

1 Towards a unifying framework for tuning analysis 2 precision by program transformation

3 Mila Dalla Preda 

4 Dipartimento di Informatica, University of Verona, Italy

5 mila.dallapreda@univr.it

6 — Abstract —

7 Static and dynamic program analyses attempt to extract useful information on program's behaviours.
8 Static analysis uses an abstract model of programs to reason on their runtime behaviour without
9 actually running them, while dynamic analysis reasons on a test set of real program executions. For
10 this reason, the precision of static analysis is limited by the presence of false positives (executions
11 allowed by the abstract model that cannot happen at runtime), while the precision of dynamic
12 analysis is limited by the presence of false negatives (real executions that are not in the test set).
13 Researchers have developed many analysis techniques and tools in the attempt to increase the
14 precision of program verification. Software protection is an interesting scenario where programs need
15 to be protected from adversaries that use program analysis to understand their inner working and
16 then exploit this knowledge to perform some illicit actions. Program analysis plays a dual role in
17 program verification and software protection: in program verification we want the analysis to be as
18 precise as possible, while in software protection we want to degrade the results of analysis as much
19 as possible. Indeed, in software protection researchers usually recur to a special class of program
20 transformations, called code obfuscation, to modify a program in order to make it more difficult to
21 analyse while preserving its intended functionality. In this setting, it is interesting to study how
22 program transformations that preserve the intended behaviour of programs can affect the precision
23 of both static and dynamic analysis. While some works have been done in order to formalise the
24 efficiency of code obfuscation in degrading static analysis and in the possibility of transforming
25 programs in order to avoid or increase false positives, less attention has been posed to formalise the
26 relation between program transformations and false negatives in dynamic analysis. In this work we
27 are setting the scene for a formal investigation of the syntactic and semantic program features that
28 affect the presence of false negatives in dynamic analysis. We believe that this understanding would
29 be useful for improving the precision of existing dynamic analysis tools and in the design of program
30 transformations that complicate the dynamic analysis.

31 *To Maurizio on its 60th birthday!*

32
33 **2012 ACM Subject Classification** Security and privacy → Software and application security →
34 Software reverse engineering

35 **Keywords and phrases** Program analysis, analysis precision, program transformation, software
36 protection, code obfuscation

37 **Digital Object Identifier** 10.4230/OASICS.Gabbrielli.2020.4

38 **Funding** The research has been partially supported by the project “Dipartimenti di Eccellenza
39 2018-2020” funded by the Italian Ministry of Education, Universities and Research (MIUR).

40 **1** Introduction

41 Program analysis refers, in general, to any examination of programs that attempts to extract
42 useful information on program's behaviours (semantics). As known from the Rice theorem,
43 all nontrivial extensional properties of program's semantics are undecidable in the general
44 case. This means that any automated reasoning on software has to involve some kind of
45 approximation. Programs can be analysed either statically or dynamically. Static program



© Mila Dalla Preda;

licensed under Creative Commons License CC-BY

Recent Developments in the Design and Implementation of Programming Languages.

Editors: Frank S. de Boer and Jacopo Mauro; Article No. 4; pp. 4:1–4:22

OpenAccess Series in Informatics



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

46 analysis reasons about the behaviour of programs without actually running them. Typically,
47 static analysis builds an abstract model that over-approximates the possible program's
48 behaviours to examine program properties. This guarantees soundness: what can be derived
49 from the analysis of the abstract model holds also on the concrete execution of the program.
50 The converse does not hold in general due to the presence of *false positives*: spurious
51 behaviours allowed by the abstract model that do not correspond to any real program
52 execution. Static analysis has proved its usefulness in many fields of computer science like
53 in optimising compilers for producing efficient code, for automatic error detection and for
54 the automatic verification of desired program properties (e.g., functional properties and
55 security properties) [21]. Many different static analysis approaches exist, as for example
56 model checking [7], deductive verification [33] and abstract interpretation [12]. In particular,
57 abstract interpretation provides a formal framework for reasoning on behavioural program
58 properties where many static analysis techniques can be formalised. In the rest of this
59 paper we focus on those static analysis that can be formalised in the abstract interpretation
60 framework. Dynamic program analyses, such as program testing [1], runtime monitoring
61 and verification [4], consider an under-approximation of program behaviour as they focus
62 their analysis on a specific subset of possible program executions. In this paper when we
63 speak of dynamic analysis we mainly refer to program testing. Testing techniques start
64 by concretely executing programs on an input set and the so obtained test set of concrete
65 executions is inspected in order to reason on program's behaviour (e.g., reveal failures or
66 vulnerabilities). It is well known that dynamic analysis can precisely detect the presence of
67 failures but cannot guarantee their absence, due to the presence of *false negatives*: concrete
68 program behaviours that do not belong to the test set. There is a famous quote by Dijkstra
69 that states that "Program testing can be used to show the presence of bugs, but never to
70 show their absence!". Since it is not possible to guarantee the absence of failures we have
71 to accept the fact that whenever we use software we incur in some risk. Software testing
72 is widely used to reveal possible software failures, to reduce the risk related to the use of
73 software and to increase the quality of software by deciding if the behaviour of software is
74 acceptable in terms of reliability, safety, maintainability, security, and efficiency [1].

75 Static analysis computes an over-approximation of program semantics, while dynamic
76 analysis under-approximates program semantics. In both cases, we have a decidable evaluation
77 of the semantic property of interest on an approximation of program semantics. For this
78 reason what we can automatically conclude regarding the behavioural properties of programs
79 has to take into account false positives for static analysis and false negatives for dynamic
80 analysis. Static analysis is precise when it is *complete* (no false positives) and this relates to
81 the well studied notion of completeness in abstract interpretation [12, 14, 23]. The intuition
82 is that static analysis is complete when the details lost by the abstract model are not relevant
83 for reasoning on the semantic property of interest. Dynamic analysis is precise when it is
84 *sound* (no false negatives) and this happens when the executions in the test set exhibit all
85 the behaviours of the program that are relevant with respect to the semantic property of
86 interest. This means that the under-approximation of program semantics considered by the
87 dynamic analysis allows us to precisely observe the behavioural property of interest. The
88 essential problem with dynamic analysis is that it is impossible to test with all inputs since
89 the input space is generally infinite. In this context, *coverage criteria* provide structured,
90 practical ways to search the input space and to decide which input set to use. The rationale
91 behind coverage criteria is to partition the input space in order to maximise the executions
92 present in the tests set that are relevant for the analysis of the semantic property of interest.
93 Coverage criteria are useful in supporting the automatic generation of input sets and in


94 providing useful rules for deciding when to terminate the generation of the test set [1].

95 Program analysis has been originally developed for verifying the correctness of programs
96 and researchers have put a great deal of effort in developing efficient and precise analysis
97 techniques and tools that try to reduce false positives and false negatives as much as possible.
98 Indeed, analysis precision relates to the ability of identifying failures and vulnerabilities that
99 may lead to unexpected behaviours, or that may be exploited by an adversary for malicious
100 purposes. For this reason the main goal of researchers has been to improve the precision and
101 efficiency of both static and dynamic analysis tools.

102 Software protection is another interesting scenario where program analysis plays a central
103 role but in a dual way. Today, software and the assets embedded in it are constantly under
104 attack. This is particularly critical for those software applications that run in an untrusted
105 environment in a scenario known as MATE (Man-At-The-End) attacks. In this setting,
106 attackers have full control over, and white-box access to, the software and the systems on
107 which the software is running. Attackers can use a wide range of analysis tools such as
108 disassemblers, code browsers, debuggers, emulators, instrumentation tools, fuzzers, symbolic
109 execution engines, customised OS features, pattern matchers, etc. to inspect, analyse and
110 alter software and its assets. In such scenarios, software protection becomes increasingly
111 important to protect the assets, even against MATE attacks. For industry, in many cases the
112 deployment of software-based defense techniques is crucial for the survival of their businesses
113 and eco-systems. In the software protection scenario, program analysis can be used by
114 adversaries to reverse engineer proprietary code and then illicitly reuse portions of the code
115 or tamper with the code in some unauthorised way. Here, in order to protect the intellectual
116 property and integrity of programs we have to force the analysis to be imprecise or so
117 expensive to make it impractical for the adversary to mount an attack.

118 To address this problem, researchers have developed software-based defense techniques,
119 called *code obfuscations*, that transform programs with the explicit intent of complicating
120 and degrading program analysis [9]. The idea of code obfuscation techniques is to transform
121 a program into a functionally equivalent one that is more difficult (ideally impossible) for
122 an analyst to understand. As well as for program analysis also for code obfuscation we
123 have an important negative result from Barak et al. [3] that proves the impossibility of
124 code obfuscation. Note that, this result states the impossibility of an ideal obfuscator that
125 obfuscates every program by revealing only the properties that can be derived from its I/O
126 semantics. Besides the negative result of Barak et al., in recent decades, we have seen a
127 big effort in developing and implementing new and efficient obfuscation strategies [8]. Of
128 course, these obfuscating techniques introduce a kind of practical obfuscators weakening the
129 ideal obfuscator of Barak et al. in different ways, and which can be effectively used in real
130 application protection in the market. For example, these obfuscators may work only for a
131 certain class of programs, or may be able to hide only certain properties of programs (e.g.,
132 control flow). Indeed, the attention on code obfuscation poses the need to deeply understand
133 what we can obfuscate, namely which kind of program properties we can hide by inducing
134 imprecision in their automatic analysis.

135 A recent survey on the existing code obfuscation techniques shows the efficiency of
136 code obfuscation in degrading the results of static analysis, while existing code obfuscation
137 techniques turn out to be less effective against dynamic analysis [31]. Consider, for example,
138 the well known control flow obfuscation based on the insertion of opaque predicates. An
139 opaque predicate is a predicate whose constant value is known to the obfuscation, while it is
140 difficult for the analyst to recognise such constant value [9]. Consider the program whose
141 control flow graph is depicted on the left of Figure 1 where we have three blocks of sequential



■ **Figure 1** Code obfuscation

142 instructions A, B and C executed in the order specified by the arrows $A \rightarrow B \rightarrow C$. Let
 143 OP^T denote a true opaque predicate, namely a predicate that always evaluates to *true*. In
 144 the middle of Figure 1 we can see what happens to the control flow graph when we insert
 145 a true opaque predicate: block D has to be considered in the static analysis of the control
 146 flow even if it is never executed at runtime. Thus, $A \rightarrow OP^T \rightarrow D \rightarrow C$ is a false positive
 147 path added by the obfuscating transformation to the static analysis, while no imprecision is
 148 added to dynamic analysis since all real executions follow the path $A \rightarrow OP^T \rightarrow B \rightarrow C$.
 149 On the right of Figure 1 we have the control flow graph of the program obtained inserting
 150 an unknown opaque predicate. An unknown opaque predicate $OP^?$ is a predicate that
 151 sometimes evaluates to *true* and sometimes evaluates to *false*. These predicates are used
 152 to diversify program execution by inserting in the true and false branches sequences of
 153 instructions that are syntactically different but functionally equivalent (e.g. blocks B and
 154 B_1) [9]. Observe that this transformation adds confusion to dynamic analysis: a dynamic
 155 analyser has now to consider more execution traces in order to cover all the paths of the
 156 control flow graph. Indeed, if the dynamic analysis observes only traces that follow the
 157 original path $A \rightarrow OP^? \rightarrow B \rightarrow C$ it may not be sound as it misses the traces that follow
 158 $A \rightarrow OP^? \rightarrow B_1 \rightarrow C$ (false negative).

159 The abstract interpretation framework has been used to formalise, prove and compare
 160 the efficiency of code obfuscation techniques in confusing static analysis [17, 25] and to
 161 derive strategies for the design of obfuscating techniques that hamper a specific analysis
 162 [19]. The general idea is that code obfuscation confuses static analysis by exploiting its
 163 conservative nature, and by modifying programs in order to increase its imprecision (adding
 164 false positives) while preserving the program intended behaviour. Observe that, in general,
 165 the imprecision added by these obfuscating transformations to confuse a static analyser is
 166 not able to confuse a dynamic attacker that cannot be deceived by false positives. This is
 167 the reason why common deobfuscation approaches often recur to dynamic analysis to reverse
 168 engineer obfuscated code [5, 10, 32, 34].

169 It is clear that to complicate dynamic analysis we need to develop obfuscation techniques
 170 that exploit the Achilles heel of dynamic analysis, namely false negatives. In the literature,
 171 there are some defense techniques that focus on hampering dynamic analysis [2, 27, 28, 30].
 172 What is still missing is a general framework where it is possible to formalise, prove and discuss
 173 the efficiency of these transformations in complicating dynamic analysis in terms of the
 174 imprecision (false negatives) that they introduce. As discussed above the main challenge for
 175 dynamic analysis is the identification of a suitable input set for testing program's behaviour.
 176 In order to automatically build a suitable input set, the analysts either design an input
 177 generation tool or an input recogniser tool. In both cases, they need a coverage criterion
 178 that defines the inputs to be considered and when to terminate the definition of the input
 179 set. Ideally, the coverage criterion is chosen in order to guarantee that the test set precisely

180 reveals the semantic property under analysis (no false negatives). However, to the best of
 181 our knowledge, there is no formal guarantee that a coverage criterion ensures the absence
 182 of false negatives with respect to a certain analysis. If hampering static analysis means to
 183 increase the presence of false positives, hampering dynamic analysis means to complicate
 184 the automatic construction of a suitable input set for a given coverage criterion. In order to
 185 formally reason on the effects that code obfuscation has on the precision of dynamic analysis
 186 it is important to develop a general framework, analogous to the one based on program
 187 semantics and abstract interpretation that formalises the relation between dynamic analysis
 188 and code obfuscation. Thus, we need to develop a framework where we can (1) formally
 189 specify the relation between the coverage criterion used and the semantic property that we
 190 are testing, (2) define when a program transformation complicates the construction of an
 191 input set that has to satisfy a given coverage criterion, (3) derive guidelines for the design of
 192 obfuscating transformations that hamper the dynamic analysis of a given program property.
 193 This formal investigation will allow us to better understand the potential and limits of code
 194 obfuscation against dynamic program analysis.

195 In the following we provide a unifying view of static and dynamic program analysis and of
 196 the approaches that researchers use to tune the precision of these analysis. From this unifying
 197 overview it turns out that while the relation between the precision of static program analysis
 198 and program transformations has been widely studied, both in the software verification and
 199 in the software protection scenario, less attention has been posed to the formal investigation
 200 of the effects that code transformations have on the precision of program testing. We start
 201 to face this problem by showing how it is possible to formally compare and relate coverage
 202 criterion, semantic property under testing and false negatives for a specific class of program
 203 properties. This discussion leads us to the identification of important and interesting new
 204 research directions that would lead to the development of the above mentioned formal
 205 framework for reasoning about the effects of program transformations on the precision of
 206 dynamic analysis. We believe that this formal reasoning would find interesting applications
 207 both in the software verification and in the software protection scenario.

208 Structure of the paper: In Section 2 we provide some basic notions. In Section 3 we
 209 discuss possible techniques for improving the precision of the analysis: Section 3.1 revise the
 210 existing and ongoing work in transforming properties and programs toward completeness of
 211 static analysis, while Section 3.2 provides the basis for a formal framework for reasoning on
 212 possible property and program transformations to achieve soundness in dynamic analysis,
 213 these are preliminary results some of which have been recently published in [18]. Section 4
 214 shows how the techniques used to improve analysis precision could be used in the software
 215 protection scenario to prove the efficiency of software protection techniques. The use of this
 216 formal reasoning for proving the efficiency of software protection techniques against static
 217 analysis is known, while it is novel for dynamic analysis. The paper ends with a discussion
 218 on the open research challenges that follow from this work.

219 **2 Preliminaries**

220 Given two sets S and T , we denote with $\wp(S)$ the powerset of S , with $S \times T$ the Cartesian
 221 product of S and T , with $S \subset T$ strict inclusion, with $S \subseteq T$ inclusion, with $S \subseteq_F T$ the
 222 fact that S is a finite subset of T . $\langle C, \leq_C, \vee_C, \wedge_C, \top_C, \perp_C \rangle$ denotes a complete lattice on
 223 the set C , with ordering \leq_C , least upper bound (*lub*) \vee_C , greatest lower bound (*glb*) \wedge_C ,
 224 greatest element (top) \top_C , and least element (bottom) \perp_C (the subscript C is omitted when
 225 the domain is clear from the context). Let C and D be complete lattices. Then, $C \xrightarrow{m} D$ and

226 $C \xrightarrow{c} D$ denote, respectively, the set and the type of all monotone and (Scott-)continuous
 227 functions from C to D . Recall that $f \in C \xrightarrow{c} D$ if and only if f preserves *lub*'s of (nonempty)
 228 chains if and only if f preserves *lub*'s of directed subsets. Let $f : C \rightarrow C$ be a function on a
 229 complete lattice C , we denote with $lfp(f)$ the least fix-point, when it exists, of function f on
 230 C . The well-known Knaster-Tarski's theorem states that any monotone operator $f : C \xrightarrow{m} C$
 231 on a complete lattice C admits a least fix point. It is known that if $f : C \xrightarrow{c} C$ is continuous
 232 then $lfp(f) = \bigvee_{i \in \mathbb{N}} f^i(\perp_C)$, where, for any $i \in \mathbb{N}$ and $x \in C$, the i -th power of f in x is
 233 inductively defined as follows: $f^0(x) = x$; $f^{i+1}(x) = f(f^i(x))$.

Program Semantics: Let us consider the set *Prog* of possible programs and the set Σ of possible program states. A program state $s \in \Sigma$ provides a snapshot of the program and memory content during the execution of the program. Given a program P we denote $Init_P$ the set of its initial states. We use Σ^* to denote the set of all finite and infinite sequences or traces of states ranged over by σ . Given a trace $\sigma \in \Sigma^*$ we denote with $\sigma_0 \in \Sigma$ the first element of sequence σ and with σ_f the final state of σ if σ is finite. Let $\tau \subseteq \Sigma \times \Sigma$ denote the transition relation between program states, thus $(s, s') \in \tau$ means that state s' can be obtained from state s in one computational step. The *trace semantics* of a program P is defined, as usual, as the least fix-point computation of function $\mathcal{F}_P : \wp(\Sigma^*) \rightarrow \wp(\Sigma^*)$ [11]:

$$\mathcal{F}_P(X) \stackrel{\text{def}}{=} Init_P \cup \{ \sigma s_i s_{i+1} \mid (s_i, s_{i+1}) \in \tau, \sigma s_i \in X \}$$

234 The trace semantics of P is $\llbracket P \rrbracket \stackrel{\text{def}}{=} lfp(\mathcal{F}_P) = \bigcup_{i \in \mathbb{N}} \mathcal{F}_P^i(\perp_C)$. $Den\llbracket P \rrbracket$ denotes the denotational
 235 (finite) semantics of program P which abstracts away the history of the computation by
 236 observing only the input-output relation of finite traces: $Den\llbracket P \rrbracket \stackrel{\text{def}}{=} \{ \sigma \in \Sigma^+ \mid \exists \eta \in \llbracket P \rrbracket : \eta_0 = \sigma_0, \eta_f = \sigma_f \}$.

237 Concrete domains are collections of computational objects where the concrete semantics
 238 is computed, while abstract domains are collections of approximate objects, representing
 239 properties of concrete objects in a domain-like structure. It is possible to interpret the
 240 semantics of programs on abstract domains thus approximating the computation with respect
 241 to the property expressed by the abstract domain. The relation between concrete and
 242 abstract domains can be equivalently specified in terms of Galois connections (GC) or upper
 243 closure operators in the abstract interpretation framework [12, 13]. The two approaches
 244 are equivalent, modulo isomorphic representations of the domain object. A GC is a tuple
 245 (C, α, γ, A) where C is the concrete domain, A is the abstract domain and $\alpha : C \rightarrow A$
 246 and $\gamma : A \rightarrow C$ are respectively the abstraction and concretisation maps that give rise
 247 to an adjunction: $\forall a \in A, c \in C : \alpha(c) \leq_A a \Leftrightarrow c \leq_C \gamma(a)$. Abstract domains can be
 248 compared with respect to their relative degree of precision: if A_1 and A_2 are abstractions
 249 of a common concrete domain C , A_1 is more precise than A_2 , denoted $A_1 \sqsubseteq A_2$ when
 250 $\forall a_2 \in A_2, \exists a_1 \in A_1 : \gamma_1(a_1) = \gamma_2(a_2)$, namely if $\gamma_2(A) \subseteq \gamma_1(A)$. An upper closure operator
 251 on a complete lattice C is an operator $\rho : C \rightarrow C$ that is monotone, idempotent, and
 252 extensive ($\forall x \in C : x \leq_C \rho(x)$). Closures are uniquely determined by their fix-points $\rho(C)$.
 253 If (C, α, γ, A) is a GC then $\rho = \gamma \circ \alpha$ is the closure associated to A , such that $\rho(C)$ is a
 254 complete lattice isomorphic to A . The closure $\gamma \circ \alpha$ associated to the abstract domain A can
 255 be thought of as the logical meaning of A in C , since this is shared by any other abstract
 256 representation for the objects of A . Thus, the closure operator approach is convenient when
 257 reasoning about properties of abstract domains independently from the representation of
 258 their objects. We denote with $uco(C)$ the set of upper closure operators over C . If C is a
 259 complete lattice then $uco(C)$ is a complete lattice where closure are ordered with respect
 260 to their relative precision $\rho_1 \sqsubseteq \rho_2 \Leftrightarrow \rho_2(C) \subseteq \rho_1(C)$ which corresponds to the ordering of
 261 abstract domains.
 262

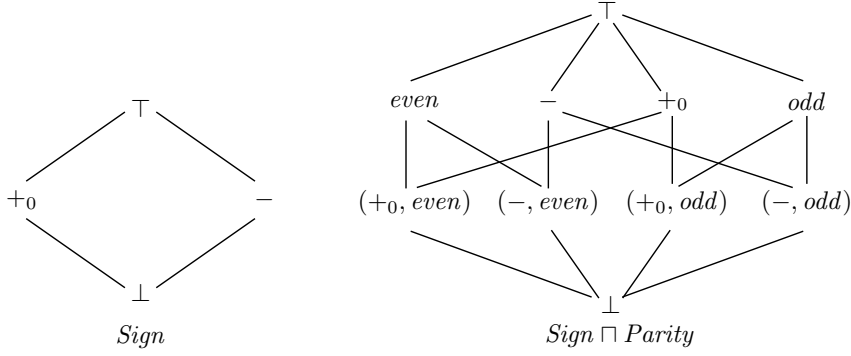
263 The abstract semantics of a program P on the abstract domain $\rho \in uco(\wp(\Sigma^*))$, denoted
 264 as $\llbracket P \rrbracket^\rho$, is defined as the fix-point computation of function $\mathcal{F}_P^\rho : \rho(\wp(\Sigma^*)) \rightarrow \rho(\wp(\Sigma^*))$ where
 265 $\mathcal{F}_P^\rho \stackrel{\text{def}}{=} \rho \circ \mathcal{F}_P \circ \rho$ is the best correct approximation of function \mathcal{F}_P on the abstract domain
 266 $\rho(\wp(\Sigma^*))$, namely $\llbracket P \rrbracket^\rho \stackrel{\text{def}}{=} \text{lfp}(\mathcal{F}_P^\rho) = \bigcup_{i \in \mathbb{N}} \mathcal{F}_P^\rho(\perp_{\rho(C)})$. Given the equivalence between GC
 267 and closures, the abstract semantics can be equivalently specified in terms of abstract traces
 268 in the corresponding abstract domain and in the following we denote the abstract semantics
 269 either with $\llbracket P \rrbracket^\rho$ or with $\llbracket P \rrbracket^A$ where (C, α, γ, A) is a GC and $\rho = \gamma \circ \alpha$.

270 *Equivalence Relations:* Let \mathcal{R} be a binary relation $\mathcal{R} \subseteq C \times C$ on a set C , given $x, y \in C$
 271 we denote with $(x, y) \in \mathcal{R}$ the fact that x is in relation \mathcal{R} with y . $\mathcal{R} \subseteq C \times C$, is an *equivalence*
 272 *relation* if \mathcal{R} is reflexive $\forall x \in C : (x, x) \in \mathcal{R}$, symmetric $\forall x, y \in C : (x, y) \in \mathcal{R} \Rightarrow (y, x) \in \mathcal{R}$
 273 and transitive $\forall x, y, z \in C : (x, y) \in \mathcal{R} \wedge (y, z) \in \mathcal{R} \Rightarrow (x, z) \in \mathcal{R}$. Given a set C equipped with
 274 an equivalence relation \mathcal{R} , we consider for each element $x \in C$ the subset $[x]_{\mathcal{R}}$ of C containing
 275 all the elements of C in equivalence relation with x , i.e., $[x]_{\mathcal{R}} = \{y \in C \mid (x, y) \in \mathcal{R}\}$. The sets
 276 $[x]_{\mathcal{R}}$ are called equivalence classes of C wrt relation \mathcal{R} and they induce a partition of the set C ,
 277 namely $\forall x, y \in C : [x]_{\mathcal{R}} = [y]_{\mathcal{R}} \vee [x]_{\mathcal{R}} \cap [y]_{\mathcal{R}} = \emptyset$ and $\cup\{[x]_{\mathcal{R}} \mid x \in C\} = C$. The partition of
 278 C induced by the relation \mathcal{R} is denoted by C/\mathcal{R} . Let $Eq(C)$ be the set of equivalence relations
 279 on the set C . The set of equivalence relations on C form a lattice $\langle Eq(C), \preceq, \sqcap_{Eq}, \sqcup_{Eq}, id, top \rangle$
 280 where id is the relation that distinguishes all the elements in C , top is the relation that cannot
 281 distinguish any element in C , and: $\mathcal{R}_1 \preceq \mathcal{R}_2$ iff $\mathcal{R}_1 \subseteq \mathcal{R}_2$ iff $(x, y) \in \mathcal{R}_1 \Rightarrow (x, y) \in \mathcal{R}_2$,
 282 $\mathcal{R}_1 \sqcap_{Eq} \mathcal{R}_2 = \mathcal{R}_1 \cap \mathcal{R}_2$, namely $(x, y) \in \mathcal{R}_1 \sqcap_{Eq} \mathcal{R}_2$ iff $(x, y) \in \mathcal{R}_1 \wedge (x, y) \in \mathcal{R}_2$; $\mathcal{R}_1 \sqcup_{Eq} \mathcal{R}_2$
 283 it is such that $(x, y) \in \mathcal{R}_1 \sqcup_{Eq} \mathcal{R}_2$ iff $(x, y) \in \mathcal{R}_1 \vee (x, y) \in \mathcal{R}_2$. When $\mathcal{R}_1 \preceq \mathcal{R}_2$ we say that
 284 \mathcal{R}_1 is a refinement of \mathcal{R}_2 . Given a subset $S \subseteq C$, we denote with $\mathcal{R}|_S \in Eq(S)$ the restriction
 285 of relation \mathcal{R} to the domain S .

286 The relation between closure operators and equivalence relations has been studied in
 287 [29]. Each closure operator $\rho \in uco(\wp(C))$ induces an equivalence relation $\mathcal{R}^\rho \in Eq(C)$
 288 where $(x, y) \in \mathcal{R}^\rho$ iff $\rho(\{x\}) = \rho(\{y\})$ and viceversa, each equivalence relation $\mathcal{R} \in Eq(C)$
 289 induces a closure operator $\rho^{\mathcal{R}} \in uco(\wp(C))$ where $\rho^{\mathcal{R}}(\{x\}) = [x]_{\mathcal{R}}$ and $\rho^{\mathcal{R}}(X) = \bigcup_{x \in X} [x]_{\mathcal{R}}$.
 290 Of course, there are many closures that induce the same partition on traces and these
 291 closures carry additional information other than the underlying state partition, and this
 292 additional information that allows us to distinguish them is lost when looking at the induced
 293 partition. Indeed, it holds that given $\mathcal{R} \in Eq(C)$ the corresponding closure is such that
 294 $\rho^{\mathcal{R}} = \sqcap\{\rho \mid \mathcal{R}^\rho = \mathcal{R}\}$. The closures in $uco(\wp(C))$ defined form a partition $\mathcal{R} \in Eq(C)$
 295 are called *partitioning* and they identify a subset of $uco(\wp(C))$: $\{\rho^{\mathcal{R}} \in uco(\wp(C)) \mid \mathcal{R} \in$
 296 $Eq(C)\} \subseteq uco(\wp(C))$ [29].

297 **3 On the precision of program analysis**

298 As argued above program analysis has been originally developed for program verification,
 299 namely to ensure that programs will actually behave as expected. Besides the impossibility
 300 result of the Rice theorem, a multitude of analysis strategies have been proposed [21]. Indeed,
 301 by tuning the precision of the behavioural feature that we want to analyse it is possible
 302 to derive an analysable semantic property that, while loosing some details of program's
 303 behaviour, may still be of practical interest [12, 14]. We are interested in semantic program
 304 properties, namely in properties that deal with the behaviour of programs, but the possibility
 305 of precisely analysing such properties depends also on the way in which programs are written.
 306 This means that there are programs that are easier to analyse than others with respect to a
 307 certain property [6]. Thus, program transformations that preserve the program's intended
 308 functionality can affect the precision of the results of the same analysis on the original and



■ **Figure 2** Abstract domain of *Sign* and *Sign* \sqcap *Parity*

309 transformed program.

310 3.1 Static Analysis

311 Precision in static program analysis means completeness, namely absence of false positives.
 312 This means that the noise introduced by the abstract model used for static program analysis
 313 does not introduce imprecision with respect to the property under analysis. Consider for
 314 example program P on the left of Figure 3 that, given an integer value a , returns its absolute
 315 value and it does it by adding some extra controls on the parity of variable a that have no
 316 effect on the result of computation¹. The semantics of program P is:

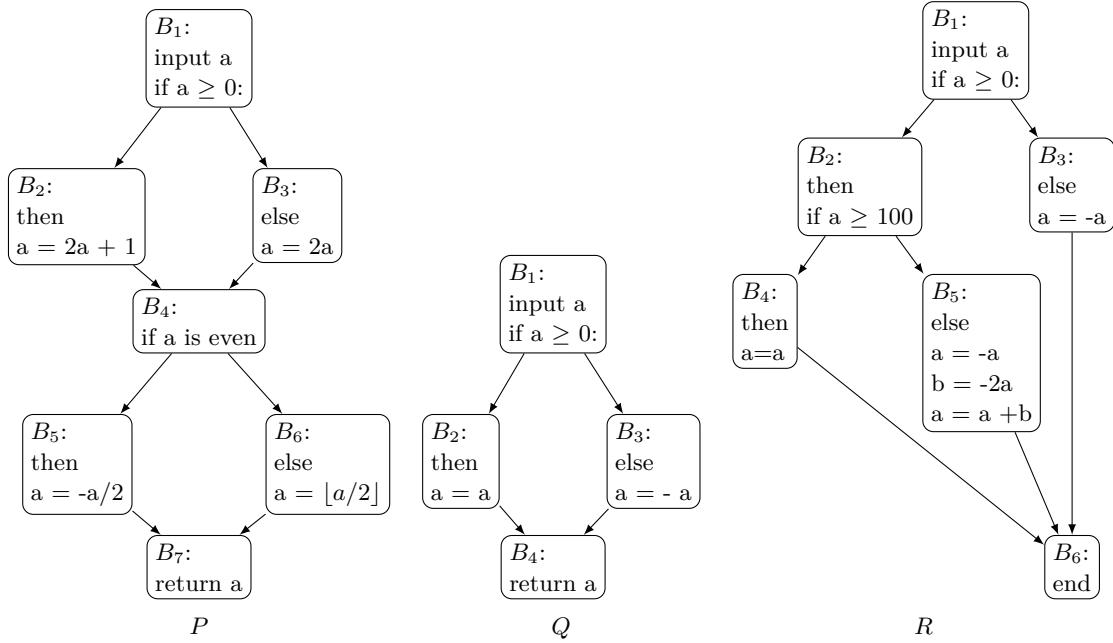
$$\begin{aligned}
 317 \quad \llbracket P \rrbracket &= \{ \langle B_1 : \perp \rangle \langle B_2 : v_1 \rangle \langle B_4 : 2 * v_1 + 1 \rangle \langle B_6 : 2 * v_1 + 1 \rangle \langle B_7 : v_1 \rangle \mid v_1 \geq 0 \} \cup \\
 318 &\quad \{ \langle B_1 : \perp \rangle \langle B_3 : v_1 \rangle \langle B_4 : 2 * v_1 \rangle \langle B_5 : 2 * v_1 \rangle \langle B_7 : -v_1 \rangle \mid v_1 < 0 \}
 \end{aligned}$$

319 where $\langle B_i, val \rangle$ denotes the program state specifying the value val of variable a when entering
 320 block B_i and \perp denotes the undefined value. Assume that we are interested in the analysis on
 321 the abstract domain *Sign* depicted on the left of Figure 2. The $Sign = \{\perp, +0, -, \top\}$ abstract
 322 domain observes the sign of integer values and it is possible to define a GC between $\wp(\mathbb{Z})$
 323 and *Sign* where the abstract element $+0$ represents all positive values plus 0, the abstract
 324 element $-$ represents all negative values, while \top represents all integer values and \perp the
 325 emptyset. We denote with $\llbracket P \rrbracket^{Sign} \in \wp(\Sigma^*)$ the abstract interpretation of program P on the
 326 domain of *Sign*, where the values of variable a are interpreted on *Sign*.

$$\begin{aligned}
 327 \quad \llbracket P \rrbracket^{Sign} &= \{ \langle B_1 : \perp \rangle \langle B_2 : +0 \rangle \langle B_4 : +0 \rangle \langle B_6 : +0 \rangle \langle B_7 : +0 \rangle, \\
 328 &\quad \langle B_1 : \perp \rangle \langle B_2 : +0 \rangle \langle B_4 : +0 \rangle \langle B_5 : +0 \rangle \langle B_7 : - \rangle [false\ positive] \\
 329 &\quad \langle B_1 : \perp \rangle \langle B_3 : - \rangle \langle B_4 : - \rangle \langle B_5 : - \rangle, \langle B_7 : +0 \rangle \\
 330 &\quad \langle B_1 : \perp \rangle \langle B_3 : - \rangle \langle B_4 : - \rangle \langle B_6 : - \rangle \langle B_7 : - \rangle [false\ positive] \}
 \end{aligned}$$

331 Each abstract trace corresponds to infinitely many concrete traces. So for example the
 332 abstract trace $\langle B_1 : \perp \rangle \langle B_2 : +0 \rangle \langle B_4 : +0 \rangle \langle B_6 : +0 \rangle \langle B_7 : +0 \rangle$ corresponds to the infinite set of
 333 concrete traces: $\{ \langle B_1 : \perp \rangle \langle B_2 : v_1 \rangle \langle B_4 : v_2 \rangle \langle B_6 : v_3 \rangle \langle B_7 : v_3 \rangle \mid v_1, v_2, v_3, v_4 \geq 0 \}$. Observe
 334 that the second and fourth abstract traces are false positives that the abstract analysis has
 335 to consider but that cannot happen during computation. This is because the guard at B_4

¹ The notation $\lfloor a/2 \rfloor$ refers to the integer division that rounds the non-integer results towards the lower integer value.



■ **Figure 3** P , Q and R are functionally equivalent programs

336 cannot be precisely evaluated on $Sign$ and therefore both branches are seen as possible. This
 337 happens because the abstract domain of $Sign$ is not complete for the analysis of program P
 338 and we have $\llbracket P \rrbracket \subset \llbracket P \rrbracket^{Sign}$. This induces imprecision in the analysis on the abstract domain
 339 $Sign$ that it is not able to conclude that the value of variable a is always positive at the end
 340 of execution. Let us denote with $\llbracket P \rrbracket(B_i)$ and with $\llbracket P \rrbracket^{Sign}(B_i)$ the possible values that can
 341 be assumed by variable a at block B_i when reasoning on the concrete and abstract semantics
 342 respectively. In this case we have that $Sign(\llbracket P \rrbracket(B_7)) = Sign(\{v \mid v \geq 0\}) = +_0$ and this is
 343 more precise than $\llbracket P \rrbracket^{Sign}(B_7) = \sqcup_{Sign} \{+_0, -\} = \top$.

344 Transforming properties towards completeness

345 It is well known that completeness is a domain property and that abstract domains can be
 346 refined in order to become complete for the analysis of a given program [23]. The idea is
 347 that in order to make the analysis complete we need to add to the abstract domain those
 348 elements that are necessary to reach completeness. In this case, if we consider the abstract
 349 domain that observes the sign and parity of integer values we reach completeness. Thus, let
 350 us consider the domain $Sign \sqcap Parity$ depicted on the right of Figure 2, where $even$ represents
 351 all the even integer values and odd represents all the odd integer values.

$$\begin{aligned}
 \llbracket P \rrbracket^{Sign \sqcap Parity} &= \{ \langle B_1 : \perp \rangle \langle B_2 : (+_0, even) \rangle \langle B_4 : (+_0, odd) \rangle \langle B_6 : (+_0, odd) \rangle \langle B_7 : +_0 \rangle \\
 &\quad \langle B_1 : (+_0, \perp) \rangle \langle B_2 : (+_0, odd) \rangle \langle B_4 : (+_0, odd) \rangle \langle B_6 : (+_0, odd) \rangle \langle B_7 : +_0 \rangle \\
 &\quad \langle B_1 : (-, \perp) \rangle \langle B_3 : (-, even) \rangle \langle B_4 : (-, even) \rangle \langle B_5 : (-, even) \rangle \langle B_7 : +_0 \rangle \\
 &\quad \langle B_1 : (-, \perp) \rangle \langle B_3 : (-, odd) \rangle \langle B_4 : (-, even) \rangle \langle B_5 : (-, even) \rangle \langle B_7 : +_0 \rangle \}
 \end{aligned}$$

356 As we can see all the abstract traces are able to precisely observe that variable a is positive at
 357 the end of the execution and that it can be either even or odd. Indeed, we have completeness
 358 with respect to the $Sign \sqcap Parity$ property $Sign \sqcap Parity(\llbracket P \rrbracket(B_7)) = \llbracket P \rrbracket^{Sign \sqcap Parity}(B_7) = +_0$.

359 Thus, a possible way for tuning the precision of static analysis is to transform the property
 360 that we want to analyse in order to reach completeness, there exists a systematic methodology
 361 that allows us to add the minimal amount of elements to the abstract domain in order to
 362 make the analysis complete for a given program [23].

363 Transforming programs towards completeness

The way in which programs are written affects the precision of the analysis. For example we can easily write a program functionally equivalent to P but for which the analysis on $Sign$ is complete. Consider, for example, program Q as the one in the middle of Figure 3, we have that:

$$\begin{aligned} \llbracket Q \rrbracket &= \{ \langle B_1 : \perp \rangle \langle B_2 : v \rangle \langle B_4 : v \rangle \mid v \geq 0 \} \cup \{ \langle B_1 : \perp \rangle \langle B_3 : v \rangle \langle B_4 : -v \rangle \mid v < 0 \} \\ \llbracket Q \rrbracket^{Sign} &= \{ \langle B_1 : \perp \rangle \langle B_2 : +0 \rangle \langle B_4 : +0 \rangle, \langle B_1 : \perp \rangle \langle B_3 : - \rangle \langle B_4 : +0 \rangle \} \end{aligned}$$

364 This makes it clear how the abstract computation loses information regarding the modulo of
 365 the value of variable a , while it precisely observes the positive value of a at the end of execution.
 366 Indeed, in this case we have that: $Sign(\llbracket Q \rrbracket(B_7)) = Sign(\{v \mid v \geq 0\}) = +0 = \llbracket Q \rrbracket^{Sign}(B_7)$.
 367 It is worth studying the possibility of transforming programs in order to make a certain
 368 analysis complete. In a recent work [6] the authors introduced the notions of complete
 369 clique $\mathbb{C}(P, A)$ and incomplete clique $\bar{\mathbb{C}}(P, A)$ that represent the set of all programs that
 370 are functionally equivalent to P and for which the analysis on the abstract domain A is
 371 respectively complete and incomplete. They prove that there are infinitely many abstractions
 372 for which the systematic removal of false positives for all programs is impossible. Moreover,
 373 they observe that false positives are related to the evaluation of boolean predicates that the
 374 abstract domain is not able to evaluate precisely (as we have seen in our earlier example). The
 375 authors claim that their investigation together with the poof system in [24] should be used
 376 as a starting point to reason on a code refactoring strategy that aims at modifying a given
 377 program in order to gain precision with respect to a predefined analysis. Given an abstract
 378 domain A , the final goal would be to derive a program transformation $\mathcal{T}_A : Prog \rightarrow Prog$ that
 379 given a program $P \in \bar{\mathbb{C}}(P, A)$ for which the analysis \mathcal{A} is incomplete, namely $\mathcal{A}(\llbracket P \rrbracket) \neq \llbracket P \rrbracket^A$,
 380 transforms it into a program $\mathcal{T}(P) \in \mathbb{C}(P, A)$ for which the analysis is complete, namely
 381 $\mathcal{A}(\llbracket P \rrbracket) = \llbracket P \rrbracket^A$.

382 These recent promising works suggest how to proceed in the investigation of program
 383 transformations as a mean for gaining precision in static program analysis.

384 3.2 Dynamic Analysis

385 Testing is typically used to discover failures (or bugs), namely an incorrect program behaviour
 386 with respect to the requirements or the description of the expected program behaviour.
 387 Precision in program testing is expressed in terms of soundness: the ideal situation where no
 388 false negatives are present. When speaking of failures, this happens when the executions
 389 considered in the test set exhibit at least one behaviour for each one of the failures present
 390 in the program. Indeed, when this happens, testing allows us to detect all the failures in the
 391 program. It is clear that the choice of the input set to use for testing is fundamental in order
 392 to minimise the number of false negatives. What we have just said holds when testing aims
 393 at detecting failures as well as for the analysis of any property of traces (as for example the
 394 order in which memory cells are accessed, the target of jumps, etc.). Let us denote with \mathbb{I}_P
 395 the input space of the possible input values needed to complete an execution of program

396 P under testing². Dynamic analysis considers a finite subset of the input space, called the
 397 input set $InSet \subseteq_F \mathbb{I}_P$, that identifies the input values that are used for execution. The
 398 execution traces generated by the input set define the test set, which is the finite set of traces
 399 used by dynamic analysis to reason on program behaviour. Given an input value $x \in \mathbb{I}_P$ we
 400 denote with $P(x) \in \llbracket P \rrbracket$ the execution of program P when fed with input x .

401 As argued above, the main source of imprecision in testing is that the number of potential
 402 inputs for most programs is so large as to be effectively infinite. Since we cannot test with
 403 all inputs, researchers typically recur to the use of coverage criteria in order to decide which
 404 test inputs to use. A coverage criterion C induces a partition on the input space and in order
 405 to minimise the false negatives the input set should contain at least one element for each
 406 class of the partition. In the left part of Figure 4 we consider a typical coverage criterion,
 407 called path coverage, for the testing of program Q in Figure 3. Path coverage criterion is
 408 satisfied when for each path in the control flow graph of the program there exists at least one
 409 execution in the test set that follows that path. When considering program Q it is immediate
 410 to derive from the coverage criterion the partition of the input space: the class of positive
 411 integer values (that follow the path $B_1 \rightarrow B_2 \rightarrow B_4$) and the class of negatives integer values
 412 (that follow the path $B_1 \rightarrow B_3 \rightarrow B_4$). In this case the coverage criterion is satisfied by every
 413 input set that contains at least one positive integer value and one negative integer value.

414 Since it is the coverage criterion that determines the input set and therefore the executions
 415 that are considered by the dynamic analysis, it is very important to select a *good* coverage
 416 criterion. However, it is not clearly stated or formally defined what makes a coverage criterion
 417 good [1], and this may be one of the reasons why many coverage criteria have been developed
 418 by researchers. Generally speaking, there are some features that it is important to consider
 419 when speaking of coverage criterion such as:


- 420 ■ the difficulty of deriving the rules to partition the input space with respect to the coverage
 421 criterion;
- 422 ■ the difficulty of generating an input set that satisfies the coverage criterion, namely that
 423 contains at least one input for each one of the classes in which the input space has been
 424 partitioned;
- 425 ■ how well a test set that satisfies the coverage criterion guarantees the absence of false
 426 negatives.

427 To the best of our knowledge there is no general framework that formalises the relation
 428 between coverage criterion, partition of the input space and false negatives in the dynamic
 429 analysis of a semantic program property. Indeed, while the soundness of dynamic analysis
 430 may not be possible in general, we think that it would be interesting to study the soundness
 431 of dynamic analysis of a program with respect to a specific semantic property (as usually
 432 done when reasoning about completeness in static analysis). We believe that this formal
 433 investigation would help in better understanding the cause of false negatives and would be
 434 useful in reducing them.

435 3.2.1 Towards a formal framework for dynamic analysis

436 We formalise the splitting of the input space induced by a coverage criterion C in terms of
 437 an equivalence relation $\mathcal{R}_I^C \in Eq(\mathbb{I}_P)$, and this allows us to formally define when an input
 438 set satisfies a coverage criterion.

² In this work, for simplicity but with no loss of generality, we speak of input values while in the general case we may need collections of values in order to complete an execution of the software under test.



■ **Figure 4** Path coverage criterion on program Q of Figure 3, and soundness conditions

439 ► **Definition 1.** Given a program P , an input set $InSet \subseteq_F \mathbb{I}_P$ and a coverage criterion
 440 C , we say that $InSet$ satisfies C , denoted $InSet \models C$, iff: $\forall [x]_{\mathcal{R}_I^C} \in \mathbb{I}_P / \mathcal{R}_I^C$ we have that
 441 $InSet \cap [x]_{\mathcal{R}_I^C} \neq \emptyset$.

442 We have seen this in Figure 4 when considering the partition induced in the input space
 443 of program Q and observing that an input set satisfies the path coverage criterion when it
 444 contains at least one positive and one negative integer value. When considering coverage
 445 criteria we need to take into account infeasible requirements: for example when considering
 446 coverage criteria related to the paths of the control flow graph we have to handle infeasible
 447 paths as it is not possible to define input values that follow these paths (as for example paths
 448 $B_1 \rightarrow B_2 \rightarrow B_4 \rightarrow B_5 \rightarrow B_7$ and $B_1 \rightarrow B_3 \rightarrow B_4 \rightarrow B_5 \rightarrow B_7$ of program P). This is a
 449 known challenging problem in dynamic analysis and testing as the detection of infeasible test
 450 requirements is undecidable for most coverage criteria [1]. This means that some preliminary
 451 analysis is needed in order to ensure the feasibility of the coverage criteria, namely to ensure
 452 that it is possible to generate an input set that satisfies a given coverage criterion. Otherwise,
 453 we need to somehow quantify how much the input set satisfies the coverage criterion, for
 454 example considering the percentage of equivalence classes that are covered by the input set.
 455 In this work we do not address this problem and we assume the feasibility of the coverage
 456 criteria.

457 Observe that the equivalence relation $\mathcal{R}_I^C \in Eq(\mathbb{I}_P)$ naturally induces an equivalence
 458 relation on traces $\mathcal{R}^C \in Eq(\llbracket P \rrbracket)$ where $(\sigma_1, \sigma_2) \in \mathcal{R}^C$ iff $\exists x_1, x_2 \in \mathbb{I}_P : P(x_1) = \sigma_1$,
 459 $P(x_2) = \sigma_2$ and $(x_1, x_2) \in \mathcal{R}_I^C$. Thus, we can say that a given coverage criterion, and
 460 therefore any test set that satisfies that coverage criterion, can be associated to a partition
 461 of program trace semantics. Our idea is that the partition of the program trace semantics
 462 induced by the coverage criterion could be used to reason on the class of semantic program
 463 properties for which the coverage criterion can ensure soundness. To this end, we need to
 464 represent semantic program properties in a way that can be compared with partitions on
 465 traces.

466 Properties of traces are naturally modelled as abstract domains, namely as closure
 467 operators in $uco(\wp(\Sigma^*))$. A semantic property $\rho \in uco(\wp(\Sigma^*))$ maps an execution trace
 468 (or a set of execution traces) to the minimal set of traces that cannot be distinguished
 469 by the considered property. Each closure operator $\rho \in uco(\wp(\Sigma^*))$ induces an equivalence
 470 relation $\mathcal{R}^\rho \in Eq(\Sigma^*)$: $\sigma_1 \mathcal{R}^\rho \sigma_2$ iff $\rho(\{\sigma_1\}) = \rho(\{\sigma_2\})$, where traces are grouped together
 471 if they are mapped in the same element by abstraction ρ . In the following, we model the
 472 properties of traces as equivalence relations over traces or equivalently as partitioning closures
 473 in $uco(\wp(\Sigma^*))$, and we denote these properties as $\mathcal{A} \in Eq(\Sigma^*)$. According to [29] there is
 474 more than one closure that maps to the same equivalence relations, thus considering the
 475 partitions induced by closure operators corresponds to focusing on the set of partitioning
 476 closures (which is a proper subset of closure operators over $\wp(\Sigma^*)$). This allows us to express

477 properties of the single traces but not relational properties that have to take into account
 478 more than one trace. This means that we can use equivalence relations in $Eq(\Sigma^*)$ to express
 479 properties such as: the order of successive accesses to memory, the order of execution of
 480 instructions, the location of the first instruction of a function, the target of jumps, function
 481 location, possible data values at certain program points, the presence of a bad states in the
 482 trace, and so on. These are properties of practical interest in dynamic program analysis.
 483 What we cannot express are properties on sets of traces, the so called hyper-properties,
 484 that express relational properties among traces, like non-interference. The extension of the
 485 framework to closures that are not partitioning is left as future work. This allows us to
 486 formally model the soundness of dynamic analysis.

487 ► **Definition 2.** *Given a program P and a property $\mathcal{A} \in Eq(\Sigma^*)$, the dynamic analysis \mathcal{A}
 488 on input set $InSet \subseteq_F \mathbb{I}_P$ is sound, denoted $InSet \overset{s}{\rightsquigarrow} \mathcal{A}(P)$, if $\forall [\sigma]_{\mathcal{A}} \in \llbracket P \rrbracket / \mathcal{A}$ we have that
 489 $[\sigma]_{\mathcal{A}} \cap InSet \neq \emptyset$.*

490 This precisely captures the fact that dynamic analysis needs to observe the different behaviours
 491 of the program with respect to the property of interest in order to be sound. Indeed, when
 492 considering a program P and a property \mathcal{A} it is enough to observe a single trace in an
 493 equivalence class $[\sigma]_{\mathcal{A}} \subseteq \llbracket P \rrbracket$ in order to observe how property \mathcal{A} behaves in all the traces of
 494 program P that belong to that equivalence class. If we consider program Q in Figure 3 we
 495 have that in order to precisely observe the evolution of the sign property along the execution
 496 we have to consider at least one trace that follows the path $B_1 \rightarrow B_2 \rightarrow B_4$ and one trace
 497 that follows the path $B_1 \rightarrow B_3 \rightarrow B_4$ as depicted in Figure 4.

498 Modelling program properties as equivalence relations makes it easy to compare them
 499 with the coverage criteria and to reason on soundness.


500 ► **Theorem 3.** *Given a program P , a coverage criterion C , an input set $InSet \subseteq_F \mathbb{I}_P$ and a
 501 property $\mathcal{A} \in Eq(\Sigma^*)$, we have that if $\mathcal{R}^C \preceq \mathcal{R}_{\llbracket P \rrbracket}^{\mathcal{A}}$ and $InSet \models C$, then $InSet \overset{s}{\rightsquigarrow} \mathcal{A}(P)$.*

502 **Proof.** $InSet \models C$ therefore $\forall [x]_{\mathcal{R}^C} \in \mathbb{I}_P / \mathcal{R}^C$ we have that $InSet \cap [x]_{\mathcal{R}^C} \neq \emptyset$. Since
 503 $\mathcal{R}^C \preceq \mathcal{A}_{\llbracket P \rrbracket}$ we have that for each equivalence class $[\sigma]_{\mathcal{R}^C}$ there exists an equivalence class
 504 $[\sigma]_{\mathcal{A}_{\llbracket P \rrbracket}}$ that $[\sigma]_{\mathcal{R}^C} \subseteq [\sigma]_{\mathcal{A}_{\llbracket P \rrbracket}}$. This implies that for every $[\sigma]_{\mathcal{A}_{\llbracket P \rrbracket}} \in \llbracket P \rrbracket / \mathcal{A}$ we have that
 505 $[\sigma]_{\mathcal{A}_{\llbracket P \rrbracket}} \cap InSet \neq \emptyset$ and therefore $InSet \overset{s}{\rightsquigarrow} \mathcal{A}(P)$. ◀

506 In Figure 4 on the right we provide a graphical representation of the above theorem. Traces
 507 in Σ^* exhibit different attributes with respect to property \mathcal{A} and this is represented by the
 508 different shapes: circle, triangle, square and star. Trace partition is then represented by the
 509 thick lines that group together traces that are undistinguishable with respect to property
 510 \mathcal{A} . Dotted lines are used to represent a trace partition induced by coverage criterion C
 511 on the traces of P and that ensures the absence of false negatives in the analysis. Indeed, from
 512 the graphical representation it is clear that when $InSet \models C$ then $InSet$ contains at least a
 513 trace for each equivalence class of \mathcal{R}^C , and this implies that it contains at least a trace for
 514 each one of the possible attributes (circle, triangle and square) that traces in $\llbracket P \rrbracket$ can exhibit
 515 with respect to property \mathcal{A} . This allows us to characterise the set of properties for which a
 516 given coverage criterion can ensure soundness.

517 ► **Definition 4.** *Given a coverage criterion C on a program P , we define the set of properties
 518 $\Pi(C) \stackrel{def}{=} \{\mathcal{A} \in Eq(\Sigma^*) \mid \mathcal{R}^C \preceq \mathcal{A}_{\llbracket P \rrbracket}\}$ that are coarsest than the equivalence relation induced
 519 by the coverage criterion.*

520 It follows that any input set that satisfies a coverage criterion C on a program P would lead
 521 to a sound dynamic analysis on any property in $\Pi(C)$.



■ **Figure 5** Comparing \mathcal{R}^C and \mathcal{A} for soundness

522 ► **Corollary 5.** *Given a coverage criterion C on a program P , and input set $InSet \subseteq_F \mathbb{I}_P$*
 523 *such that $InSet \models C$, then $\forall \mathcal{A} \in \Pi(C)$ we have that $InSet \overset{s}{\rightsquigarrow} \mathcal{A}(P)$.*

524 In Figure 5 we summarise the relation between coverage criteria and soundness of a particular
 525 program property. Given a program P , Figure 5 depicts the domain of equivalence relations
 526 over $\llbracket P \rrbracket$ where id denotes the most fine equivalence relation that corresponds to the identity
 527 relation, $\forall \sigma_1, \sigma_2 \in \llbracket P \rrbracket : (\sigma_1, \sigma_2) \in id$ iff $\sigma_1 = \sigma_2$, and top denotes the coarser equivalence
 528 relation that sees every trace as equivalent $\forall \sigma_1, \sigma_2 \in \llbracket P \rrbracket$ it holds that $(\sigma_1, \sigma_2) \in top$. As
 529 stated in Theorem 3 whenever $\mathcal{R}^C \preceq \mathcal{A}_{\llbracket P \rrbracket}$ then the coverage criterion C can be used to
 530 ensure soundness of the analysis of property \mathcal{A} on program P . As stated by Corollary 5 a
 531 coverage criterion C can ensure soundness for all those properties in $\Pi(C)$.

532 Following our reasoning, the most natural coverage criterion for a given semantic property
 533 \mathcal{A} is the one for which $\mathcal{R}^C = \mathcal{A}$, namely the coverage criterion whose partition on states
 534 corresponds to the property under analysis. In the literature there exists many different
 535 coverage criteria and some of them turn out to be equivalent when compared with respect
 536 to the partition that they induce on the input space. It has been observed that all existing
 537 test coverage criteria can be formalised on four mathematical structures: input domains,
 538 graphs, logic expressions, and syntax descriptions (grammars) [1]. Even if these coverage
 539 criteria are not explicitly related to the properties being analysed they have probably been
 540 designed while having in mind the kind of properties of interest. For example, some of the
 541 most widely known coverage criteria are based on graph features and are typically used for
 542 the analysis of properties related to a graphical representation of programs, like control flow
 543 or data flow properties of code or variables that can be verified on the control flow graph of
 544 a program, or function calls that can be verified on the call graph or a program, and so on.
 545 For example code coverage requires the execution of all the basic blocks of a control flow
 546 graph and wants to ensure that all the reachable instructions of a program are considered at
 547 least in one execution of the test set.

548 What we have stated so far allows us to begin to answer the question regarding how
 549 well the coverage criterion behaves with respect to the analysis of a given semantic property
 550 (when this can be modelled as a partitioning closure on the powerset of program traces). The
 551 design of an automatic or systematic strategy for the generation of an input set that covers
 552 a given coverage criterion remains an open challenge that deserves further investigation.

553 **Transforming properties towards soundness**

554 There are two questions that naturally arise from our reasoning and that would be interesting
 555 to investigate regarding the systematic transformation of the property under analysis or the
 556 coverage criterion towards soundness.

- 557 1. Consider a program P , a coverage criterion C that induces a partition $\mathcal{R}^C \in Eq(\llbracket P \rrbracket)$ on
 558 the traces of program P and a trace property \mathcal{A}_1 for which the coverage criterion C cannot
 559 ensure soundness. We wonder if it is possible to design a systematic transformation of
 560 property \mathcal{A}_1 that, by grouping some of its equivalence classes, returns a trace property
 561 for which we have soundness when C is satisfied by the input set. It would be interesting
 562 to understand when this transformation is possible without reaching *top*, i.e., while still
 563 being able to distinguish trace properties. This is depicted by the arrow labeled with T_1
 564 in the upper part of Figure 5.
- 565 2. Consider a program P , a coverage criterion C_1 that induces a partition $\mathcal{R}^{C_1} \in Eq(\llbracket P \rrbracket)$ on
 566 the traces of program P and a trace property \mathcal{A} for which the coverage criterion C_1 cannot
 567 ensure soundness. We wonder if it is possible to design a systematic transformation of
 568 \mathcal{R}^{C_1} that, by further splitting its equivalence classes, returns a partition of the program
 569 traces, and therefore a coverage criterion, that when satisfied by the input set ensures
 570 soundness for the analysis of property \mathcal{A} . In this case it is interesting to investigate when
 571 this refinement is possible without ending up with the identity relation, namely without
 572 collapsing to *id* where all program traces needs to be considered for coverage. This is
 573 depicted by the arrow labeled with T_2 in the bottom part of Figure 5.

574 **Transforming programs towards soundness**

575 As for static analysis also for dynamic analysis the way in which programs are written
 576 influences the precision of the analysis either because they expand the input set that satisfies
 577 a given coverage criterion, thus requiring the observation of more program runs, or because
 578 they complicate the automatic/systematic extraction of an input set that satisfies a given
 579 coverage criterion. We focus on the first case since we still have to formally investigate the
 580 extraction of input sets for a given coverage criterion, namely the input generation and input
 581 recogniser procedure.

582 Let us consider program R on the right of Figure 3 that computes the absolute value of
 583 an integer value and does it by adding some extra control on the range of the input integer
 584 value in order to proceed with the computation of the modulo in some syntactically different,
 585 but semantically equivalent ways. Indeed, in this example it is easy to observe that blocks B_4
 586 and B_5 are equivalent, but we can think about more sophisticated ways to write equivalent
 587 code in such a way that it would be difficult for the analyst to automatically recognise that
 588 they are equivalent. If we consider again the path coverage criterion we can observe that in
 589 order to cover the control flow graph of program R we need at least three input values: a
 590 negative integer, a positive integer smaller than 100 and a positive integer greater than or
 591 equal to 100. Of course what is done in block B_2 can be replicated many times, as far as we
 592 are able to write blocks that are syntactically different but semantically equivalent to B_4 or
 593 B_3 . According to our framework, path coverage is more complicated to reach on program
 594 R than on program Q . Indeed, in this case, every input set that satisfies path coverage for
 595 program R also satisfies path coverage for program Q while the converse does not hold in
 596 general. This reasoning is limited to the amount of traces that we need to satisfy a given
 597 coverage criterion and does not take into account the difficulty of generating such traces. Of
 598 course both aspects would need to be taken into account by our formal framework.

Moreover, as done for static analysis in [6], it would be interesting to define the notions of sound clique $\mathbb{S}(P, InSet, \mathcal{A})$ and of unsound clique $\bar{\mathbb{S}}(P, InSet, \mathcal{A})$ that represent the sets of all programs that are functionally equivalent to P and for which the dynamic analysis of property \mathcal{A} on input set $InSet \subseteq \mathbb{I}_P$ is respectively sound and not sound:

$$\mathbb{S}(P, InSet, \mathcal{A}) \stackrel{\text{def}}{=} \{Q \in Prog \mid Den(\llbracket P \rrbracket) = Den(\llbracket Q \rrbracket), InSet \overset{s}{\rightsquigarrow} \mathcal{A}(P)\}$$

$$\bar{\mathbb{S}}(P, InSet, \mathcal{A}) \stackrel{\text{def}}{=} \{Q \in Prog \mid Den(\llbracket P \rrbracket) = Den(\llbracket Q \rrbracket), Q \notin \mathbb{S}(P, InSet, \mathcal{A})\}$$

599 We plan to study the existence of transformations from $\bar{\mathbb{S}}(P, InSet, \mathcal{A})$ to $\mathbb{S}(P, InSet, \mathcal{A})$ in
 600 order to rewrite a program toward soundness. It is interesting to understand which are the
 601 properties for which this can be done in a systematic way and what is the key for reaching
 602 soundness. The intuition is that for reaching soundness with respect to a property \mathcal{A} on an
 603 input set $InSet$ we should choose programs whose variations of property \mathcal{A} are all considered
 604 by the input set as stated in Theorem 3. Thus, in general, if we reduce variations of the
 605 considered property by merging traces that are functionally equivalent even if they have
 606 diversified \mathcal{A} properties we would probably facilitate soundness. This needs to be formally
 607 understood, proved and validated on some existing dynamic analysis.

608 **4 Software protection: a new perspective**

609 In the software protection scenario we are interested in preventing program analysis while
 610 preserving the intended behaviour of programs. To face this problem Collberg et al. [9]
 611 introduced the notion of code obfuscation: program transformations designed with the explicit
 612 intent of complicating and degrading program analysis while preserving program functionality.
 613 Few years later Barak et al. [3] proved that it is not possible to obfuscate everything but
 614 the input-output behaviour for all programs with limited penalty in performances. However,
 615 it is possible to relax some of the requirements of Barak et al. and design obfuscating
 616 techniques that are able to complicate certain analysis of programs. This is witnessed by the
 617 great amount of obfuscation tools and techniques that researchers, both from academia and
 618 industry, have been developing in the last twenty years [8]. What it means for a program
 619 transformation to complicate program analysis is something that needs to be formally
 620 stated and proved when defining new obfuscating transformations. The extent to which an
 621 obfuscating technique complicates, and therefore protects, the analysis of certain program
 622 properties is referred to as *potency* of the obfuscation. A formal proof of the quality of
 623 obfuscation in terms of its potency is very important in order to compare the efficiency of
 624 different obfuscation techniques and in order to understand the degree of protection that they
 625 guarantee. Unfortunately, a unifying methodology for the quantitative evaluation of software
 626 protection techniques is still an open challenge, as witnessed by the recent Dagstuhl Seminar
 627 on this topic [20]. What we have are specific measurements done when new techniques are
 628 proposed, or formal proofs that reduce the analysis of obfuscated programs to well known
 629 complex analysis tasks (like alias analysis, shape analysis, etc.).

630 In our framework, complicating program analysis means inducing imprecision in the
 631 results of the analysis of the obfuscated program with respect to the results of the analysis
 632 of the original program. This means that code obfuscation should induce false positives in
 633 static program analysis and false negatives in dynamic program analysis.

634 **4.1 Program transformations against static program analysis**


635 The abstract interpretation framework has been used to reason on the semantic properties
 636 that code obfuscation transformations are able to protect and the ones that they can still be

637 analysed on the obfuscated program. It has been observed that a program property expressed
 638 by an abstract domain \mathcal{A} is obfuscated (protected) by an obfuscation $\mathcal{O} : Prog \rightarrow Prog$ on
 639 a program P whenever $\llbracket P \rrbracket^{\mathcal{A}} \leq_{\mathcal{A}} \llbracket \mathcal{O}(P) \rrbracket^{\mathcal{A}}$, namely when the analysis \mathcal{A} on the obfuscated
 640 program returns a less precise result with respect to the analysis of the same property on
 641 the original program P . The spurious information added to the analysis by the obfuscation
 642 is the noise that confuses the analyst, thus making the analysis more complicated. The
 643 relation between potency of code obfuscation and the notion of (in)completeness in abstract
 644 interpretation has been proven, as obfuscating a property means to induce incompleteness
 645 in its analysis [22]. So, for example, the insertion of a true opaque predicate O^T (see the
 646 program in the middle of Figure 1) would confuse all those analyses that are not able to
 647 precisely evaluate such a predicate and have to consider both branches as possible. No
 648 confusion is added for those analyses that are able to precisely evaluate the opaque predicate
 649 and consider only the true branch as possible, namely those analyses that are complete for
 650 the evaluation of the predicate value. Following this idea, a formal framework based on
 651 program semantics and abstract interpretation has been developed, where it is possible to
 652 formally prove that a property is obfuscated by a given program transformation, compare
 653 the efficiency of different obfuscating techniques in protecting a given property, define a
 654 systematic strategy for the design of a code obfuscation technique for protecting a given
 655 program property [17, 19, 22, 25]. This semantic understanding of the effects that code
 656 obfuscation has on the semantics and semantic properties of programs as shown its usefulness
 657 also in the malware detection scenario where malware writers use code obfuscation to evade
 658 automatic detection [15, 16].

659 Thus we can say that the effects of functionality preserving program transformations on
 660 program semantics and on the precision of the results of static analysis has been extensively
 661 studied and a mature formal framework has been provided [15, 16, 17, 19, 22, 25].

662 4.2 Program transformations against dynamic program analysis

663 To the best of our knowledge, the effects of functionality preserving program transformations
 664 on the precision of dynamic analysis have not been fully investigated yet. Following our
 665 reasoning, the general idea is that dynamic analysis is complicated by program transformations
 666 that induce false negatives while preserving program's functionality. Let $\mathcal{A} \in Eq(\Sigma^*)$ denote
 667 a property of interest for dynamic analysis. Inducing false negatives for the analysis of
 668 a property \mathcal{A} can be done by exploiting the partial observation of program's executions
 669 innate in the test set, and thus adding traces that do not belong to the test set and have
 670 a different \mathcal{A} property. Thus, the key for software protection against dynamic analysis is
 671 software *diversification* with respect to the property under analysis. The ideal obfuscation
 672 against the dynamic analysis of property \mathcal{A} should specialise programs with respect to every
 673 input in such a way that every input exhibits a different behaviour for property \mathcal{A} . Namely,
 674 an ideal obfuscation against \mathcal{A} is a program transformation $\mathcal{O} : Prog \rightarrow Prog$ such that
 675 $\forall \sigma_1, \sigma_2 \in \llbracket \mathcal{O}(P) \rrbracket$ we have that $\mathcal{A}(\sigma_1) = \mathcal{A}(\sigma_2) \Leftrightarrow \sigma_1 = \sigma_2$. In this ideal situation in order
 676 to avoid false negatives the analyst should consider every possible execution trace of $\mathcal{O}(P)$
 677 since each trace exhibits a different aspects of property \mathcal{A} , so missing a trace would mean
 678 to miss such an aspect. This intuition is confirmed in a preliminary work in this direction
 679 where it is shown how diversification is the basis of existing software protection techniques
 680 against dynamic analysis [18]. This work provides a topological characterisation of the
 681 soundness of the dynamic analysis of properties expressed as equivalence relations (as we
 682 have done in Section 3.2.1). This formal characterisation is then used to define the notion of
 683 transformation potency for dynamic analysis.



■ **Figure 6** Transformation Potency

684 ► **Definition 6.** A functionality preserving program transformation $\mathcal{O} : \text{Prog} \rightarrow \text{Prog}$ is
 685 potent for the analysis of $\mathcal{A} \in \text{Eq}(\Sigma^*)$ of program P if:

686 ■ $\forall \sigma_1, \sigma_2 \in [[\mathcal{O}(P)]] : [\sigma_1]_{\mathcal{A}} = [\sigma_2]_{\mathcal{A}}, \forall \nu_1, \nu_2 \in [[P]] : \text{Den}(\nu_1) = \text{Den}(\sigma_1), \text{Den}(\nu_2) =$
 687 $\text{Den}(\sigma_2)$ then $[\nu_1]_{\mathcal{A}} = [\nu_2]_{\mathcal{A}}$

688 ■ $\exists \nu_1, \nu_2 \in [[P]] : [\nu_1]_{\mathcal{A}} = [\nu_2]_{\mathcal{A}}$ for which $\exists \sigma_1, \sigma_2 \in [[\mathcal{O}(P)]] : \text{Den}(\nu_1) = \text{Den}(\sigma_1), \text{Den}(\nu_2) =$
 689 $\text{Den}(\sigma_2)$ such that $[\sigma_1]_{\mathcal{A}} \neq [\sigma_2]_{\mathcal{A}}$

690 Figure 6 provides a graphical representation of the notion of potency. On the left we have the
 691 traces of the original program P partitioned according to the equivalence relation \mathcal{A} induced
 692 by the property of interest, while on the right we have the traces of the transformed program
 693 $\mathcal{O}(P)$ partitioned according to \mathcal{A} . Traces that are denotationally equivalent have the same
 694 shape (triangle, square, circle, oval), but different dimension since they are in general different
 695 traces. The first condition means that the traces of $\mathcal{O}(P)$ that property \mathcal{A} maps to the
 696 same equivalence class (circle and square), are denotationally equivalent to traces of P that
 697 property \mathcal{A} maps to the same equivalence class. This means that what is grouped together
 698 by \mathcal{A} on $[[\mathcal{O}(P)]]$ was grouped together by \mathcal{A} on $[[P]]$, modulo the denotational equivalence
 699 of traces. The second condition requires that there are traces of P (triangle and star) that
 700 property \mathcal{A} maps to the same equivalence class and whose denotationally equivalent traces in
 701 $\mathcal{O}(P)$ are mapped by \mathcal{A} to different equivalence classes. This means that a defense technique
 702 against dynamic analysis with respect to a property \mathcal{A} is successful when it transforms a
 703 program into a functionally equivalent one for which property \mathcal{A} is more diversified among
 704 execution traces. This implies that it is necessary to collect more execution traces in order
 705 for the analysis to be precise. At the limit we have an optimal defense technique when \mathcal{A}
 706 varies at every execution trace.

707 The above definition of transformation potency for dynamic analysis has been validated
 708 by modelling in the proposed framework some existing software defence strategies against
 709 dynamic analysis for the extraction of the control flow graph of programs like Range Dividers
 710 [2] and Gadget diversification [30]. In both cases it is possible to show that the proposed
 711 transformations complicate the dynamic extraction of the control flow graph by adding new
 712 diversified paths to the control flow graph, as stated in Definition 6. In the following we
 713 report a simple example from [18] that shows how the key for obfuscating properties of data
 714 values for dynamic analysis is diversification.

715 ► **Example 7.** Consider the following programs P and Q that compute the sum of natural
 716 numbers from $x \geq 0$ to 49 (we assume that the inputs values for x are natural numbers).

<pre style="margin: 0;"> 717 P input x; sum := 0; while x < 50 • { X = [0, 49] } sum := sum + x; x := x + 1; </pre>	<pre style="margin: 0;"> Q input x; n := select(N,x) x := x * n; sum := 0; while x < 50 * n • { X = [0, n * 50 - 1] } sum := sum + x/n; x := x + n; x := x/n; </pre>
----------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Consider a dynamic analysis that observes the maximal value assumed by x at program point \bullet . For every possible execution of program P we have that the maximal value assumed by x at program point \bullet is 49. Consider a state $s \in \Sigma$ as a tuple $\langle pp, [val_x, val_{sum}] \rangle$, where pp denotes the current program point, val_x and val_{sum} denote the current values of variables x and sum respectively. We define a function $\tau : \Sigma \rightarrow \mathbb{N}$ that observes the value assumed by x at state s when s refers to program point \bullet , and function $max : \Sigma^* \rightarrow \mathbb{N}$ that observes the maximal value assumed by x at \bullet along an execution trace:

$$\tau(s) \stackrel{\text{def}}{=} \begin{cases} val_x & \text{if } pp = \bullet \\ \emptyset & \text{otherwise} \end{cases} \quad max(\sigma) \stackrel{\text{def}}{=} max(\{\tau(s) \mid s \in \sigma\})$$

718 This allows us to define the equivalence relation $\mathcal{A}_{max} \in Eq(\Sigma^*)$ that observes traces with
719 respect to the maximal value assumed by x at \bullet , as $(\sigma, \sigma') \in \mathcal{A}_{max}$ iff $max(\sigma) = max(\sigma')$.
720 We can observe that all the execution traces of P belong to the same equivalence class of
721 \mathcal{A}_{max} . In this case, a dynamic analysis of property \mathcal{A}_{max} on P is sound whenever the test
722 set contains at least one execution trace of P . This happens because the property that we
723 are looking for is an invariant property of program executions and it can be observed on any
724 execution trace.

725 Let us now consider program Q equivalent to P , i.e., $Den[[P]] = Den[[Q]]$, where the
726 value of x is diversified by multiplying it by the parameter n . The guard and the body
727 of the **while** are adjusted in order to preserve the functionality of the program. When
728 observing property \mathcal{A}_{max} on Q , we have that the maximal value assumed by x at program
729 point \bullet is determined by the parameter n generated in the considered trace. The statement
730 $n := \text{select}(N, x)$ assigns to n a value in the range $[0, N]$ depending on the input value x . We
731 have that the traces of program Q are grouped by \mathcal{A}_{max} depending on the value assumed by
732 n . Thus, $\mathcal{A}([Q])$ contains an equivalence class for every possible value assumed by n during
733 execution. This means that the transformation that rewrites P into Q is potent according
734 to Definition 6. Dynamic analysis of property \mathcal{A}_{max} on program Q is sound if the test set
735 contains at least one execution trace for each of the possible values of n generated during
736 execution.

5 Open research directions

738 We have provided an unifying view of the relations between properties and program transfor-
739 mations and the precision of static and dynamic analysis in the standard analysis scenario
740 and in the software protection scenario. Researchers have proposed possible ways for tuning
741 the precision of static analysis while less attention has been posed to the formal investigation
742 of dynamic analysis. In this context it is worth to mention the recent work of O'Hearn [26]
743 that defines a formalism called incorrectness logic, which is similar to Hoare's logic, and
744 allows us to prove the presence of bugs but not their absence, thus capturing the essence

745 of program testing. The incorrectness logic is based on a under-approximation triple that
746 plays a dual role when compared to the standard over-approximation triple that we are
747 used to see in Hoare’s logic. Indeed, while logic and symbolic reasoning are useful since
748 they can cover many states or program paths at once, they do not allow in general to cover
749 all paths and this makes it difficult to prove the absence of errors. The author claims the
750 necessity and usefulness of incorrectness logic that formalises under-approximate reasoning
751 in order to provide a logical proof of the presence of bugs. Such reasoning should of course
752 be combined with standard correctness proof in order to obtain a global view of program’s
753 runtime behaviour. The incorrectness logic of O’Hearn does not try to gain soundness,
754 namely to avoid or reduce false negatives, but provides formal proofs for what can be derived
755 in an unsound context. Our idea is to investigate the extent to which it is possible to induce
756 or force soundness by modifying either the program, the property to be analysed or the
757 coverage criterion. Once we have understood when and how soundness can be forced we
758 should see how this interacts with incorrectness logic.

759 The preliminary work done in the investigation of program and properties transformations
760 towards sound dynamic analysis have pointed out many interesting aspects that need to be
761 studied and that we list below as future research directions.

762 The preliminary results that relate program properties, coverage criteria and the soundness
763 of the analysis should be generalised and extended to properties that cannot be modelled
764 as partitioning closures. Soundness of the analysis and transformation potency should
765 be redefined probably in terms of join-irreducible elements instead of equivalence classes.
766 This further investigation would probably lead to a classification of the properties usually
767 considered by dynamic analysis based on the domain model needed to express them: properties
768 of traces, properties of sets of traces, relational properties, hyper-properties. For each class
769 of properties it would then be interesting to derive a suitable obfuscation strategy. This
770 unifying framework would provide a common ground where to interpret and compare the
771 potency of different software protection techniques in harming dynamic analysis.

772 As regarding the transformation of properties towards soundness, we plan to verify if
773 and when it is possible to refine the coverage criterion C in order to ensure soundness with
774 respect to a given property \mathcal{A} , or when it is possible to further abstract the semantic property
775 \mathcal{A} in order to make it sound for a given coverage criterion C . This should be done starting
776 with properties that can be expressed as partitioning closures and then generalised to the
777 other classes of properties.

778 As regarding the transformation of programs towards soundness, it is important to
779 investigate when it is possible to transform a program P for which the dynamic analysis of a
780 given property \mathcal{A} is sound (resp. unsound) into a different program P' which is functionally
781 equivalent to P and for which the dynamic analysis of property \mathcal{A} is unsound (resp. sound).

782 It would also be important to extend the framework in order to take into account the
783 feasibility of the considered coverage criterion, maybe defining some constraints that a
784 program has to satisfy in order to guarantee the feasibility of a given coverage criterion, or
785 by modelling and measuring situations when full coverage is not possible.

786 References

- 787 1 Paul Ammann and Jeff Offutt. *Introduction to software testing*. Cambridge University Press,
788 2016.
- 789 2 Sebastian Banescu, Christian Collberg, Vijay Ganesh, Zack Newsham, and Alexander
790 Pretschner. Code obfuscation against symbolic execution attacks. In *Proceedings of the*
791 *32nd Annual Conference on Computer Security Applications*, pages 189–200, 2016.

- 792 3 B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. P. Vadhan, and K. Yang. On
793 the (im)possibility of obfuscating programs. In *CRYPTO '01: Proceedings of the 21st Annual*
794 *International Cryptology Conference on Advances in Cryptology*, pages 1–18. Springer-Verlag,
795 2001.
- 796 4 Ezio Bartocci and Yliès Falcone, editors. *Lectures on Runtime Verification - Introductory and*
797 *Advanced Topics*, volume 10457 of *Lecture Notes in Computer Science*. Springer, 2018.
- 798 5 Tim Blazytko, Moritz Contag, Cornelius Aschermann, and Thorsten Holz. Syntia: Synthesizing
799 the semantics of obfuscated code. In *26th USENIX Security Symposium, USENIX Security*
800 *2017, Vancouver, BC, Canada, August 16-18, 2017*, pages 643–659. USENIX Association,
801 2017.
- 802 6 Roberto Bruni, Roberto Giacobazzi, Roberta Gori, Isabel Garcia-Contreras, and Dusko
803 Pavlovic. Abstract extensionality: on the properties of incomplete abstract interpretations.
804 *Proc. ACM Program. Lang.*, 4(POPL):28:1–28:28, 2020.
- 805 7 Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors.
806 *Handbook of Model Checking*. Springer, 2018.
- 807 8 C. Collberg and J. Nagra. *Surreptitious Software: Obfuscation, Watermarking, and Tamper-*
808 *proofing for Software Protection*. Addison-Wesley Professional, 2009.
- 809 9 C. Collberg, C. Thomborson, and D. Low. Manufacturing cheap, resilient, and stealthy opaque
810 constructs. In *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of*
811 *programming languages (POPL '98)*, pages 184–196. ACM Press, 1998.
- 812 10 Kevin Coogan, Gen Lu, and Saumya K. Debray. Deobfuscation of virtualization-obfuscated
813 software: a semantics-based approach. In *Proceedings of the 18th ACM Conference on Computer*
814 *and Communications Security, CCS 2011, Chicago, Illinois, USA, October 17-21, 2011*, pages
815 275–284. ACM, 2011.
- 816 11 P. Cousot. Constructive design of a hierarchy of semantics of a transition system by abstract
817 interpretation. *Theor. Comput. Sci.*, 277(1-2):47–103, 2002.
- 818 12 P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis
819 of programs by construction or approximation of fixpoints. In *Conference Record of the 4th*
820 *ACM Symposium on Principles of Programming Languages (POPL '77)*, pages 238–252. ACM
821 Press, 1977.
- 822 13 P. Cousot and R. Cousot. A constructive characterization of the lattices of all retractions,
823 preclosure, quasi-closure and closure operators on a complete lattice. *Portug. Math.*, 38(2):185–
824 198, 1979.
- 825 14 P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In *Conference*
826 *Record of the 6th ACM Symposium on Principles of Programming Languages (POPL '79)*,
827 pages 269–282. ACM Press, 1979.
- 828 15 M. Dalla Preda, M. Christodorescu, S. Jha, and S. Debray. A semantics-based approach
829 to malware detection. In *POPL '07: Proceedings of the 34th annual ACM SIGPLAN-*
830 *SIGACT symposium on Principles of programming languages*, pages 377–388. ACM Press,
831 2007. doi:<http://doi.acm.org/10.1145/1190216.1190270>.
- 832 16 Mila Dalla Preda, Mihai Christodorescu, Somesh Jha, and Saumya K. Debray. A semantics-
833 based approach to malware detection. *ACM Trans. Program. Lang. Syst.*, 30(5):25:1–25:54,
834 2008.
- 835 17 Mila Dalla Preda and Roberto Giacobazzi. Semantics-based code obfuscation by abstract
836 interpretation. *Journal of Computer Security*, 17(6):855–908, 2009.
- 837 18 Mila Dalla Preda, Roberto Giacobazzi, and Niccoló Marastoni. Formal framework for reasoning
838 about the precision of dynamic analysis. In *Static Analysis, 16th International Symposium,*
839 *SAS 2020*, page to appear, 2020.
- 840 19 Mila Dalla Preda and Isabella Mastroeni. Characterizing a property-driven obfuscation strategy.
841 *J. Comput. Secur.*, 26(1):31–69, 2018.

- 842 20 Bjorn De Sutter, Christian S. Collberg, Mila Dalla Preda, and Brecht Wyseur. Software
843 protection decision support and evaluation methodologies (dagstuhl seminar 19331). *Dagstuhl*
844 *Reports*, 9(8):1–25, 2019.
- 845 21 Hanne Riis Nielson Flemming Nielson and Chris Hankin. *Principles of Program Analysis*.
846 Springer, 1999.
- 847 22 R. Giacobazzi. Hiding information in completeness holes - new perspectives in code obfus-
848 cation and watermarking. In *Proc. of The 6th IEEE International Conferences on Software*
849 *Engineering and Formal Methods (SEFM'08)*, pages 7–20. IEEE Press., 2008.
- 850 23 R. Giacobazzi, F. Ranzato, and F. Scozzari. Making abstract interpretation complete. *Journal*
851 *of the ACM*, 47(2):361–416, March 2000.
- 852 24 Roberto Giacobazzi, Francesco Logozzo, and Francesco Ranzato. Analyzing program analyses.
853 In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of*
854 *Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, pages 261–273,
855 2015.
- 856 25 Roberto Giacobazzi, Isabella Mastroeni, and Mila Dalla Preda. Maximal incompleteness as
857 obfuscation potency. *Formal Asp. Comput.*, 29(1):3–31, 2017.
- 858 26 Peter W. O’Hearn. Incorrectness logic. *Proc. ACM Program. Lang.*, 4(POPL):10:1–10:32,
859 2020.
- 860 27 Mathilde Ollivier, Sébastien Bardin, Richard Bonichon, and Jean-Yves Marion. How to kill
861 symbolic deobfuscation for free (or: unleashing the potential of path-oriented protections). In
862 *Proceedings of the 35th Annual Computer Security Applications Conference*, pages 177–189,
863 2019.
- 864 28 Andre Pawlowski, Moritz Contag, and Thorsten Holz. Probfuscation: an obfuscation approach
865 using probabilistic control flows. In *International Conference on Detection of Intrusions and*
866 *Malware, and Vulnerability Assessment*, pages 165–185. Springer, 2016.
- 867 29 Francesco Ranzato and Francesco Tapparo. Strong preservation as completeness in abstract
868 interpretation. In *Programming Languages and Systems, 13th European Symposium on*
869 *Programming, ESOP 2004, Held as Part of the Joint European Conferences on Theory and*
870 *Practice of Software, ETAPS 2004, Barcelona, Spain, March 29 - April 2, 2004, Proceedings*,
871 volume 2986 of *Lecture Notes in Computer Science*, pages 18–32. Springer, 2004.
- 872 30 Sebastian Schrittwieser and Stefan Katzenbeisser. Code obfuscation against static and dynamic
873 reverse engineering. In *International workshop on information hiding*, pages 270–284. Springer,
874 2011.
- 875 31 Sebastian Schrittwieser, Stefan Katzenbeisser, Johannes Kinder, Georg Merzdovnik, and
876 Edgar R. Weippl. Protecting software through obfuscation: Can it keep pace with progress in
877 code analysis? *ACM Comput. Surv.*, 49(1):4:1–4:37, 2016.
- 878 32 Monirul I. Sharif, Andrea Lanzi, Jonathon T. Giffin, and Wenke Lee. Automatic reverse
879 engineering of malware emulators. In *30th IEEE Symposium on Security and Privacy (S&P*
880 *2009), 17-20 May 2009, Oakland, California, USA*, pages 94–109. IEEE Computer Society,
881 2009.
- 882 33 Bernhard Steffen and Gerhard J. Woeginger, editors. *Computing and Software Science - State*
883 *of the Art and Perspectives*, volume 10000 of *Lecture Notes in Computer Science*. Springer,
884 2019.
- 885 34 Babak Yadegari, Brian Johannsmeyer, Ben Whitely, and Saumya Debray. A generic approach
886 to automatic deobfuscation of executable code. In *2015 IEEE Symposium on Security and*
887 *Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 674–691. IEEE Computer
888 Society, 2015.