



HAL
open science

Hypercollecting Semantics and its Application to Static Analysis of Information Flow

Mounir Assaf, David A Naumann, Julien Signoles, Eric Totel, Frédéric Tronel

► **To cite this version:**

Mounir Assaf, David A Naumann, Julien Signoles, Eric Totel, Frédéric Tronel. Hypercollecting Semantics and its Application to Static Analysis of Information Flow. POPL 2017 - ACM Symposium on Principles of Programming Languages, Jan 2017, Paris, France. pp.874-887, 10.1145/3009837.3009889 . hal-01618360

HAL Id: hal-01618360

<https://inria.hal.science/hal-01618360v1>

Submitted on 17 Oct 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Hypercollecting Semantics and its Application to Static Analysis of Information Flow

Mounir Assaf

Stevens Institute of Technology,
Hoboken, US
first.last@stevens.edu

David A. Naumann

Stevens Institute of Technology,
Hoboken, US
first.last@stevens.edu

Julien Signoles

Software Reliability and Security Lab,
CEA LIST, Saclay, FR
first.last@cea.fr

Éric Totel

CIDRE, CentraleSupélec,
Rennes, FR
first.last@centralesupelec.fr

Frédéric Tronel

CIDRE, CentraleSupélec,
Rennes, FR
first.last@centralesupelec.fr

Abstract

We show how static analysis for secure information flow can be expressed and proved correct entirely within the framework of abstract interpretation. The key idea is to define a Galois connection that directly approximates the hyperproperty of interest. To enable use of such Galois connections, we introduce a fixpoint characterisation of hypercollecting semantics, i.e. a “set of sets” transformer. This makes it possible to systematically derive static analyses for hyperproperties entirely within the calculational framework of abstract interpretation. We evaluate this technique by deriving example static analyses. For qualitative information flow, we derive a dependence analysis similar to the logic of Amtoft and Banerjee (SAS’04) and the type system of Hunt and Sands (POPL’06). For quantitative information flow, we derive a novel cardinality analysis that bounds the leakage conveyed by a program instead of simply deciding whether it exists. This encompasses problems that are hypersafety but not k -safety. We put the framework to use and introduce variations that achieve precision rivalling the most recent and precise static analyses for information flow.

Categories and Subject Descriptors D.2.4 [Software Engineering]: Software/Program Verification—Assertion checkers; D.3 [Programming Languages]; F.3.1 [Logics and meanings of programs]: Semantics of Programming Language

Keywords static analysis, abstract interpretation, information flow

1. Introduction

Most static analyses tell something about all executions of a program. This is needed, for example, to validate compiler optimizations. Functional correctness is also formulated in terms of a predicate on observable behaviours, i.e. more or less abstract execution traces: A program is correct if all its traces satisfy the predicate.

By contrast with such *trace properties*, extensional definitions of dependency involve more than one trace. To express that the final value of a variable x may depend only on the initial value of a variable y , the requirement—known as *noninterference* in the security literature (?)—is that any two traces with the same initial value for y result in the same final value for x . Sophisticated information flow policies allow dependency subject to quantitative bounds—and their formalisations involve more than two traces, sometimes unboundedly many.

For secure information flow formulated as decision problems, the theory of *hyperproperties* classifies the simplest form of noninterference as *2-safety* and some quantitative flow properties as *hypersafety properties* (?). A number of approaches have been explored for analysis of dependency, including type systems, program logics, and dependency graphs. Several works have used abstract interpretation in some way. One approach to 2-safety is by forming a *product program* that encodes execution pairs (???), thereby reducing the problem to ordinary safety which can be checked by abstract interpretation (?) or other means. Alternatively, a 2-safety property can be checked by dedicated analyses which may rely in part on ordinary abstract interpretations for trace properties (?).

The theory of abstract interpretation serves to specify and guide the design of static analyses. It is well known that effective application of the theory requires choosing an appropriate notion of observable behaviour for the property of interest (???). Once a notion of “trace” is chosen, one has a program semantics and “all executions” can be formalized in terms of *collecting semantics*, which can be used to define a trace property of interest, and thus to specify an abstract interpretation (???).

The foundation of abstract interpretation is quite general, based on Galois connections between semantic domains on which collecting semantics is defined. ? formalize the notion of hyperproperty in a very general way, as a set of sets of traces. Remarkably, prior works using abstract interpretation for secure information flow do not directly address the set-of-sets dimension and instead involve various ad hoc formulations. This paper presents a new approach of deriving information flow static analyses within the calculational framework of abstract interpretation.

First contribution. We lift collecting semantics to sets of trace sets, dubbed *hypercollecting semantics*, in a fixpoint formulation which is not simply the lifted direct image. This can be composed with Galois connections that specify hyperproperties beyond 2-

safety, without recourse to ad hoc additional notions. On the basis of this foundational advance, it becomes possible to derive static analyses entirely within the calculational framework of abstract interpretation (??).

Second contribution. We use hypercollecting semantics to derive an analysis for ordinary dependency. This can be seen as a rational reconstruction of both the type system of ?? and the logic of ?. They determine, for each variable x , a conservative approximation of the variables y whose initial values influence the final value of x .

Third contribution. We derive a novel analysis for quantitative information flow. This shows the benefit of taking hyperproperties seriously by means of abstract interpretation. For noninterference, once the variables y on which x depends have fixed values, there can be only one final value for x . For quantitative information flow, one is interested in measuring the extent to which other variables influence x : for a given range of variation for the “high inputs”, what is the range of variation for the final values of x ? We directly address this question as a hyperproperty: given a set of traces that agree only on the low inputs, what is the cardinality of the possible final values for x ? Using the hypercollecting semantics, we derive a novel *cardinality abstraction*. We show how it can be used for analysis of quantitative information problems including a bounding problem which is not k -safety for any k .

The calculational approach disentangles key design decisions and it enabled us to identify opportunities for improving precision. We assess the precision of our analyses and provide a formal characterisation of precision for a quantitative information flow analysis vis a vis qualitative. Versions of our analyses rival state of the art analyses for qualitative and quantitative information flow.

Our technical development uses the simplest programming language and semantic model in which the ideas can be exposed. One benefit of working entirely within the framework of abstract interpretation is that a wide range of semantics and analyses are already available for rich programming languages.

Outline. Following the background (??), we introduce domains and Galois connections for hyperproperties (??) and hypercollecting semantics (??). Hyperproperties for information flow are defined in ???. We use the framework to derive the static analyses in ?? and ??. ??. uses examples to evaluate the precision of the analyses, and shows how existing analyses can be leveraged to improve precision. We discuss related work (??) and conclude. *An accompanying technical report (?) contains detailed proofs for all results, as well as a table of symbols.*

2. Background: Collecting Semantics, Galois Connections

The formal development uses deterministic imperative programs over integer variables. Let n range over literal integers \mathbb{Z} , x over variables, and \oplus (resp. cmp) over some arithmetic (resp. comparison) operators.

$\mathbf{c} ::= \text{skip} \mid x := e \mid c_1; c_2 \mid \text{if } b \text{ then } c_1 \text{ else } c_2 \mid \text{while } b \text{ do } c$
 $\mathbf{e} ::= n \mid x \mid e_1 \oplus e_2 \mid b$
 $\mathbf{b} ::= e_1 \text{ cmp } e_2$

Different program analyses may consider different semantic domains as needed to express a given class of program properties. For imperative programs, the usual domains are based on states $\sigma \in \mathbf{States}$ that map each variable to a value (?). Some program properties require the use of traces that include intermediate states; others can use more abstract domains. For information flow properties involving intermediate outputs, or restricted to explicit data flow (?), details about intermediate steps are needed. By contrast, bounding the range of variables can be expressed in terms of final

$$\mathcal{P}(\mathbf{States}^*) \mathcal{P}(\mathbf{Trc}) \mathcal{P}(\mathbf{States}) \searrow$$

Figure 1. Fragment of the hierarchy of semantic domains ($\xrightarrow{\text{abstraction}}$)

states. As another example, consider determining which variables are left unchanged: To express this, we need both initial and final states.

In this paper we use the succinct term *trace* for elements of \mathbf{Trc} defined by $\mathbf{Trc} \triangleq \mathbf{States} \times \mathbf{States}$, interpreting $t \in \mathbf{Trc}$ as an initial and final state. In the literature, these are known as *relational traces*, by contrast with *maximal trace* semantics using the set \mathbf{States}^* of finite sequences. A uniform framework describes the relationships and correspondences between these and many other semantic domains using Galois connections (?). Three of these domains are depicted in ??.

Given partially ordered sets \mathcal{C}, \mathcal{A} , the monotone functions $\alpha \in \mathcal{C} \rightarrow \mathcal{A}$ and $\gamma \in \mathcal{A} \rightarrow \mathcal{C}$ comprise a *Galois connection*, a proposition we write $(\mathcal{C}, \leq) \xleftrightarrow[\alpha]{\gamma} (\mathcal{A}, \sqsubseteq)$, provided they satisfy $\alpha(c) \sqsubseteq a$ iff $c \leq \gamma(a)$ for all $c \in \mathcal{C}, a \in \mathcal{A}$.

For example, to specify an analysis that determines which variables are never changed, let \mathcal{A} be sets of variables. Define $\alpha \in \mathcal{P}(\mathbf{Trc}) \rightarrow \mathcal{P}(\mathbf{Vars})$ by $\alpha(T) = \{x \mid \forall(\sigma, \sigma') \in T, \sigma(x) = \sigma'(x)\}$ and $\gamma(X) = \{(\sigma, \sigma') \mid \forall x \in X, \sigma(x) = \sigma'(x)\}$. Then $(\mathcal{P}(\mathbf{Trc}), \sqsubseteq) \xleftrightarrow[\alpha]{\gamma} (\mathcal{P}(\mathbf{Var}), \supseteq)$.

For the hierarchy of usual domains, depicted in ??, the connections are defined by an “element-wise abstraction”. Define $\text{elt} \in \mathbf{States}^* \rightarrow \mathbf{Trc}$ by $\text{elt}(\sigma_0 \sigma_1 \dots \sigma_n) \triangleq (\sigma_0, \sigma_n)$. This lifts to an abstraction $\mathcal{P}(\mathbf{States}^*) \rightarrow \mathcal{P}(\mathbf{Trc})$.

Lemma 1 Element-wise abstraction. Let $\text{elt} \in \mathcal{C} \rightarrow \mathcal{A}$ be a function between sets. Let $\alpha_{\text{elt}}(C) \triangleq \{\text{elt}(c) \mid c \in C\}$ and $\gamma_{\text{elt}}(A) \triangleq \{c \mid \text{elt}(c) \in A\}$. Then $(\mathcal{P}(\mathcal{C}), \sqsubseteq) \xleftrightarrow[\alpha_{\text{elt}}]{\gamma_{\text{elt}}} (\mathcal{P}(\mathcal{A}), \sqsubseteq)$.

The domain $\mathcal{P}(\mathbf{States})$, which suffices to describe the final reachable states of a program, is an abstraction of the relational domain $\mathcal{P}(\mathbf{Trc})$, by $\text{elt}(\sigma, \tau) \triangleq \tau$. In this paper we focus on the domain \mathbf{Trc} because it is the simplest that can express dependency.

Program semantics. We define both the denotational semantics $\llbracket c \rrbracket \in \mathbf{Trc}_\perp \rightarrow \mathbf{Trc}_\perp$ of commands and the denotational semantics $\llbracket e \rrbracket \in \mathbf{Trc} \rightarrow \mathbf{Val}$ of expressions. Here $\mathbf{Val} \triangleq \mathbb{Z}$ and \mathbf{Trc}_\perp adds bottom element \perp using the flat ordering.

Standard semantics of commands		$\llbracket c \rrbracket \in \mathbf{Trc}_\perp \rightarrow \mathbf{Trc}_\perp$
$\llbracket c \rrbracket \perp \triangleq \perp$	$\llbracket x := e \rrbracket(\sigma, \tau) \triangleq (\sigma, \tau[x \mapsto \llbracket e \rrbracket(\sigma, \tau)])$	
$\llbracket c_1; c_2 \rrbracket t \triangleq \llbracket c_2 \rrbracket \circ \llbracket c_1 \rrbracket t$	$\llbracket \text{skip} \rrbracket t \triangleq t$	
$\llbracket \text{if } b \text{ then } c_1 \text{ else } c_2 \rrbracket t \triangleq \begin{cases} \llbracket c_1 \rrbracket t & \text{if } \llbracket b \rrbracket t = 1 \\ \llbracket c_2 \rrbracket t & \text{if } \llbracket b \rrbracket t = 0 \end{cases}$		
$\llbracket \text{while } b \text{ do } c \rrbracket t \triangleq (\text{lfp}_{(\lambda t. \perp)}^{\llbracket c \rrbracket} \mathcal{F})(t)$	where $\mathcal{F}(w)(t) \triangleq \begin{cases} t & \text{if } \llbracket b \rrbracket t = 0 \\ w \circ \llbracket c \rrbracket t & \text{otherwise} \end{cases}$	

Let t be a trace (σ, τ) . The denotation $\llbracket e \rrbracket t$ evaluates e in the “current state”, τ . (In ?? we also use $\llbracket e \rrbracket_{\text{pre}} t$ which evaluates e in the initial state, σ .) The denotation $\llbracket c \rrbracket t$ is (σ, τ') where execution of c in τ leads to τ' . The denotation is \perp in case c diverges from τ . Boolean

expressions evaluate to either 0 or 1. We assume programs do not go wrong. We denote by \preceq the point-wise lifting to $\mathbf{Trc}_\perp \rightarrow \mathbf{Trc}_\perp$ of the approximation order \preceq on \mathbf{Trc}_\perp .

The terminating computations of c can be written as its image on the initial traces: $\{\llbracket c \rrbracket t \mid t \in \mathbf{IniTrc} \text{ and } \llbracket c \rrbracket t \neq \perp\}$ where

$$\mathbf{IniTrc} \triangleq \{(\sigma, \sigma) \mid \sigma \in \mathbf{States}\}$$

To specify properties that hold for all executions we use *collecting semantics* which lifts the denotational semantics to arbitrary sets $T \in \mathcal{P}(\mathbf{Trc})$ of traces. The idea is that $\llbracket c \rrbracket T$ is the direct image of $\llbracket c \rrbracket$ on T . To be precise, in this paper we focus on termination-insensitive properties, and thus $\llbracket c \rrbracket T$ is the set of non- \perp traces t' such that $\llbracket c \rrbracket t = t'$ for some $t \in T$. Later we also use the collecting semantics of expressions: $\llbracket e \rrbracket T \triangleq \{\llbracket e \rrbracket t \mid t \in T\}$.

Importantly, the collecting semantics $\llbracket c \rrbracket \in \mathcal{P}(\mathbf{Trc}) \rightarrow \mathcal{P}(\mathbf{Trc})$ can be defined compositionally using fixpoints (2, Sec. 7). For conditional guard b , write $\llbracket \text{grd}^b \rrbracket$ for the filter defined by $\llbracket \text{grd}^b \rrbracket T \triangleq \{t \in T \mid \llbracket b \rrbracket t = 1\}$.

Collecting semantics	$\llbracket c \rrbracket \in \mathcal{P}(\mathbf{Trc}) \rightarrow \mathcal{P}(\mathbf{Trc})$
$\llbracket x := e \rrbracket T \triangleq \{\llbracket x := e \rrbracket t \mid t \in T\}$	
$\llbracket c_1; c_2 \rrbracket T \triangleq \llbracket c_2 \rrbracket \circ \llbracket c_1 \rrbracket T$	$\llbracket \text{skip} \rrbracket T \triangleq T$
$\llbracket \text{if } b \text{ then } c_1 \text{ else } c_2 \rrbracket T \triangleq \llbracket c_1 \rrbracket \circ \llbracket \text{grd}^b \rrbracket T \cup \llbracket c_2 \rrbracket \circ \llbracket \text{grd}^{\neg b} \rrbracket T$	
$\llbracket \text{while } b \text{ do } c \rrbracket T \triangleq \llbracket \text{grd}^{\neg b} \rrbracket \left(\text{lfp}_T^{\subseteq} \llbracket \text{if } b \text{ then } c \text{ else skip} \rrbracket \right)$	

The clause for while loops uses the denotation of a constructed conditional command as a definitional shorthand—its denotation is compositional.

Given a Galois connection $(\mathcal{P}(\mathbf{Trc}), \subseteq) \xleftrightarrow[\alpha]{\gamma} (\mathcal{A}, \sqsubseteq)$, such as the one for unmodified variables, the desired analysis is specified as $\alpha \circ \llbracket c \rrbracket \circ \gamma$. Since it is not computable in general, we only require an approximation $f^\sharp \in \mathcal{A} \rightarrow \mathcal{A}$ that is *sound* in this sense:

$$\alpha \circ \llbracket c \rrbracket \circ \gamma \sqsubseteq f^\sharp \quad (1)$$

where \sqsubseteq denotes the point-wise lifting of the partial order \sqsubseteq .

To explain the significance of this specification, suppose one wishes to prove program c satisfies a trace property $T \in \mathcal{P}(\mathbf{Trc})$, i.e. to prove that $\llbracket c \rrbracket(\mathbf{IniTrc}) \subseteq T$. Given ?? it suffices to find an abstract value a that approximates \mathbf{IniTrc} , i.e. $\mathbf{IniTrc} \subseteq \gamma(a)$, and show that

$$\gamma(f^\sharp(a)) \subseteq T \quad (2)$$

?? is equivalent to $\llbracket c \rrbracket \circ \gamma \sqsubseteq \gamma \circ f^\sharp$ by a property of Galois connections. So ?? implies $\llbracket c \rrbracket(\gamma(a)) \subseteq T$ which (by monotonicity of $\llbracket c \rrbracket$) implies $\llbracket c \rrbracket(\mathbf{IniTrc}) \subseteq \llbracket c \rrbracket(\gamma(a)) \subseteq T$.

The beauty of specification ?? is that f^\sharp can be obtained as an abstract interpretation $\llbracket c \rrbracket^\sharp$, derived systematically for all c by calculating from the left side of ?? as shown by ?.

3. Domains and Galois Connections for Hyperproperties

To express hyperproperties, we need Galois connections for domains that involve sets of sets of observable behaviours. This section spells out how such powerset domains form a hierarchy as illustrated along the top of ?. We describe how dependency and quantitative information flow can be formulated as Galois connections. We spell out a methodology whereby the standard notions and techniques of abstract interpretation can be applied to specify and derive—in the same form as ??—static analyses for hyperproperties.




Figure 2. Extended hierarchy of semantic domains ($\xrightarrow{\text{abstraction}}$)

As a first example, consider the condition: the final value of x depends only on the initial value of y . Its expression needs, at least, two traces: If two traces, denoted by (σ, σ') and (τ, τ') , agree on the initial value of y then they agree on the final value of x . That is, $\sigma(y) = \tau(y)$ implies $\sigma'(x) = \tau'(x)$. This must hold for any two traces of the program. This is equivalent to the following: For all sets T of traces, if traces in T all agree on the initial value of y then they all agree on the final value of x . Later we extend this example to an analysis that infers which dependences hold.

Consider the problem of quantifying dependence with min-capacity (?). For a program on two integer variables h, l , the problem is to infer how much information is conveyed via l about h : considering some traces that agree on the initial value of l , how many final values are possible for l . For example, the program $l := (h \text{ mod } 2) + l$ has two final values for l , for each initial l , though there are many possible initial values for h . This cardinality problem generalizes prior work on quantitative flow analysis, where typically low inputs are not considered.

Whereas the simple dependency problem can be formulated in terms of 2-element sets, the cardinality problem involves sets of unbounded size. In the terminology of hyperproperties, it is not a k -safety hyperproperty for any k (? , Sec. 3), although it is hyper-safety (?). For a fixed k , the problem “variable l has at most $k - 1$ final values” is k -safety, which means it can be formulated in terms of sets with at most k traces.

It turns out that by using Galois connections on sets of sets, we can develop a general theory that encompasses many hyperproperties and which enables derivation of interesting abstract interpreters. For our applications, we use relational traces as the notion of observable behavior, and thus $\mathcal{P}(\mathcal{P}(\mathbf{Trc}))$. The approach works as well for other notions, so there is a hierarchy of domains as shown at the top of ??, in parallel with the ordinary hierarchy shown along the bottom.

The abstractions of this hierarchy are obtained by lifting each abstraction between two standard collecting semantics (?) to their hypercollecting versions, by element-wise abstraction (??). For instance, ?? justifies the abstraction between $\mathcal{P}(\mathcal{P}(\mathbf{Trc}))$ and $\mathcal{P}(\mathcal{P}(\mathbf{States}))$, by lifting the abstraction between $\mathcal{P}(\mathbf{Trc})$ and $\mathcal{P}(\mathbf{States})$ (? , Sec. 8). Additionally, the diagonal lines in ?? represent abstractions between hypercollecting semantics defined over some form of observations and the corresponding collecting semantics defined over the same observations.

Lemma 2 . Let \mathcal{C} be a set. Define $\alpha_{\text{hpp}}(\mathcal{C}) \triangleq \bigcup_{C \in \mathcal{C}} C$ and $\gamma_{\text{hpp}}(\mathcal{C}) \triangleq \mathcal{P}(\mathcal{C})$. These form a Galois connection:

$$(\mathcal{P}(\mathcal{P}(\mathcal{C})), \subseteq) \xleftrightarrow[\alpha_{\text{hpp}}]{\gamma_{\text{hpp}}} (\mathcal{P}(\mathcal{C}), \subseteq)$$

It is noted by ? that any trace property can be lifted to a unique hyperproperty; this lifting is exactly the concretisation γ_{hpp} of ?? . Although the model of ? is quite general, it does focus on infinite traces. But hyperproperties can be formulated in terms of other notions of observation, as illustrated in ?? .

Cardinality abstraction. To lay the groundwork for our quantitative information flow analysis, we consider abstracting a set of values by its cardinality. Cardinality is one ingredient in many quantitative information flow analyses estimating the amount of sensitive

information a program may leak (?????). The lattice of abstract representations we consider is the set

$$[0, \infty] \triangleq \mathbb{N} \cup \{\infty\}$$

where ∞ denotes an infinite cardinal number. We use the natural order \leq , and \max as a join. Consider the abstraction operator $\text{crdval} \in \mathcal{P}(\mathbf{Val}) \rightarrow [0, \infty]$ computing cardinality and given by $\text{crdval}(V) \triangleq |V|$. This operator crdval is not *additive*, i.e. it does not preserve joins; e.g. $\text{crdval}(\{1, 2\} \cup \{2, 3\}) \neq \max(\text{crdval}(\{1, 2\}), \text{crdval}(\{2, 3\}))$. Thus, there exists no associated concretisation f for which crdval is the lower adjoint in a Galois connection. Yet, we can lift the abstraction operator crdval to a Galois connection over $\mathcal{P}(\mathcal{P}(\mathbf{Val}))$ through what is called a supremus abstraction (?, p.52).

Lemma 3 Supremus abstraction. Let $\text{elt} \in \mathcal{C} \rightarrow \mathcal{A}$ be a function from a set \mathcal{C} , with codomain forming a complete lattice $(\mathcal{A}, \sqsubseteq)$. Let $\alpha_{\text{elt}}(C) \triangleq \sqcup_{c \in C} \text{elt}(c)$ and $\gamma_{\text{elt}}(a) \triangleq \{c \in \mathcal{C} \mid \text{elt}(c) \sqsubseteq a\}$. Then

$$(\mathcal{P}(\mathcal{C}), \sqsubseteq) \xleftrightarrow[\alpha_{\text{elt}}]{\gamma_{\text{elt}}} (\mathcal{A}, \sqsubseteq)$$

For example, define $\alpha_{\text{crdval}}(\mathbb{V}) \triangleq \max_{V \in \mathbb{V}} \text{crdval}(V)$ and $\gamma_{\text{crdval}}(n) \triangleq \{V \mid \text{crdval}(V) \leq n\}$. Thus we obtain a Galois connection $(\mathcal{P}(\mathcal{P}(\mathbf{Val})), \sqsubseteq) \xleftrightarrow[\alpha_{\text{crdval}}]{\gamma_{\text{crdval}}} ([0, \infty], \leq)$.

As another example let us consider, in simplified form, an ingredient in dependency or noninterference analysis. For program variable x , $\text{agree}_x \in \mathcal{P}(\mathbf{States}) \rightarrow \{\text{tt}, \text{ff}\}$ determines whether a set of states contains only states that all agree on x 's value:

$$\text{agree}_x(\Sigma) \triangleq (\forall \sigma, \sigma' \in \Sigma, \llbracket x \rrbracket \sigma = \llbracket x \rrbracket \sigma')$$

Function agree_x is not additive, so it is not part of a Galois connection from $\mathcal{P}(\mathbf{States})$ to $\{\text{tt}, \text{ff}\}$. The same problem arises with agreements on multiple variables, and with more concrete domains like the finite maximal trace semantics $\mathcal{P}(\mathbf{States}^*)$.

We lift the operator agree_x to a Galois connection over $\mathcal{P}(\mathcal{P}(\mathbf{States}))$. A supremus abstraction yields

$$\begin{aligned} \alpha_{\text{agree}_x}(\mathbb{S}) &\triangleq (\forall \Sigma \in \mathbb{S}, \text{agree}_x(\Sigma)) \\ \gamma_{\text{agree}_x}(\text{bv}) &\triangleq \{\Sigma \mid \text{agree}_x(\Sigma) \leftarrow \text{bv}\} \end{aligned}$$

so that $(\mathcal{P}(\mathcal{P}(\mathbf{States})), \sqsubseteq) \xleftrightarrow[\alpha_{\text{agree}_x}]{\gamma_{\text{agree}_x}} (\{\text{tt}, \text{ff}\}, \leftarrow)$.

These examples are consistent with the many formulations of noninterference (e.g. (?????)) that motivated the characterisation of information-flow security requirements as hyperproperties (?). Concretising an abstract value a can be seen as defining the denotation of a type expression (as in, for instance, ?, Sec. 3.3.1 and ?), i.e. defining the set of objects that satisfy the description a . Thus, concretising tt , when tt is interpreted as “satisfies a property requirement”, naturally yields a set of traces. Concretising tt , where tt is interpreted as “satisfies a security requirement”, yields a set of sets of traces.

Intuitively, the most abstract denotation/concretisation of a property requirement is defined in terms of a set of traces. The most abstract concretisation/denotation of a security requirement yields a set of sets of traces, namely a hyperproperty. Hints of this intuition appear in the literature (?????); e.g. security policies “are predicates on sets of traces (i.e. they are higher order)” (?, p.2). However, only recently has a comprehensive framework proposed a sharp characterisation of security policies as hyperproperties (??).

Abstract interpretation of hyperproperties. The basic methodology for the verification of a hyperproperty HP, may be described as follows:

Step 1. Design approximate representations forming a complete lattice \mathcal{A} , choose a collecting semantics \mathcal{C} among the extended hierarchy (set of sets domains, e.g. $\mathcal{P}(\mathcal{P}(\mathbf{Trc}))$), and define α, γ for a Galois connection $(\mathcal{C}, \leq) \xleftrightarrow[\alpha]{\gamma} (\mathcal{A}, \sqsubseteq)$.

Step 2. Compute an approximation $a \in \mathcal{A}$ of the semantics $C \in \mathcal{C}$ of the program P of interest.

Step 3. Prove that the inferred approximation a implies that P satisfies HP. The concretisation $\gamma(a)$ is a set of trace sets, of which the program’s trace set is a *member*—by contrast to approximations of trace properties, which infer a single trace set of which the program trace set is a *subset*. Then, it suffices to prove $\gamma(a) \subseteq \text{HP}$.

Step 1 is guided by the need to have $\gamma(a) \subseteq \text{HP}$, i.e. a describes a hyperproperty that implies HP. The calculational design (?) of abstract domains greatly systematises Step 2, by relying on the Galois connection defined in Step 1. Collecting semantics can be adapted to the additional structure of sets, as we show in ??.

4. Hypercollecting Semantics

In the following, we introduce a hypercollecting semantics defined over sets $\mathbb{T} \in \mathcal{P}(\mathcal{P}(\mathbf{Trc}))$ of sets of traces. This is used in subsequent sections to derive static analyses.

Here is Step 2 of the methodology, spelled out in detail. Given a Galois connection $(\mathcal{P}(\mathcal{P}(\mathbf{Trc})), \sqsubseteq) \xleftrightarrow[\alpha]{\gamma} (\mathcal{A}, \sqsubseteq^{\#})$ built by the supremus abstraction, and an approximation a of the initial traces (i.e. **IniTrc** is in $\gamma(a)$), find an approximation $a' \in \mathcal{A}$ of the analysed program c , i.e. $\llbracket c \rrbracket \mathbf{IniTrc}$ is in $\gamma(a')$. Then prove that the program satisfies the hyperproperty HP of interest, i.e. $\gamma(a') \subseteq \text{HP}$. In order to compute a' , we define a hypercollecting semantics $\llbracket c \rrbracket \in \mathcal{P}(\mathcal{P}(\mathbf{Trc})) \rightarrow \mathcal{P}(\mathcal{P}(\mathbf{Trc}))$. That will serve to derive—in the manner of ??—a static analysis that is correct by construction.

Hypercollecting semantics $\llbracket c \rrbracket \in \mathcal{P}(\mathcal{P}(\mathbf{Trc})) \rightarrow \mathcal{P}(\mathcal{P}(\mathbf{Trc}))$

$$\llbracket x := e \rrbracket \mathbb{T} \triangleq \{\llbracket x := e \rrbracket T \mid T \in \mathbb{T}\}$$

$$\llbracket c_1 ; c_2 \rrbracket \mathbb{T} \triangleq \llbracket c_2 \rrbracket \circ \llbracket c_1 \rrbracket \mathbb{T} \quad \llbracket \text{skip} \rrbracket \mathbb{T} \triangleq \mathbb{T}$$

$$\begin{aligned} \llbracket \text{if } b \text{ then } c_1 \text{ else } c_2 \rrbracket \mathbb{T} &\triangleq \\ &\{\llbracket c_1 \rrbracket \circ \llbracket \text{grd}^b \rrbracket T \cup \llbracket c_2 \rrbracket \circ \llbracket \text{grd}^{-b} \rrbracket T \mid T \in \mathbb{T}\} \end{aligned}$$

$$\llbracket \text{while } b \text{ do } c \rrbracket \mathbb{T} \triangleq \llbracket \text{grd}^{-b} \rrbracket \left(\text{lfp}_{\mathbb{T}}^{\subseteq} (\llbracket \text{if } b \text{ then } c \text{ else skip} \rrbracket) \right)$$

$$\llbracket \text{grd}^b \rrbracket \mathbb{T} \triangleq \{\llbracket \text{grd}^b \rrbracket T \mid T \in \mathbb{T}\}$$

Recall from ?? that standard collecting semantics is a fixpoint-based formulation that captures the direct image on sets of the underlying program semantics – this is proved, for example, by ??. The fixpoint formulation at the level of sets-of-sets we use is not simply the direct image of the standard collecting semantics. The direct image of the standard collecting semantics would yield a set of (inner) fixpoints over sets of traces, whereas an outer fixpoint over sets of sets of traces enables straightforward application of the fixpoint transfer theorem.

Theorem 1. For all c and all $T \in \mathcal{P}(\mathbf{Trc})$, $\llbracket c \rrbracket T$ is in $\llbracket c \rrbracket \{T\}$.

For a singleton $\{T\}$, the set $\llbracket c \rrbracket \{T\} \in \mathcal{P}(\mathcal{P}(\mathbf{Trc}))$ is not necessarily a singleton set containing only the element $\llbracket c \rrbracket T$. If c is a loop, $\llbracket c \rrbracket \{T\}$ yields a set of sets R of traces, where each set R of traces contains only traces that exit the loop after less than k iterations, for $k \in \mathbb{N}$. We prove this theorem as corollary of the following:

$$\forall \mathbb{T} \in \mathcal{P}(\mathcal{P}(\mathbf{Trc})), \{\llbracket c \rrbracket T \mid T \in \mathbb{T}\} \subseteq \llbracket c \rrbracket \mathbb{T}$$

This is proved by structural induction on commands. For loops, there is a secondary induction on iterations of the loop body.

In summary, suppose one wishes to prove program c satisfies hyperproperty $\text{HP} \in \mathcal{P}(\mathcal{P}(\mathbf{Trc}))$, i.e. one wishes to prove that $\llbracket c \rrbracket(\mathbf{IniTrc}) \in \text{HP}$. Suppose we have an approximation $f^\#$ of the hypercollecting semantics, similarly to ??, i.e.

$$\alpha \circ \llbracket c \rrbracket \circ \gamma \sqsubseteq^\# f^\# \quad (3)$$

Given ?? it suffices to find an abstract value a that approximates \mathbf{IniTrc} , i.e. $\mathbf{IniTrc} \in \gamma(a)$, and show that:

$$\gamma(f^\#(a)) \subseteq \text{HP} \quad (4)$$

Why? ?? is equivalent to $\llbracket c \rrbracket \circ \gamma \sqsubseteq^\# \gamma \circ f^\#$ by a property of Galois connections. So we have $\llbracket c \rrbracket(\mathbf{IniTrc}) \in \llbracket c \rrbracket(\gamma(a)) \subseteq \gamma(f^\#(a)) \subseteq \text{HP}$ using $\mathbf{IniTrc} \in \gamma(a)$, the Theorem, and ??.

5. Information Flow

This section gives a number of technical definitions which build up to the definition of Galois connections with which we specify information flow policies explicitly as hyperproperties.

When a fixed main program is considered, we refer to it as P and its variables as Var_P . Our analyses are parametrised by the program P to analyse, and an initial **typing context** $\Gamma \in \text{Var}_P \rightarrow \mathcal{L}$ mapping each variable to a security level $l \in \mathcal{L}$ for its initial value. We assume $(\mathcal{L}, \sqsubseteq, \sqcup, \sqcap)$ is a finite lattice. In the most concrete case, \mathcal{L} may be defined as the **universal flow lattice**, i.e. the powerset of variables $\mathcal{P}(\text{Var}_P)$, from which all other information flow types can be inferred through a suitable abstraction (?, Sec. 6.2); the initial typing context is then defined as $\lambda x. \{x\}$.

Initial l -equivalence and variety. A key notion in information flow is **l -equivalence**. Two states are l -equivalent iff they agree on the values of variables having security level at most l . We introduce the same notion over a set of traces, requiring that the **initial states** are l -equivalent. Let us first denote by $\llbracket e \rrbracket_{\text{pre}} \in \mathbf{Trc} \rightarrow \mathbf{Val}$ the evaluation of expression e in the initial state σ of a trace $(\sigma, \tau) \in \mathbf{Trc}$ —unlike $\llbracket e \rrbracket \in \mathbf{Trc} \rightarrow \mathbf{Val}$ which evaluates expression e in the final state τ . Then, we denote by $T \models_\Gamma l$ the judgement that all traces in a set $T \subseteq \mathbf{Trc}$ are **initially l -equivalent**, i.e. they all initially agree on the value of variables up to a security level $l \in \mathcal{L}$.

For example, in the case that \mathcal{L} is the universal flow lattice, $T \models_\Gamma \{x, y\}$ means $\forall t_1, t_2 \in T, \llbracket x \rrbracket_{\text{pre}} t_1 = \llbracket x \rrbracket_{\text{pre}} t_2 \wedge \llbracket y \rrbracket_{\text{pre}} t_1 = \llbracket y \rrbracket_{\text{pre}} t_2$.

Initial l-equivalence	$T \models_\Gamma l$
$T \models_\Gamma l$ iff. $\forall t_1, t_2 \in T, \forall x \in \text{Var}_P,$	
$\Gamma(x) \sqsubseteq l \implies \llbracket x \rrbracket_{\text{pre}} t_1 = \llbracket x \rrbracket_{\text{pre}} t_2$	

The notion of variety (?) underlies most definitions of qualitative and quantitative information flow security. Information is transmitted from a to b over execution of program P if by “varying the initial value of a (exploring the variety in a), the resulting value in b after P ’s execution will also vary (showing that variety is conveyed to b)” (?). We define the **l -variety** of expression e , as the set of sets of values e may take, when considering only initially l -equivalent traces. The variety is defined first as a function $\mathcal{O}^l \llbracket e \rrbracket \in \mathcal{P}(\mathbf{Trc}) \rightarrow \mathcal{P}(\mathcal{P}(\mathbf{Val}))$ on trace sets, from which we obtain a function $\mathcal{O}^l \langle e \rangle \in \mathcal{P}(\mathcal{P}(\mathbf{Trc})) \rightarrow \mathcal{P}(\mathcal{P}(\mathbf{Val}))$, on sets of trace sets. Intuitively, l -variety of expression e is the variety that is conveyed to e by varying only the input values of variables having a security level l' such that $\neg(l' \sqsubseteq l)$.

l-variety	$\mathcal{O}^l \llbracket e \rrbracket$	$\mathcal{O}^l \langle e \rangle$
-------------------------------	---	-----------------------------------

$$\begin{aligned} \mathcal{O}^l \llbracket e \rrbracket &\in \mathcal{P}(\mathbf{Trc}) \rightarrow \mathcal{P}(\mathcal{P}(\mathbf{Val})) \\ \mathcal{O}^l \llbracket e \rrbracket T &\triangleq \{ \llbracket e \rrbracket R \mid R \subseteq T \text{ and } R \models_\Gamma l \} \\ \mathcal{O}^l \langle e \rangle &\in \mathcal{P}(\mathcal{P}(\mathbf{Trc})) \rightarrow \mathcal{P}(\mathcal{P}(\mathbf{Val})) \\ \mathcal{O}^l \langle e \rangle \mathbb{T} &\triangleq \bigcup_{T \in \mathbb{T}} \mathcal{O}^l \llbracket e \rrbracket T \end{aligned}$$

Each set $V \in \mathcal{O}^l \llbracket e \rrbracket T$ of values results from initially l -equivalent traces ($R \models_\Gamma l$ for $R \subseteq T$). Thus, expression e does not leak sensitive information to attackers having a security clearance $l \in \mathcal{L}$ if $\mathcal{O}^l \llbracket e \rrbracket T$ is a set of singleton sets. Indeed, sensitive data for attackers with security clearance $l \in \mathcal{L}$ is all data having a security level l' for which attackers do not have access (i.e. $\neg(l' \sqsubseteq l)$ (?)). Thus, if $\mathcal{O}^l \llbracket e \rrbracket T$ is a set of singleton sets, this means that no matter how sensitive information varies, this variety is not conveyed to expression e .

Besides a pedagogical purpose, we define l -variety $\mathcal{O}^l \llbracket e \rrbracket$ (resp. $\mathcal{O}^l \langle e \rangle$) instead of simply lifting the denotational semantics $\llbracket e \rrbracket$ of expressions to sets of traces (resp. sets of sets of traces) since we want to build modular abstractions of traces by relying on underlying abstractions of values. Thus, l -variety enables us to pass information about initially l -equivalent traces to the underlying domain of values by keeping disjoint values that originate from traces that are not initially l -equivalent.

Specifying information flow. We now have the ingredients needed to describe information flow for command c , with respect to typing context $\Gamma \in \text{Var}_P \rightarrow \mathcal{L}$. A quantitative security metric, introduced by ??, relies on min-entropy and min-capacity (?) in order to estimate the leakage of a program. Let us assume a program P that is characterized by a set $T_P \in \mathcal{P}(\mathbf{Trc})$ of traces, i.e. $T_P \triangleq \llbracket P \rrbracket \mathbf{IniTrc}$. For simplicity, assume attackers only observe the value of a single variable $x \in \text{Var}_P$. (The generalization to multiple variables is straightforward.) The leakage of P , as measured by **min-capacity**, to attackers having security clearance $l \in \mathcal{L}$ is defined by

$$\mathcal{ML}_l \triangleq \log_2 \circ \alpha_{\text{crdval}} \circ \mathcal{O}^l \llbracket x \rrbracket T_P$$

(The definition of α_{crdval} follows ??.) For our purposes, it suffices to know that this quantity aims to measure, in bits, the remaining uncertainty about sensitive data for attackers with security clearance l . Refer to the original work (?) for more details.

Leaving aside the logarithm in the definition of \mathcal{ML}_l , a quantitative security requirement may enforce a limit on the amount of information leaked to attackers with security clearance $l \in \mathcal{L}$, by requiring that the l -cardinality of variable x is less than or equal to some non-negative integer k . We denote by $\text{SR}(l, k, x)$ the hyperproperty that characterises this security requirement, i.e. the set of program denotations satisfying it:

$$\text{SR}(l, k, x) \triangleq \{ T \in \mathcal{P}(\mathbf{Trc}) \mid \alpha_{\text{crdval}} \circ \mathcal{O}^l \llbracket x \rrbracket T \leq k \}$$

Note that SR implicitly depends on the choice of initial typing Γ , as does $\mathcal{O}^l \llbracket x \rrbracket T$.

The termination-insensitive noninterference policy “the final value of x depends only on the initial values of variables labelled at most l ” corresponds to the hyperproperty $\text{SR}(l, 1, x)$. Therefore, the program P satisfies $\text{SR}(l, 1, x)$ if $\alpha_{\text{crdval}} \circ \mathcal{O}^l \llbracket x \rrbracket T_P \leq 1$. Let $\mathbb{T} = \langle P \rangle \{ \mathbf{IniTrc} \}$. Since T_P is in \mathbb{T} (??), then P satisfies $\text{SR}(l, 1, x)$ if $\alpha_{\text{crdval}} \circ \mathcal{O}^l \langle x \rangle \mathbb{T} \leq 1$, by monotony of α_{crdval} and by $\mathcal{O}^l \llbracket x \rrbracket T_P \subseteq \mathcal{O}^l \langle x \rangle \mathbb{T}$ from the definition of $\mathcal{O}^l \langle - \rangle$.

6. Dependences

We rely on abstract interpretation to derive a static analysis similar to existing ones inferring dependences (????).

Recall that our analyses are parametrised on a security lattice \mathcal{L} and program P . We denote by $l \rightsquigarrow x$ an atomic dependence constraint, with $l \in \mathcal{L}$ and $x \in \text{Var}_P$, read as “agreement up to security level l leads to agreement on x ”. It is an atomic pre-post contract expressing that the final value of x must only depend on initial values having at most security level l . Said otherwise, $l \rightsquigarrow x$ states the noninterference of variable x from data that is sensitive for attackers with security clearance l , i.e. all inputs having security level l' such that $\neg(l' \sqsubseteq l)$.

Dependencies are similar to information flow types (?) and are the dual of independences assertions (?). Both interpretations are equivalent (? , Sec. 5).

Lattice of dependence constraints Dep $\mathcal{D} \in \text{Dep}$

Given a lattice \mathcal{L} and program P , define

$$\begin{aligned} \text{Dep} &\triangleq \mathcal{P}(\{l \rightsquigarrow x \mid l \in \mathcal{L}, x \in \text{Var}_P\}) \\ \mathcal{D}_1 \sqsubseteq^{\natural} \mathcal{D}_2 &\triangleq \mathcal{D}_1 \supseteq \mathcal{D}_2 \quad \mathcal{D}_1 \sqcup^{\natural} \mathcal{D}_2 \triangleq \mathcal{D}_1 \cap \mathcal{D}_2 \end{aligned}$$

In the rest of this section, \mathcal{L} and P are fixed, together with a typing context $\Gamma \in \text{Var}_P \rightarrow \mathcal{L}$.

The semantic characterisation of dependences is tightly linked to variety. An atomic constraint $l \rightsquigarrow x$ holds if no variety is conveyed to x when the inputs up to security level l are fixed. We use this intuition to define the Galois connections linking the hypercollecting semantics and the lattice Dep , by instantiating the supremus abstraction in ??.

The agreements abstraction approximates a set $\mathbb{V} \in \mathcal{P}(\mathcal{P}(\mathbf{Val}))$ by determining whether it contains variety.

Agreements abstraction agree $\alpha_{\text{agree}} \gamma_{\text{agree}}$

$$\begin{aligned} \text{agree} &\in \mathcal{P}(\mathbf{Val}) \rightarrow \{\text{tt}, \text{ff}\} \\ \text{agree}(V) &\triangleq (\forall v_1, v_2 \in V, v_1 = v_2) \\ \alpha_{\text{agree}} &\in \mathcal{P}(\mathcal{P}(\mathbf{Val})) \rightarrow \{\text{tt}, \text{ff}\} \\ \alpha_{\text{agree}}(\mathbb{V}) &\triangleq \bigwedge_{V \in \mathbb{V}} \text{agree}(V) \\ \gamma_{\text{agree}} &\in \{\text{tt}, \text{ff}\} \rightarrow \mathcal{P}(\mathcal{P}(\mathbf{Val})) \\ \gamma_{\text{agree}}(\text{bv}) &\triangleq \{V \in \mathcal{P}(\mathbf{Val}) \mid \text{agree}(V) \leftarrow \text{bv}\} \end{aligned}$$

$$(\mathcal{P}(\mathcal{P}(\mathbf{Val})), \sqsubseteq) \xleftrightarrow[\alpha_{\text{agree}}]{\gamma_{\text{agree}}} (\{\text{tt}, \text{ff}\}, \leftarrow)$$

Note that $\gamma_{\text{agree}}(\text{tt})$ is $\{V \in \mathcal{P}(\mathbf{Val}) \mid \text{agree}(V)\}$ and $\gamma_{\text{agree}}(\text{ff})$ is $\mathcal{P}(\mathbf{Val})$. Also, $\text{agree}(V)$ iff $|V| \leq 1$.

The dependences abstraction approximates a set $\mathbb{T} \in \mathcal{P}(\mathcal{P}(\mathbf{Trc}))$ by a dependence constraint $\mathcal{D} \in \text{Dep}$. Recall that $\mathcal{O}^l \llbracket x \rrbracket T$ is the set of final values for variable x in traces $t \in T$ that agree on inputs of level at most l . So $\alpha_{\text{agree}}(\mathcal{O}^l \llbracket x \rrbracket T)$ holds just if there is at most one final value.

Dependencies abstraction deptr $\alpha_{\text{deptr}} \gamma_{\text{deptr}}$

$$\begin{aligned} \text{deptr} &\in \mathcal{P}(\mathbf{Trc}) \rightarrow \text{Dep} \\ \text{deptr}(T) &\triangleq \{l \rightsquigarrow x \mid l \in \mathcal{L}, x \in \text{Var}_P, \alpha_{\text{agree}}(\mathcal{O}^l \llbracket x \rrbracket T)\} \\ \alpha_{\text{deptr}} &\in \mathcal{P}(\mathcal{P}(\mathbf{Trc})) \rightarrow \text{Dep} \\ \alpha_{\text{deptr}}(\mathbb{T}) &\triangleq \sqcup^{\natural}_{T \in \mathbb{T}} \text{deptr}(T) \\ \gamma_{\text{deptr}} &\in \text{Dep} \rightarrow \mathcal{P}(\mathcal{P}(\mathbf{Trc})) \\ \gamma_{\text{deptr}}(\mathcal{D}) &\triangleq \{T \mid \text{deptr}(T) \sqsubseteq^{\natural} \mathcal{D}\} \end{aligned}$$

$$(\mathcal{P}(\mathcal{P}(\mathbf{Trc})), \sqsubseteq) \xleftrightarrow[\alpha_{\text{deptr}}]{\gamma_{\text{deptr}}} (\text{Dep}, \sqsubseteq^{\natural})$$

Note that $\text{deptr}(T)$ is the set of dependences $l \rightsquigarrow x$ for which $\alpha_{\text{agree}}(\mathcal{O}^l \llbracket x \rrbracket T)$ holds. For instance, the initial typing context $\Gamma \in \text{Var}_P \rightarrow \mathcal{L}$ determines the initial dependences of a program:

$$\begin{aligned} \alpha_{\text{deptr}}(\{\mathbf{IniTrc}\}) &= \{l \rightsquigarrow x \mid l \in \mathcal{L}, x \in \text{Var}_P \text{ and } \alpha_{\text{agree}}(\mathcal{O}^l \llbracket x \rrbracket \mathbf{IniTrc})\} \\ &= \{l \rightsquigarrow x \mid l \in \mathcal{L}, x \in \text{Var}_P \text{ and } \Gamma(x) \sqsubseteq l\} \end{aligned}$$

We derive an approximation¹ $\mathcal{O}_D^l(e)^{\natural}$ of l -variety $\mathcal{O}^l(e)$. This approximation $\mathcal{O}_D^l(e)^{\natural} \in \text{Dep} \rightarrow \{\text{tt}, \text{ff}\}$, called l -agreement of expression e , determines whether a set \mathcal{D} of dependence constraints guarantees that no variety is conveyed to expression e when the inputs up to security level l are fixed.

l -agreement of expressions $\mathcal{O}_D^l(e)^{\natural} \in \text{Dep} \rightarrow \{\text{tt}, \text{ff}\}$

$$\begin{aligned} \mathcal{O}_D^l(n)^{\natural} \mathcal{D} &\triangleq \text{tt} & \mathcal{O}_D^l(x)^{\natural} \mathcal{D} &\triangleq (l \rightsquigarrow x \in \mathcal{D}) \\ \mathcal{O}_D^l(e_1 \oplus e_2)^{\natural} \mathcal{D} &\triangleq \mathcal{O}_D^l(e_1)^{\natural} \mathcal{D} \wedge \mathcal{O}_D^l(e_2)^{\natural} \mathcal{D} \\ \mathcal{O}_D^l(e_1 \text{ cmp } e_2)^{\natural} \mathcal{D} &\triangleq \mathcal{O}_D^l(e_1)^{\natural} \mathcal{D} \wedge \mathcal{O}_D^l(e_2)^{\natural} \mathcal{D} \end{aligned}$$

Deriving the clauses defining $\mathcal{O}_D^l(-)^{\natural}$ amounts to a constructive proof of the following.

Lemma 4. $\mathcal{O}_D^l(e)^{\natural}$ is sound:

$$\forall e, \forall l, \forall \mathcal{D}, \quad \alpha_{\text{agree}} \circ \mathcal{O}^l(e) \circ \gamma_{\text{deptr}}(\mathcal{D}) \leftarrow \mathcal{O}_D^l(e)^{\natural} \mathcal{D} .$$

Dependencies abstract semantics. We derive a dependences abstract semantics $\llbracket c \rrbracket^{\natural}$ by approximating the hypercollecting semantics $\llbracket c \rrbracket$. This abstract semantics $\llbracket c \rrbracket^{\natural} \in \text{Dep} \rightarrow \text{Dep}$ over-approximates the dependence constraints that hold after execution of a command c , on inputs satisfying initial dependence constraints.

We assume a static analysis approximating the variables that a command modifies.

Modifiable variables Mod $\in \text{Com} \rightarrow \mathcal{P}(\text{Var})$

For all c, x , if there exists $t, t' \in \mathbf{Trc}$ such that $\llbracket c \rrbracket t = t'$ and $\llbracket x \rrbracket_{\text{pre}t'} \neq \llbracket x \rrbracket_{t'}$, then $x \in \text{Mod}(c)$.

¹ We use symbol \natural here, for contrast with similar notation using \sharp in later sections

Dependences abstract semantics $\langle c \rangle^{\sharp} \in \text{Dep} \rightarrow \text{Dep}$

$$\langle \text{skip} \rangle^{\sharp} \mathcal{D} \triangleq \mathcal{D} \quad \langle c_1; c_2 \rangle^{\sharp} \mathcal{D} \triangleq \langle c_2 \rangle^{\sharp} \circ \langle c_1 \rangle^{\sharp} \mathcal{D}$$

$$\langle x := e \rangle^{\sharp} \mathcal{D} \triangleq \{l \rightsquigarrow y \in \mathcal{D} \mid y \neq x\} \cup \{l \rightsquigarrow x \mid l \in \mathcal{L}, \mathcal{O}'_D \langle e \rangle^{\sharp} \mathcal{D}\}$$

$$\begin{aligned} \langle \text{if } b \text{ then } c_1 \text{ else } c_2 \rangle^{\sharp} \mathcal{D} &\triangleq \\ \text{let } \mathcal{D}_1 &= \langle c_1 \rangle^{\sharp} \mathcal{D} \text{ in} \\ \text{let } \mathcal{D}_2 &= \langle c_2 \rangle^{\sharp} \mathcal{D} \text{ in} \\ \text{let } W &= \text{Mod}(\text{if } b \text{ then } c_1 \text{ else } c_2) \text{ in} \\ \bigcup_{l \in \mathcal{L}} &\left\{ \begin{array}{ll} \pi^l(\mathcal{D}_1) \sqcup^{\sharp} \pi^l(\mathcal{D}_2) & \text{if } \mathcal{O}'_D \langle b \rangle^{\sharp} \mathcal{D} \\ \{l \rightsquigarrow x \in \pi^l(\mathcal{D}) \mid x \notin W\} & \text{otherwise} \end{array} \right. \end{aligned}$$

$$\langle \text{while } b \text{ do } c \rangle^{\sharp} \mathcal{D} \triangleq \text{lfp}_{\mathcal{D}}^{\square^{\sharp}} \langle \text{if } b \text{ then } c_1 \text{ else } c_2 \rangle^{\sharp}$$

$$\pi^l(\mathcal{D}) \triangleq \{l \rightsquigarrow x \in \mathcal{D} \mid x \in \text{Var}_P\}$$

The abstract semantics of assignments $x := e$ discards all atomic constraints related to variable x in the input set \mathcal{D} of constraints, and adds atomic constraints $l \rightsquigarrow x$ if \mathcal{D} guarantees l -agreement for expression e . For conditionals, for each security level l , if the input set \mathcal{D} guarantees l -agreement of the conditional guard, the abstract semantics computes the join over the dependences of both conditional branches, after projecting to only those atomic constraints related to l (notation $\pi^l(-)$). If \mathcal{D} does not guarantee l -agreement of the conditional guard, atomic constraints related to both l and variables possibly modified are discarded. Intuitively, if \mathcal{D} guarantees l -agreement of the conditional guard, then l -agreement over some variable x in both branches guarantees l -agreement over x after the conditional command. Otherwise, the only l -agreements that are guaranteed after the conditional are those that hold before the conditional for variables that are not modified.

Theorem 2. The dependences semantics is sound:

$$\alpha_{\text{deptr}} \circ \langle c \rangle \circ \gamma_{\text{deptr}} \sqsubseteq^{\sharp} \langle c \rangle^{\sharp}.$$

We denote by \sqsubseteq^{\sharp} the point-wise lifting of the partial order \sqsubseteq^{\sharp} . We can derive this abstract semantics by directly approximating the relational hypercollecting semantics $\langle c \rangle$ through the dependences Galois connection $(\alpha_{\text{deptr}}, \gamma_{\text{deptr}})$. The derivation is by structural induction on commands. It leverages mathematical properties of Galois connections. We start with the specification of the best abstract transformer $\alpha_{\text{deptr}} \circ \langle c \rangle \circ \gamma_{\text{deptr}} \in \text{Dep} \rightarrow \text{Dep}$, and successively approximate it to finally obtain the definition of the dependence abstract semantics for each form of command. The derivation is the proof, and the obtained definition of the abstract semantics is correct by construction.

Let us showcase the simplest derivation for a sequence of commands in order to illustrate this process:

$$\begin{aligned} &\alpha_{\text{deptr}} \circ \langle c_1; c_2 \rangle \circ \gamma_{\text{deptr}} \\ &= \text{\{By definition of the hypercollecting semantics\}} \\ &\alpha_{\text{deptr}} \circ \langle c_2 \rangle \circ \langle c_1 \rangle \circ \gamma_{\text{deptr}} \\ &\sqsubseteq^{\sharp} \text{\{By } \gamma_{\text{deptr}} \circ \alpha_{\text{deptr}} \text{ is extensive \}} \\ &\alpha_{\text{deptr}} \circ \langle c_2 \rangle \circ \gamma_{\text{deptr}} \circ \alpha_{\text{deptr}} \circ \langle c_1 \rangle \circ \gamma_{\text{deptr}} \\ &\sqsubseteq^{\sharp} \text{\{By induction hypothesis } \alpha_{\text{deptr}} \circ \langle c \rangle \circ \gamma_{\text{deptr}} \sqsubseteq^{\sharp} \langle c \rangle^{\sharp} \}} \\ &\langle c_2 \rangle^{\sharp} \circ \langle c_1 \rangle^{\sharp} \\ &\triangleq \text{\{Take this last approximation as the definition.\}} \\ &\langle c_1; c_2 \rangle^{\sharp} \end{aligned}$$

Alternatively, we can leverage Galois connections to give the analysis as an approximation of the cardinality analysis. We work this out by ????, introduced in ??.

Comparison with previous analyses. Our dependence analysis is similar to the logic of ? as well as the flow-sensitive type system of ?. The relationship between our sets $\mathcal{D} \in \text{Dep}$ of dependence constraints and the type environments $\Delta \in \text{Var}_P \rightarrow \mathcal{L}$ of Hunt and Sands can be formalised by the abstraction:

$$\begin{aligned} \alpha_{\text{hs}} &\in \text{Dep} \rightarrow \text{Var}_P \rightarrow \mathcal{L} \\ \alpha_{\text{hs}}(\mathcal{D}) &\triangleq \lambda x. \sqcap \{l \mid l \rightsquigarrow x \in \mathcal{D}\} \\ \gamma_{\text{hs}} &\in (\text{Var}_P \rightarrow \mathcal{L}) \rightarrow \text{Dep} \\ \gamma_{\text{hs}}(\Delta) &\triangleq \{l \rightsquigarrow x \mid x \in \text{Var}_P, l \in \mathcal{L}, \Delta(x) \sqsubseteq l\} \end{aligned}$$

This is in fact an isomorphism because of the way we interpret dependences. Indeed, if $l \rightsquigarrow x$ holds, then also $l' \rightsquigarrow x$ for all $l' \in \mathcal{L}$ such that $l \sqsubseteq l'$ (cf. report (?)). This observation suggests reformulating the sets $\mathcal{D} \in \text{Dep}$ of dependence constraints to contain only elements with minimal level, but we refrain from doing so for simplicity of presentation.

Our dependences analysis is at least as precise as the type system of Hunt and Sands. To state this result, we denote by $\perp_{\mathcal{L}}$ the bottom element of the lattice \mathcal{L} . We also assume that the modified variables is precise enough to simulate the same effect as the program counter used in the type system: $\text{Mod}(c)$ is a subset of the variables that are targets of assignments in c .

Theorem 3. For all $c, \mathcal{D}_0, \mathcal{D} \in \text{Dep}$, $\Delta_0, \Delta \in \text{Var}_P \rightarrow \mathcal{L}$, where $\perp_{\mathcal{L}} \vdash \Delta_0 \{c\} \Delta$, and $\mathcal{D} = \langle c \rangle^{\sharp} \mathcal{D}_0$, it holds that:

$$\alpha_{\text{hs}}(\mathcal{D}_0) \sqsubseteq^{\sharp} \Delta_0 \implies \alpha_{\text{hs}}(\mathcal{D}) \sqsubseteq^{\sharp} \Delta.$$

7. Cardinality Abstraction

Dependence analysis is only concerned with whether variety is conveyed. We refine this analysis by deriving a cardinality abstraction that enumerates variety.

We denote by $l \rightsquigarrow x \# n$ an atomic cardinality constraint where $l \in \mathcal{L}$, $x \in \text{Var}_P$ and $n \in [0, \infty]$, read as “agreement up to security level l leads to a variety of at most n values in variable x ”.

Lattice of cardinality constraints $\text{Card} \quad \mathcal{C} \in \text{Card}$

For a program P and lattice \mathcal{L} , we say \mathcal{C} is a **valid set of constraints** iff $\forall x \in \text{Var}_P, \forall l \in \mathcal{L}, \exists! n \in [0, \infty], l \rightsquigarrow x \# n \in \mathcal{C}$.

Let Card be the set of valid sets of constraints.

It is a complete lattice:

$$\begin{aligned} \mathcal{C}_1 \sqsubseteq^{\sharp} \mathcal{C}_2 &\text{ iff } \forall l \rightsquigarrow x \# n_1 \in \mathcal{C}_1, \exists n_2, \\ &\quad l \rightsquigarrow x \# n_2 \in \mathcal{C}_2 \wedge n_1 \leq n_2 \\ \mathcal{C}_1 \sqcup^{\sharp} \mathcal{C}_2 &\triangleq \{l \rightsquigarrow x \# \max(n_1, n_2) \mid \\ &\quad l \rightsquigarrow x \# n_1 \in \mathcal{C}_1, l \rightsquigarrow x \# n_2 \in \mathcal{C}_2\} \end{aligned}$$

In the rest of this section, \mathcal{L} and P are fixed, together with a typing context $\Gamma \in \text{Var}_P \rightarrow \mathcal{L}$.

A valid constraint set is essentially a function from l and x to n . So \sqsubseteq^{\sharp} is essentially a pointwise order on functions, and we ensure that \sqsubseteq^{\sharp} is antisymmetric.

The cardinality abstraction relies on the abstraction α_{cardval} , introduced in ??, in order to approximate l -variety of a variable into a cardinality $n \in [0, \infty]$.

Cardinality abstraction crdtr α_{crdtr} γ_{crdtr}

$$\begin{aligned}
\text{crdtr} &\in \mathcal{P}(\mathbf{Trc}) \rightarrow \text{Card} \\
\text{crdtr}(T) &\triangleq \{l \rightsquigarrow x\#n \mid l \in \mathcal{L}, x \in \text{Var}_{\mathbf{P}}, \\
&\quad n = \alpha_{\text{crdval}}(\mathcal{O}^l\{x\}T)\} \\
\alpha_{\text{crdtr}} &\in \mathcal{P}(\mathcal{P}(\mathbf{Trc})) \rightarrow \text{Card} \\
\alpha_{\text{crdtr}}(\mathbb{T}) &\triangleq \sqcup_{T \in \mathbb{T}} \text{crdtr}(T) \\
\gamma_{\text{crdtr}} &\in \text{Card} \rightarrow \mathcal{P}(\mathcal{P}(\mathbf{Trc})) \\
\gamma_{\text{crdtr}}(\mathcal{C}) &\triangleq \{T \mid \text{crdtr}(T) \sqsubseteq^{\#} \mathcal{C}\}
\end{aligned}$$

$$(\mathcal{P}(\mathcal{P}(\mathbf{Trc})), \subseteq) \xleftrightarrow[\alpha_{\text{crdtr}}]{\gamma_{\text{crdtr}}} (\text{Card}, \sqsubseteq^{\#})$$

The cardinality abstraction enables us to derive an approximation $\mathcal{O}_C^l(e)^{\#}$ of l -variety $\mathcal{O}^l(e)$. This approximation $\mathcal{O}_C^l(e)^{\#} \in \text{Card} \rightarrow [0, \infty]$, called l -cardinality of expression e , enumerates the l -variety conveyed to expression e assuming a set $\mathcal{C} \in \text{Card}$ of cardinality constraints holds. Note that the infinite cardinal ∞ is absorbing, i.e. $\forall n, \infty \times n \triangleq \infty$.

l -cardinality of expressions $\mathcal{O}_C^l(e)^{\#} \in \text{Card} \rightarrow [0, \infty]$

$$\begin{aligned}
\mathcal{O}_C^l(n)^{\#} \mathcal{C} &\triangleq 1 & \mathcal{O}_C^l(x)^{\#} \mathcal{C} &\triangleq n \text{ where } l \rightsquigarrow x\#n \in \mathcal{C} \\
\mathcal{O}_C^l(e_1 \oplus e_2)^{\#} \mathcal{C} &\triangleq \mathcal{O}_C^l(e_1)^{\#} \mathcal{C} \times \mathcal{O}_C^l(e_2)^{\#} \mathcal{C} \\
\mathcal{O}_C^l(e_1 \text{ cmp } e_2)^{\#} \mathcal{C} &\triangleq \min\left(2, \mathcal{O}_C^l(e_1)^{\#} \mathcal{C} \times \mathcal{O}_C^l(e_2)^{\#} \mathcal{C}\right)
\end{aligned}$$

Lemma 5. $\mathcal{O}_C^l(e)^{\#}$ is sound:

$$\forall e, \forall l, \quad \alpha_{\text{crdval}} \circ \mathcal{O}^l(e) \circ \gamma_{\text{crdtr}} \leq \mathcal{O}_C^l(e)^{\#}$$

We now derive a cardinality abstract semantics by approximating the relational hypercollecting semantics of \mathbf{Trc} . It uses definitions to follow.

Cardinality abstract semantics $\llbracket c \rrbracket^{\#} \in \text{Card} \rightarrow \text{Card}$

$$\begin{aligned}
\llbracket \text{skip} \rrbracket^{\#} \mathcal{C} &\triangleq \mathcal{C} & \llbracket c_1; c_2 \rrbracket^{\#} \mathcal{C} &\triangleq \llbracket c_2 \rrbracket^{\#} \circ \llbracket c_1 \rrbracket^{\#} \mathcal{C} \\
\llbracket x := e \rrbracket^{\#} \mathcal{C} &\triangleq \\
&\{l \rightsquigarrow y\#n \in \mathcal{C} \mid y \neq x\} \\
&\cup \{l \rightsquigarrow x\#n \mid l \in \mathcal{L}, x \in \text{Var}_{\mathbf{P}}, n = \mathcal{O}_C^l(e)^{\#} \mathcal{C}\}
\end{aligned}$$

$$\begin{aligned}
\llbracket \text{if } b \text{ then } c_1 \text{ else } c_2 \rrbracket^{\#} \mathcal{C} &\triangleq \\
&\text{let } \mathcal{C}_1 = \llbracket c_1 \rrbracket^{\#} \mathcal{C} \text{ in} \\
&\text{let } \mathcal{C}_2 = \llbracket c_2 \rrbracket^{\#} \mathcal{C} \text{ in} \\
&\text{let } W = \text{Mod}(\text{if } b \text{ then } c_1 \text{ else } c_2) \text{ in} \\
&\bigcup_{l \in \mathcal{L}} \begin{cases} \pi^l(\mathcal{C}_1) \sqcup^{\#} \pi^l(\mathcal{C}_2) & \text{if } \mathcal{O}_C^l(b)^{\#} \mathcal{C} = 1 \\ \pi^l(\mathcal{C}_1) \sqcup_{\text{add}(W, \pi^l(\mathcal{C}))}^{\#} \pi^l(\mathcal{C}_2) & \text{otherwise} \end{cases}
\end{aligned}$$

$$\llbracket \text{while } b \text{ do } c \rrbracket^{\#} \mathcal{C} \triangleq \text{lfp}_{\mathcal{C}}^{\#} \llbracket \text{if } b \text{ then } c_1 \text{ else } c_2 \rrbracket^{\#}$$

$$\begin{aligned}
\pi^l(\mathcal{C}) &\triangleq \{l \rightsquigarrow x\#n \in \mathcal{C} \mid x \in \text{Var}_{\mathbf{P}}, n \in [0, \infty]\} \\
C_1 \sqcup_{\text{add}(W, C_0)}^{\#} C_2 &\triangleq \bigcup_{x \in \text{Var}_{\mathbf{P}} \setminus W} \{l \rightsquigarrow x\#n \in C_0\} \\
&\quad \cup \bigcup_{x \in W} \{l \rightsquigarrow x\#(n_1+n_2) \mid \\
&\quad \quad l \rightsquigarrow x\#n_j \in C_j, j = 1, 2\}
\end{aligned}$$

The abstract semantics of assignments $x := e$ is similar in spirit to the one for dependences: discard atomic constraints related to x , and add new ones by computing l -cardinality of expression e . The

abstract semantics of conditionals is also similar to dependences: if the conditional guard does not convey l -variety, then all initially l -equivalent traces follow the same execution path and the join operator (defined as max over cardinality) over both conditional branches over-approximates the l -cardinality after the conditional. Otherwise, the l -cardinality over both conditional branches have to be summed—for the variables that may be modified in the conditional branches—to soundly approximate the l -cardinality after the conditional.

Theorem 4. The cardinality abstract semantics is sound:

$$\alpha_{\text{crdtr}} \circ \llbracket c \rrbracket^{\#} \circ \gamma_{\text{crdtr}} \sqsubseteq^{\#} \llbracket c \rrbracket^{\#}$$

The lattice Card is complete, although not finite. We may define a widening operator $\nabla \in \text{Card} \times \text{Card} \rightarrow \text{Card}$ to ensure convergence of the analysis $(?)^{\#}(\?)^{\#}$ (Sec. 4).

$$\begin{aligned}
\mathcal{C}_1 \nabla \mathcal{C}_2 &\triangleq \{l \rightsquigarrow x\#n \mid l \rightsquigarrow x\#n_1 \in \mathcal{C}_1, l \rightsquigarrow x\#n_2 \in \mathcal{C}_2, \\
&\quad n = n_1 \nabla n_2\} \\
n_1 \nabla n_2 &\triangleq \text{if } (n_2 \leq n_1) \text{ then } n_1 \text{ else } \infty
\end{aligned}$$

The occurrence of widening depends on the iteration strategy employed by the static analyser. Widening accelerates or forces the convergence of fixpoint computations. In the simplest setting, the analyser passes as arguments to the widening operator the old set \mathcal{C}_1 of cardinality as well as the new set \mathcal{C}_2 that is computed. For each atomic cardinality constraint, the widening operator then compares the old cardinality n_1 to the new cardinality n_2 . If the cardinality is still strictly increasing ($n_2 > n_1$), the widening forces the convergence by setting it to ∞ . If the cardinality is decreasing, the widening operator sets it to the maximum cardinality n_1 in order to force convergence and ensure the sequence of computed cardinalities is stationary.

Min-capacity leakage. So far, we showed how one can derive static analyses of hyperproperties—the abstract representations themselves are interpreted as hyperproperties—by approximating hypercollecting semantics. Let us now recall the security requirement $\text{SR}(l, k, x)$ introduced in \mathbf{Trc} in order to illustrate how these analyses may prove that a program satisfies a hyperproperty, i.e. Step 3 of the methodology in \mathbf{Trc} (see also \mathbf{Trc}).

Consider a program P characterised by a set $T_P \in \mathcal{P}(\mathbf{Trc})$ of traces, i.e. T_P is $\llbracket P \rrbracket^{\#} \mathbf{IniTrc}$. How do we prove that P satisfies the hyperproperty $\text{SR}(l, k, x)$? We can use the cardinality analysis to prove that variable x has a l -cardinality that is at most k . Indeed, if \mathcal{C} approximates T_P (i.e. $\alpha_{\text{crdtr}}(\{T_P\}) \sqsubseteq^{\#} \mathcal{C}$) then $\alpha_{\text{crdval}} \circ \mathcal{O}^l\{x\}T_P \leq \mathcal{O}_C^l(x)^{\#} \mathcal{C}$. Thus, if the inferred l -cardinality of \mathcal{C} is at most k then program P is guaranteed to satisfy the hyperproperty $\text{SR}(l, k, x)$. We have $\{T_P\} \subseteq \gamma_{\text{crdtr}}(\mathcal{C})$ since \mathcal{C} approximates T_P (i.e. $\alpha_{\text{crdtr}}(\{T_P\}) \sqsubseteq^{\#} \mathcal{C}$). And we have $\gamma_{\text{crdtr}}(\mathcal{C}) \subseteq \text{SR}(l, k, x)$ by assumption $\mathcal{O}_C^l(x)^{\#} \mathcal{C} \leq k$. Hence $T_P \in \text{SR}(l, k, x)$.

The hyperproperty $\text{SR}(l, k, x)$ is a $(k+1)$ -safety hyperproperty $(?)$, i.e. it requires exhibiting at most $k+1$ traces in order to prove that a program does not satisfy $\text{SR}(l, k, x)$. For example, termination-insensitive noninterference for security level l , which corresponds to the hyperproperty $\text{SR}(l, 1, x)$, is 2-safety. A k -safety hyperproperty of a program can be reduced to a safety property of a k -fold product program $(\mathbf{Trc})^k$.

Various quantitative information flow properties are not k -safety. For example, the bounding problem that the cardinality analysis targets, namely min-capacity leakage is not a k -safety hyperproperty for any k (Sec. 3). Instead, this bounding problem is hyper-safety $(?)$.

Cardinality vs. dependence. Just as quantitative security metrics are the natural generalisations of qualitative metrics such as noninterference, the cardinality abstraction is a natural generalisation of

dependence analysis. Instead of deciding if variety is conveyed, the cardinality analysis enumerates this variety. In other words, dependences are abstractions of cardinalities. We can factor the Galois connections, e.g. $(\alpha_{\text{agree}}, \gamma_{\text{agree}})$ is $(\alpha_{\text{lqone}} \circ \alpha_{\text{crdval}}, \gamma_{\text{crdval}} \circ \gamma_{\text{lqone}})$ for suitable $(\alpha_{\text{lqone}}, \gamma_{\text{lqone}})$.

Lemma 6. $(\alpha_{\text{agree}}, \gamma_{\text{agree}})$ is the composition of two Galois connections $(\alpha_{\text{crdval}}, \gamma_{\text{crdval}})$ and $(\alpha_{\text{lqone}}, \gamma_{\text{lqone}})$:

$$(\mathcal{P}(\mathcal{P}(\mathbf{Val})), \sqsubseteq) \xleftarrow[\alpha_{\text{crdval}}]{\gamma_{\text{crdval}}} ([0, \infty], \leq) \xleftarrow[\alpha_{\text{lqone}}]{\gamma_{\text{lqone}}} (\{\text{tt}, \text{ff}\}, \Leftarrow)$$

with:

$$\alpha_{\text{lqone}}(n) \triangleq \begin{cases} \text{tt} & \text{if } n \leq 1 \\ \text{ff} & \text{otherwise.} \end{cases}, \text{ and}$$

$$\gamma_{\text{lqone}}(\text{bv}) \triangleq \begin{cases} 1 & \text{if } \text{bv} = \text{tt} \\ \infty & \text{otherwise.} \end{cases}$$

Lemma 7. $(\alpha_{\text{deptr}}, \gamma_{\text{deptr}})$ is the composition of two Galois connections $(\alpha_{\text{crdtr}}, \gamma_{\text{crdtr}})$ and $(\alpha_{\text{lqonecc}}, \gamma_{\text{lqonecc}})$:

$$(\mathcal{P}(\mathcal{P}(\mathbf{Trc})), \sqsubseteq) \xleftarrow[\alpha_{\text{crdtr}}]{\gamma_{\text{crdtr}}} (\text{Card}, \sqsubseteq^{\sharp}) \xleftarrow[\alpha_{\text{lqonecc}}]{\gamma_{\text{lqonecc}}} (\text{Dep}, \sqsubseteq^{\sharp})$$

with:

$$\alpha_{\text{lqonecc}}(\mathcal{C}) \triangleq \{l \rightsquigarrow x \mid l \rightsquigarrow x \# n \in \mathcal{C} \text{ and } \alpha_{\text{lqone}}(n)\}$$

$$\gamma_{\text{lqonecc}}(\mathcal{D}) \triangleq \bigcup_{l \in \mathcal{L}, x \in \text{VarP}} \{l \rightsquigarrow x \# n \mid n = \gamma_{\text{lqone}}(l \rightsquigarrow x \in \mathcal{D})\}$$

We use $????$ to abstract further the cardinality abstract semantics and derive the correct by construction dependence analysis of $??$. This derivation, which can be found in (?), proves $??$ and $??$ stated earlier.

As a corollary and by $??$, this also proves the precision of the cardinality analysis relative to Amtoft and Banerjee’s logic (?) as well as Hunt and Sands’ type system (??).

Corollary 1 No leakage for well-typed programs. For all $c, \mathcal{C}_0, \mathcal{C} \in \text{Card}$, $\Delta_0, \Delta \in \text{VarP} \rightarrow \mathcal{L}$, where $\perp_{\mathcal{L}} \vdash \Delta_0\{c\}\Delta$, and $\mathcal{C} = \langle c \rangle^{\sharp} \mathcal{C}_0$, it holds that:

$$\alpha_{\text{hs}} \circ \alpha_{\text{lqonecc}}(\mathcal{C}_0) \sqsubseteq^{\sharp} \Delta_0 \implies$$

$$(\forall x \in \text{VarP}, l \in \mathcal{L}, \Delta(x) \sqsubseteq l \implies \mathcal{O}_C^l(x)^{\sharp} \leq 1)$$

The cardinality analysis determines that there is no leakage for programs that are “well-typed” by the flow-sensitive type system of Hunt and Sands. By “well-typed”, we mean that the final typing environment that is computed by the type system allows attackers with security clearance $l \in \mathcal{L}$ to observe a variable $x \in \text{VarP}$.

To the best of our knowledge, the cardinality abstraction is the first approximation-based analysis for quantitative information flow that provides a formal precision guarantee wrt. traditional analyses for qualitative information flow. This advantage makes the cardinality analysis appealing even when interested in proving a qualitative security policy such as non-interference, since the cardinality abstraction provides quantitative information that may assist in making better informed decisions if declassification is necessary. Nonetheless, we need further experimentation to compare to other quantitative analyses —see $??$.

8. Towards More Precision

This section introduces examples to evaluate the precision of the analyses, and shows how existing analyses can be leveraged to improve precision. For simplicity, we consider a two point lattice $\{L, H\}$ and an initial typing context where variables y_i are the only low variables ($\Gamma(y_i) = L$). As is usual, low may flow to high ($L \sqsubseteq H$).

Consider the following program.

```

1 | if (y1 ≥ secret) then
2 |   x := y2
3 | else
4 |   x := y3

```

Listing 1. Leaking 1 bit of secret

The cardinal abstraction determines that x has at most 2 values after the execution of the program in $??$, for initially L -equivalent traces. For fixed low inputs, x has one value in the then branch and one value in the else branch, and these cardinalities get summed after the conditional guard since the conditional guard may evaluate to 2 different values. Thus, the cardinality abstraction proves that this example program satisfies the hyperproperty $\text{SR}(L, 2, x)$.

Stronger trace properties. Another way of proving a hyperproperty is by proving a stronger trace property. If a program is proven to satisfy a trace property $T \in \mathcal{P}(\mathbf{Trc})$, then proving that T is stronger than hyperproperty $H \in \mathcal{P}(\mathcal{P}(\mathbf{Trc}))$ —in the sense that $\gamma_{\text{hpp}}(T) \subseteq H$ —guarantees the program satisfies the hyperproperty H . For instance, by proving for some program that an output variable x ranges over an interval of integer values whose size is k , we can prove that program satisfies $\text{SR}(L, k, x)$.

However, approximating a hyperproperty by a trace property may be too coarse for some programs, as we can illustrate with an interval analysis (?) on the example program in $??$. Such an interval analysis loses too much precision in the initial state of this program, since it maps all low input variables y_1, y_2 and y_3 to $[-\infty, +\infty]$. After the conditional, it determines that x belongs to the interval $[-\infty, +\infty]$, which is a coarse over-approximation. Also, a polyhedron (?) does not capture the disjunction that is needed for this example program ($x = y_2$ or $x = y_3$). Both abstract domains and many more existing ones are not suitable for the task of inferring cardinalities or dependences because they are convex. Using them as a basis to extract counting information delivers an over-approximation of the leakage, but a very coarse one, especially in the presence of low inputs.

A disjunction of two polyhedra (through powerset domains, disjunctive postconditions, or trace partitioning) is as precise as the cardinality analysis for this example. However, disjunctions are not tractable in general. As soon as one fixes a maximum number of disjunctive elements (as in the quantitative information flow analysis of $??$) or defines a widening operator to guarantee convergence, one loses the relative precision wrt. classical dependence analyses (??) that the cardinality analysis guarantees (Cf. ??).

Consider the following program.

```

1 | if (y1 ≥ secret) then x := y2 else x := y3;
2 | o := y4 * x

```

Listing 2. Leaking x

The cardinal abstraction determines that variable o leaks the two possible values of x ; for fixed low inputs, x has two possible values whereas y_4 has one possible value. Relational abstract domains such as polyhedra (?) or octogons (?) do not support non-linear expressions, and therefore are unable to compute a precise bound of the leakage for variable o . An analysis with a disjunction $\{x =$

$y_2 \vee x = y_3$ of polyhedra and with linearisation over intervals (?) will compute either $\{o = y_2 * [-\infty, +\infty] \vee y_3 * [-\infty, +\infty]\}$ if the linearisation happens for the right expression, or $\{o = [-\infty, +\infty] * x, y_3 * [-\infty, +\infty] * x\}$ if the linearisation happens for the left expression; none will deduce that variable o has at most 2 values.

Scaling to richer languages. We can rely on existing abstract domains to support richer language constructs, e.g. pointers and aliasing. Consider the following variation of ??.

```

if (y1 ≥ secret) then
  p := &y2
else
  p := &y3
o := *p

```

Listing 3. Leaking 1 bit of secret

The cardinality abstraction determines that initially L -equivalent memories lead to a variety of at most 2 in the pointer p after the conditional, whereas both y_2 and y_3 have a variety of 1. Assuming an aliasing analysis determines that p may point to y_2 or y_3 , the cardinality analysis determines that variable o has a variety of at most 2, for initially L -equivalent memories.

Improving precision. To improve precision of the cardinality abstraction, we can augment it with existing abstract domains. One shortcoming of the cardinality analysis is the fact that it is not relational. Assuming attackers with security clearance L observe both variable x and o after the execution of ??, the cardinality abstraction leads us to compute a leakage of two bits: four different possible values, instead of only 2 possible values for initially L -equivalent memories. Relying on a relational domain with linearisation (?) over cardinality (instead of intervals) captures the required constraints $\{L \rightsquigarrow o \# 1 * x, L \rightsquigarrow x \# 2\}$ to compute a leakage of only one bit; these constraints are to be interpreted as “initially L -equivalent memories result in x having at most 2 values, and o being equal to one fixed integer times x ”.

We leave these extensions of the cardinality abstraction for future work. In the following, we focus on one particular improvement to both previous analyses in order to gain more precision. We uncovered this case while deriving the analysis by relying on the calculational framework of abstract interpretation. Indeed, notice that the following holds:

$$\alpha_{\text{crdval}} \circ \mathcal{O}^l(\{x_1\}) \circ (\text{grd}^{x_1==x_2}) \circ \gamma_{\text{crdtr}}(\mathcal{C}) \leq \mathcal{O}_C^l(\{x_2\}) \# \mathcal{C}$$

$$\alpha_{\text{crdval}} \circ \mathcal{O}^l(\{x_2\}) \circ (\text{grd}^{x_1==x_2}) \circ \gamma_{\text{crdtr}}(\mathcal{C}) \leq \mathcal{O}_C^l(\{x_1\}) \# \mathcal{C}$$

Therefore, we can deduce that:

$$\alpha_{\text{crdtr}} \circ (\text{grd}^{x_1==x_2}) \circ \gamma_{\text{crdtr}}(\mathcal{C})$$

$$\sqsubseteq^\sharp \{l \rightsquigarrow x \# n \in \mathcal{C} \mid x \neq x_1, x \neq x_2\}$$

$$\cup \{l \rightsquigarrow x_1 \# \min(n_1, n_2), l \rightsquigarrow x_2 \# \min(n_1, n_2) \mid$$

$$l \rightsquigarrow x_1 \# n_1 \in \mathcal{C}, l \rightsquigarrow x_2 \# n_2 \in \mathcal{C}\}$$

$$\triangleq (\text{grd}^{x_1==x_2}) \# \mathcal{C}$$

For other comparison operators, we use as before $(\text{grd}^b) \# \mathcal{C} \triangleq \mathcal{C}$.

We can now also improve the dependences abstraction:

$$\alpha_{\text{qonecc}} \circ (\text{grd}^{x_1==x_2}) \# \circ \gamma_{\text{qonecc}}(\mathcal{D})$$

$$\sqsubseteq^\sharp \alpha_{\text{qonecc}}(\{l \rightsquigarrow x \# n \in \gamma_{\text{qonecc}}(\mathcal{D}) \mid x \neq x_1, x \neq x_2\})$$

$$\cup \alpha_{\text{qonecc}}(\{l \rightsquigarrow x_1 \# \min(n_1, n_2), l \rightsquigarrow x_2 \# \min(n_1, n_2) \mid$$

$$l \rightsquigarrow x_1 \# n_1 \in \gamma_{\text{qonecc}}(\mathcal{D}), l \rightsquigarrow x_2 \# n_2 \in \gamma_{\text{qonecc}}(\mathcal{D})\})$$

$$\sqsubseteq^\sharp \{l \rightsquigarrow x \in \mathcal{D} \mid x \neq x_1, x \neq x_2\}$$

$$\cup \{l \rightsquigarrow x_1, l \rightsquigarrow x_2 \mid l \rightsquigarrow x_1 \in \mathcal{D} \text{ or } l \rightsquigarrow x_2 \in \mathcal{D}\}$$

$$\triangleq (\text{grd}^{x_1==x_2}) \# \mathcal{D}$$

For other comparison operators, we also use $(\text{grd}^b) \# \mathcal{D} \triangleq \mathcal{D}$.

With these new definitions, we can update the abstract semantics of conditionals and loops, for both dependences and cardinalities, to leverage the transfer functions $(\text{grd}^-) \#$ and $(\text{grd}^-) \#$.

Improved dependences abstract semantics $(c) \# \in \text{Dep} \rightarrow \text{Dep}$

$$(\text{if } b \text{ then } c_1 \text{ else } c_2) \# \mathcal{D} \triangleq$$

$$\text{let } \mathcal{D}_1 = (\text{grd}^b) \# \circ (c_1) \# \mathcal{D} \text{ in}$$

$$\text{let } \mathcal{D}_2 = (\text{grd}^{-b}) \# \circ (c_2) \# \mathcal{D} \text{ in}$$

$$\text{let } W = \text{Mod}(\text{if } b \text{ then } c_1 \text{ else } c_2) \text{ in}$$

$$\bigcup_{l \in \mathcal{L}} \begin{cases} \pi^l(\mathcal{D}_1) \sqcup^\sharp \pi^l(\mathcal{D}_2) & \text{if } \mathcal{O}_D^l(b) \# \mathcal{D} \\ \{l \rightsquigarrow x \in \pi^l(\mathcal{D}) \mid x \notin W\} & \text{otherwise} \end{cases}$$

$$(\text{while } b \text{ do } c) \# \mathcal{D} \triangleq (\text{grd}^{-b}) \# \circ \text{lfp}_{\mathcal{D}}^\# (\text{if } b \text{ then } c_1 \text{ else } c_2) \#$$

Improved cardinality abs. semantics $(c) \# \in \text{Card} \rightarrow \text{Card}$

$$(\text{if } b \text{ then } c_1 \text{ else } c_2) \# \mathcal{C} \triangleq$$

$$\text{let } \mathcal{C}_1 = (\text{grd}^b) \# \circ (c_1) \# \mathcal{C} \text{ in}$$

$$\text{let } \mathcal{C}_2 = (\text{grd}^{-b}) \# \circ (c_2) \# \mathcal{C} \text{ in}$$

$$\text{let } W = \text{Mod}(\text{if } b \text{ then } c_1 \text{ else } c_2) \text{ in}$$

$$\bigcup_{l \in \mathcal{L}} \begin{cases} \pi^l(\mathcal{C}_1) \sqcup^\sharp \pi^l(\mathcal{C}_2) & \text{if } \mathcal{O}_C^l(b) \# \mathcal{C} = 1 \\ \pi^l(\mathcal{C}_1) \sqcup_{\text{add}(W, \pi^l(\mathcal{C}))}^\sharp \pi^l(\mathcal{C}_2) & \text{otherwise} \end{cases}$$

$$(\text{while } b \text{ do } c) \# \mathcal{C} \triangleq (\text{grd}^{-b}) \# \circ \text{lfp}_{\mathcal{C}}^\# (\text{if } b \text{ then } c_1 \text{ else } c_2) \#$$

To illustrate the benefits of this improvement, consider the following example.

```

1 while (secret != y3) do {
2   x := x+1;
3   secret := secret - 1;
4 }
5 o := secret;

```

Listing 4. Improved precision

The cardinality analysis determines that initially L -equivalent memories result in x having an infinity of values: the L -cardinality of x grows until it is widened to ∞ . In contrast, the cardinalities also determines that variables o and $secret$ have only 1 value, assuming L -equivalent memories. This is because of the reduction that concerns variable $secret$ after the while loop, specifically $(\text{grd}^{secret==y_3}) \#$. Similarly, the improved dependences analysis also determines that both variables $secret$ and o are low.

Remarkably, this has been overlooked by many previous analyses. In fact, this simple improvement makes our dependence analysis strictly more precise than ?'s and ??'s analyses and incomparable to the more recent dependence analysis of ?.

Combination with intervals. Consider now the following example inspired from ?.

```

1 if (secret == 0) then {
2   x := 0;
3   y := y + 1;
4 }
5 else {
6   x := 0;
7 }

```

Listing 5. Example program from ?

The analysis of ? determines that x is low, whereas the cardinality abstraction determines that L -equivalent memories result in at most 2 values for variable x , because it does not track the actual values of the variables. We can combine cardinality with an interval analysis to be more precise in such cases, through a reduced product (??).

Assume a set StInt of interval environments provided with the usual partial order that we denote by $\leq^{\#, \text{Int}}$. Assume also a Galois connection $(\alpha^{\text{Int}}, \gamma^{\text{Int}})$ enabling the derivation of an interval analysis as an approximation of a standard collecting semantics defined over $\mathcal{P}(\mathbf{Trc})$. We can lift this Galois connection to $\mathcal{P}(\mathcal{P}(\mathbf{Trc}))$ to obtain a Galois connection by composing with $(\alpha_{\text{hpp}}, \gamma_{\text{hpp}})$, to obtain $(\alpha', \gamma') \triangleq (\alpha^{\text{Int}} \circ \alpha_{\text{hpp}}, \gamma^{\text{Int}} \circ \gamma_{\text{hpp}})$ with:

$$(\mathcal{P}(\mathcal{P}(\mathbf{Trc})), \subseteq) \xleftarrow[\alpha_{\text{hpp}}]{\gamma_{\text{hpp}}} (\mathcal{P}(\mathbf{Trc}), \subseteq) \xleftarrow[\alpha^{\text{Int}}]{\gamma^{\text{Int}}} (\text{StInt}, \leq^{\#, \text{Int}})$$

A Granger's reduced product ? for the cardinality abstraction and an interval analysis may be defined as a pair of functions $\text{toint} \in \text{Card} \times \text{StInt} \rightarrow \text{StInt}$ and $\text{tocard} \in \text{Card} \times \text{StInt} \rightarrow \text{Card}$ verifying the following conditions:

1. soundness:

$$\begin{aligned} \gamma'(\text{toint}(\mathcal{C}, \iota)) \cap \gamma_{\text{crdtr}}(\mathcal{C}) &= \gamma'(\iota) \cap \gamma_{\text{crdtr}}(\mathcal{C}) \\ \gamma'(\iota) \cap \gamma_{\text{crdtr}}(\text{tocard}(\mathcal{C}, \iota)) &= \gamma'(\iota) \cap \gamma_{\text{crdtr}}(\mathcal{C}) \end{aligned}$$

2. reduction:

$$\begin{aligned} \text{toint}(\mathcal{C}, \iota) &\leq^{\#, \text{Int}} \iota \\ \text{tocard}(\mathcal{C}, \iota) &\sqsubseteq^{\#} \mathcal{C} \end{aligned}$$

Let us denote by size the function that returns the size of an interval. One such Granger's reduced product can be defined as:

$$\begin{aligned} \text{tocard} &\in \text{Card} \times \text{StInt} \rightarrow \text{Card} \\ \text{tocard}(\mathcal{C}, \iota) &\triangleq \{l \rightsquigarrow x\#n' \mid l \rightsquigarrow x\#n \in \mathcal{C} \text{ and} \\ &\quad n' = \min(n, \text{size } \iota(x))\} \\ \text{toint} &\in \text{Card} \times \text{StInt} \rightarrow \text{Card} \\ \text{toint}(\mathcal{C}, \iota) &\triangleq \iota \end{aligned}$$

Once enhanced with this reduced product, the cardinality analysis determines for the program in ??, that L -equivalent memories result in at most one possible value for variable x .

The dependences analysis can be improved similarly, with a reduction function defined as follows:

$$\begin{aligned} \text{todep} &\in \text{Dep} \times \text{StInt} \rightarrow \text{Dep} \\ \text{todep}(\mathcal{D}, \iota) &\triangleq \mathcal{D} \cup \{l \rightsquigarrow x \mid l \in \mathcal{L} \text{ and } \text{size } \iota(x) = 1\} \end{aligned}$$

Once extended with a reduced product with intervals, the dependence analysis is also able to determine that variable x is low for the program in ??.

More reduced products. As a final example, let us consider ??, inspired by ?, program 7, that we annotate with the result of the improved cardinality abstraction. To the best of our knowledge, no

```

0 //L ~ h#∞, L ~ y1#1, L ~ y2#1, L ~ y3#1
1 y1 := 1; //L ~ y1#1
2 if (h == y1) then {
3   skip; //L ~ h#1, L ~ y1#1, L ~ y2#1
4 }
5 else {
6   y2 := 5; //L ~ y1#1, L ~ y2#1
7   while (y2 != 1) do {
8     y2 := y2-1; //L ~ y2#1
9     y1 := y2; //L ~ y1#1
10  } //L ~ y1#1, L ~ y2#1
11 }
12 //L ~ h#∞, L ~ y1#2, L ~ y2#2, L ~ y3#1
13 o := y1 * y3; //L ~ o#2

```

Listing 6. No leakage for variable o

existing automated static analysis determines that variable o is low at the end of this program. Also, no prior monitor but the one recently presented by ? accepts all executions of this program, assuming attackers with clearance L can observe variable o .

For initially L -equivalent memories, the cardinality abstraction determines that variables y_1 , y_2 and o have at most two values. This result is precise for y_2 , but not precise for y_1 and o . As a challenge, let us see what is required to gain more precision to determine that both variables y_1 and o have at most 1 possible value – in other words, they are low.

To tackle this challenge, we need to consider cardinality combined with an interval analysis and a simple relational domain tracking equalities. With the equality $y_1 = y_2$ at the exit of the loop, both y_1 and y_2 will be reduced to the singleton interval $[1, 1]$. After the conditional, we still deduce that y_2 has at most 2 different values thanks to the cardinality abstraction. Using intervals, we deduce that variable y_1 has only one value (singleton interval $[1, 1]$). And finally, at the last assignment the cardinalities abstraction determines that variable o has only one possible value. Similarly, this same combination of analyses can be put to use to let the dependence analysis reach the desired precision.

9. Related Work

Although noninterference has important applications, for many security requirements it is too strong. That is one motivation for research in quantitative information flow analysis. In addition, a number of works investigate weakenings of noninterference and downgrading policies that are conditioned on events or data values (??). ? introduce *abstract noninterference*, which generalizes noninterference by means of abstract interpretations that specify, for example, limits on the attacker's power and the extent of partial releases (declassification). The survey by ? further generalizes the notion and highlights, among other things, its applicability to a range of underlying semantics. The Galois connections in this work are at the level of trace sets, not sets of sets. Abstract noninterference retains the explicit 2-run formulation (??): from two related initial states, two executions lead to related final states. The relations are defined in terms of abstract interpretations of the individual states/executions. ? show how to infer indistinguishability relations—modelling attackers' observations—to find the best abstract noninterference policy that holds. The inference algorithm iteratively refines the relation by using counter-examples and abstract domain completion (?).

Set-of-sets structures occur in work on abstraction for nondeterministic programs, but in those works one level of sets are powerdomains for nondeterminacy; the properties considered are trace properties (??). ? develop a binding time analysis and a strictness

analysis (?) based on partial equivalence relations: Their concretisations are sets of equivalence classes. ? point out that this analysis could be achieved by a collecting semantics over sets-of-sets, defined simply as a direct image. To the best of our knowledge this has not been explored further in the literature, except in unpublished work on which this paper builds (??).

?? extend temporal logic with means to quantify over multiple traces in order to express hyperproperties, and provide model checking algorithms for finite space systems. ? introduce a technique for runtime verification of k -safety properties.

The dependences analysis we derive is similar to the information flow logic of ? and the equivalent flow-sensitive type system of ?. Amtoft and Banerjee use the domain $\mathcal{P}(\mathbf{Trc})$ and on the basis of a relational logic they validate a forward analysis. In effect their interpretation of “independences” is a Galois connection with sets of sets, but the analysis is not formulated or proved correct as an abstract interpretation. To deal with dynamically allocated state, ? augment the relational assertions of information flow logic with region assertions, which can be computed by abstract interpretation. This is used both to express agreement relations between the two executions and to approximate modifiable locations. This approach is generalized in ? to a relational Hoare logic for object-based programs that encompasses information flow properties with conditional downgrading (?).

? give a backwards analysis that infers dependencies and is proved strictly more precise than (??). This is achieved by product construction that facilitates inferring relations between variables in executions that follow different control paths. Correctness of the analysis is proved by way of a relational Hoare logic. The variations of our proposed analyses, in ??, rivals theirs in terms of precision—they are incomparable.

Our dependence analysis relies on an approximation of the modifiable variables, to soundly track implicit flows due to control flow, instead of labelling a program counter variable pc to account for implicit flows (?). ? also derive a similar analysis through a syntactic Galois connection—a syntactic assignment $z := x * y$ is abstracted into a propositional formula $x \rightarrow z \wedge y \rightarrow z$ denoting an information flow from variables x and y to variable z . The soundness of this analysis wrt. a semantic property such as noninterference requires more justification, though it is remarkable that the concretisation of propositional formula yields, roughly speaking, a set of program texts. ? also provides an abstract interpretation account of a flow-insensitive type system (?) enforcing noninterference by guaranteeing a stronger safety property, namely that sensitive locations should not influence public locations ?.

? explicitly formulate termination-insensitive noninterference as an abstract interpretation, namely the “merge over all twin computations” that makes explicit both the 2-safety aspect and the need for an analysis to relate some aligned intermediate states. Their analysis, like many others, is based on reducing the problem to a safety property of product programs. ? implement an algorithm that automates reasoning in a Hoare logic for k -safety, implicitly constructing product programs; the performance compares favorably with explicit

construction of product programs. Program dependency graphs are another approach to dependency, is shown to be correct for noninterference by ?, by way of slicing and a simulation argument.

?, Chap. 5 proposes the first quantitative measure of a program’s leakage in terms of Shannon entropy (?). Other quantitative metrics emerge in the literature (?????). These quantitative security metrics model different scenarios suitable for different policies. Most existing static analyses for quantitative information flow leverage existing model checking tools and abstract domains for safety; they prove that a program satisfies a quantitative security requirement by proving a stronger safety property. In contrast, the cardinal abstraction proves a hyperproperty by inferring a stronger hyperproperty satisfied by the analysed program. This is key to target quantitative information flow in multilevel security lattices, beyond the 2-point lattice $\{L, H\}$.

? synthesize equivalence classes induced by outputs over low equivalent memories by relying on software model checkers, in order to bound various quantitative metrics. ? also rely on a similar technique to quantify information flow for database queries. ? note that the exact computation of information-theoretic characteristics is prohibitively hard, and propose to rely on approximation-based analyses, among which are randomisation techniques and abstract interpretation ones. They also propose to rely on a self-composed product program to model a scenario where attackers may refine their knowledge by influencing the low inputs. ? relies on similar techniques to handle programs with low inputs, and uses polyhedra to synthesize linear constraints (?) over variables. ? decide whether answering a query on sensitive data augments attackers’ knowledge beyond a certain threshold, by using probabilistic polyhedra.

10. Conclusion

Galois connection-based semantic characterisations of program analyses provide new perspectives and insights that lead to improved techniques. We have extended the framework to fully encompass hyperproperties, through a remarkable form of hypercollecting semantics that enables calculational derivation of analyses. This new foundation raises questions too numerous to list here.

One promising direction is to combine dependency and cardinality analysis with existing abstract domains, e.g. through advanced symbolic methods (?), and partitioning (??).

Static analysis of secure information flow has yet to catch up with recent advances in dynamic information flow monitoring (????). We discussed, in ??, how existing static analyses may be of use to statically secure information flow. It seems likely that hypercollecting semantics will also be of use for dynamic analyses.

Acknowledgments

Thanks to Anindya Banerjee and the anonymous reviewers for thoughtful comments and helpful feedback. This work was partially supported by NSF awards CNS-1228930 and CCF-1649884, ANR project AnaStaSec ANR-14-CE28-0014 and a CFR CEA Phd Fellowship.