” notion of non-interference since it basically says that a specific execution is (abstract) non-interferent with any other execution.

Finally, we can quantify both free variables. Again, by definition of  $\pi_{\text{ANI}}$ , if one of the two variables is existentially quantified, then when  $\phi \neq \top$  or  $\eta \neq \text{id}$ , the property would be trivially satisfied by any program, resulting in a not particular meaningful notion. Instead, if we universally quantify both variables, i.e., by considering the confidentiality property  $\forall x_1. \forall x_2. \pi_{\text{ANI}}$ , we obtain a generalized notion of Abstract Non-Interference.

*Definition 3.2 (Generalized Abstract Non-Interference).* Let  $\phi, \eta$  in  $\text{uco}(\wp(\mathbb{I}))$  be abstractions on input and  $\rho$  in  $\text{uco}(\wp(\mathbb{O}))$  be an abstraction on output. A function  $f : \mathbb{I} \rightarrow \mathbb{O}$  satisfies *Generalized Abstract Non-Interference (GANI in short)* w.r.t.  $\phi, \eta$  in input and  $\rho$  in output, written  $f \models \text{GANI}[\langle \phi^{\#}, \eta^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ , if:

$$\forall x_1. \forall x_2. (\phi(x_1) = \phi(x_2) \wedge \eta(x_1) \neq \eta(x_2) \Rightarrow \rho f(x_1) = \rho f(x_2))$$

Let us show now that this notion precisely captures all the flavors of Abstract Non-Interference presented in the background. In particular, narrow (NANI) and declassified (DANI) Abstract Non-Interference can be obtained by checking Definition 3.2 for specific input observations:  $\phi = \top$  in the case of DANI, not admitting any form of declassification; and  $\eta = \text{id}$  in the case of NANI, not fixing any form of protection.

**PROPOSITION 3.3.** *Let  $\phi, \eta$  in  $\text{uco}(\wp(\mathbb{I}))$  be abstractions on input and  $\rho$  in  $\text{uco}(\wp(\mathbb{O}))$  be an abstraction on output. The following hold for any function  $f : \mathbb{I} \rightarrow \mathbb{O}$ .*

- $f \models \text{DANI}[\eta, \rho]$  iff  $f \models \text{GANI}[\langle \top^{\#}, \eta^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$
- $f \models \text{NANI}[\phi, \rho]$  iff  $f \models \text{GANI}[\langle \phi^{\#}, \text{id}^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$

**PROOF.** The proof directly follows from Definitions 3.2, 2.7, and 2.6. □

As far as Abstract Non-Interference (ANI) of Definition 2.4 is concerned, we can capture such a notion by using Definition 3.2 and changing the function under consideration: from the function  $f$  computing on concrete values (i.e., on  $\mathbb{I}$ ) to the function  $f \circ \eta$  computing on abstract values (i.e., on  $\eta(\mathbb{I})$ ). This is necessary since ANI is defined as abstract program executions instead of concrete ones (as all other confidentiality notions).

**PROPOSITION 3.4.** *Let  $\phi, \eta$  in  $\text{uco}(\wp(\mathbb{I}))$  be abstractions on input and  $\rho$  in  $\text{uco}(\wp(\mathbb{O}))$  be an abstraction on output. For any function  $f : \mathbb{I} \rightarrow \mathbb{O}$  we have:*

$$f \models \text{ANI}[\langle \phi, \eta \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}] \text{ iff } f \circ \eta \models \text{GANI}[\langle \phi^{\#}, \text{id}^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$$

PROOF. The proof follows from Definitions 3.2 and 2.4.  $\square$

*Example 3.5.* Consider the following *Program 3* computing the remainder of the division by 2. That is, the program computes the function  $f_5 : \mathbb{Z} \rightarrow \mathbb{Z}$  such that  $f_5(x) = x \bmod 2$ .

```

/* Program 3 */
remainder := number;
if (number < 0) {
  remainder := -(remainder);
}
while (remainder >= 2) {
  remainder := remainder - 2;
}
out remainder;

```

Consider  $\mathbb{I} = \mathbb{O} = \mathbb{Z}$  and the following abstractions  $\phi, \eta, \rho \in \text{uco}(\wp(\mathbb{Z}))$ :  $\phi = \{\emptyset, \mathbb{Z}_{2k}, \mathbb{Z}_{2k+1}, \mathbb{Z}\}$  (Parity);  $\eta = \{\emptyset, \mathbb{Z}^-, \mathbb{Z}^{+0}, \mathbb{Z}\}$  (Sign); and  $\rho = \{\emptyset, \mathbb{Z}\} \cup \{\{z\} \mid z \in \mathbb{Z}\}$  (Constants). We have that  $f_5 \models \text{GANI}[\langle \phi^-, \eta^\# \rangle, \langle \rho \rangle_{\mathbb{O}}]$  holds. To see that, we have to check the pairs of inputs values with the same parity but with different sign, namely the pairs of the form:

- a:**  $v_1 \in \mathbb{Z}_{2k} \cap \mathbb{Z}^-$  and  $v_2 \in \mathbb{Z}_{2k} \cap \mathbb{Z}^{+0}$  e.g.,  $(-4, 0)$  or  $(-2, 6)$   
**b:**  $v_1 \in \mathbb{Z}_{2k+1} \cap \mathbb{Z}^-$  and  $v_2 \in \mathbb{Z}_{2k+1} \cap \mathbb{Z}^{+0}$  e.g.,  $(-3, 5)$

For the pairs of the form **a**, the function  $f_5$  returns 0, e.g.,  $f_5(-4) = 0 = f_5(0)$ , and  $f_5(-2) = 0 = f_5(6)$ ; while for those of the form **b**, it returns 1, e.g.,  $f_5(-3) = 1 = f_5(5)$ . When we observe constants in output, we cannot see any difference from the application of  $f_5$  on any pair of the two forms. Indeed, for each pair we have that  $f_5(v_1) = f_5(v_2)$ , hence  $\rho f_5(v_1) = \rho f_5(v_2)$  (in particular, it is either the constant  $\{0\}$  or  $\{1\}$ ). In this specific example, GANI holds for any  $\eta$ , thus, it is equivalent to  $\text{NANI } f_5 \models \text{NANI}[\phi, \rho]$  (Definition 2.6).

*End Example.*

### 3.2 Generalizing Opacity

To find common ground to compare confidentiality notions, specifically Abstract Non-Interference and Opacity, we now characterize (and indeed generalize) Opacity as a confidentiality property. Similarly to the previous subsection, we define the execution requirement imposed by Opacity (Definition 2.1). To be more generic, we define the requirement directly by adopting closure operators instead of equivalence relations<sup>9</sup>. Doing so, the property that has to be *opaque*, usually modeled as a binary predicate (Definition 2.1), may induce more than two partitions on values.

Similar to Abstract Non-Interference, the Opacity requirement stipulates a constraint on pairs of executions. Such a requirement intuitively says that a variation of the input values, up to an input abstraction  $\eta$ , should not result in a variation of the output values, up to an output abstraction  $\rho$ , when considering executions yielding from input values indistinguishable by another input abstraction  $\phi$ . This is tantamount to saying that from the output observation  $\rho$ , we cannot deduce the input property  $\eta$  of values. Indeed, we always have the same output, up to  $\rho$ , stemming from input values having different abstractions in  $\eta$ . Formally:

$$\pi_{\text{AOp}} \triangleq \phi(x_1) = \phi(x_2) \wedge \eta(x_1) \neq \eta(x_2) \wedge \rho f(x_1) = \rho f(x_2)$$

We can now fix two specific executions, instantiating  $x_1$  and  $x_2$  with two values. But, as it happens for Abstract Non-Interference, this would simply say that the two specific executions are (abstract) opaque. More interesting is when we fix one computation and quantify the other. In this case, the conjunctions in  $\pi_{\text{AOp}}$  make an outer universal quantification meaningless: the only function satisfying such a confidentiality property would be the one such that  $\rho f$  is a constant, since it would

<sup>9</sup>Abstract domains strictly generalize equivalence relations, and therefore binary predicates [26].

be the case that  $\phi = \top$  and  $\eta = id$ . Nevertheless, we obtain a meaningful confidentiality property when considering an outer existential quantification, i.e., when considering  $\exists x_1 . \pi_{\text{AOP}}[\mathbb{V}/x_2]$ , for a fixed input value  $v$  in  $\mathbb{I}$ .

*Definition 3.6 (Single Abstract Opacity).* Let  $\phi, \eta$  in  $uco(\wp(\mathbb{I}))$  be abstractions on input and  $\rho$  in  $uco(\wp(\mathbb{O}))$  be an abstraction on output. Given an input value  $v$  in  $\mathbb{I}$ , the property  $\eta$  is *abstract opaque* for a function  $f : \mathbb{I} \rightarrow \mathbb{O}$  on  $v$  w.r.t.  $\phi$  in input and  $\rho$  in output, written  $f \models \text{AOP}(v) [\eta : \langle \phi \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ , if:

$$\exists x . (\phi(x) = \phi(v) \wedge \eta(x) \neq \eta(v) \wedge \rho f(x) = \rho f(v))$$

Going further, we can quantify both free variables. Again, since  $\pi_{\text{AOP}}$  is composed only of conjunctions, when both variables are universally quantified, we actually impose that  $\phi = \top$  and  $\eta = id$ , implying that  $\rho f$  is a constant, and resulting in a meaningless confidentiality property. A similar reasoning can be done if we existentially quantify  $x_1$  and we universally quantify  $x_2$ . Instead, doing the opposite, namely universally quantifying  $x_1$  and existentially quantifying  $x_2$ , we obtain the confidentiality property  $\forall x_1 . \exists x_2 . \pi_{\text{AOP}}$ , that generalizes the notion of Opacity (Definition 2.1).

*Definition 3.7 (Abstract Opacity).* Let  $\phi, \eta$  in  $uco(\wp(\mathbb{I}))$  be abstractions input and  $\rho$  in  $uco(\wp(\mathbb{O}))$  be an abstraction on output. The property  $\eta$  is *abstract opaque* for a function  $f : \mathbb{I} \rightarrow \mathbb{O}$  w.r.t.  $\phi$  in input and  $\rho$  in output, written  $f \models \text{AOP}[\eta : \langle \phi \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ , if:

$$\forall x_1 . \exists x_2 . (\phi(x_1) = \phi(x_2) \wedge \eta(x_1) \neq \eta(x_2) \wedge \rho f(x_1) = \rho f(x_2))$$

The above definition indeed generalizes the classic notion of Opacity. Consider  $P \in \text{Op}(f, \overset{\sim}{\sim}, \overset{\sim}{\sim})$ , where  $P \subseteq \mathbb{I}$  is a binary predicate on input values, while  $\overset{\sim}{\sim}$  in  $\text{Eq}(\mathbb{I})$  and  $\overset{\sim}{\sim}$  in  $\text{Eq}(\mathbb{O})$  are equivalence relations on input and output values, respectively. The corresponding partitioning closure for the binary predicate  $P \subseteq \mathbb{I}$  is  $\eta_P \triangleq \{\emptyset, P, \bar{P}, \mathbb{I}\}$ . For the equivalence relations  $\overset{\sim}{\sim}$  and  $\overset{\sim}{\sim}$ , the corresponding partitioning closures are  $\phi_i \triangleq \text{Clo}^{\overset{\sim}{\sim}}$  and  $\rho_o \triangleq \text{Clo}^{\overset{\sim}{\sim}}$ , respectively. In the following, to simplify the notation, when we write  $\phi_i$  we mean the closure induced by the equivalence relation  $\overset{\sim}{\sim}$  (the same applies for  $\rho_o$  with  $\overset{\sim}{\sim}$ ).

**PROPOSITION 3.8.** *Let  $\overset{\sim}{\sim}$  in  $\text{Eq}(\mathbb{I})$  and  $\overset{\sim}{\sim}$  in  $\text{Eq}(\mathbb{O})$  be equivalence relations on input and output values, respectively, and  $P \subseteq \mathbb{I}$  be a binary predicate on input values. Then,  $P \in \text{SOp}(f, \overset{\sim}{\sim}, \overset{\sim}{\sim})$  iff  $f \models \text{AOP}[\eta_P : \langle \phi_i \rangle_{\mathbb{I}}, \langle \rho_o \rangle_{\mathbb{O}}]$ .*

**PROOF.** The proof trivially follows from observing that: by definition of  $\eta_P$ , we have  $\eta_P(x_1) \neq \eta_P(x_2)$  iff  $x_1 \in P$  and  $x_2 \in \bar{P}$  (or vice versa); by definition of  $\phi_i$ , we have  $\phi_i(x_1) = \phi_i(x_2)$  iff  $x_1 \overset{\sim}{\sim} x_2$ ; by definition of  $\rho_o$ , we have  $\rho_o f(x_1) = \rho_o f(x_2)$  iff  $f(x_1) \overset{\sim}{\sim} f(x_2)$ .  $\square$

When dealing with partitioning closures, what Abstract Opacity stipulates is that every equivalence class of a given equivalence relation is opaque, as per classic Opacity (Definition 2.1). Indeed, it is easy to see that when we consider a generic equivalence relation  $R$  in  $\text{Eq}(\mathbb{O})$ , we have that  $f \models \text{AOP}[\eta_R : \langle \phi_i \rangle_{\mathbb{I}}, \langle \rho_o \rangle_{\mathbb{O}}]$  holds when  $[v]_R \in \text{Op}(f, \overset{\sim}{\sim}, \overset{\sim}{\sim})$  for every input value  $v$  in  $\mathbb{I}$ . The particular case when the equivalence relation stems from a predicate  $P \subseteq \mathbb{I}$  is exactly the classic definition of Symmetric Opacity (Definition 2.1). Indeed, the equivalence classes of the equivalence relation induced by  $P$  are precisely  $P$  and  $\bar{P}$ .

*Example 3.9.* Suppose to have a system that deals with the balance of bank accounts. Let  $\mathbb{I} = \mathbb{O} = \mathbb{Z}$ , where  $\mathbb{Z}$  is the set of integer numbers. Consider the input observation  $\phi_i \triangleq \{\emptyset, [-50, 50], \mathbb{Z} \setminus [-50, 50], \mathbb{Z}\}$  (where  $\overset{\sim}{\sim} = \{[-50, 50], \mathbb{Z} \setminus [-50, 50]\}$  is the corresponding equivalence relation on  $\mathbb{I}$ ) and the input property to protect  $\eta_R \triangleq \{\emptyset, [-\infty, -10], [-9, 9], [10, \infty], \mathbb{Z}\}$  (where the corresponding equivalence relation is  $R = \{[-\infty, -10], [-9, 9], [10, \infty]\}$ ). The property to protect may represent investment risk factors:  $[10, \infty]$ , for low risk;  $[-9, 9]$ , for normal risk; and  $[-\infty, -10]$ ,




Fig. 2. Graphical explanation of Abstract Opacity for *Program 4*.

for high risk. Instead, the input observation can be used to profile users:  $[-50, 50]$ , for a common market profile; and  $\mathbb{Z} \setminus [-50, 50]$ , for corner case profiles (either too poor or too rich). Finally, let  $\rho_o \triangleq \{\emptyset, \mathbb{Z}^-, \mathbb{Z}^{+0}, \mathbb{Z}\}$  (Sign domain) be the output observation (where  $\overset{\circ}{=} = \{[-\infty, -1], [0, +\infty]\}$  is the corresponding equivalence relation). Now consider a *Program 4* that aims to hide low and high risk balances.

```

/* Program 4 */
if (-9 <= balance and balance <= 9) { out balance; }
else {
  if (abs(balance) >= 50) {
    if (balance < 0) { out balance + 51; }
    else { out balance - 52; }
  } else {
    if (balance < 0) { out balance + 11; }
    else { out balance - 11; }
  }
}

```

The program computes the function  $f_4 \in \mathbb{Z} \rightarrow \mathbb{Z}$  defined as

$$f_4(x) \triangleq \begin{cases} x + 51 & \text{if } x \leq -50 \\ x - 52 & \text{if } x \geq 50 \\ x + 11 & \text{if } -50 < x \leq -10 \\ x - 11 & \text{if } 10 \leq x < 50 \\ x & \text{otherwise} \end{cases}$$

The program protects the property  $\eta_R$ , since every equivalence class of  $R$  is made opaque by the program, namely  $f_4 \models \text{AOp}[\eta_R : \langle \phi_i \rangle_{\mathbb{I}}, \langle \rho_o \rangle_{\mathbb{O}}]$ . Indeed, for each value (e.g., 0) we can find another value (e.g., 11) in the relation  $\overset{\circ}{\sim}$  (i.e.,  $0 \overset{\circ}{\sim} 11$ ) but not in the relation  $R$  (i.e.,  $0 \not\sim 11$ ) with the same result for  $\rho_o f_4$  (i.e.,  $f_4(0) \overset{\circ}{\sim} f_4(11)$ ). This is summarized in Figure 2, where we report a graphical representation of Abstract Opacity for *Program 4*. On the left of the figure, we have the regions of values satisfying the input conjuncts of Abstract Opacity, that is  $\phi_i(v_1) = \phi_i(v_2) \wedge \eta_R(v_1) \neq \eta_R(v_2)$ , while on the right we have the regions of values satisfying the output conjunct of Abstract Opacity, that is  $\rho_o(v_1) = \rho_o(v_2)$ . *Program 4* (i.e.,  $f_4$ ) maps input regions to output regions, satisfying Abstract Opacity. Note that, if we do not apply the program (e.g., by considering  $f$  as the identity function)

we have that the equivalence class  $[-9, 9]$  is not opaque. Thus, by observing the sign of output values, we can infer the investment risk factor of a given account.

*End Example.*

By rewriting Opacity in terms of abstractions, we can also state a correlation between the former and Generalized Abstract Non-Interference. Let us consider the single case first.

**THEOREM 3.10.** *Let  $\phi, \eta$  in  $uco(\wp(\mathbb{I}))$  be abstractions on input and  $\rho$  in  $uco(\wp(\mathbb{O}))$  be an abstraction on output. Given an input value  $v$  in  $\mathbb{I}$ , if there exists another input value  $v'$  in  $\mathbb{I}$  such that  $\phi(v) = \phi(v')$  and  $\eta(v) \neq \eta(v')$ , then:*

$$f \models \text{GANI}_{(v)}[\langle \phi^{\#}, \eta^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}] \text{ implies } f \models \text{AOp}_{(v)}[\eta : \langle \phi \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$$

**PROOF.** Since, by assumption there exists  $v'$  such that  $\phi(v) = \phi(v')$  and  $\eta(v) \neq \eta(v')$ , when GANI on  $v$  holds (by Definition 3.1) we have that  $\rho f(v) = \rho f(v')$ . Hence, we have that  $\exists x. (\phi(x) = \phi(v) \wedge \eta(x) \neq \eta(v) \wedge \rho f(x) = \rho f(v))$  holds by substituting  $x$  with  $v'$ . This implies (by Definition 3.6)  $f \models \text{AOp}_{(v)}[\eta : \langle \phi \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ .  $\square$

Such a correlation can be trivially extended to the non-single versions of Abstract Opacity and Generalized Abstract Non-Interference.

**COROLLARY 3.11.** *Let  $\phi, \eta$  in  $uco(\wp(\mathbb{I}))$  be abstractions on input and  $\rho$  in  $uco(\wp(\mathbb{O}))$  be an abstraction on output. If for each input value  $v$  in  $\mathbb{I}$  there exists another input value  $v'$  in  $\mathbb{I}$  such that  $\phi(v) = \phi(v')$  and  $\eta(v) \neq \eta(v')$ , then:*

$$f \models \text{GANI}[\langle \phi^{\#}, \eta^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}] \text{ implies } f \models \text{AOp}[\eta : \langle \phi \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$$

**PROOF.** The proof follows from Theorem 3.10 and from observing that:  $f \models \text{GANI}[\langle \phi^{\#}, \eta^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$  iff for every  $v$  in  $\mathbb{I}$  we have  $f \models \text{GANI}_{(v)}[\langle \phi^{\#}, \eta^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ ; and  $f \models \text{AOp}[\eta : \langle \phi \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$  iff for every  $v$  in  $\mathbb{I}$  we have  $f \models \text{AOp}_{(v)}[\eta : \langle \phi \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ .  $\square$

### 3.3 Generalizing Final Opacity

We can also define final opacity [6] in our formalism, acting on the output of program semantics. Recall that the global version of opacity [6] is obtained by changing the denotation of program results, i.e.,  $\mathbb{O}$ . For instance, we can consider traces of values instead of single (output) values. We directly provide a generalization, introducing an input observation, as it has been done for (initial) Opacity. Note that, even if the definitions of opacity and final opacity appear quite different in [6], they are the same concept instantiated on a different predicate to make opaque. Indeed, given a predicate  $P \subseteq \mathbb{O}$  on output values, we can define a corresponding predicate  $P^{-f} \subseteq \mathbb{I}$  on input values, parametric on a given function  $f : \mathbb{I} \rightarrow \mathbb{O}$ . In particular, we define  $P^{-f} \triangleq \{v \mid f(v) \in P\}$  as the set of input values whose image under  $f$  satisfies  $P$ . To simplify the notation, we denote with  $P^{\sim}$  the predicate  $P^{-f}$  when  $f$  is clear from the context. Then, final opacity [6] for  $P$  is simply Opacity for  $P^{\sim}$ .

**Definition 3.12 (Final Opacity).** Let  $\overset{\sim}{\sim}$  in  $Eq(\mathbb{I})$  and  $\overset{\circ}{\sim}$  in  $Eq(\mathbb{O})$  be equivalence relations on input and output, respectively. A binary predicate  $P \subseteq \mathbb{O}$  is *finally opaque* for a function  $f : \mathbb{I} \rightarrow \mathbb{O}$  w.r.t.  $\overset{\sim}{\sim}$  in input and  $\overset{\circ}{\sim}$  in output, denoted  $P \in \text{FOp}(f, \overset{\sim}{\sim}, \overset{\circ}{\sim})$ , if:

$$\forall_{P^{\sim}} x_1. \exists_{\bar{P}^{\sim}} x_2. (x_1 \overset{\sim}{\sim} x_2 \wedge f(x_1) \overset{\circ}{\sim} f(x_2))$$

When  $P, \bar{P} \in \text{FOp}(f, \overset{\sim}{\sim}, \overset{\circ}{\sim})$ ,  $P$  is *symmetrically and finally opaque*, written  $P \in \text{SFOp}(f, \overset{\sim}{\sim}, \overset{\circ}{\sim})$ .

It is easy to note that  $P \in \text{FOp}(f, \overset{\sim}{\sim}, \overset{\circ}{\sim})$  if and only if  $P^{\sim} \in \text{Op}(f, \overset{\sim}{\sim}, \overset{\circ}{\sim})$ . Note that, when  $\overset{\circ}{\sim}$  is the identity, we have that  $f(P^{\sim}) \cap f(\bar{P}^{\sim}) = \emptyset$ , and therefore  $P$  cannot be finally opaque.

Similar to what we have done in the previous subsection for Opacity, we can generalize Final Opacity by adopting generic properties in place of binary predicates and by identifying the execution

requirement to define it as a confidentiality property. Final Opacity (Definition 3.12) requires that executions starting from input values indistinguishable by the input abstraction  $\phi$  yield a variation of the output values, up to an output abstraction  $\eta$ , that cannot be distinguished by another output observation  $\rho$ . That is, changes in the output observation  $\eta$  are not revealed by observing the computation output up to  $\rho$ . Formally:

$$\pi_{\text{AFOP}} \triangleq \phi(x_1) = \phi(x_2) \wedge \eta f(x_1) \neq \eta f(x_2) \wedge \rho f(x_1) = \rho f(x_2)$$

Analogously to what we have seen for Abstract Non-Interference and Opacity, we can define confidentiality properties for Final Opacity by quantifying and/or instantiating the free variables of  $\pi_{\text{AFOP}}$ . As we observed for Opacity, given the logical structure of the execution requirement, the only meaningful case is the one where we instantiate the variable  $x_2$  and we existentially quantify  $x_1$ , that is  $\exists x_1 . \pi_{\text{AFOP}}[\forall/x_2]$ , given an input value  $v$  in  $\mathbb{I}$ .

*Definition 3.13 (Single Abstract Final Opacity).* Let  $\phi$  in  $\text{uco}(\wp(\mathbb{I}))$  be an abstraction on input and  $\eta, \rho$  in  $\text{uco}(\wp(\mathbb{O}))$  be abstractions on output. Given an input value  $v$  in  $\mathbb{I}$ , the property  $\eta$  is *finally abstract opaque* for a function  $f : \mathbb{I} \rightarrow \mathbb{O}$  on  $v$  w.r.t.  $\phi$  in input and  $\rho$  in output, written  $f \models \text{AFOP}(v) [\eta : \langle \phi \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ , if:

$$\exists x . \phi(x) = \phi(v) \wedge \eta f(x) \neq \eta f(v) \wedge \rho f(x) = \rho f(v)$$

When quantifying on both variables, with a similar reasoning as for Opacity, we consider as meaningful the case where we universally quantify on  $x_1$ , and we existentially quantify on  $x_2$ , that is  $\forall x_1 . \exists x_2 . \pi_{\text{AFOP}}$ , obtaining so far an abstract notion of Final Opacity.

*Definition 3.14 (Abstract Final Opacity).* Let  $\phi$  in  $\text{uco}(\wp(\mathbb{I}))$  be an abstraction on input and  $\eta, \rho$  in  $\text{uco}(\wp(\mathbb{O}))$  be abstractions on output. The property  $\eta$  is *finally abstract opaque* for a function  $f : \mathbb{I} \rightarrow \mathbb{O}$  w.r.t.  $\phi$  in input and  $\rho$  in output, written  $f \models \text{AFOP}[\eta : \langle \phi \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ , if:

$$\forall x_1 . \exists x_2 . \phi(x_1) = \phi(x_2) \wedge \eta f(x_1) \neq \eta f(x_2) \wedge \rho f(x_1) = \rho f(x_2)$$

The above definition indeed generalizes the classic notion of Final Opacity. Consider  $P \in \text{FOp}(f, \overset{\sim}{\sim}, \overset{\sim}{\sim})$ , where  $P \subseteq \mathbb{O}$  is a binary predicate on output values, while  $\overset{\sim}{\sim}$  in  $\text{Eq}(\mathbb{I})$  and  $\overset{\sim}{\sim}$  in  $\text{Eq}(\mathbb{O})$  are equivalence relations on input and out values, respectively. The corresponding partitioning closure for the binary predicate  $P \subseteq \mathbb{I}$  is  $\eta_P \triangleq \{\emptyset, P, \bar{P}, \mathbb{O}\}$ . For the equivalence relations  $\overset{\sim}{\sim}$  and  $\overset{\sim}{\sim}$ , the corresponding partitioning closures are  $\phi_i \triangleq \text{Clo}_{\overset{\sim}{\sim}}$  and  $\rho_o \triangleq \text{Clo}_{\overset{\sim}{\sim}}$  respectively.

**PROPOSITION 3.15.** *Let  $\overset{\sim}{\sim}$  in  $\text{Eq}(\mathbb{I})$  and  $\overset{\sim}{\sim}$  in  $\text{Eq}(\mathbb{O})$  be equivalence relations on input and output, respectively, and  $P \subseteq \mathbb{O}$  a binary predicate on output values. Then,  $P \in \text{FOp}(f, \overset{\sim}{\sim}, \overset{\sim}{\sim})$  iff  $f \models \text{AFOP}[\eta_P : \langle \phi_i \rangle_{\mathbb{I}}, \langle \rho_o \rangle_{\mathbb{O}}]$ .*

**PROOF.** The proof trivially follows from observing that: by definition of  $\eta_P$ , we have  $\eta_P f(x_1) \neq \eta_P f(x_2)$  iff  $f(x_1) \in P$  and  $f(x_2) \in \bar{P}$ ; by definition of  $P^\sim$ , we have  $f(x_1) \in P$  iff  $x_1 \in P^\sim$  and  $f(x_2) \in \bar{P}$  iff  $x_2 \in P^\sim$ ; by definition of  $\phi_i$ , we have  $\phi_i(x_1) = \phi_i(x_2)$  iff  $x_1 \overset{\sim}{\sim} x_2$ ; by definition of  $\rho_o$ , we have  $\rho_o f(x_1) = \rho_o f(x_2)$  iff  $f(x_1) \overset{\sim}{\sim} f(x_2)$ .  $\square$

#### 4 Confidentiality Notions for Code Protection

As noted in the previous section, Final Opacity can be seen as Opacity w.r.t. the inverse predicate/property to protect. However, in this section, we will provide a different viewpoint of Final Opacity, proving that it can be used to precisely define a novel confidentiality property useful in code protection, precisely in code *obfuscation* [10].

Indeed, by using both Final Opacity and Generalized Abstract Non-Interference, we can characterize a confidentiality property determining whether a program (already) obfuscates a desired

semantic property, in such a case we may avoid to transform (i.e., obfuscate) the program. Precisely as it happens in language-based security, where we have a confidentiality property to satisfy (e.g., Abstract Non-Interference or Opacity) and the corresponding enforcing techniques (e.g., security type systems), also for code protection, we can consider the existing obfuscation techniques as the code transformations enforcing a desired *code protection property*.

In Reference [15], the authors described how obfuscation strategies can be modeled regarding complete program transformations. In particular, obfuscation should achieve two simultaneous goals: *conceal*, i.e., hide as much as possible, a given aspect of the original program; and *preserve*, i.e., do not distort, another given aspect of the original program. In Reference [15], the authors proved that concealment corresponds to making incomplete a program analysis, aiming at retrieving the program aspect that obfuscation is supposed to hide. Recall that, a function  $f : \mathbb{I} \rightarrow \mathbb{O}$  is *complete* [24] for an abstraction  $\phi$  on input values and an abstraction  $\eta$  on *output* values when the following semantic property holds:  $\forall x. (\eta f(x) = \eta f\phi(x))$ . Moreover, recall that, as in the rest of the article, the *output* domain  $\mathbb{O}$  of the semantic function  $f$  is not necessarily the program output, but is in general the denotation modeling the program meaning (e.g., output data, concrete execution trace log, etc.).

We can formally define code obfuscation as considered in Reference [15] in terms of program (in-)completeness and Generalized Abstract Non-Interference (the latter for the preservation part). In the following definition,  $\phi$  represents the observation capability on input values.

*Definition 4.1 (Weak Abstract Obfuscation).* Let  $\phi$  in  $uco(\wp(\mathbb{I}))$  be an abstraction on input and  $\eta, \rho$  in  $uco(\wp(\mathbb{O}))$  be abstractions on output. A function  $f : \mathbb{I} \rightarrow \mathbb{O}$  satisfies *Weak Abstract Obfuscation* under  $\phi$  in input, w.r.t.  $\eta$  to *conceal* and  $\rho$  to *preserve*, written  $f \models \text{WAObf}[\langle \phi \rangle_{\mathbb{I}}, \langle \eta^c, \rho^p \rangle_{\mathbb{O}}]$ , if:

- (weak-conceal)  $f$  weak conceals  $\eta$  under  $\phi$ , that is,  $f$  is incomplete for  $\eta$  and  $\phi$ ; and
- (preserve)  $f$  preserves  $\rho$  under  $\phi$ , that is,  $f \models \text{GANI}[\langle \phi^{\#}, id^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ .

In de-obfuscation by static analysis, the abstractions  $\eta$  and  $\phi$  usually coincide. Indeed,  $\eta f \eta$  represents a static analysis of (the program)  $f$  aiming at retrieving the property  $\eta$  of  $f$  hidden by the obfuscation transformation. When (weak-conceal) holds the analysis is incomplete, i.e., there exists an input  $v$  such that  $\eta f(v) \neq \eta f \eta(v)$ , indicating a de-obfuscation failure, at least on that input. The preservation part given by GANI indicates that the obfuscation transformation has not distorted the property  $\rho$ .

In Reference [36], it has been proven that Completeness and a formulation of Abstract Non-Interference corresponding to GANI of Definition 3.2 are equivalent. That is, given a function  $f : \mathbb{I} \rightarrow \mathbb{O}$ , an abstraction  $\phi$  on input values and an abstraction  $\eta$  on output values, we have [36]:

$$\forall x. (\eta f(x) = \eta f\phi(x)) \text{ iff } f \models \text{GANI}[\langle \phi^{\#}, id^{\#} \rangle_{\mathbb{I}}, \langle \eta \rangle_{\mathbb{O}}]$$

Due to this correspondence, the (weak-conceal) requirement of Definition 4.1 is equivalent to the fact that  $f \models \text{GANI}[\langle \phi^{\#}, id^{\#} \rangle_{\mathbb{I}}, \langle \eta \rangle_{\mathbb{O}}]$  does not hold, thus modeling (weak) obfuscation in terms of Generalized Abstract Non-Interference only. Under this different viewpoint, we can note that Definition 4.1 is a too permissive notion, potentially failing to adequately model concealment. Indeed, the negation of GANI says that it should exist a pair of inputs, indistinguishable by  $\phi$ , such that  $f$  conceals (i.e., protect)  $\eta$ . Formally:  $\exists x_1. \exists x_2. (\phi(x_1) = \phi(x_2) \wedge \eta f(x_1) \neq \eta f(x_2))$ , by definition of GANI. Furthermore, Definition 4.1 says that  $f$  must preserve  $\rho$  for any possible pair of inputs indistinguishable by  $\phi$ . Formally:  $\forall x_1. \forall x_2. (\phi(x_1) = \phi(x_2) \Rightarrow \rho f(x_1) = \rho f(x_2))$ , again, by definition of GANI. We deemed this notion of obfuscation as weak since the concealment part is indeed too permissive. The problem is that being Completeness (and, indeed, GANI) a universal property, its negation only requires the existence of some (potentially few) concealing

executions (for all the other input values the program aspect may not be concealed). Instead, the preservation part is a universal property, hence (correctly) requiring that all executions are preserving.

*Example 4.2.* Consider the following program  $P_{\circ}$  computing the 2nd power of a given number. When the number is odd, the program also inverts the sign of the result.

```
/* Original Program */
if (num % 2 == 0) {
  res := num * num;
} else {
  res := -(num * num);
}
out res;
```

Consider the classic parity  $\text{Par}$  and sign  $\text{Sgn}$  abstractions on numeric values reported in Figure 1 (left-center). Suppose that the parity of values is observed in input (i.e.,  $\phi = \text{Par}$ ); the parity of values should be *preserved* on output (i.e.,  $\rho = \text{Par}$ ); and the sign of values should be *concealed* on the semantic observation (i.e.,  $\eta = \text{Sgn}$ ). Consider an analysis  $f^{\#} = \llbracket P_{\circ} \rrbracket^{\text{ParSgn}}$  on the abstraction  $\text{ParSgn}$  reported in Figure 1 (right), in this case, the observation is the final abstract value obtained by the analysis computed on the  $\text{ParSgn}$  domain. Specifically, the analysis can retrieve the combined information about the parity and sign of numeric values.

The program  $P_{\circ}$  is not obfuscated w.r.t. the considered abstractions, since  $f^{\#} \models \text{WAObf}[\langle \text{Par} \rangle_i, \langle \text{Sgn}^c, \text{Par}^p \rangle_o]$  does not hold. In this case,  $\text{Par}$  is preserved, but  $\text{Sgn}$  is not concealed. Indeed, we have completeness for  $f^{\#}$ . Each pair of even numbers, independently from their sign, is mapped to  $\mathbb{Z}_{2k}^{+0}$  by  $f^{\#}$ , hence indistinguishable by both  $\text{Par}$  and  $\text{Sgn}$ , meaning that both are preserved (none is concealed). Indeed,  $\rho(\mathbb{Z}_{2k}^{+0}) = \text{Par}(\mathbb{Z}_{2k}^{+0}) = \mathbb{Z}_{2k}$ . Similarly, each pair of odd numbers, independently from the numbers' sign, is mapped to  $\mathbb{Z}_{2k+1}^{-}$  by  $f^{\#}$ , hence, again, indistinguishable by both  $\text{Par}$  and  $\text{Sgn}$ . Indeed,  $\eta(\mathbb{Z}_{2k+1}^{-}) = \text{Sgn}(\mathbb{Z}_{2k+1}^{-}) = \mathbb{Z}^{-}$ .

Therefore, by Definition 4.1, we realize that the *program must be obfuscated*. Consider the following obfuscation  $P_{\bullet}$  of  $P_{\circ}$ , where multiplication is implemented by using iterative sums [9, 18, 19]. Since iterative sums are slower than direct multiplication, we may perform obfuscation on even numbers only, for performance needs. The semantics of the original program is not affected by the obfuscation, namely  $\llbracket P_{\circ} \rrbracket = \llbracket P_{\bullet} \rrbracket$ .

```
/* Obfuscation 1 */
if (num % 2 == 0) {
  if (num < 0) { t := -num; }
  else { t := num; }
  res := 0;
  while (t != 0) {
    res := res + num;
    t := t - 1;
  }
  if (num < 0) { res := -res; }
} else {
  res = -(num * num);
}
out res;
```

The program  $P_{\bullet}$  is weakly obfuscated, since we have that  $f^{\#} \models \text{WAObf}[\langle \text{Par} \rangle_i, \langle \text{Sgn}^c, \text{Par}^p \rangle_o]$  holds. In this case,  $\text{Par}$  is preserved, and  $\text{Sgn}$  is only weakly concealed. Indeed, we have incompleteness for  $f^{\#}$ . Specifically, positive even numbers are mapped to  $\mathbb{Z}_{2k}^{+0}$  by  $f^{\#}$ , while negative even numbers are mapped to  $\mathbb{Z}_{2k}$  by  $f^{\#}$ . In the latter case, when computing the abstract value of  $\text{res}$ , the first iteration  $0 + \text{num}$  of the loop is abstracted into  $\mathbb{Z}_{2k}^{+0} + \mathbb{Z}_{2k}^{-} = \mathbb{Z}_{2k}$ , so the final abstract value for  $\text{res}$  is also  $\mathbb{Z}_{2k}$ . Hence, they are indistinguishable by  $\text{Par}$  (preserved)

but not by Sgn (concealed). Indeed,  $\rho(\mathbb{Z}_{2k}^{+0}) = \text{Par}(\mathbb{Z}_{2k}^{+0}) = \mathbb{Z}_{2k} = \text{Par}(\mathbb{Z}_{2k}) = \rho(\mathbb{Z}_{2k})$ , while  $\eta(\mathbb{Z}_{2k}^{+0}) = \text{Sgn}(\mathbb{Z}_{2k}^{+0}) = \mathbb{Z}^{+0} \neq \mathbb{Z} = \text{Sgn}(\mathbb{Z}_{2k}) = \eta(\mathbb{Z}_{2k})$ . However, each pair of odd numbers, independently from the sign, is mapped to  $\mathbb{Z}_{2k+1}^-$  by  $f^\sharp$ , hence indistinguishable by both Par (preserved) and Sgn (not concealed). Hence, *not all inputs are obfuscated*.

*End Example.*

We can make the definition of obfuscation stronger by considering Final Opacity in place of GANI (and, thus, Completeness) for the concealment part. Consider an (input) property  $\eta$  being finally abstract opaque for  $f$  w.r.t. the input abstraction  $\phi$  and the output abstraction  $\tau$ . Then, we can introduce a stronger notion of concealment:

(conceal)  $f$  conceals  $\eta$  under  $\phi$ , that is,  $f \models \text{AFOp}[\eta : \langle \phi \rangle_{\mathbb{I}}, \langle \tau \rangle_{\mathbb{O}}]$

This version of concealment says that for every possible input value, it must exist another input value, indistinguishable by  $\phi$ , such that  $f$  conceals (i.e., protect)  $\eta$ . Formally:  $\forall x_1. \exists x_2. \phi(x_1) = \phi(x_2) \wedge \eta f(x_1) \neq \eta f(x_2)$ , by definition of Abstract Final Opacity. It is trivial to observe that (conceal) is stronger than (weak-conceal). Indeed, it stipulates that for *any* input value we have to conceal the program aspect to protect. We can then define obfuscation as a combination of Abstract Final Opacity and Generalized Abstract Non-Interference.

*Definition 4.3 (Abstract Obfuscation).* Let  $\phi$  in  $\text{uco}(\wp(\mathbb{I}))$  be an abstraction on input and  $\eta, \rho$  in  $\text{uco}(\wp(\mathbb{O}))$  be an abstraction on output. A function  $f : \mathbb{I} \rightarrow \mathbb{O}$  satisfies *Abstract Obfuscation* under  $\phi$  in input, w.r.t.  $\eta$  to conceal and  $\rho$  to preserve, written  $f \models \text{AObf}[\langle \phi \rangle_{\mathbb{I}}, \langle \tau \rangle_{\mathbb{O}}]$ , if:

(conceal)  $f$  conceals  $\eta$  under  $\phi$ , that is,  $f \models \text{AFOp}[\eta : \langle \phi \rangle_{\mathbb{I}}, \langle \tau \rangle_{\mathbb{O}}]$ ; and  
 (preserve)  $f$  preserves  $\rho$  under  $\phi$ , that is,  $f \models \text{GANI}[\langle \phi^\# \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ .

In the context of de-obfuscation by static analysis, the abstractions  $\eta$  and  $\phi$  usually coincide. But, by replacing (weak-conceal) with (conceal), we have that the static analysis of (the program)  $f$  is incomplete on *all* inputs, thus making de-obfuscation infeasible. The preservation part given by GANI still indicates that the obfuscation transformation has not distorted the property  $\rho$ .

*Example 4.4.* Consider the obfuscation  $P_\bullet$  of the original program  $P_\circ$  from Example 4.2. Now suppose to apply data-obfuscation on odd numbers, by inserting semantics-preserving operations into expressions (e.g., adding zero into mathematical expressions) and yielding the following obfuscation  $P_\bullet$ . Again, the semantics of the original program is not affected by the obfuscation, namely  $\llbracket P_\circ \rrbracket = \llbracket P_\bullet \rrbracket$ .

```

/* Obfuscation 2 */
if (num % 2 == 0) {
  if (num < 0) { t := -num; }
  else { t := num; }
  res := 0;
  while (t != 0) {
    res := res + num;
    t := t - 1;
  }
  if (num < 0) { res := -res; }
} else {
  res := -(num + 0)*(num + 0);
}
out res;

```

The program  $P_\bullet$  is (fully) obfuscated, as per Definition 4.3, since we have that  $f^\sharp \models \text{AObf}[\langle \text{Par} \rangle_{\mathbb{I}}, \langle \text{Sgn}^c, \text{Par}^p \rangle_{\mathbb{O}}]$  holds when considering the analysis on ParSgn for the obfuscated program, i.e.,  $f^\sharp = \llbracket P_\bullet \rrbracket^{\text{ParSgn}}$ . In this case, Par is preserved, and Sgn is concealed. Indeed, positive even numbers are mapped to  $\mathbb{Z}_{2k}^{+0}$  by  $f^\sharp$ , while negative even numbers are mapped to  $\mathbb{Z}_{2k}$  by  $f^\sharp$ . Hence, they are indistinguishable by Par (preserved) but not by Sgn (concealed). Similarly, positive

odd numbers are mapped to  $\mathbb{Z}_{2k+1}^-$  by  $f^\sharp$ , while negative odd numbers are mapped to  $\mathbb{Z}_{2k+1}$  by  $f^\sharp$ . In the latter case, when computing the abstract value of `res`, the expression `num + 0` is abstracted into  $\mathbb{Z}_{2k+1}^- + \mathbb{Z}_{2k}^{+0} = \mathbb{Z}_{2k+1}$ , so the final abstract value for `res` is also  $\mathbb{Z}_{2k+1}$ . Hence, they are indistinguishable by `Par` (preserved) but not by `Sgn` (concealed). Indeed,  $\rho(\mathbb{Z}_{2k+1}^-) = \text{Par}(\mathbb{Z}_{2k+1}^-) = \mathbb{Z}_{2k+1} = \text{Par}(\mathbb{Z}_{2k+1}) = \rho(\mathbb{Z}_{2k+1})$ , while  $\eta(\mathbb{Z}_{2k+1}^-) = \text{Sgn}(\mathbb{Z}_{2k+1}^-) = \mathbb{Z}^- \neq \mathbb{Z} = \text{Sgn}(\mathbb{Z}_{2k+1}) = \eta(\mathbb{Z}_{2k+1})$ . Hence, for each input, there exists another input (one with the opposite sign and the same parity), allowing us to conceal `Sgn`.

Therefore, by Definition 4.3, we can realize that the *program does not need to be obfuscated*, as the program aspects of interest are already preserved or concealed, respectively.

*End Example.*

## 5 Verifying Confidentiality Notions

In this section, we show how the verification mechanism for non-interference proposed in Reference [34] can be generalized to other, more permissive, confidentiality notions, such as Generalized Abstract Non-Interference and Opacity.

### 5.1 A Verification Mechanism for Generalized Abstract Non-Interference

Let us start with Generalized Abstract Non-Interference (GANI). Being a hyperproperty, this confidentiality notion requires comparing multiple program executions. As explained in Section 2, this means that GANI is defined on collections of sets of program executions, namely it is an element of  $\wp(\wp(\mathbb{I} \times \mathbb{O}))$ . We recall that here we consider as program executions the program input-output pairs (see Section 2.6 for details).

Note that, the notion of Abstract Non-Interference (ANI) introduced in Section 2 cannot be defined as a hyperproperty of (concrete) program executions, that is, as a hyperproperty of computations over data values. This is not surprising since ANI already considers abstract computations instead of concrete ones: the computation starts from abstract values in  $\eta$  and not from concrete values in  $\mathbb{I}$ . In other words, ANI specifies a relation between the input observation ( $\phi$ ) and the output observation ( $\rho$ ) of *abstract* computations ( $f\eta$ ), namely computations executed on abstract values. So, we cannot define ANI, in general, on the same denotation domain of concrete executions. Things are different when the input abstraction  $\eta$  is the identity, meaning that no abstraction is indeed applied. In that case, ANI with  $\eta = id$  specifies a relation between the input observation ( $\phi$ ) and the output one ( $\rho$ ) of *concrete* computations ( $f$ ). The former reasoning applies to GANI as well, since no abstraction is performed in input to concrete computations, by definition. Thus, GANI can be expressed as a hyperproperty of computations over data values.

By a straightforward application of Definition 3.2, we can formulate GANI as a set of sets of the program *concrete* executions, i.e., as a hyperproperty, as follows:

$$\text{HP}^{\text{GANI}} \triangleq \left\{ X \subseteq \mathbb{I} \times \mathbb{O} \mid \forall (v_1, v'_1), (v_2, v'_2) \in X. \right. \\ \left. \phi(v_1) = \phi(v_2) \wedge \eta(v_1) \neq \eta(v_2) \Rightarrow \rho(v'_1) = \rho(v'_2) \right\}$$

We now show how it is possible to define an abstract interpretation-based verification mechanism for  $\text{HP}^{\text{GANI}}$ , and, thus, for Generalized Abstract Non-Interference. Recalling Section 2.6, a program  $P$  satisfies GANI, w.r.t.  $\phi, \eta$  and  $\rho$ , namely  $f_P \models \text{GANI}[\langle \phi^\sharp, \eta^\sharp \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ , iff  $\{\llbracket P \rrbracket\} \subseteq \text{HP}^{\text{GANI}}$ , where here  $\llbracket P \rrbracket \subseteq \mathbb{I} \times \mathbb{O}$  is the input-output semantics of  $P$ . This check is undecidable, so we can exploit the abstract interpretation framework [12] in order to make the verification feasible, similarly to what has been done for non-interference in [34]. GANI is a 2-bounded hyperproperty [33], which means that we can check it on the subsets of the semantics with cardinality 2 (i.e., on pairs of executions). Hence, given  $\llbracket P \rrbracket \triangleq \{X \subseteq \llbracket P \rrbracket \mid |X| = 2\}$ , we have that  $f_P \models \text{GANI}[\langle \phi^\sharp, \eta^\sharp \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$  iff

$\langle P \rangle \subseteq \text{HP}^{\text{GANI}}$ . Furthermore,  $\langle P \rangle$  can be partitioned as follows:

$$\begin{aligned} \langle P \rangle_{\phi, \eta} &\triangleq \{ \{ (v_1, v'_1), (v_2, v'_2) \} \in \langle P \rangle \mid \phi(v_1) = \phi(v_2) \wedge \eta(v_1) \neq \eta(v_2) \} \\ \langle P \rangle_{\overline{\phi, \eta}} &\triangleq \{ \{ (v_1, v'_1), (v_2, v'_2) \} \in \langle P \rangle \mid \phi(v_1) \neq \phi(v_2) \vee \eta(v_1) = \eta(v_2) \} \end{aligned}$$

Since  $\langle P \rangle_{\overline{\phi, \eta}} \subseteq \text{HP}^{\text{GANI}}$  is trivially always true, we just have to check whether  $\langle P \rangle_{\phi, \eta} \subseteq \text{HP}^{\text{GANI}}$ . Following [34], we can simplify even further the verification process considering the Galois connection<sup>10</sup>:

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

where, without loss of generality, we assume that input and output domains coincide, i.e.,  $\mathbb{D} \triangleq \mathbb{O} = \mathbb{I}$ . The non-relational abstraction  $\alpha_{nr}$  [34] first forgets the relation between the input and the output of the computation (i.e., we only keep the relation between the whole set of inputs and the whole set of outputs), then it forgets the relation between sets of executions (i.e., we only keep the relation between the set of all possible input sets of and the set of all output sets).

At this point, we have that  $f_P \models \text{GANI}[\langle \phi^{\bar{\cdot}}, \eta^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$  iff  $\alpha_{nr}(\langle P \rangle_{\phi, \eta}) \subseteq \alpha_{nr}(\text{HP}^{\text{GANI}})$ . Then, it is easy to note that  $\alpha_{nr}(\langle P \rangle_{\phi, \eta}) \subseteq \alpha_{nr}(\text{HP}^{\text{GANI}})$  iff<sup>11</sup>  $\alpha_{nr}(\langle P \rangle_{\phi, \eta})_{\cdot}$  contains only sets whose elements all agree on  $\rho$ . We denote with  $\text{equiv}_{\rho}$  the set  $\{ X \subseteq \mathbb{D} \mid \forall v_1, v_2 \in X. \rho(v_1) = \rho(v_2) \}$  of all sets of elements in  $\mathbb{D}$  that agree on  $\rho$ . We have the following equivalence:

$$f_P \models \text{GANI}[\langle \phi^{\bar{\cdot}}, \eta^{\#} \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}] \text{ iff } \alpha_{nr}(\langle P \rangle_{\phi, \eta})_{\cdot} \subseteq \text{equiv}_{\rho}. \quad (3)$$

*Example 5.1.* Consider the abstractions and the *Program 3* defined in Example 3.5. We have that  $\langle P_3 \rangle = \{ \{ (v_1, v'_1), (v_2, v'_2) \} \mid v'_1 = v_1 \bmod 2 \text{ and } v'_2 = v_2 \bmod 2 \}$ . Furthermore,  $\langle P_3 \rangle_{\phi, \eta}$  is:

$$\begin{aligned} &\{ \{ (v_1, v'_1), (v_2, v'_2) \} \in \langle P_3 \rangle \mid v_1 \in \mathbb{Z}_{2k} \cap \mathbb{Z}^- \text{ and } v_2 \in \mathbb{Z}_{2k} \cap \mathbb{Z}^{+0} \} \cup && \text{set } A \\ &\{ \{ (v_1, v'_1), (v_2, v'_2) \} \in \langle P_3 \rangle \mid v_1 \in \mathbb{Z}_{2k+1} \cap \mathbb{Z}^- \text{ and } v_2 \in \mathbb{Z}_{2k+1} \cap \mathbb{Z}^{+0} \} && \text{set } B \end{aligned}$$

For all elements  $\{ (v_1, v'_1), (v_2, v'_2) \}$  in the set *A*, by definition of  $P_3$ , we have that  $v'_1 = v'_2 = 0$ . Instead, for all elements  $\{ (v_1, v'_1), (v_2, v'_2) \}$  in the set *B*, by definition of  $P_3$ , we have that  $v'_1 = v'_2 = 1$ . So, we have that  $\alpha_{nr}(\langle P \rangle_{\phi, \eta})_{\cdot} = \{ \{0\}, \{1\} \}$ , which is a subset of  $\text{equiv}_{\rho} = \{ \{n\} \mid n \in \mathbb{Z} \}$ , as expected.

*End Example.*

Now, let  $\mathbb{D}_2 \triangleq \{ \{v_1, v_2\} \subseteq \mathbb{D} \mid \phi(v_1) = \phi(v_2) \wedge \eta(v_1) \neq \eta(v_2) \}$  and  $\{\!\{P\}\!\} : \wp(\mathbb{D}) \rightarrow \wp(\mathbb{D})$  be the collecting semantics of  $P$ . It is easy to note that  $\alpha_{nr}(\langle P \rangle_{\phi, \eta})_{\cdot}$  coincides with the set  $\{ \{\!\{P\}\!\}(X) \mid X \in \mathbb{D}_2 \}$ . By defining a correct collecting hypersemantics  $\{\!\{P\}\!\}_h : \wp(\wp(\mathbb{D})) \rightarrow \wp(\wp(\mathbb{D}))$  for  $P$  (see again Subsection 2.6 for details), we have that  $\{ \{\!\{P\}\!\}(X) \mid X \in \mathbb{D}_2 \} \subseteq \{\!\{P\}\!\}_h(\mathbb{D}_2)$ . Hence, the verification of GANI boils down to check  $\{\!\{P\}\!\}_h(\mathbb{D}_2) \subseteq \text{equiv}_{\rho}$ , which implies  $\alpha_{nr}(\langle P \rangle_{\phi, \eta})_{\cdot} \subseteq \text{equiv}_{\rho}$ .

This check is undecidable since it still involves the computation of the concrete semantics. To perform verification, we need to abstract the computation of the concrete collecting hypersemantics. We can instantiate the *hyperlevel constants* domain of Reference [33] on  $\wp(\mathbb{D})$ , since  $\rho$  is an abstract domain of  $\mathbb{D}$  (i.e., it is a closure on  $\wp(\mathbb{D})$ ). In particular, following the construction of [33], we have that  $\rho^{\text{bc}} = \wp(\text{Atom}(\rho)) \cup \{\rho\}$ . Applying the lifting transformer of [33] we obtain  $\mathcal{L}(\rho) \triangleq \lambda X. \{ \rho(X) \mid X \in \mathcal{X} \}$ . Composing the two, we obtain  $\alpha_{gani} \triangleq \rho^{\text{bc}} \circ \mathcal{L}(\rho)$ , that forms, together with its left adjoint  $\gamma_{gani} = \alpha_{ani}^-$ , the Galois connection [33]:

$$\langle \wp(\wp(\mathbb{D})), \subseteq \rangle \xleftrightarrow[\alpha_{gani}]{\gamma_{gani}} \langle \rho^{\text{bc}}, \subseteq \rangle$$

<sup>10</sup>We abuse notation denoting with  $\subseteq$  both set-inclusion and its component-wise lift to pairs, i.e.,  $(X, Y) \subseteq (\mathcal{T}, \mathcal{Z}) \triangleq (X \subseteq \mathcal{T} \wedge Y \subseteq \mathcal{Z})$ .

<sup>11</sup>Given a generic pair  $(a, b)$ , we denote with  $(a, b)_{\cdot}$  its projection on the second element.

By construction,  $Atom(\rho) = \text{equiv}_\rho$ , which belongs to  $\rho^{\text{bc}}$ . We finally have that if  $\alpha_{gani}(\llbracket P \rrbracket_h(\mathbb{D}_2)) \subseteq \text{equiv}_\rho$  then  $\llbracket P \rrbracket_h(\mathbb{D}_2) \subseteq \text{equiv}_\rho$ . In turn, this implies that  $\{\llbracket P \rrbracket\} \subseteq \text{HP}^{\text{GANI}}$ . Hence, we can soundly approximate the verification of GANI by computing the approximated (abstract) hypersemantics on the hyper domain  $\rho^{\text{bc}}$ , checking whether all pairs of computations have constant value in  $\rho$ . With a sound abstract hypersemantics  $\llbracket P \rrbracket_h^\rho$ , namely a semantics such that for every  $\mathcal{X}$  in  $\wp(\wp(\mathbb{D}))$  it holds  $\alpha_{gani}(\llbracket P \rrbracket_h(\mathcal{X})) \subseteq \llbracket P \rrbracket_h^\rho(\alpha_{gani}(\mathcal{X}))$ , we can move the verification completely on the abstract domain.

**THEOREM 5.2 (GANI VERIFICATION).** *Given  $\mathbb{D}_2^\rho$  in  $\rho^{\text{bc}}$  such that  $\mathbb{D}_2 \subseteq \gamma_{gani}(\mathbb{D}_2^\rho)$ , if  $\llbracket P \rrbracket_h^\rho(\mathbb{D}_2^\rho) \subseteq \text{equiv}_\rho$  then  $P$  satisfies GANI, that is,  $f_P \models \text{GANI}[\langle \phi^-, \eta^\# \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$ .*

**PROOF.** The proof is a trivial generalization of the proof of Theorem 4.6 provided in [34].  $\square$

*Note on implementation.* GANI is indeed a collection of hyperproperties, since it is parametric on the abstractions  $\phi$ ,  $\eta$  and  $\rho$ . The abstract hypersemantics is strictly related to these abstractions. Hence, an actual abstract hypersemantics for GANI cannot be systematically defined. Once the abstractions are fixed, an abstract hypersemantics can be defined, for instance, following the approach used for non-interference described in [34].

## 5.2 A Verification Mechanism for Opacity

Being also Opacity a hyperproperty, its verification requires comparing different program executions. Again, we recall that here we consider as program executions just the program input-output pairs (see Section 2.6 for details). Let us observe that Symmetric Opacity (of Definition 2.1 and of [41]) can be formulated as the hyperproperty  $\text{HP}^{\text{SOP}}(\text{P}) \triangleq \text{HP}^{\text{OP}}(\text{P}) \cap \text{HP}^{\text{OP}}(\bar{\text{P}})$ , where:

$$\text{HP}^{\text{OP}}(\text{P}) \triangleq \{X \subseteq \mathbb{I} \times \mathbb{O} \mid \forall (v_1, v'_1) \in X (v_1 \in \text{P}) \exists (v_2, v'_2) \in X (v_2 \notin \text{P}). v_1 \stackrel{L}{\sim} v_2, \wedge v'_1 \stackrel{O}{\sim} v'_2 \}$$

It is easy to note that  $\text{HP}^{\text{SOP}}(\text{P})$  is not subset-closed<sup>12</sup> [7, 32]. Hence, it is quite a general hyperproperty that is hard to verify using static analysis. Nevertheless, we can exploit the connection between GANI and Abstract Opacity, as introduced in Definition 3.7, to verify Symmetric Opacity.

Indeed, we have already probed the equivalence between Abstract Opacity and Symmetric Opacity, and that GANI implies Abstract Opacity under some assumptions.

**THEOREM 5.3.** *Let  $\stackrel{L}{\sim}$  in  $\text{Eq}(\mathbb{I})$  and  $\stackrel{O}{\sim}$  in  $\text{Eq}(\mathbb{O})$  be equivalence relations on input and output values, respectively, and  $\text{P} \subseteq \mathbb{I}$  be a binary predicate on input values. If for each input value  $v$  in  $\mathbb{I}$  there exists another input value  $v'$  in  $\mathbb{I}$  such that  $\phi_{\mathbb{I}}(v) = \phi_{\mathbb{I}}(v')$  and  $\rho_{\mathbb{O}}(v) \neq \rho_{\mathbb{O}}(v')$ , then:*

$$f \models \text{GANI}[\langle \phi_{\mathbb{I}}^-, \eta_{\mathbb{P}}^\# \rangle_{\mathbb{I}}, \langle \rho_{\mathbb{O}} \rangle_{\mathbb{O}}] \text{ implies } \text{P} \in \text{SOP}(f, \stackrel{L}{\sim}, \stackrel{O}{\sim})$$

**PROOF.** The proof trivially follows from Proposition 3.8 and Corollary 3.11.  $\square$

**Example 5.4.** Consider the abstractions and the Program 3 defined in Example 3.5. Now, let  $\stackrel{L}{\sim} = \text{Rel}^{\phi} = \{(z_1, z_2) \mid z_1, z_2 \in \mathbb{Z}_{2k} \text{ or } z_1, z_2 \in \mathbb{Z}_{2k+1}\}$ . Note that the abstraction  $\eta$  from Example 3.5 is the partitioning closure associated with the binary predicate  $\text{P} \triangleq \{z \mid z \in \mathbb{Z}^{+0}\}$ . The premises of Theorem 5.3 are trivially fulfilled, namely for each  $v_1$  in  $\mathbb{I}$  there exists  $v_2$  in  $\mathbb{I}$  such that  $\phi(v_1) = \phi(v_2)$ , or, equivalently,  $v_1 \stackrel{L}{\sim} v_2$  (i.e., having the same parity) and  $\eta(v_1) \neq \eta(v_2)$  (i.e., having different sign). For instance, if  $v_1$  is even and greater than or equal to zero, then we can take another even number lower than zero. We have already seen that  $f \models \text{GANI}[\langle \phi^-, \eta^\# \rangle_{\mathbb{I}}, \langle \rho \rangle_{\mathbb{O}}]$  holds. By the same case analysis of Example 3.5, we also have that  $\rho f(v_1) = \rho f(v_2)$  (i.e., the same constant is observed

<sup>12</sup>Subsets of opaque semantics may not be opaque due to the existential quantifier in the Opacity definition.

in output). Hence, given  $\overset{\circ}{\sim} = Rel^\rho = \{(z, z) \mid z \in \mathbb{Z}\}$ , we have that  $f(v_1) \overset{\circ}{\sim} f(v_2)$ . Then, we can conclude that  $P \in \text{SOP}(f, \overset{\circ}{\sim}, \overset{\circ}{\sim})$ .

*End Example.*

This means that we can reuse enforcement and verification mechanisms developed for GANI in Subsection 5.1 to verify Opacity. Indeed, due to Theorem 5.3, Symmetric Opacity  $P \in \text{SOP}(f_P, \overset{\circ}{\sim}, \overset{\circ}{\sim})$  is implied by  $f_P \models \text{GANI}[\langle \phi_i^{\circ}, \eta_P^{\circ} \rangle_{\mathbb{I}}, \langle \rho_o \rangle_{\circ}]$  where  $\phi_i = Clo^{\overset{\circ}{\sim}}$ ,  $\rho_o = Clo^{\overset{\circ}{\sim}}$ , and  $\eta_P = \{\emptyset, P, \bar{P}, \mathbb{I}\}$ ; and assuming that for all  $v_1$  in  $P$  there exists  $v_2$  in  $\bar{P}$  such that  $\phi_i(v_1) = \phi_i(v_2)$ . The latter check can be easily performed by looking at the structure of  $\phi_i$ . Therefore, to show that a predicate  $P$  is symmetrically opaque for  $f_P$ , we may just check GANI for  $P$ . In particular, the check boils down to  $\alpha_{nr}(\langle P \rangle_{\phi_i, \eta_P})_{\circ} \subseteq \text{equiv}_{\rho_o}$ , where

$$\text{equiv}_{\rho_o} \triangleq \{X \subseteq \mathbb{D} \mid \forall v_1, v_2 \in X. \rho_o(v_1) = \rho_o(v_2)\}$$

is the set of all sets of elements in  $\mathbb{D}$  that agree on  $\rho_o$ , and

$$\langle P \rangle_{\phi_i, \eta_P} \triangleq \{ \{(v_1, v'_1), (v_2, v'_2)\} \in \langle P \rangle \mid \phi_i(v_1) = \phi_i(v_2) \text{ and } \eta_P(v_1) \neq \eta_P(v_2) \}$$

Similarly to the GANI case, if  $\alpha_{nr}(\langle P \rangle_{\phi_i, \eta_P})_{\circ} \subseteq \text{equiv}_{\rho_o}$  holds, then  $P \in \text{SOP}(f_P, \overset{\circ}{\sim}, \overset{\circ}{\sim})$ . Again, it is easy to note that  $\alpha_{nr}(\langle P \rangle_{\phi_i, \eta_P})_{\circ}$  coincides with the set  $\{\langle P \rangle(X) \mid X \in \mathbb{D}_2\}$ , where  $\mathbb{D}_2 \triangleq \{\{v_1, v_2\} \subseteq \mathbb{I} \mid \phi_i(v_1) = \phi_i(v_2) \text{ and } \eta_P(v_1) \neq \eta_P(v_2)\}$ . By definition of correct hypersemantics, we have that  $\{\langle P \rangle(X) \mid X \in \mathbb{D}_2\} \subseteq \langle P \rangle_h(\mathbb{D}_2)$ . So, we just need to check  $\langle P \rangle_h(\mathbb{D}_2) \subseteq \text{equiv}_{\rho_o}$ , in turn implying  $\alpha_{nr}(\langle P \rangle_{\phi_i, \eta_P})_{\circ} \subseteq \text{equiv}_{\rho_o}$ .

This check is also undecidable. We can then again exploit the hyperlevel constants domain (parametric in  $\rho_o$  this time)  $\rho_o^{\text{bc}} \triangleq \wp(\text{Atom}(\rho_o)) \cup \{\rho_o\}$  to compute an abstract collecting hypersemantics approximating the check. The Galois connection is given by the abstraction function  $\alpha_{sop} \triangleq \rho_o^{\text{bc}} \circ \mathcal{L}(\rho_o)$  and its corresponding concretization  $\gamma_{sop} = \alpha_{sop}^-$ . By construction,  $\text{Atom}(\rho_o) = \text{equiv}_{\rho_o}$ , which belongs to  $\rho_o^{\text{bc}}$ . We then have that  $\alpha_{sop}(\langle P \rangle_h(\mathbb{D}_2)) \subseteq \text{equiv}_{\rho_o}$  implies  $\langle P \rangle_h(\mathbb{D}_2) \subseteq \text{equiv}_{\rho_o}$ . Finally, with a sound abstract hypersemantics  $\langle P \rangle_h^{\rho_o}$ , we can move the verification of Symmetric Opacity totally on the abstract domain.

**THEOREM 5.5 (OPACITY VERIFICATION).** *Given  $\mathbb{D}_2^{\rho_o}$  in  $\rho_o^{\text{bc}}$  such that  $\mathbb{D}_2 \subseteq \gamma_{sop}(\mathbb{D}_2^{\rho_o})$ , if  $\langle P \rangle_h^{\rho_o}(\mathbb{D}_2^{\rho_o}) \subseteq \text{equiv}_{\rho_o}$  then  $P$  satisfies Symmetric Opacity, that is,  $P \in \text{SOP}(f_P, \overset{\circ}{\sim}, \overset{\circ}{\sim})$ .*

**PROOF.** The proof follows from Theorems 5.3 and 5.2.  $\square$

The parametricity of Theorem 5.5 in the predicate  $P$  to check is given by  $\mathbb{D}_2$  (and, hence, by  $\mathbb{D}_2^{\rho_o}$ ), that is defined in terms of  $P$ . Moreover, as it happens for GANI, Opacity is also a collection of hyperproperties, parametric on the equivalence relations  $\overset{\circ}{\sim}$  and  $\overset{\circ}{\sim}$ , and the predicate  $P$  to check. Hence, again, an actual abstract hypersemantics for Opacity can be defined only when the abstractions are fixed.

*Note on precision.* Being Opacity a hyperproperty containing a  $\forall \exists$  quantifier alternation, it is theoretically very challenging to statically verify, and even to dynamically test [11]. The proposed verification analysis for Opacity trades soundness for an inevitable precision loss. We indeed had to resort to approximating Opacity verification with ANI verification, which is theoretically simpler. Thus, our analysis can lose precision when used to verify Opacity. Nevertheless, while possibly yielding false positives, our analysis is sound and practically fast, being the abstract domain adopted limited in size.

## 6 Conclusion

In this article, we investigated the relation between two fundamental confidentiality notions, i.e., opacity [39, 41] and abstract non-interference [21], showing how they represent a different

perspective on the same security issue. To do so, we provided a formal framework for describing the different flavors of abstract non-interference and (abstract) opacity in the same setting, allowing us to compare the two notions formally. In particular, we proved that abstract non-interference is stronger than opacity under specific assumptions. It should be mentioned that the proposed formalism is not specific to abstract non-interference and opacity considered in the article, but it can be used to model generic confidentiality notions.

Another outcome of understanding the connection between opacity and abstract non-interference has been the characterization of a novel confidentiality notion modeling code obfuscation, which can be useful to check whether a system really needs to be obfuscated. This represents a more precise formalization than the one proposed in References [15, 18, 19, 22] based on the incompleteness of program analyses. Indeed, being an existential property (i.e., its refutation requires a single counterexample), incompleteness is too weak for being realistic in the context of code obfuscation.

The proven relation between opacity and abstract non-interference also led to effective verification of opacity by extending what has been done for non-interference in Reference [34]. Specifically, adapting the abstract interpretation-based verification mechanism of Reference [34] to abstract non-interference resulted in effective verification of opacity by exploiting the connection between the two notions discovered in the current article.

Additionally, by leveraging the introduced formalism, we were able to model in the same setting also (abstract) final opacity [6], checking opacity on outputs instead of inputs. Due to its “opposite” nature, the verification mechanism existing for non-interference cannot be seamlessly extended to cope with final opacity. As future work, we plan to investigate this issue in order to make feasible the verification of confidentiality notions focusing on outputs rather than inputs.

From a practical point of view, as future work, we aim at extending the implementation for non-interference [34] to concretely provide a static analysis for verifying the proposed confidentiality properties when abstractions have been specified. From a theoretical point of view, we can observe that from the proposed confidentiality properties, we can derive their dual versions. Intuitively, a dual confidentiality property checks whether a program *does not* satisfy the execution requirement (as per Section 3) of the original confidentiality property. For instance, the dual of Abstract Non-Interference (Definition 2.4) describes a form of *abstract dependency* between input and output data-properties. In contrast, the dual of Abstract Opacity (Definition 4.3) describes a form of *abstract transparency* of an input data-property. Since the notion of abstract dependency is not new in the literature [31, 37], we believe that a deeper understanding of such notions, in particular with respect to the existing literature, deserves further investigation.

Another line of future work could be to consider *active attackers* in our framework. Active attackers introduce the capability of an attacker to alter code, rather than simply observing code execution. A path to follow when considering such an attack model is to check program *robustness*, in addition to confidentiality compliance. In this context, robustness, as introduced in Reference [42], asserts that an active attacker, who can modify program code in some fixed points (holes), is not able to disclose more private information than a passive attacker, who merely observes public data. The problem of checking abstract non-interference under robustness assumptions (i.e., considering active attackers) has been tackled in [2] by exploiting a weakest precondition semantics. In particular, this semantics simulates the analysis an attacker can perform backward, i.e., from the public output towards the private input. However, such a work has been developed on the specific notion of ANI originally proposed for security (allowing a public/private partition of data only), and not on GANI, which is more general. Hence, extending GANI to deal with active attackers within the definition itself is surely an interesting future work. A possibility here could be to introduce “holes” in programs as done in [2, 42], and quantify over the code that can fill such holes. Such an

extension will impact our verification mechanism. The analysis proposed in Section 5 builds on an abstract interpretation-based hyperproperty check, which is inductively defined on the syntax of a *single program*. To the best of our knowledge, there are no abstract interpretation-based static analyses considering *multiple programs* at the same time, and it could be interesting to consider this problem as well. A possibility here could be to take inspiration from recent developments in program synthesis [27], where simple program properties are verified for infinite sets of programs by exploiting a Hoare logic-like analysis.

## 6.1 Related Work

The original notion of opacity [6] defined for Labeled Transition Systems (LTSs) only considers an observation on the output of the computation. More recently, the same notion defined for programs [41] has been extended by considering also an input observation. In Reference [41], it is shown how this notion is crucial in several privacy scenarios, seeking to answer fundamental questions on understanding and enforcing opacity. Moreover, it is shown that a program satisfies non-interference if it is opaque for the possible input predicates. Unfortunately, while adding the input abstraction, in Reference [41], the relation between the observed input data property and the predicate to protect has not been investigated (it should be clear that if the abstraction implies the predicate, then it cannot be opaque, independently of the analyzed computation). In our approach, where the observation capability and the data property to protect are modeled in the same way, this relation becomes clear.

Concerning opacity verification, Balun and Masopust [3] provide an algorithmic approach to verify specific flavors of *state-based opacity*, that is, weak and strong  $k$ -step opacity, on Discrete Event Systems (DESs). Here,  $k$ -step opacity is an opacity notion considered in the context of DESs that characterizes a system's ability to maintain confidential behavior for a bounded time window under passive attackers. In essence, a system is  $k$ -step opaque if a passive attacker, who only sees a limited set of events, cannot definitively determine that the system was in a secret state at any time instant within the last  $k$  observable steps leading up to the current observation. The weak version of  $k$ -step opacity says that the attacker cannot determine the exact time when the system was in a secret state within the last  $k$  observable steps leading up to the current observation. Bourouis et al. [4] consider the same flavors of opacity on finite LTSs. Then, they present efficient algorithms for verifying opacity in all these forms with an on-the-fly, dynamic approach. Similarly, Falcone and Marchand [17] show how to model-check, verify, and enforce  $k$ -step opacity at runtime on LTSs. Klein et al. [28] provide an algorithmic approach to verify current-state opacity, corresponding to  $k$ -opacity with no time window (i.e.,  $k = 0$ ) on timed automata. Saadaoui et al. [40] discuss the formalization and verification of *language-based opacity* when considering DESs represented by partially observable Petri nets. The difference between state-based and language-based opacity is that the latter considers the whole computation, while the former considers only the current system state. Specifically, they identify two opacity notions, called *consistency* and *non-secrecy*, and exploit the mathematical characterization of a Petri net system to separately check each notion. The combined verification allows for the ultimately verification of language-based opacity.

Studying the correlation between the previously mentioned flavors of opacity and the classic notions of opacity considered in the current article would be interesting, but we have not yet investigated this aspect. Since our formalization can be instantiated with different denotations of program executions, we believe that modeling  $k$ -step opacity and similar notions as instantiations of Abstract (Final) Opacity would be possible. From the point of view of verification, all previous approaches consider DESs, while ours considers actual programs. Hence, the applicability of our approach is more direct, since it does not require a program to DSE encoding.

Finally, Liu et al. [29] introduce a quantifiable measure of current-state opacity that considers the likelihood of satisfying opacity for stochastic control systems modeled as general Markov Decision Processes (gMDPs). They also propose verification methods tailored to this novel quantitative notion of opacity for finite gMDPs by using value iteration techniques.

Our formalization of opacity considers abstractions on program execution denotations, modeled in terms of data properties. Thus, our approximation of the system behavior is qualitative and orthogonal to the quantitative approximation considered by Liu et al. Nevertheless, it would be interesting to study the correlation between these two apparently different notions of opacity in a formal setting.

An initial investigation of the relation between opacity and abstract non-interference has been carried out in our previous work [35], for which the current article represents a follow-up work. In the current article, we propose a formalism and a methodology to reason about confidentiality properties in general terms, for which opacity and abstract non-interference are particular cases. By exploiting the formal framework proposed in the current article, we were able to generalize and extend the results from [35]. Indeed, differently from [35], the relation proved in the current article does not require any assumption on the abstractions adopted. Moreover, by incrementally reasoning on confidentiality notions, starting from execution requirements and then adding execution quantifiers, in the current article, we were able to draw simpler connections between opacity and abstract non-interference, and to characterize confidentiality properties on single executions. In addition, we investigated the connection between opacity and abstract non-interference with an apparently unrelated concept like code obfuscation. This investigation led us to provide a way for modeling more realistic obfuscation properties, enabling the identification of programs for which obfuscation is not necessary since the data-property to protect is already concealed. As far as verification is concerned, the approach provided in the current article is a simple adaptation of the one proposed in Reference [35] to the new formalism introduced.

## References

- [1] 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.
- [2] M. Balliu and I. Mastroeni. 2010. A weakest precondition approach to robustness. *LNCS Transactions on Computational Science* 10 (2010), 261–297. DOI : <https://doi.org/10.1007/978-3-642-17499-5>
- [3] J. Balun and T. Masopust. 2023. Verifying weak and strong k-step opacity in discrete-event systems. *Automatica* 155 (2023), 111153. DOI : <https://doi.org/10.1016/j.automatica.2023.111153>
- [4] Amina Bourouis, Kais Klai, Yamen El Touati, and Nejib Ben Hadj-Alouane. 2015. Checking opacity of vulnerable critical systems on-the-fly. *Int. J. Inf. Technol. Web Eng.* 10, 1 (Jan. 2015), 1–30. DOI : <https://doi.org/10.4018/ijitwe.2015010101>
- [5] R. Bruni, R. Giacobazzi, R. Gori, and F. Ranzato. 2021. A logic for locally complete abstract interpretations. In *Symposium on Logic in Computer Science, LICS*. IEEE, New York City, US, 1–13.
- [6] J. W. Bryans, M. Koutny, L. Mazaré, and P. Y. A. Ryan. 2008. Opacity generalised to transition systems. *Int. J. Inf. Sec.* 7, 6 (2008), 421–435.
- [7] M. R. Clarkson and F. B. Schneider. 2010. Hyperproperties. *Journal of Computer Security* 18, 6 (2010), 1157–1210.
- [8] E. Cohen. 1977. Information transmission in computational systems. *SIGOPS Oper. Syst. Rev.* 11, 5 (1977), 133–139.
- [9] C. Collberg and J. Nagra. 2009. *Surreptitious Software: Obfuscation, Watermarking, and Tamperproofing for Software Protection*. Addison-Wesley Professional.
- [10] C. Collberg, C. Thomborson, and D. Low. 1997. *A Taxonomy of Obfuscating Transformations*. Technical Report 148. Department of Computer Sciences, The University of Auckland. Retrieved from [http://www.cs.auckland.ac.nz/~sim\\$collberg/Research/Publications/CollbergThomborsonLow97a/index.html](http://www.cs.auckland.ac.nz/~sim$collberg/Research/Publications/CollbergThomborsonLow97a/index.html)
- [11] A. Correnson, T. Nießen, B. Finkbeiner, and G. Weissenbacher. 2024. Finding  $\forall\exists$  hyperbugs using symbolic execution. *Proc. ACM Program. Lang.* 8, OOPSLA2, Article 321 (Oct. 2024), 26 pages. DOI : <https://doi.org/10.1145/3689761>
- [12] 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.
- [13] P. Cousot and R. Cousot. 1979. Systematic design of program analysis frameworks. In *Proc. of Conf. Record of the 6th ACM Symp. on Principles of Programming Languages (POPL'79)*. ACM Press, New York, 269–282.

- [14] M. Dalla Preda and Roberto Giacobazzi. 2005. Semantic-based code obfuscation by abstract interpretation. In *Proc. of ICALP*. 1325–1336.
- [15] M. Dalla Preda and I. Mastroeni. 2018. Characterizing a property-driven obfuscation strategy. *J. Comput. Secur.* 26, 1 (2018), 31–69.
- [16] D. Devriese and F. Piessens. 2010. Noninterference through secure multi-execution. In *2010 IEEE Symposium on Security and Privacy*. 109–124. DOI : <https://doi.org/10.1109/SP.2010.15>
- [17] Y. Falcone and H. Marchand. 2015. Enforcement and validation (at runtime) of various notions of opacity. *Discrete Event Dynamic Systems* 25, 4 (Dec. 2015), 531–570. DOI : <https://doi.org/10.1007/s10626-014-0196-4>
- [18] R. Giacobazzi. 2008. Hiding information in completeness holes—new perspectives in code obfuscation and watermarking. In *Proc. of The 6th IEEE International Conferences on Software Engineering and Formal Methods (SEFM'08)*. IEEE Press, 7–20.
- [19] R. Giacobazzi, N. D. Jones, and I. Mastroeni. 2012. Obfuscation by partial evaluation of distorted interpreters. In *Proc. of the ACM SIGPLAN Symp. on Partial Evaluation and Semantics-Based Program Manipulation (PEPM'12)*, O. Kiselyov and S. Thompson (Eds.). ACM Press, 63–72.
- [20] R. Giacobazzi and I. Mastroeni. 2010. Adjoining classified and unclassified information by abstract interpretation. *Journal of Computer Security* 18, 5 (2010), 751–797.
- [21] R. Giacobazzi and I. Mastroeni. 2018. Abstract non-interference: A unifying framework for weakening information-flow. *ACM Trans. Priv. Secur.* 21, 2 (2018), 1–31.
- [22] Roberto Giacobazzi, Isabella Mastroeni, and Mila Dalla Preda. 2017. Maximal incompleteness as obfuscation potency. *Formal Aspects Comput.* 29, 1 (2017), 3–31.
- [23] R. Giacobazzi, F. Ranzato, and F. Scozzari. 2000. Making abstract interpretation complete. *Journal of the ACM* 47, 2 (2000), 361–416.
- [24] Roberto Giacobazzi, Francesco Ranzato, and Francesca Scozzari. 2000. Making abstract interpretations complete. *J. ACM* 47, 2 (2000), 361–416. DOI : <https://doi.org/10.1145/333979.333989>
- [25] D. J. D. Hughes and V. Shmatikov. 2004. Information hiding, anonymity and privacy: A modular approach. *J. Comput. Secur.* 12, 1 (2004), 3–36.
- [26] S. Hunt and I. Mastroeni. 2005. The PER model of abstract non-interference. In *Proc. of SAS*. 171–185.
- [27] Jinwoo Kim, Shaan Nagy, Thomas Reps, and Loris D'Antoni. 2025. Semantics of sets of programs. *Proc. ACM Program. Lang.* 9, OOPSLA1, Article 110 (April 2025), 27 pages. DOI : <https://doi.org/10.1145/3720515>
- [28] Julian Klein, Paul Kogel, and Sabine Glesner. 2024. Verifying opacity of discrete-timed automata. In *Proceedings of the 2024 IEEE/ACM 12th International Conference on Formal Methods in Software Engineering (FormalISE) (FormalISE'24)*. Association for Computing Machinery, New York, NY, USA, 55–65. DOI : <https://doi.org/10.1145/3644033.3644376>
- [29] Siyuan Liu, Xiang Yin, Dimos V. Dimarogonas, and Majid Zamani. 2025. On approximate opacity of stochastic control systems. *IEEE Trans. Automat. Control* 70, 6 (2025), 3846–3861. DOI : <https://doi.org/10.1109/TAC.2024.3516202>
- [30] I. Mastroeni and A. Banerjee. 2011. Modelling declassification policies using abstract domain completeness. *Mathematical Structures in Computer Science* 21, 6 (2011), 1253–1299.
- [31] I. Mastroeni and D. Nikolic. 2010. An abstract unified framework for (abstract) program slicing. In *12th International Conference on Formal Engineering Methods, ICFEM 201 (Lecture Notes in Computer Science)*. Springer-Verlag, 452–467.
- [32] I. Mastroeni and M. Pasqua. 2017. Hyperhierarchy of semantics - A formal framework for hyperproperties verification. In *Proc. of SAS*. 232–252.
- [33] I. Mastroeni and M. Pasqua. 2018. Verifying bounded subset-closed hyperproperties. In *Proc. of SAS*. 1–20.
- [34] I. Mastroeni and M. Pasqua. 2019. Statically analyzing information flows: An abstract interpretation-based hyperanalysis for non-interference. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing, SAC 2019*. 2215–2223. DOI : <https://doi.org/10.1145/3297280.3297498>
- [35] I. Mastroeni and M. Pasqua. 2022. Verifying opacity by abstract interpretation. In *SAC'22: The 37th ACM/SIGAPP Symposium on Applied Computing, Virtual Event, April 25 - 29, 2022*, Jiman Hong, Miroslav Bures, Juw Won Park, and Tomáš Cerný (Eds.). ACM, 1817–1826. DOI : <https://doi.org/10.1145/3477314.3507119>
- [36] I. Mastroeni and M. Pasqua. 2023. Domain precision in galois connection-less abstract interpretation. In *Static Analysis*, Manuel V. Hermenegildo and José F. Morales (Eds.). Springer Nature Switzerland, Cham, 434–459.
- [37] I. Mastroeni and D. Zanardini. 2017. Abstract program slicing: An abstract interpretation-based approach to program slicing. *ACM Trans. Comput. Logic* 18, 1, Article 7 (feb 2017), 58 pages. DOI : <https://doi.org/10.1145/3029052>
- [38] F. Ranzato and F. Tapparo. 2004. Strong preservation as completeness in abstract interpretation. In *Proc. of ESOP*. 18–32.
- [39] P.Y.A. Ryan and T. Peacock. 2006. Opacity - Further Insights on an Information Flow Property. (2006). Technical Report CS-TR-958, University of Newcastle upon Tyne.

- [40] I. Saadaoui, A. Labed, Z. Li, A. M. El-Sherbeeney, and H. Du. 2023. Language-based opacity verification in partially observed petri nets through linear constraints. *Mathematics* 11, 18 (2023). DOI: <https://doi.org/10.3390/math11183880>
- [41] D. Schoepe and A. Sabelfeld. 2015. Understanding and enforcing opacity. In *Proc. of CSF*. 539–553.
- [42] S. Zdancewic and A. C. Myers. 2001. Robust declassification. In *Proceedings of the 14th IEEE Workshop on Computer Security Foundations (CSFW'01)*. IEEE Computer Society, USA, 5.

Received 22 April 2025; revised 26 September 2025; accepted 1 December 2025