

On Time-Sensitive Control Dependencies

MARTIN HECKER, SIMON BISCHOF, and GREGOR SNETLING, Karlsruhe Institute of Technology

We present efficient algorithms for *time-sensitive* control dependencies (CDs). If statement y is time-sensitively control dependent on statement x , x not only decides whether y is executed, but also how many time steps after x . If y is not standard control dependent on x , but time-sensitively control dependent, then y will always be executed after x , but the execution time between x and y varies. This allows to discover e.g. timing leaks in security-critical software.

We systematically develop properties and algorithms for time-sensitive CDs, as well as for nontermination-sensitive CDs. These do not only work for standard control flow graphs (CFGs), but also for CFGs lacking a unique exit node (e.g. reactive systems). We show that Cytron's efficient algorithm for dominance frontiers [10] can be generalized to allow efficient computation not just of classical CDs, but also of time-sensitive and nontermination-sensitive CDs. We then use time-sensitive CDs and time-sensitive slicing to discover cache timing leaks in an AES implementation. Performance measurements demonstrate scalability of the approach.

CCS Concepts: • **Software and its engineering** → **Automated static analysis**; • **Theory of computation** → **Program analysis**; • **Security and privacy** → **Logic and verification**.

Additional Key Words and Phrases: control dependency, program slicing, timing dependency, timing leak

ACM Reference Format:

Martin Hecker, Simon Bischof, and Gregor Snelting. 2021. On Time-Sensitive Control Dependencies. *ACM Trans. Program. Lang. Syst.* 1, 1 (August 2021), 39 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION AND OVERVIEW

Timing Leaks are a major source of software security problems today. Attacks based on timing leaks such as *Spectre* [22] have become known to the general public. Yet there are not many program analysis tools that detect timing leaks in software.

In this article we describe a new kind of dependency between program statements, the *time-sensitive control dependency*. It is able to discover timing leaks, and can be implemented as an automatic program analysis. We will explain time-sensitive dependencies, provide efficient algorithms, provide a soundness proof, and apply it to discover timing leaks in an implementation of the AES cryptographic standard.

The construction of time-sensitive control dependencies starts with classical control dependencies. We will thus begin by sketching the research path from CDs to timing

Authors' address: Martin Hecker; Simon Bischof; Gregor Snelting Karlsruhe Institute of Technology, gregor.snelting@kit.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

0164-0925/2021/8-ART \$15.00

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

dependencies, and provide introductory examples. Later in the article, we will provide formal definitions, proofs, and algorithms.

Control dependencies (CDs), originally introduced by [11, 33], are a fundamental building block in program analysis. CDs have many applications: they can, for example, be used for program optimizations such as code scheduling, loop fusion, or code motion (see e.g. [25]); or for program transformations such as partial evaluation (e.g. [20]) or refactoring (e.g. [5]). CDs are in particular fundamental for program dependence graphs (PDGs) and program slicing [11, 19, 23]. Intuitively, a program statement y is control dependent on another statement x , written $x \rightarrow_{cd} y$, if x – typically an **if** or **while** statement – decides whether y will be executed or not. Classical CDs are defined via postdominators; in fact classical CDs are essentially the postdominance frontiers in the control flow graph (CFG) [10]. Postdominance frontiers can be computed by an efficient algorithm due to Cytron [10].

Unfortunately, the classic CD definition is limited to CFGs with unique exit node, and thus assumes that all programs can terminate. In 2007 Ranganath et al. [29] generalized control dependence to CFGs without unique exits and nonterminating programs (e.g. reactive systems); providing the first algorithm for nontermination-sensitive CDs. Later, Amtoft [3] provided definitions and algorithms for nontermination-insensitive CDs, which allow for sound analysis and slicing of nonterminating programs. But these algorithms could no longer be based on the efficient Cytron algorithm for postdominance frontiers.

In this contribution, we not only present new efficient algorithms for the Ranganath-Amtoft CD definitions. We will also provide definitions and algorithms for *time-sensitive* CDs and time-sensitive slicing. Time-sensitive CD, written $x \rightarrow_{tscd} y$, holds if x decides whether y is executed, or if x decides *when* y is executed (even if y is always executed after x) – that is, how many time units after execution of x . Intuitively, $x \rightarrow_{tscd} y$ while *not* $x \rightarrow_{cd} y$ means that y will always be executed after x , but the execution time between x and y varies. The latter property is important to discover timing leaks in security-critical software.

We systematically develop theoretical properties and efficient algorithms for \rightarrow_{tscd} , and evaluate their performance. It turns out that Cytron’s efficient algorithm for dominance frontiers can be generalized to an abstract notion of dominance, which then can be used for the efficient computation of both the Ranganath/Amtoft CD, as well as our new time-sensitive CD. We then apply \rightarrow_{tscd} to (models of) hardware microarchitectures, and use it to find cache timing leaks in an AES implementation.

Many of the theorems in this article have been formalized and machine-checked using the machine prover *Isabelle*. Such theorems are marked with a 🍌 sign. The Isabelle proofs can be found in the electronic appendix of this article. For some theorems in section 5, such an Isabelle proof has not yet been completed. Manual proofs are available, but are not presented in this article. Consequently, such theorems are called “observations”.

1.1 Overview

The main part of this article will present time-sensitive CDs and algorithms in a rather technical manner. Before we embark on this, we present an informal overview of our research path and results. We begin with a discussion of classical control dependencies, and compare these to our new notion of time-sensitive control dependencies.

1.1.1 Control Dependence. Informally, a control dependence in a CFG, written $x \rightarrow_{cd} y$, means that x decides whether y is executed or not. In structured programs, x is typically an **if** or **while** statement. Figure 1 presents two examples: In the first example (left), node (5) is control dependent on node (1); node (3) is not control dependent on (1) but on

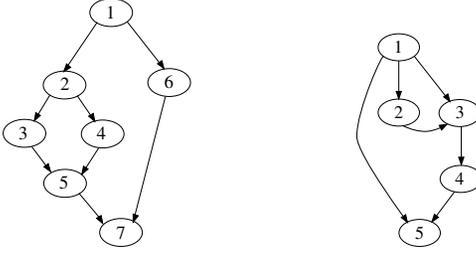


Fig. 1. Two simple control flow graphs illustrating control dependence

(2). In the second example (right)¹, nodes (2), (3), and (4) are control dependent on node (1). Technically, CD is based on the notion of *postdomination* in CFGs. y postdominates x (written $y \sqsubseteq_{\text{POST}} x$) if any path from x to the *exit* node must pass through y . Several formal CD definitions exist; as this may be confusing we will relate the most popular definitions to the examples in Figure 1. The original definition of CD in [11] is as follows:

$$x \rightarrow_{cd} y \iff \neg(y \sqsubseteq_{\text{POST}} x) \wedge \exists \text{ path } \pi : x \rightarrow^* y \text{ such that } \forall z \in \pi \setminus \{x, y\} : y \sqsubseteq_{\text{POST}} z$$

The condition that y is not a postdominator for x means that from x there is a second path (not containing y) to the exit node. That is, there is a conditional branch at x . The next condition demands that there is a path from x to y , and that y is a postdominator for all nodes z between x and y . Thus there is no side branch from any z to the exit node; hence x is directly controlling whether y is executed.

In Figure 1 (left), node (5) postdominates all nodes on paths between (1) and (5), but (5) does not postdominate (1); hence $(1) \rightarrow_{cd} (5)$. But (3) does not postdominate (2) (this node being the only one between (1) and (3)), hence $\neg((1) \rightarrow_{cd} (3))$. In Figure 1 (right), node (4) is control dependent on node (1). Since we have $(1) \rightarrow (5)$, node (4) does not postdominate (1). The path $(1) \rightarrow (3) \rightarrow (4)$ only contains the additional $z = (3)$, and (4) postdominates (3), so the second condition is satisfied. But what about the path $(1) \rightarrow (2) \rightarrow (3) \rightarrow (4)$? It is irrelevant, as the CD definition only demands there *exists* a path where for all z etc; it does not demand the z condition for all paths. Likewise, (2) as well as (3) are control dependent on (1).

An alternate, more compact CD definition was provided in [33], and is used in this article. Here x is a branch node with direct successors n and m , where the control-dependent y postdominates one but not the other:

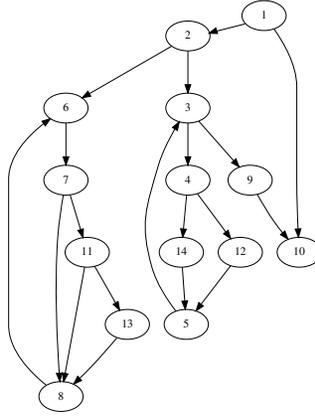
$$x \rightarrow_{cd} y \iff \exists n, m : x \rightarrow n, x \rightarrow m, y \sqsubseteq_{\text{POST}} n, \neg(y \sqsubseteq_{\text{POST}} m)$$

Lemma 1.1. 🌟 The above definitions for $x \rightarrow_{cd} y$ are equivalent whenever $x \neq y$.

Applied to Figure 1 (right), again we conclude that (4) is CD on (1). Choose $n = (2)$ ($n = (3)$ also works), $m = (5)$, then (4) postdominates (2) (and (3)), but (4) does not postdominate (5). Note that both n and m in the definition are existentially quantified. Thus the definition neither demands nor inhibits that (4) postdominates (3).

Lemma 1.2. 🌟 In Figure 1 (right), we have $(4) \sqsubseteq_{\text{POST}} (2)$, $(4) \sqsubseteq_{\text{POST}} (3)$, $\neg((4) \sqsubseteq_{\text{POST}} (5))$, and thus $(1) \rightarrow_{cd} (4)$.

¹We thank one reviewer for suggesting this example.

Fig. 2. A CFG G

1.1.2 *CFGs Without Unique Exit.* CFGs without unique exit, in particular with no exit or unreachable exits, are important for modern language constructs, for example event handlers or loops in reactive systems. Ranganath and Amtoft had generalized postdominance and CD for such CFGs. The resulting postdominance relations are called max- and sink-postdominance, and will be explained in section 2.2. If these are used in CD definitions, one obtains *non-termination-sensitive* control dependence, written $\rightarrow_{\text{ntscd}}$; and *non-termination-insensitive* control dependence, written $\rightarrow_{\text{nticd}}$. $\rightarrow_{\text{nticd}}$ is identical to \rightarrow_{cd} , but also works for graphs without unique exit. $\rightarrow_{\text{ntscd}}$ is identical to \rightarrow_{wcd} (weak control dependence, see section 2.1), but also works for graphs without unique exit.

$\rightarrow_{\text{nticd}}$ and $\rightarrow_{\text{ntscd}}$ are important building blocks for $\rightarrow_{\text{tscd}}$. A comparison between $\rightarrow_{\text{nticd}}$, $\rightarrow_{\text{ntscd}}$, and $\rightarrow_{\text{tscd}}$ is given in the next section.

1.1.3 *Time-Sensitive Control Dependence.* Time-sensitive CD, written $x \rightarrow_{\text{tscd}} y$, holds if x decides *when or whether* y is executed. This dependence is more relaxed than standard CD. Intuitively, $x \rightarrow_{\text{tscd}} y$ while *not* $x \rightarrow_{\text{cd}} y$ means that y will always be executed after x , but the *execution time between x and y varies*.

The latter property is important to discover timing leaks in security-critical software. A typical situation is as follows: y is not control dependent on x , but there are at least two paths from x to y . Then the run time between x and y varies: $x \rightarrow_{\text{tscd}} y$. If this variation depends on secret data, and can be measured by an attacker, a timing leak has been born. $x \rightarrow_{\text{tscd}} y$ will uncover this leak.

In our work time is discrete; a unit of time coincides with a transition in the CFG. Since steps of an abstraction of real programs and hardware are timed, this is therefore a “weakly timing-sensitive” model in the sense of [21].

Now, let us illustrate the differences between the different kinds of control dependences. A node y is non-termination sensitively control dependent on node x , written $x \rightarrow_{\text{ntscd}} y$, if x decides whether y will be executed. In Figure 2, we have $(1) \rightarrow_{\text{ntscd}} (2)$, because we will execute (2) when choosing (2) as the successor of (1) but not if we choose (10). Also, due to the loop at (3), we have $(3) \rightarrow_{\text{ntscd}} (10)$: By choosing (9) as the successor of (3) we are guaranteed to reach (10). But if we choose (4) as the successor, we might avoid reaching (10) by staying in the loop forever, so (3) decides if (10) will be executed.

197 y is non-termination insensitively control dependent on x , written $x \rightarrow_{\text{nticd}} y$, if x decides
 198 whether y will be executed, *assuming we eventually exit all loops that can be exited*. In
 199 Figure 2, we still have $(1) \rightarrow_{\text{nticd}} (2)$, with the same reasoning as above. But now we have
 200 $\neg((3) \rightarrow_{\text{ntscd}} (10))$: Since we assume that we always exit the loop at (3), we are guaranteed
 201 to reach (10), no matter which successor we choose at (3).

202 y is timing sensitively control dependent on x , written $x \rightarrow_{\text{tscd}} y$, if x decides *when* y will
 203 be executed. In Figure 2, we have $(7) \rightarrow_{\text{tscd}} (8)$, because we will execute (8) after one step
 204 when choosing (8) as the successor of (7) but not if we choose (11), when it takes two or
 205 three steps. Therefore, the choice taken at (7) influences the timing of (8). On the contrary,
 206 $\neg((4) \rightarrow_{\text{tscd}} (5))$, because no matter how we choose, we will always reach (5) in two steps.
 207 An interesting case is $(1) \rightarrow_{\text{tscd}} (2)$: If we choose (2) as successor of (1), we will reach (2)
 208 in exactly one step, but we will not if we choose (10) because we then will not reach (2).
 209

210
 211 **1.1.4 Applications for Software Security.** As indicated, $\rightarrow_{\text{tscd}}$ may help to discover timing
 212 leaks. More generally, $\rightarrow_{\text{tscd}}$ is useful for *Information Flow Control* (IFC). IFC uses program
 213 analysis techniques to discover leaks in software. Technically, *noninterference* is a property
 214 which guarantees that a program does not leak secret data. *Probabilistic* noninterference
 215 guarantees that there are no internal timing leaks, which arise if secret data influence
 216 scheduling or other measurable timing properties. For an introduction to IFC, see e.g. [31].

217 Indeed $\rightarrow_{\text{tscd}}$ was developed as an instrument to improve the precision of probabilistic
 218 noninterference analysis. We will report on applications of $\rightarrow_{\text{tscd}}$ for IFC in a separate article.
 219 In the current article, we focus on algorithms for $\rightarrow_{\text{tscd}}$; and use a different security example:
 220 in section 4, we will analyse an implementation of the AES cryptographic standard, and
 221 discover cache leaks in this implementation. These infamous cache leaks have been known
 222 for some time [4], but so far no program analysis tool was able to discover such leaks.
 223

224 **1.1.5 Algorithms.** The major part of this contribution is concerned with efficient algorithms
 225 for $\rightarrow_{\text{tscd}}$. For the classical \rightarrow_{cd} , Cytron's efficient algorithm for dominance frontiers can be
 226 used; but this algorithm was not employed by Ranganath/Amtoft.

227 We discovered that a generalized version of Cytron's algorithm can not only be used for
 228 both $\rightarrow_{\text{nticd}}$ and $\rightarrow_{\text{ntscd}}$, but also for $\rightarrow_{\text{tscd}}$. Thus we have been able to obtain efficient
 229 implementations for all these dependence notions. The algorithms are described in section 5.
 230 Performance evaluations are described in section 6.
 231

232 2 CONTROL DEPENDENCE IN GRAPHS WITHOUT UNIQUE EXIT

233 Our work was strongly motivated by earlier results of Ranganath et al. [29] and Amtoft
 234 [3]. These authors extended the classical notion of CD and slicing to CFGs which do not
 235 contain a unique exit node. As multiple exit nodes can trivially be handled by adding a new
 236 "global" exit node, Ranganath's and Amtoft's work is in fact concerned with CFGs which do
 237 not have a single, unique exit node. A typical example is a CFG with an infinite loop from
 238 which an exit node cannot be reached. Such CFGs are relevant, because modern programs
 239 need not necessarily terminate through exit nodes. One paramount example are reactive
 240 systems, which are assumed to run forever; and thus have no exit node at all. Another
 241 example are event handlers, which may shutdown a thread while the thread has no explicit
 242 exit. Thus Ranganath and Amtoft opened the door to apply CD and slicing to modern
 243 program structures.
 244

245

Time-sensitive CD will also work on graphs without unique exit. It is therefore necessary to recall Ranganath’s and Amtoft’s work. We begin with fundamental definitions of CDs and postdomination for CFGs with no unique exit.

2.1 Classical Control Dependence and Weak Control Dependence

CFGs are a standard representation of programs e.g. in compilers, and many tools are available which extract CFGs from source code.² Thus let $G = (N, \rightarrow_G)$ be the CFG of a program. In this article, we once and for all assume a fixed CFG G and therefore omit the sub- or superscript G whenever possible; e.g. we write $n \rightarrow m$ instead of $n \rightarrow_G m$. In the classical case of a unique exit node, there is $exit \in N$ such that $n \rightarrow^* exit$ for all $n \in N$, and $exit \rightarrow n$ for no node $n \in N$.

Node m postdominates n ($m \sqsubseteq_{\text{POST}} n$) iff $m \in \pi$ for every path π from n to $exit$. Node m *strongly* postdominates n ($m \sqsubseteq_{\text{SPOST}} n$) iff $m \sqsubseteq_{\text{POST}} n$, and there exists some $k \geq 1$ such that $m \in \pi$ for every path π starting in n with length $\geq k$ [27]. In contrast to $m \sqsubseteq_{\text{POST}} n$, $m \sqsubseteq_{\text{SPOST}} n$ does *not* hold if there is an infinite loop between m and n : Assume such a loop exists, then there will be paths π starting at n of arbitrary length k which never pass through m : $\forall k \exists \pi : \text{len}(\pi) = k \wedge m \notin \pi$. If this happens, $\sqsubseteq_{\text{SPOST}}$ is *not* supposed to hold; hence the negation of the latter condition must hold for $\sqsubseteq_{\text{SPOST}}$.

Classical (nontermination-insensitive) CD, denoted \rightarrow_{cd} , is defined in terms of postdominance. Formally (as already explained above),

$$x \rightarrow_{cd} y \iff \exists n, m \in N : x \rightarrow n, x \rightarrow m, y \sqsubseteq_{\text{POST}} n, \neg(y \sqsubseteq_{\text{POST}} m)$$

This CD definition can be modified in order to react sensitively to infinite loops. This *nontermination sensitive* form of CD, called “weak control dependence” and written $x \rightarrow_{wcd} y$, was introduced in [27]; and is defined in terms of *strong* postdominance. The formal definition is identical to the above CD definition; with $\sqsubseteq_{\text{SPOST}}$ instead of $\sqsubseteq_{\text{POST}}$. Even if $x \rightarrow_{cd} y$ does not hold, $x \rightarrow_{wcd} y$ might still hold if there is an infinite loop between x and y . Note that weak control dependence does not imply that this infinite loop is in fact executed.

2.2 Postdominance in Graphs Without Unique Exit

In order to understand how the above definitions are generalized to arbitrary graphs with no unique exit node, consider the example in Figure 2. It has no unique exit node, since the only candidate node 10 is unreachable from, e.g., node 6. Thus the classical definitions for $\sqsubseteq_{\text{POST}}$ and $\sqsubseteq_{\text{SPOST}}$ cannot be applied. Instead in [29], Ranganath et al. proposed control dependence for arbitrary graphs based on the notions of *maximal* and *sink* paths.

A maximal path is a path which cannot be extended (i.e.: is infinite, or ends in some node n without successor). On the other hand, a (control-) sink is a strongly connected component S such that no edge leaves S .³ Specifically, all nodes n without successor (in particular $n = exit$) form a (trivial) sink. A sink path then is a path π such that $s \in S$ for some node $s \in \pi$ and some sink S , and if S is nontrivial (i.e. not a singleton), then all nodes in S appear in π infinitely often. In programming terms, S would be an infinite loop in the CFG, and a sink path corresponds to an execution which infinitely loops in S .

Definition 2.1 (Implicit in [29]). A node $m \in N$ nontermination-sensitively postdominates a node $n \in N$ (written $m \sqsubseteq_{\text{MAX}} n$) iff $m \in \pi$ for all maximal paths π starting in n . Similarly, a

²All examples and measurements in this article are based on CFGs which were produced using the JOANA system. JOANA is a system for IFC and can in particular check probabilistic noninterference for full Java with arbitrary threads [6, 13, 14].

³In a strongly connected component (SCC) S , there is a path between all $x, y \in S$. Every cycle is an SCC.

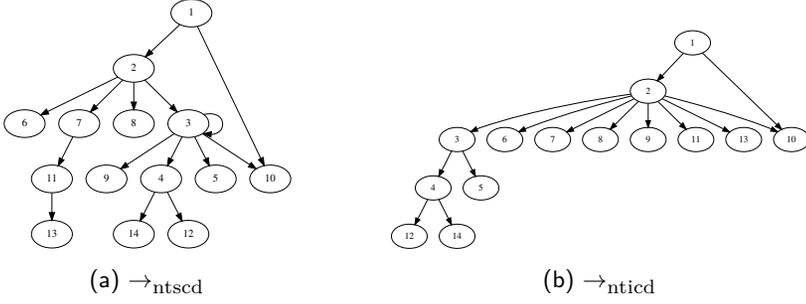


Fig. 3. Nontermination-(in)sensitive CDs for CFG from Figure 2

node m nontermination-insensitively postdominates a node n (written $m \sqsubseteq_{\text{SINK}} n$) iff $m \in \pi$ for all sink paths π starting in n .

Since every sink path is a maximal path, $m \sqsubseteq_{\text{MAX}} n$ implies $m \sqsubseteq_{\text{SINK}} n$. $m \sqsubseteq_{\text{SINK}} n$ while $\neg(m \sqsubseteq_{\text{MAX}} n)$ means that on reaching n , m will later be executed unless an infinite loop is entered.

The following definition is equivalent to those in [29] whenever $n \neq m$.

Definition 2.2. A node $y \in N$ is *non-termination sensitively (resp. insensitively) control-dependent on* $x \in N$, written $x \rightarrow_{\text{ntscd}} y$ (resp. $x \rightarrow_{\text{nticd}} y$), if there exist edges $x \rightarrow n$, $x \rightarrow m$ such that $y \sqsubseteq_{\text{MAX}} n$ (resp. $y \sqsubseteq_{\text{SINK}} n$), but $\neg(y \sqsubseteq_{\text{MAX}} m)$ (resp. $\neg(y \sqsubseteq_{\text{SINK}} m)$).

Note that this definition is identical to the original CD definition, with \sqsubseteq_{MAX} resp. $\sqsubseteq_{\text{SINK}}$ instead of $\sqsubseteq_{\text{POST}}$. In fact, for graphs with unique exit node, we have $\rightarrow_{\text{ntscd}} = \rightarrow_{\text{wcd}}$ and $\rightarrow_{\text{nticd}} = \rightarrow_{\text{cd}}$ [29]. Figure 3 shows $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{nticd}}$ for the CFG from Figure 2.

Like many program analysis problems, postdominance and CD can be characterized as a fixpoint computation. Our first new insight is that both \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ can be characterized as a greatest resp. least fix point of *one* rule set D . This surprising fact is the basis for our generalization of Cytron's algorithm. Note that the rule set can be interpreted as a functional which transforms a set of dominance relationships $\{x \sqsubseteq y\}$ into a new set $D(\{(x \sqsubseteq y)\})$. If such a functional is monotone, it has a least as well as a greatest fixpoint.

Theorem 2.1.⁴ Let D be the following rule system, and let D also denote its implicit functional; write μ for the least fixpoint, and ν the greatest fixpoint. Then D is a monotone functional in the (finite) lattice $(2^{N \times N}, \subseteq)$, and $\mu D = \sqsubseteq_{\text{MAX}}$, and $\nu D = \sqsubseteq_{\text{SINK}}$.

$$\text{Rule system } D : \quad \frac{}{n \sqsubseteq n} D^{\text{self}} \quad \frac{\forall n \rightarrow x : m \sqsubseteq x \quad n \rightarrow^* m}{m \sqsubseteq n} D^{\text{succ}}$$

The reachability side-condition $n \rightarrow^* m$ is in most cases redundant for the least fixed point μD , but essential for the greatest fixed point νD .⁵

Of course, algorithms for $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{nticd}}$ are needed, and indeed Ranganath et al proposed such algorithms. The algorithm for $\rightarrow_{\text{ntscd}}$ from [29] can in principle be thought of

⁴Lemmas and Theorems marked with  have been formalized and proved in the machine prover Isabelle. The proof explanations and scripts can be found in the electronic appendix of this article.

⁵We mention in passing that for graphs with unique exit node, replacing this condition with $n \neq \text{exit}$ results in a similar rule system P on which the algorithm from [8] is based. We will not describe P in detail, but note that $\sqsubseteq_{\text{POST}} = \nu P$ [15].

344 as a simple least fixed point computation of the set of nodes m such that $m \sqsubseteq_{\text{MAX}} n$, but
 345 only for nodes n that are successors of *branching* nodes.

346 We however discovered a more general and systematic algorithmic approach, which exploits
 347 the above fix-point theorem. It is based on the insight that Cytron’s efficient algorithm for
 348 dominance frontiers can be generalized to an abstract notion of “dominance”; and thus can
 349 be used for $\sqsubseteq_{\text{POST}}$, $\sqsubseteq_{\text{SPOST}}$, \sqsubseteq_{MAX} , $\sqsubseteq_{\text{SINK}}$, $\rightarrow_{\text{ntscd}}$, $\rightarrow_{\text{nticd}}$, and in particular for $\rightarrow_{\text{tscd}}$.
 350 We will present all algorithms in a separate section (section 5), as they demand a rather
 351 heavy technical machinery.

352 In conclusion of this section, we recall another notion from Ranganath et al., which will
 353 also be helpful to characterise $\rightarrow_{\text{tscd}}$.

354 **Definition 2.3** ([29]). *Decisive order dependence*, written $n \rightarrow_{\text{dod}} (m_1, m_2)$ is a ternary relation
 355 which means that n controls the *order* in which m_1, m_2 are executed.
 356

357 We omit the formal definition, but provide an intuition: In [29], the necessity of \rightarrow_{dod} was
 358 motivated by an irreducible⁶ graph, such as the graph shown in Figure 6a. Here, neither m_1
 359 nor m_2 is nontermination sensitively control dependent on n : $\neg(n \rightarrow_{\text{ntscd}} m_1) \wedge \neg(n \rightarrow_{\text{ntscd}}$
 360 $m_2)$. But the decision at n determines which node is executed next: Leaving n via the left
 361 branch will execute m_1 before m_2 , but leaving n via the right branch will execute m_2 before
 362 m_1 . Thus $n \rightarrow_{\text{dod}} (m_1, m_2)$ holds. Ranganath and Amtoft used $\rightarrow_{\text{ntscd}}$ and \rightarrow_{dod} to define a
 363 sound notion of *nontermination-sensitive backward slicing*. This is consistent with the fact
 364 that CDs are fundamental for slicing and PDGs.
 365

366 3 TIMING SENSITIVE CONTROL DEPENDENCE

367 3.1 Why Time-Sensitivity Matters

368 Before we formally develop timing sensitive CDs, let us motivate the usefulness of this
 369 concept for software security analysis. Known attacks exploiting timing side channels include
 370 Spectre [22] and cache attacks on implementations of the cryptographic standard AES
 371 [4]. In this kind of attacks, the attacker is able to observe the timing behaviour of certain
 372 instructions; from this observation determine whether some specific data are in the cache or
 373 not; and from this knowledge infer values of secret variables (e.g. by using the secret value
 374 as an array index), or draw conclusions about control flow.

375 Timing sensitive CDs can reveal such potential attacks, or prove that such attacks are
 376 impossible. For example, in a specific AES implementation⁷ we find the code lines
 377

```

378 (1)   for r1: [0,1,...,15]
379 (2)     r2 := state[r1]; // state depends on key and plain text
380 (3)     r3 := sbbox[r2]; // sbbox is a constant array
381 (4)     state[r1] := r3
382 (5)   end
```

383 `sbbox` is a constant array which typically spans multiple memory blocks, while `r1`, `r2`, `r3`
 384 are registers. Thus the value of `r2` in one iteration may influence whether the read of `r3` in a
 385 later iteration is served from cache (namely if, earlier, the corresponding memory block was
 386 already loaded into the cache), or from main memory. This makes a difference in execution
 387 time and can be observed by an attacker; who may thus be able to infer the value of `r2`.
 388

389 ⁶A CFG is reducible if the forward edges form a directed acyclic graph, and in any backedge $m \rightarrow n$, n
 390 dominates m . Structured programs have reducible CFGs; wild `gotos` typically lead to irreducible CFGs.

391 ⁷This implementation was presented in [4]. It assumes that all accesses to the `sbbox` array need constant time.
 392 But in fact access time is cache dependent.

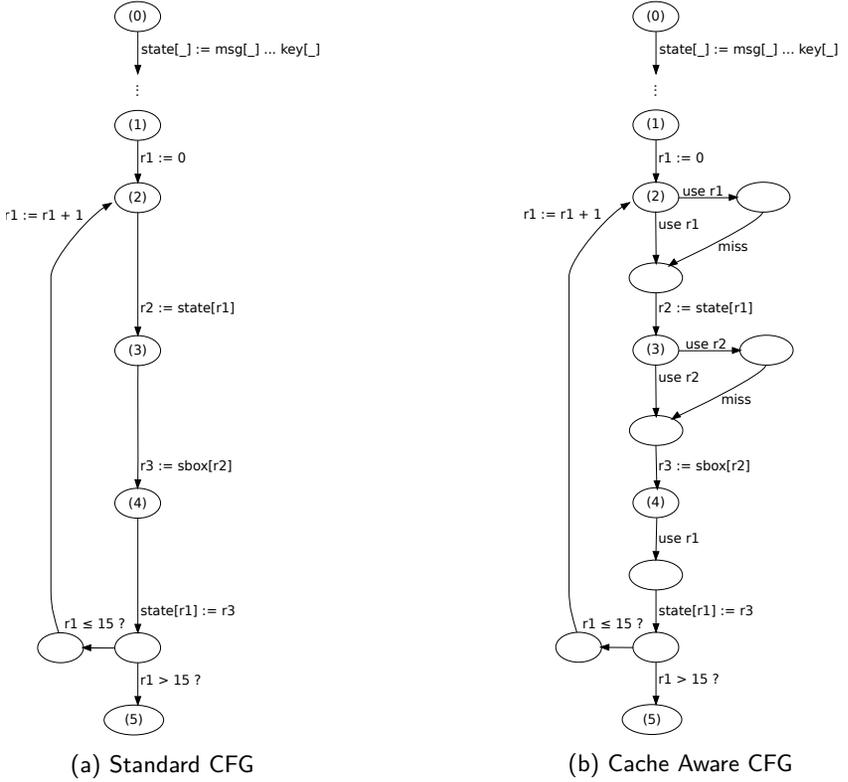


Fig. 4. Control Flow Graphs for AES Sbox substitution.

Such leaks can be discovered by timing-sensitive CDs; provided the CFG not just describes the code, but additionally models relevant hardware features such as cache behaviour.

Figure 4 shows the standard CFG for the AES code fragment, as well as the *micro-architectural CFG* which models timing differences due to cache hits and misses. The latter CFG indicates that array access $r3 := sbox[r2]$ in line (3) may either result in a cache hit or cache miss. In the control flow, this is modeled via two paths leaving node (3) that are joined after following a *different* number of edges, and take a different amount of time to execute. Specifically, the time at which execution reaches the exit node (5) depends on which paths are taken at (3): node (5) is time sensitively control dependent on node (3). The edge annotations *use r2* indicate that the value register $r2$ determines which array index, and hence which cache line is accessed at (3). Furthermore, due to the previous assignment $r2 := state[r1]$, node (3) is data-dependent on the initialization of the state array from the plain text message and the key, as indicated in node (0).⁸ Thus we obtain the following dependency chain (where \rightarrow_{dd} denotes data dependency): (secret input) \rightarrow_{dd} *state* \rightarrow_{dd} $r2 \rightarrow_{dd}$ (3) \rightarrow_{tscd} (5). That means: the time until (5) is reached depends on secret input. Thus time-sensitive CDs reveal that *sbox* access is not constant time (in contrast to the AES specification); opening a door to cache-leak based attacks.

⁸Besides CDs, data dependencies are important for security analysis. This is described in section 3.4. For the current AES example, the reader may assume all data dependencies are available as necessary.



(a) A CFG with external timing leak

(b) A CFG without timing leak

Fig. 5. Dependence of execution time of m_x on n .

If a $\rightarrow_{\text{tscd}}$ dependency is not (indirectly) data dependent on secret data, it does *not* generate a timing leak. E.g. in the AES CFG, (2) $\rightarrow_{\text{tscd}}$ (5) also holds. But the value of $r1$ (and thus (2)) is not data dependent on any secret value, so no timing leak arises via (2). We will discuss this in more depth in section 3.4; and come back to AES and micro-architectural CFGs in section 4.

3.2 Timing Sensitive Control Dependence

Consider Figure 5a: m_x is guaranteed to be executed, no matter which branch is taken at n , so we have $\neg(n \rightarrow_{\text{ntscd}} m_x)$. But let us assume that we could measure execution times. Now, n can control at which time m_x will be executed, namely 4 or 2 steps after executing n . We will say that m_x is timing sensitively control-dependent on n , or $n \rightarrow_{\text{tscd}} m_x$. In Figure 5b, however, m_x will always be executed 4 steps after n , so there we have $\neg(n \rightarrow_{\text{tscd}} m_x)$.

We will now formally define $\rightarrow_{\text{tscd}}$. Specifically, we will

- (1) Propose a notion $\sqsubseteq_{\text{TIME}}$ of *timing sensitive* postdominance.
- (2) Give a least fixed point characterization of $\sqsubseteq_{\text{TIME}}$.
- (3) Propose a notion $\rightarrow_{\text{tscd}}$ of *timing sensitive* control dependence. It will be based on $\sqsubseteq_{\text{TIME}}$ the same way that $\rightarrow_{\text{ntscd}}$ is based on \sqsubseteq_{MAX} .
- (4) Prove soundness and minimality of $\rightarrow_{\text{tscd}}$.

To start with, remember that \sqsubseteq_{MAX} was defined via

$$m \sqsubseteq_{\text{MAX}} n \iff \forall \pi \in {}_n\Pi_{\text{MAX}}. m \in \pi$$

where ${}_n\Pi_{\text{MAX}}$ is the set of maximal paths starting in n . For example, in Figure 5a it holds that $m_x \sqsubseteq_{\text{MAX}} n$, because any maximal path starting in any successor of n must contain m_x (i.e.: both $m_x \sqsubseteq_{\text{MAX}} n'$ and $m_x \sqsubseteq_{\text{MAX}} n''$), and so must any maximal path starting in n .

Now for time-sensitive postdominance we additionally want to express that in Figure 5a m_x can be reached via two different paths, *with varying execution time*. In order to account for the different *timing* of the (first) occurrence of m_x in maximal paths starting in n , the following auxiliary definition is needed.

Definition 3.1. Given any path $\pi = m_0, m_1, m_2, \dots$ we say that m appears in π at position k iff $m = m_k$, and write $m \in^k \pi$. If additionally, $m_i \neq m$ for all $i < k$, we say that m *first* appears in π at position k , and write $m \in_{\text{FIRST}}^k \pi$.

Using this notation, we can define time-sensitive postdominance as follows.

Definition 3.2. (a) m timing-sensitively postdominates n at position $k \in \mathbb{N}$, written $m \sqsubseteq_{\text{TIME}}^k n$, iff on all maximal paths starting in n , m first appears at position k . Formally

$$m \sqsubseteq_{\text{TIME}}^k n \iff \forall \pi \in {}_n\Pi_{\text{MAX}}. m \in_{\text{FIRST}}^k \pi$$

(b) m timing-sensitively postdominates n , written $m \sqsubseteq_{\text{TIME}} n$, if there exists k such that $m \sqsubseteq_{\text{TIME}}^k n$. Thus

$$m \sqsubseteq_{\text{TIME}} n \iff \exists k \in \mathbb{N} \forall \pi \in {}_n\Pi_{\text{MAX}}. m \in_{\text{FIRST}}^k \pi$$

If we compare $m \sqsubseteq_{\text{TIME}}^k n$ to $m \sqsubseteq_{\text{MAX}} n$, the difference is that in the latter, m must occur somewhere in all maximal paths from n , while in the former m must first occur at a specific position k in all maximal paths from n . Thus if $m \sqsubseteq_{\text{TIME}} n$, m must appear in all maximal paths from n *at the same position*. Therefore in Figure 5a, $m_x \sqsubseteq_{\text{TIME}} n$ does not hold, while in Figure 5b, $m_x \sqsubseteq_{\text{TIME}} n$ does hold.

Lemma 3.1. 🌀 Given m and n , the k such that $m \sqsubseteq_{\text{TIME}}^k n$ (if it exists) is unique.

Following the definitions for nontermination sensitive and insensitive control dependence $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{nticd}}$, we define the following timing sensitive notion of control dependence:

Definition 3.3. y is said to be *timing sensitively control-dependent* on x , written $x \rightarrow_{\text{tscd}} y$, if there exist edges $x \rightarrow n$ and $x \rightarrow m$ as well as some $k \in \mathbb{N}$ such that

$$y \sqsubseteq_{\text{TIME}}^k n \text{ and } \neg(y \sqsubseteq_{\text{TIME}}^k m)$$

Note that this definition is identical to the definition of $\rightarrow_{\text{ntscd}}$ resp. \rightarrow_{cd} ; with $\sqsubseteq_{\text{TIME}}^k$ instead of \sqsubseteq_{MAX} resp. $\sqsubseteq_{\text{POST}}$. Thus $\rightarrow_{\text{tscd}}$ has the same formal structure as classical CD and its later extensions.

In Figure 5a we have $m_x \sqsubseteq_{\text{TIME}}^3 n'$ but $\neg(m_x \sqsubseteq_{\text{TIME}}^3 n'')$, thus we have $n \rightarrow_{\text{tscd}} m_x$; while in Figure 5b we have $m_x \sqsubseteq_{\text{TIME}}^3 n'$ and $m_x \sqsubseteq_{\text{TIME}}^3 n''$ and thus *not* $n \rightarrow_{\text{tscd}} m_x$. For more complex examples, consider again the CFG in Figure 2. The timing sensitive postdominance for this CFG is shown in Figure 7b. Figure 7c and Figure 7d show the corresponding non-termination sensitive and timing sensitive control dependencies. Note, for example, that $7 \rightarrow_{\text{tscd}} 8$ because a choice $7 \rightarrow_G 11$ can delay node 8, but in contrast: $\neg(7 \rightarrow_{\text{ntscd}}^* 8)$, because no choice at node 7 can *prevent* node 8 from being executed. It is *not* the case that, in general, $n \rightarrow_{\text{ntscd}} m$ implies $n \rightarrow_{\text{tscd}} m$. For example: $2 \rightarrow_{\text{ntscd}} 8$, but $\neg(2 \rightarrow_{\text{tscd}} 8)$. What *does* hold here is $2 \rightarrow_{\text{tscd}}^* 8$ via $2 \rightarrow_{\text{tscd}} 7 \rightarrow_{\text{tscd}} 8$.

We will now provide a fixpoint characterization of $m \sqsubseteq_{\text{TIME}}^k n$. Remember from Theorem 2.1 that \sqsubseteq_{MAX} is the least fixed point of the rule system D

$$\frac{}{n \sqsubseteq n} \text{D}^{\text{self}} \quad \frac{\forall n \rightarrow x. m \sqsubseteq x \quad n \rightarrow^* m}{m \sqsubseteq n} \text{D}^{\text{succ}}$$

in the lattice $(2^{N \times N}, \sqsubseteq)$. Similarly, the ternary relation $m \sqsubseteq_{\text{TIME}}^k n$ is the least fixed point of the rule system T_{FIRST} in the underlying lattice $(2^{N \times N \times N}, \sqsubseteq)$.

Theorem 3.1. 🌀 Let T_{FIRST} be the rule-system

$$\frac{}{n \sqsubseteq^0 n} \text{T}_{\text{FIRST}}^{\text{self}} \quad \frac{\forall n \rightarrow x. m \sqsubseteq^k x \quad m \neq n \quad n \rightarrow^* m}{m \sqsubseteq^{k+1} n} \text{T}_{\text{FIRST}}^{\text{succ}}$$

Then $\sqsubseteq_{\text{TIME}} = \mu \text{T}_{\text{FIRST}}$.

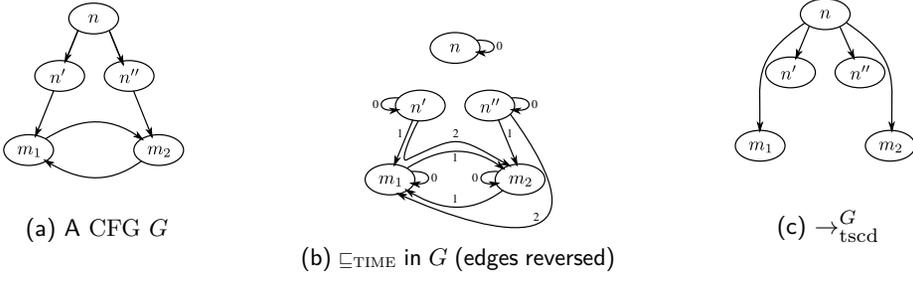


Fig. 6. The canonical irreducible graph, where neither $n \rightarrow_{\text{ntscd}} m_1$ nor $n \rightarrow_{\text{ntscd}} m_2$.

Note that the condition $n \rightarrow^* m$ is redundant for nodes n that have *some* successor x , since we consider only the least, but not the greatest, fixed point of T_{FIRST} . The condition $m \neq n$ ensures that we only consider the first occurrence of m in each path.

We will now demonstrate that $\rightarrow_{\text{tsCD}}$ is *transitively* a stricter requirement than non-termination sensitive control independence. To this end, we use the following notation.

Definition 3.4. For $M \subseteq N$ and \rightarrow a relation on N , the backward slice of M is

$$(\rightarrow)^*(M) = \{y \mid \exists x \in M : y \rightarrow^* x\}$$

This definition can be generalized to the ternary relation \rightarrow_{dod} : if $y \rightarrow_{\text{dod}}(x_1, x_2)$, $y \in (\rightarrow_{\text{dod}})^*(M)$ only if x_1 and $x_2 \in (\rightarrow_{\text{dod}})^*(M)$ [29].

Theorem 3.2. Let $M \subseteq N$. Then

$$(\rightarrow_{\text{tsCD}})^*(M) \supseteq (\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{dod}})^*(M)$$

That is, there are more transitive time-sensitive CDs than the transitive closure of even the union of $\rightarrow_{\text{ntscd}}$ and \rightarrow_{dod} . Now remember that Ranganath and Amtoft introduced $\rightarrow_{\text{ntscd}}$ and \rightarrow_{dod} in order to provide a sound notion of nontermination-sensitive backward slicing. Thus in the language of PDGs, $(\rightarrow)^*(M)$ is just the backward slice of M , and the theorem states that the timing sensitive backward slice of M contains the nontermination sensitive backward slice of M .

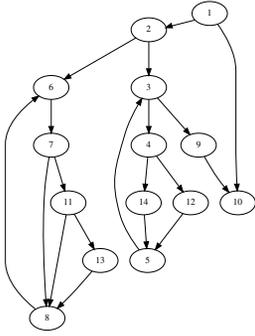
It is worth noting that the $\rightarrow_{\text{tsCD}}$ slice in Theorem 3.2 does *not* require a timing sensitive analogue of the relation \rightarrow_{dod} . As seen above, the necessity of \rightarrow_{dod} was motivated by an irreducible graph, such as the graph in Figure 6. But while in Figure 6 neither m_1 nor m_2 is nontermination sensitively control dependent on n , *both* m_1 and m_2 are timing-sensitively control dependent on n (e.g.: $n \rightarrow_{\text{tsCD}} m_1$ because $m_1 \sqsubseteq_{\text{TIME}}^1 n'$, but $\neg(m_1 \sqsubseteq_{\text{TIME}}^1 n'')$, and also: $m_1 \sqsubseteq_{\text{TIME}}^2 n''$, but $\neg(m_1 \sqsubseteq_{\text{TIME}}^2 n')$). This m_1/m_2 symmetry makes a ternary “ $\rightarrow_{\text{tsdod}}$ ” unnecessary.

3.3 Soundness and minimality of $\rightarrow_{\text{tsCD}}$

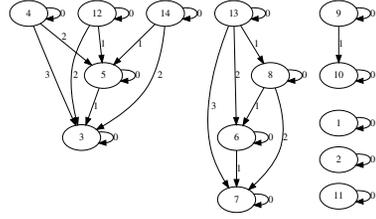
It is our ultimate goal to discover timing leaks. We thus need a soundness proof for $\rightarrow_{\text{tsCD}}$, which guarantees that $\rightarrow_{\text{tsCD}}$ will indeed discover all potential timing leaks. We will further show that $\rightarrow_{\text{tsCD}}$ is minimal, which means there are no spurious time-sensitive dependencies.

Any soundness proof makes assumptions about the possibilities of attackers; this is called the attacker model. To prove soundness of $\rightarrow_{\text{tsCD}}$ under an attacker model, we use a technique called trace equivalence. Let us thus describe our attacker model, and then define trace

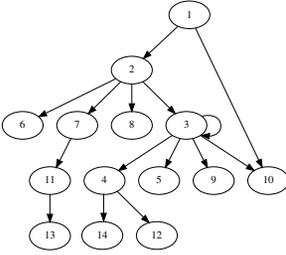
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637



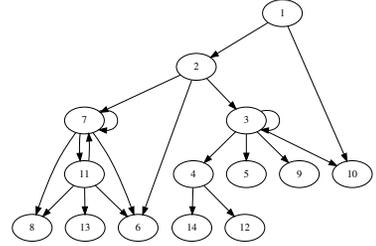
(a) CFG from Figure 2



(b) Its relation $\sqsubseteq_{\text{TIME}}$ (edges reversed)



(c) Its non-termination sensitive control dependence $\rightarrow_{\text{ntscd}}$



(d) Its timing sensitive control dependence $\rightarrow_{\text{tscd}}$

Fig. 7. Timing sensitive postdomination. Edges $n \xrightarrow{k} m$ indicate $m \sqsubseteq_{\text{TIME}}^k n$.

equivalence. We imagine an attacker who tries to infer secret values (such as $\mathbf{r2}$ in the AES example) measuring execution times for certain execution paths. But the attacker cannot observe all nodes, he can only observe certain “observable” nodes.⁹ The goal of security analysis is then to guarantee that secret information cannot flow to observable nodes, resp. that execution times measured at observable nodes will not allow the attacker to infer secret values at unobservable nodes.¹⁰ Technically, for $\rightarrow_{\text{tscd}}$ this guarantee is based on trace equivalence of *clocked traces*.

Definition 3.5. An (unlocked) trace t is a sequence of edges $(n, n') \in (\rightarrow_G) \cup (N_x \times \{\perp\})$ that is either finite with $t = (n_e, n_1), (n_1, n_2), \dots, (n_k, n_x), (n_x, \perp)$ for some exit node $n_x \in N_x$, or infinite with $t = (n_e, n_1), (n_1, n_2), \dots$. Partial edges (n, \perp) occur only at exit nodes.

Definition 3.6. A *clocked trace* is a trace where every step is additionally annotated with a time stamp. We write $t \overset{\circledast}{\circlearrowleft} [i]$ if a trace step t has time stamp i . Given a trace $t = (n_e, n_1), (n_1, n_2), \dots$, its *clocked version* is thus

$$t^{\circledast} = (n_e, n_1) \overset{\circledast}{\circlearrowleft} [0], (n_1, n_2) \overset{\circledast}{\circlearrowleft} [1], \dots$$

⁹These observable nodes are called “low” nodes in the literature on software security analysis (see e.g. [31]).

¹⁰This kind of security analysis is called *information flow control* (IFC), and is based on the technical notion of *noninterference*. We will describe technical details on the application of $\rightarrow_{\text{tscd}}$ for IFC in a separate article; here we present only the AES example and do not discuss technical details of noninterference.

Next, we assume there is a fixed set $S \subseteq N$ of observable nodes.¹¹

Definition 3.7. Let $S \subseteq N$; let a trace t be given. We define the S -observation $t|_S$ of t to be the sub-sequence of t containing only edges (n, n') with $n \in S$. Traces t_1, t_2 are called S -equivalent if $t_1|_S = t_2|_S$.

These definitions work for unlocked and clocked traces. S -observability means that we assume an attacker to observe exactly those choices made at nodes $n \in S$. Specifically, we assume that an attacker can observe neither the nodes in a subtrace between observable nodes, nor – for unlocked traces – the *time spent* between two observable nodes (i.e: the *length* of the subtrace between two observable nodes).

Now we consider traces caused by specific inputs. We write t_i for the (possibly infinite) trace caused by input i . As we want to abstract away from particular input formats or data objects, we use a nonstandard formalization of input: i is a map from CFG nodes to a (perhaps infinite) list of CFG successor nodes: $i : N \rightarrow N^*$. An input i causes t_i as follows: If e.g. an **if** node $n \in N$ is visited for the k th time during the execution with input i , the execution will continue with the k th element of $i(n)$, which is a successor node (i.e. true or false path) of n . If n is only visited finitely often, superfluous entries in $i(n)$ are ignored.

This encoding has the effect that our CFGs are *state-free*: they contain CDs and nothing else. In particular the CFG does not contain program variables or program state – these are hidden in the i encoding. From a practical viewpoint this is however no restriction, and no weakening of the soundness property: we do not constrain possible i , and the soundness theorem below holds for all i, i' . Note however that for practical discovery of timing leaks, *data dependences* are additionally needed; this is described in section 3.4.

Next, we need the notion of S -equivalent inputs. For $S \subseteq N$, $i|_S : S \rightarrow N^*$ is the restriction of the map i to nodes $n \in S$, thereby only determining the successor nodes chosen at condition nodes $\in S$. Two inputs i, i' are called S -equivalent, written $i \sim_S i'$, if $i|_S = i'|_S$. An attacker cannot distinguish S -equivalent inputs.

We will now explain why – in the absence of timing leaks – S -equivalent inputs demand S -equivalent traces. It is essential to consider *clocked* traces: even if two unlocked traces are S -equivalent, their clocked versions may be different. This is the essence of time-sensitivity! For illustration consider Figure 5a, with observable nodes $S = \{m, m_x\}$. Regardless of the choice made at n , all inputs i, i' starting in m have the same observable trace

$$t_i|_S = (m, n), (m_x, \perp) = t_{i'}|_S$$

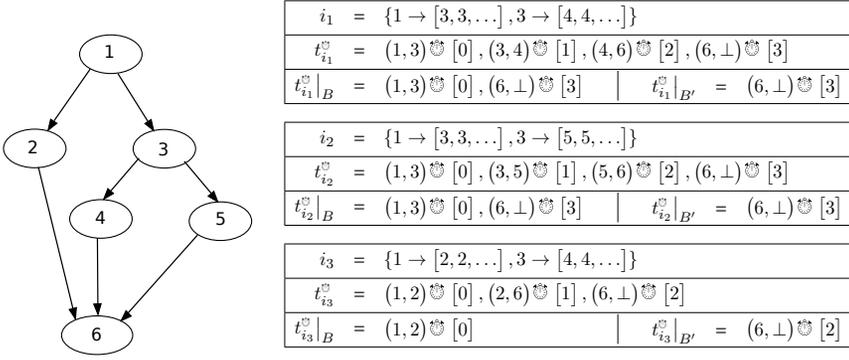
Hence t_i and $t_{i'}$ are always trace equivalent. Thus an attacker without clock cannot extract any secret information from observing traces. However, if equipped with a suitably precise clock, an attacker will observe m_x after 5 steps for the input i that chooses n' at n , but already after 3 steps for i' that chooses n'' at n , exposing a timing difference. This becomes obvious if we use the clocked versions of $t_i, t_{i'}$, and then compare their S -observation:

$$\begin{aligned} t_i^{\clubsuit}|_S &= (m, n) \clubsuit [0], (m_x, \perp) \clubsuit [5] \\ &\neq (m, n) \clubsuit [0], (m_x, \perp) \clubsuit [3] = t_{i'}^{\clubsuit}|_S \end{aligned}$$

Since the attacker cannot distinguish i and i' (they only differ in the choices for the unobservable node n), this timing difference allows the attacker to gain additional information, leading to a timing leak. On the other hand, the program in Figure 5b has no timing leak:

¹¹The assumption of a fixed, static S and batch-like execution is standard in IFC and noninterference. It can be generalised and made more realistic in various ways; which however is not a topic of this article. Likewise, technical details of noninterference will not be discussed in this article.

687
688
689
690
691
692
693
694
695
696
697
698



699 Fig. 8. In the CFG on the left, let $M = \{6\}$ be the slicing criterion. Then $B = BS(\{6\}) = \{1, 6\}$ is the
700 time-sensitive backward slice of M , because $1 \rightarrow_{\text{tscd}} 6$. $B' = \{6\}$ is a slice that is too small. Right: 3
701 different inputs with their traces and observable behaviour regarding B and B' .
702

703 even if we annotate each edge in the observable trace with its execution time, all inputs i, i'
704 starting in m have the same observable *clocked* trace

705
706

$$t_i^c|_S = (m, n) \circlearrowleft [0], (m_x, \perp) \circlearrowleft [5] = t_{i'}^c|_S$$

707

This discussion motivates the following definition of timing leaks:

708

709 **Definition 3.8.** Let $S \subseteq N$ be a set of observable (“low”) nodes. A program is free of timing
710 leaks if for all inputs i, i'

711

$$i \sim_S i' \implies t_i^c|_S = t_{i'}^c|_S$$

712

This definition is formally identical to classical noninterference definitions (cmp. e.g. [31]),
713 but is based on clocked traces.

714

To prevent a timing leak, it is necessary that *all nodes which influence the timing of*
715 *observable nodes $\in S$ are observable itself*. Otherwise, a secret node might influence the
716 timing of an observable node. For example, Figure Figure 5a contains – as described above
717 – a timing leak if we assume $S = \{m, m_x\}$. Indeed $n \rightarrow_{\text{tscd}} m_x$, but *not* $n \in S$. With
718 $S' = S \cup \{n\} = \{m, n, m_x\}$, the timing leak disappears: While the timing of m_x still differs,
719 i and i' are now distinguishable for the attacker, so this timing difference does not give
720 additional information.

721

We will now show how $\rightarrow_{\text{tscd}}$ can be used to check for timing leaks. In particular we
722 demonstrate that for any observable $M \subseteq S$, the time-sensitive backward slice $B = (\rightarrow_{\text{tscd}}$
723 $)^*(M)$ fulfills the condition of definition 3.8. This implies that B is not too small, i.e. $\rightarrow_{\text{tscd}}$
724 is sound.

725

Before we state the theorem, consider what happens if B is too small. In that case, the
726 $\rightarrow_{\text{tscd}}$ dependency would have “missing edges”. Then there could exist two inputs that agree
727 on B , but lead to different traces: $i \sim_B i'$ but $t_i^c|_B \neq t_{i'}^c|_B$. Figure 8 presents one such
728 example. For $B = BS(\{6\}) = \{1, 6\}$, we have $i_1 \sim_B i_2$ and indeed $t_{i_1}^c|_B = t_{i_2}^c|_B$. In contrast,
729 the unsound “slice” $B' = \{6\}$ leads to $i_1 \sim_{B'} i_3$ but $t_{i_1}^c|_{B'} \neq t_{i_3}^c|_{B'}$. (Note that the only
730 difference between the two slices is the timing of $(6, \perp)$, so we have $t_i^c|_{B'} = t_{i'}^c|_{B'}$ for the
731 *unclocked* traces. In fact, $\{6\}$ is a sound slice when ignoring timing and using $\rightarrow_{\text{ntscd}}$.) If
732 however $i \sim_B i'$ always implies $t_i^c|_B = t_{i'}^c|_B$, soundness is guaranteed.
733

734

735

Theorem 3.3 (Soundness of $\rightarrow_{\text{tscd}}$).  Let $M \subseteq S$. Let $B = (\rightarrow_{\text{tscd}})^*(M)$ be the timing sensitive backward slice w.r.t M . Then, for any inputs i, i' such that $i \sim_B i'$, we have

$$t_i^{\otimes} \upharpoonright_B = t_{i'}^{\otimes} \upharpoonright_B$$

Corollary 3.1.  If $BS(S) \subseteq S$, definition 3.8 holds, i.e. there is no timing leak.

As $S \subseteq BS(S)$ always holds, the corollary's premise is in fact $S = BS(S)$. If the premise is not satisfied, i.e. for some $x \in BS(S)$: $x \notin S$, x – as explained above – is a timing leak.

Minimality of slicing now shows that $B = BS(M)$ is as small as possible: Any set of nodes B' that includes the slicing criterion M can only be secure if it is a superset of B .

Theorem 3.4 (Minimality of $\rightarrow_{\text{tscd}}$).  Under the assumptions of Theorem 3.3, for any $B' \supseteq M$ with $B' \not\supseteq B$ there exist inputs i, i' such that $i \sim_{B'} i'$, but:

$$t_i^{\otimes} \upharpoonright_{B'} \neq t_{i'}^{\otimes} \upharpoonright_{B'}$$

It should be noted that the proof for both theorems relies on the non-standard, state-free input encoding of i, i' , which was described above.

3.4 The Full Time-Sensitive Backward Slice

Our nonstandard input encoding (which “factors away” all state information) is not practical for “real” programs. In such programs, time-sensitive influences through variables must be considered too. For this reason, discovery of timing leaks needs data dependences in addition to control dependences. Data dependences have in fact already been used in the AES example. For completeness and better understanding, we will thus describe the full algorithm for discovering timing leaks. Note that in this article, we do not provide a modified soundness proof for the full algorithm, as it does not contribute to $\rightarrow_{\text{tscd}}$ “as such”.

We denote data dependencies by \rightarrow_{dd} . $x \rightarrow_{dd} y$ means that a variable v which is defined (assigned) at x is used at y ; provided there is a CFG path $x \rightarrow^* y$, and v is not redefined on this path [11]. We will not describe the construction of \rightarrow_{dd} in detail, but note that for full languages with functions, objects, multithreading etc. the computation of precise data dependencies is nontrivial and requires context-sensitive summary dependencies, precise points-to analysis, may-happen-in-parallel analysis, and much more (see e.g. [14, 23, 30]).

The full algorithm for discovering timing leaks then assumes \rightarrow_{dd} , and proceeds as follows.

- (1) Compute $\rightarrow_{\text{tscd}}$. If $x \rightarrow_{\text{tscd}} y$, but *not* $x \rightarrow_{cd} y$, then there may be a timing leak at y , but only if it can be influenced by secret data.
- (2) Using \rightarrow_{dd} , the full time-sensitive backward slice is defined as

$$BS_{ts}(M) = (\rightarrow_{\text{tscd}} \cup \rightarrow_{dd})^*(M)$$

This slice contains all CFG nodes which may influence M ; other nodes which influence M cannot exist [6, 14, 18].

- (3) Now if $x \rightarrow_{\text{tscd}} y$, and $BS_{ts}(\{x\})$ contains any secret input or variables, there is a timing leak at y : the execution time between x and y varies, depending on secret data.

This procedure is fully analogous to the slicing-based noninterference check used in JOANA (see [6, 14]; these papers include soundness proofs and other details about slicing-based IFC), but with $\rightarrow_{\text{tscd}}$ instead of \rightarrow_{cd} . Note that in the current article, we consider only context-insensitive timing-dependencies (while JOANA uses context-sensitive, object-sensitive dependencies). A context-sensitive $\rightarrow_{\text{tscd}}$ is future work.

4 TIMING SENSITIVITY FOR MICROARCHITECTURAL CFGS

Already in the abstract we mentioned the infamous AES cache timing leaks which were discovered by Bernstein [4]. Some details of this attack were described in section 3.1. We will now describe in more detail how such cache leaks can be discovered resp. prevented via time-sensitive CDs in microarchitectural CFGs. Basically, the algorithm from section 3.4 is used, but the underlying CFG must be extended to model cache behaviour.

In the following, we describe this cache-modelling CFG extension in detail. The CFG edges are labeled with assignments and guards that refer to (cacheable) variables $\mathbf{a}, \mathbf{b}, \dots$, and uncacheable registers $\mathbf{r1}, \mathbf{r2}, \dots$.

We assume a simple data cache of size four, with a least recently used eviction strategy. The (micro-architectural) cache-state hence consists of a list $[x_1, x_2, x_3, x_4]$ of variables, with x_1 being the most recently used, and x_4 the next to be evicted. In Figure 9b, we show — under an abstraction that considers cache state only — all possible executions of the control flow graph, assuming an empty initial cache. For example, the abstract node $(9, [\mathbf{x}, \mathbf{d}, \mathbf{c}, \mathbf{b}])$ represents all those concrete configurations at control node 9 in which the concrete micro-architectural cache contains cached values for the variables $[\mathbf{x}, \mathbf{d}, \mathbf{c}, \mathbf{b}]$, in this order (with arbitrary concrete macro-architectural state).

In the example, executions can reach the control node $m = 15$ at cache states represented by either $[\mathbf{b}, \mathbf{y}, \mathbf{c}, \mathbf{x}]$, or by $[\mathbf{b}, \mathbf{y}, \mathbf{d}, \mathbf{x}]$. Which of these (abstract) cache states is reached is determined by the macro-architectural choice made at $n = 9$. But it is easy to see that the execution time of the read of \mathbf{y} at node $m = 15$ does *not* depend on the choice made at $n = 9$, since in both (classes of) executions that reach node $m = 15$, the cache *does* contain the variable \mathbf{y} , which is the only cacheable variable accessed by the edge $15 \xrightarrow{x_2:=y} 16$ at m .

For the read of variable \mathbf{b} at node $m = 14$, on the other hand, one class of executions reaches m in $(14, [\mathbf{y}, \mathbf{c}, \mathbf{b}, \mathbf{x}])$ (containing \mathbf{b}), while another class of executions reaches m in $(14, [\mathbf{y}, \mathbf{d}, \mathbf{x}, \mathbf{c}])$ (*not* containing \mathbf{b}). Whether the relevant variable \mathbf{b} is in the cache at $m = 14$ (and hence: the execution time of the read of \mathbf{b} at $m = 14$) or not depends here on the choice made at $n = 9$.

Now consider the read of \mathbf{c} at node $m = 21$. Does its cache state depend on the choice made right before at $n' = 16$? There are four (abstract) cache states at $m = 21$. Two contain the variable \mathbf{c} : $(21, [\mathbf{b}, \mathbf{y}, \mathbf{c}, \mathbf{x}])$ and $(21, [\mathbf{a}, \mathbf{y}, \mathbf{b}, \mathbf{c}])$. The other two do *not* contain \mathbf{c} : $(21, [\mathbf{a}, \mathbf{y}, \mathbf{b}, \mathbf{d}])$ and $(21, [\mathbf{b}, \mathbf{y}, \mathbf{d}, \mathbf{x}])$. The cache states containing \mathbf{c} are reachable from configurations at control node $n' = 16$. At the same time: cache states *not* containing \mathbf{c} are *also* reachable from configurations at control node $n' = 16$. But in fact, whether \mathbf{c} is in cache at m does *not* depend on the choice made at n' . To see this, note that node $n' = 16$ can be reached at two different cache states. The first abstract configuration is $(16, [\mathbf{y}, \mathbf{b}, \mathbf{c}, \mathbf{x}])$. But whenever $m = 21$ is reached from this abstract configuration, \mathbf{c} *is* in the cache (either $(21, [\mathbf{b}, \mathbf{y}, \mathbf{c}, \mathbf{x}])$ or $(21, [\mathbf{a}, \mathbf{y}, \mathbf{b}, \mathbf{c}])$). The second abstract configuration at which $n' = 16$ can be reached is $(16, [\mathbf{y}, \mathbf{b}, \mathbf{d}, \mathbf{x}])$. But whenever $m = 21$ is reached from that configuration, \mathbf{c} *is not* in the cache ($(21, [\mathbf{a}, \mathbf{y}, \mathbf{b}, \mathbf{d}])$ or $(21, [\mathbf{b}, \mathbf{y}, \mathbf{d}, \mathbf{x}])$).

On the other hand, the cache status of \mathbf{c} at node $m = 21$ *does* depend on the choice made earlier at $n = 9$. In this example this is necessarily so, since the node $n = 9$ is the only other macro-architectural conditional node in the control flow graph. But this is also directly evident by the structure of the graph in Figure 9b.

Note that through a small modification of the program, the cache status of \mathbf{c} at $m = 21$ could have been *independent* from the choice made earlier at $n = 9$. For example: had there been reads to two additional variables (e.g: \mathbf{e}, \mathbf{f}) right before $m = 21$, then *all* cache states

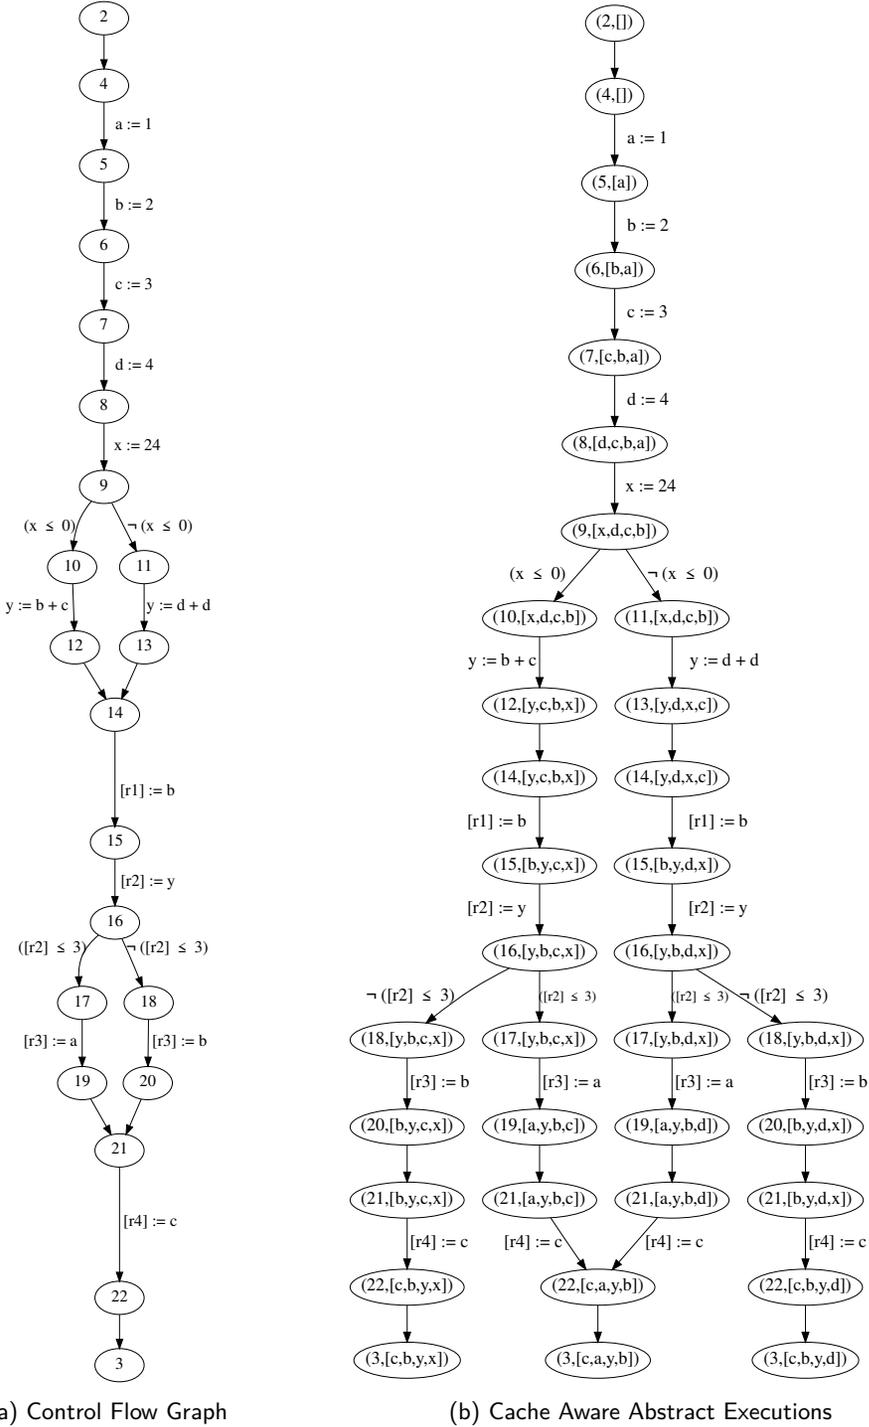


Fig. 9. A CFG and its possible cache-aware abstract executions.

883 at m would *not* have contained c . This is because these two reads would have evicted c even
 884 from $[b, y, c, x]$ (and $[a, y, b, c]$).

885 In summary, the choice made at $n = 9$ does influence the relevant (micro-architectural)
 886 cache state at $m \in \{21, 14\}$. In fact for this micro-architecture, these are the *only* micro-
 887 architectural dependencies in this CFG. The example indicates how a CFG G can be
 888 transformed into a cache-aware version. We will not present the formal definitions here (see
 889 [16]), but just present the transformed CFG for the above example.

890 Figure 10 shows the micro-architectural-aware CFG G' for Figure 9; together with an
 891 explicit timing cost model C' . A cache-miss is assumed to take 10 units of time, while a
 892 cache-hit takes 2 units¹². At node 14, the read from b takes either 2 or 10 units of time,
 893 since b there might either be in the cache, or not.¹³ Hence in G' , node 14 has *two* artificial
 894 successors: the read from b takes either 2 or 10 units of time, since b there might either
 895 be in the cache, or not. On the other hand, node 15 still has only one successor, reached
 896 with timing cost $3 = 2 + 1$ (cache access plus register access), since we found that there the
 897 variable y is always in cache.

898 In G' , we now have (as desired) that node 21 is in the backward slice of the exit node 3.
 899 Formally,

$$900 \quad 21 \in \left(\rightarrow_{\text{tscd}}^{G'[C']} \right)^* (\{3\})$$

902 Together with the microarchitectural dependence from node 9 to node 21, we conclude that
 903 the decision at node 9 may influence the execution time of node 3.

904 Note that even if G is deterministic, G' usually is not. This is no problem, because we
 905 can still use the micro-architectural dependencies $\rightarrow_{\mu d}^G$ (and data dependencies \rightarrow_{dd}) from
 906 the original graph G , and only use G' for timing sensitive control dependence $\rightarrow_{\text{tscd}}^{G'[C']}$.

907 For the AES code, the cache-sensitive graph G' has been shown in section 3.1, and we
 908 already described how cache leaks in AES have been discovered through time-sensitive
 909 backward slicing. More details can be found in [16].

912 5 ALGORITHMS

913 Our algorithms are based on the fundamental insight that Cytron's original algorithm for
 914 dominance frontiers can be generalized to CFGs with loops and multiple exit nodes; and
 915 even to the computation of time-sensitive CD. We consider this "generic" algorithm our
 916 major contribution: without it, the new $\rightarrow_{\text{tscd}}$ definition would be worthless in practice; and
 917 even Ranganath's and Amtoft's $\rightarrow_{\text{ntscd}}/\rightarrow_{\text{nticd}}$ are more efficient to compute using the new
 918 algorithms.

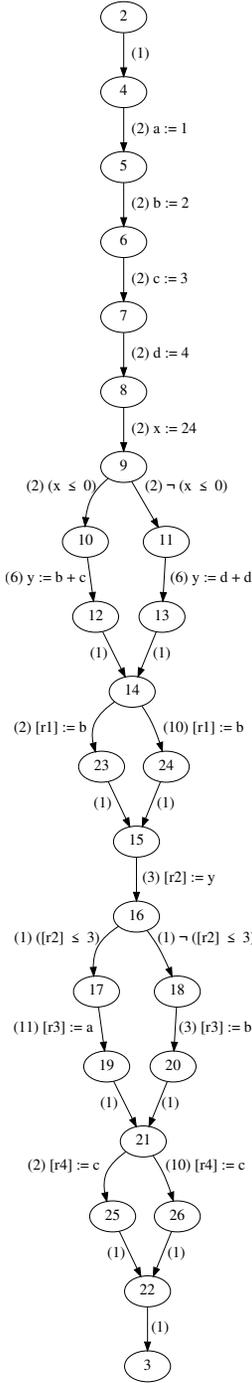
920 5.1 New algorithms for \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$

921 Let us begin with new algorithms for $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{nticd}}$. These will – in generalization of
 922 Cytron's approach – be constructed as postdominance frontiers of \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$. The
 923 efficient implementation of \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ needs some technical machinery, namely transitive
 924 reductions and pseudo-forests.

925 Both \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ will always be represented by their transitive reductions; allowing
 926 efficient construction algorithms. A transitive reduction $<$ of a transitive relation \sqsubseteq is

927 ¹²memory writes are assumed to always take 2 units of times, and register accesses take 1 unit of time

928 ¹³In the timing cost model C , the cost 11 = 10 + 1 that stems from one uncached variable access plus one
 929 register access is split into two edges. We need to do this because in our notion of graphs, there can be no
 930 multi-edges, and we require cost models C to be *strictly* positive.

Fig. 10. Micro-Architecture Aware CFG G' for Figure 9.

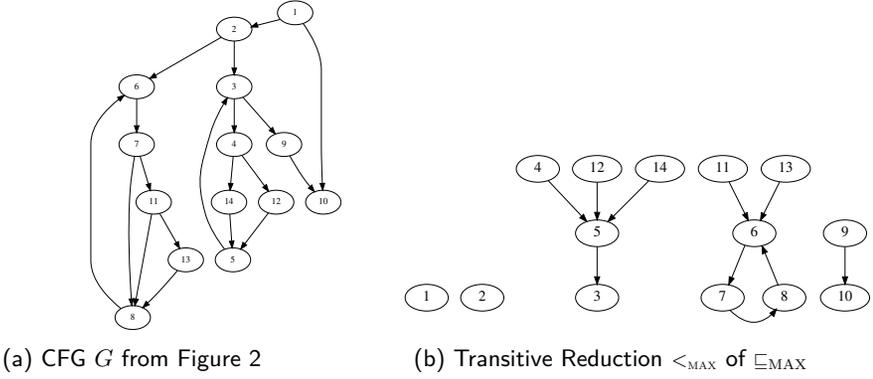


Fig. 11. Nontermination-sensitive Postdominance

a minimal subset $<$ of \sqsubseteq such that $(<)^* = \sqsubseteq$. Thus $<$ has a minimal number of edges but the same transitive closure as \sqsubseteq . Efficient algorithms for transitive reductions have long been known [2]. But remember that \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ may contain cycles (i.e. are not antisymmetric), in contrast to the classical $\sqsubseteq_{\text{POST}}$. Hence their transitive reductions may also contain cycles. Therefore the transitive reductions of \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ are not forests (i.e. sets of trees) as for $\sqsubseteq_{\text{POST}}$, but so-called *pseudo-forests*.

Definition 5.1. A pseudo-forest is a relation $<$ such that for every node $n \in N$, $m < n$ for at most one node m .

Thus in a pseudo-forest every node has at most one parent node, but in contrast to ordinary forests, pseudo-forests may contain cycles. Summarizing this discussion, we obtain

Observation 5.1. 1. Both \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ are reflexive and transitive, but not necessarily anti-symmetric.

2. Any transitive, reflexive reduction $<_{\text{MAX}}$ of \sqsubseteq_{MAX} is a pseudo-forest.

3. Any transitive, reflexive reduction $<_{\text{SINK}}$ of $\sqsubseteq_{\text{SINK}}$ is a pseudo-forest.

Figure 11 (b) shows a reduction $<_{\text{MAX}}$ of \sqsubseteq_{MAX} for the CFG in Figure 11 (a). This pseudo-forest has five trees, with roots 1, 2, 3, $\{6, 7, 8\}$ and 10.¹⁴ Node 9 does *not* \sqsubseteq_{MAX} -postdominate node 3 because the loop at 3 may not terminate. On the other hand, node 9 *does* $\sqsubseteq_{\text{SINK}}$ -postdominate node 3: a path looping forever at 3 is *not* a sink path, and any sink path starting at 3 must eventually reach the trivial sink at node 10.

We will now present new algorithms to compute \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$. The representation of both \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ by pseudo-forests is crucial, as pseudo-forests admit efficient algorithms for their computation. Based on pseudo-forests, our algorithm for \sqsubseteq_{MAX} is a standard fixpoint iteration. Beginning with the empty pseudo-forest, new edges are added to $<_{\text{MAX}}$ according to Theorem 2.1 until a fixpoint is reached. Since \sqsubseteq_{MAX} is efficiently represented by a pseudo-forest $<_{\text{MAX}}$, it is straightforward to derive an efficient algorithm for the computation of \sqsubseteq_{MAX} , see algorithm 2. In addition, we need an efficient implementation of set-intersection

¹⁴In the figure, downarrows $n \rightarrow m$ mean that $m < n$.

```

1030 Input      : A pseudo-forest  $\lt$ , represented as a map  $\text{IMDOM} : N \hookrightarrow N$  s.t.
1031              $\text{IMDOM}[n] = m$  iff  $m \lt n$ .
1032 Input      : Nodes  $n_0, m_0$ 
1033 Output     : A least common ancestor of  $n_0, m_0$ , or  $\perp$  if there is none.
1034 begin
1035   | return  $\text{lca}(n_0, m_0)$ 
1036 end
1037 Function  $\text{lca}(\pi_n, \pi_m)$ 
1038   | Input      : A  $\lt$ -path  $\pi_n = n_0, \dots, n$  ending in  $n$ 
1039   | Input      : A  $\lt$ -path  $\pi_m = m_0, \dots, m$  ending in  $m$ 
1040   | if  $m \in \pi_n$  then return  $m$ 
1041   | switch  $\text{IMDOM}[n]$  do
1042   |   | case  $\perp$  do return  $\text{lin}[\pi_n](\pi_m)$ 
1043   |   | case  $n'$  do
1044   |   |   | if  $n' \in \pi_n$  then
1045   |   |   |   | return  $\text{lin}[\pi_n](\pi_m)$ 
1046   |   |   |   end
1047   |   |   return  $\text{lca}(\pi_m, \pi_n n')$ 
1048   |   end
1049 end
1050 Function  $\text{lin}[\pi_n](\pi_m)$ 
1051   | Input      : A  $\lt$ -path  $\pi_m = m_0, \dots, m$  ending in  $m$ 
1052   | Implicit    : A  $\lt$ -path  $\pi_n = n_0, \dots, n$  ending in  $n$ 
1053   | switch  $\text{IMDOM}[m]$  do
1054   |   | case  $\perp$  do return  $\perp$ 
1055   |   | case  $m'$  do
1056   |   |   | if  $m' \in \pi_n$  then return  $m'$ 
1057   |   |   | if  $m' \in \pi_m$  then return  $\perp$ 
1058   |   |   | return  $\text{lin}[\pi_n](\pi_m m')$ 
1059   |   |   end
1060   |   end
1061 end

```

Algorithm 1: A *least common ancestor* algorithm for pseudo-forests. $N \hookrightarrow N$ denotes a partial map from N to N .

in the representation \lt , i.e.: a *least common ancestor* algorithm lca_\lt for pseudo-forests; see algorithm 1.

Algorithm 1 calculates the least common ancestor of n_0 and m_0 by alternately extending \lt -paths from n_0 and m_0 one by one. If the newly added node is already contained in the other path, it is returned as the result of $\text{lca}(n_0, m_0)$: Since this is the first time the two paths overlap, this node is not only a common ancestor but also the least one. If one path cannot be extended (because its IMDOM is \perp or starts to contain a cycle), only the other path is extended (procedure lin). When the other path cannot be extended anymore either, we do not have an lca , so we return \perp .

Algorithm 2 works in two phases: First, we establish trivial IMDOM relations for nodes with only one successor. For the graph in Figure 14 (left), these would be $\text{IMDOM}[5]=7$, $\text{IMDOM}[7]=8$, $\text{IMDOM}[8]=9$ and $\text{IMDOM}[9]=8$.

1079 Next, we calculate IMDOM for conditional nodes (nodes with more than one successor). We
 1080 keep a queue of such nodes for which IMDOM has not been calculated. For each conditional
 1081 node x , we try to calculate the lca of their successors. If this returns a node $a \neq \perp$, we
 1082 set $\text{IMDOM}[x]=a$ and remove x from the worklist, otherwise it is put back at the end. The
 1083 algorithm terminates once the worklist is empty or we have completed a full iteration through
 1084 the worklist without a change to IMDOM . The variable oldest tracks the first node after the
 1085 last change; once we visit it again, we are done.

1086 In the case of the graph in Figure 14 (left), assume we iterate in order $\text{COND}_G = [1, 2, 3, 4]$,
 1087 which becomes our first workqueue . For $x=1$, we calculate $\text{lca}(\{2, 3, 4\})$. Let's sup-
 1088 pose we try to calculate $\text{lca}(2, 3) = \text{lca}([2], [3])$ first. Since $\text{IMDOM}[2]=\perp$, we imme-
 1089 diately call $\text{lin}([2])([3])$, but since $\text{IMDOM}[3]=\perp$ as well, we return \perp as $\text{lca}(2, 3)$.
 1090 But then $\text{lca}(\{2, 3, 4\})=\perp$ as well. 1 is therefore put back into the queue, so we now
 1091 have $\text{workqueue}=[2, 3, 4, 1]$ and $\text{oldest}=1$. For $x=2$, when calculating $\text{lca}(6, 7)$, we have
 1092 $\text{IMDOM}[6]=\perp$, so we immediately call $\text{lin}([6])([7])$. During the recursion in lin , we ex-
 1093 tend $[7]$ to $[7, 8]$ and $[7, 8, 9]$ (since no new node is in $[6]$). The next step would be
 1094 $\text{IMDOM}[9]=8$. Since $8 \in [7, 8, 9]$ (which would lead to a loop), we return \perp . 2 is therefore put
 1095 back into the queue, so we now have $\text{workqueue}=[3, 4, 1, 2]$ and $\text{oldest}=1$. For $x=3$, when
 1096 calculating $\text{lca}(5, 7)$, we have $\text{IMDOM}[5]=7$ and $7 \in [7]$, so we return 7 as our lca . Since we
 1097 have an lca , we update $\text{IMDOM}[3]=7$, keep 3 out of the workqueue (so $\text{workqueue}=[4, 1, 2]$)
 1098 and set $\text{oldest}=\perp$. For $x=4$, when calculating $\text{lca}(9, 5)$, we extend both paths alternately
 1099 until the path starting in 9 would enter a loop, then only the path starting in 5 is extended.
 1100 Then we will find that 8 is our lca , since it is in both paths. In detail, $\text{lca}([9], [5]) =$
 1101 $\text{lca}([5], [9, 8]) = \text{lca}([9, 8], [5, 7]) = \text{lin}([9, 8])([5, 7]) = \text{lin}([9, 8])([5, 7, 8])$
 1102 $= 8$. We update $\text{IMDOM}[4]=8$, keep 4 out of the workqueue (so $\text{workqueue}=[1, 2]$) and keep
 1103 $\text{oldest}=\perp$. Now we are back to $x=1$. When calculating $\text{lca}(2, 3)$, we now have $\text{IMDOM}[3]=7$,
 1104 so we can extend $[3]$ until we get $[3, 7, 8, 9]$. But still, no node is contained in $[2]$, so
 1105 we still have \perp as our lca . We put 1 back into the workqueue (so $\text{workqueue}=[2, 1]$) and
 1106 set $\text{oldest}=1$. After finding for $x=2$ that $\text{lca}(6, 7)=\perp$, we put 2 back into the workqueue
 1107 (so $\text{workqueue}=[1, 2]$) and keep $\text{oldest}=1$. But now our next element in the queue is our
 1108 oldest , so we are done.

1109 The computation of $\sqsubseteq_{\text{SINK}}$ is slightly more complicated. As it is a greatest fixpoint, in
 1110 principle we must start with $N \times N$ and reduce it according to the rules; until the greatest
 1111 fixpoint is reached. But $N \times N$ cannot be represented by a pseudo-forest. Hence we need
 1112 to initialize the fixed point iteration with an approximation \sqsubseteq_0 of $\sqsubseteq_{\text{SINK}}$ (i.e.: $\sqsubseteq_0 \supseteq \sqsubseteq_{\text{SINK}}$)
 1113 that is representable by a pseudo-forest $<_0$. We can build $<_0$ by interleaving a traversal of a
 1114 preliminary pseudo forest $<$ with $\text{lca}_<$ computations. Consider the preliminary $<$ in figure
 1115 12b. We need to establish $3 < 1$, but find that $\text{lca}_<(\{2, 3\}) = \perp$ for the successors of 1. We
 1116 would like to assume *both* $3 < 1$ and $2 < 1$, the latter of which would then be invalidated in
 1117 the (downward) fixed point iteration. But then $<$ no longer would be a pseudo forest! If we
 1118 assumed just $2 < 1$, we would obtain a $<_0$ such that *not*: $<_0^* \supseteq \sqsubseteq_{\text{SINK}}$, so we need to make
 1119 the assumption $3 < 1$. This example illustrates how the fixpoint iteration must proceed. It
 1120 is based on the

1121
 1122 **Observation 5.2.** Let $<_{\text{SINK}}$ be a transitive reduction of $\sqsubseteq_{\text{SINK}}$. Then whenever $x <_{\text{SINK}} y$ and
 1123 any path starting in x is bound for a sink S (such S is necessarily unique), then any path
 1124 starting in y is bound for S as well.

1125
 1126
 1127

```

1128 Input : A CFG  $G$ 
1129 Data: A pseudo-forest  $<$  represented as a map  $\text{IMDOM} : N \leftrightarrow N$  s.t.  $\text{IMDOM}[n] = m$  iff
1130  $m < n$ 
1131 Output: A transitive reduction  $<_{\text{MAX}}$  of  $\sqsubseteq_{\text{MAX}}$ 
1132 begin
1133   for  $x \in N$ ,  $\{z \mid x \rightarrow z\} = \{z\}$ ,  $z \neq x$  do
1134      $\text{IMDOM}[x] \leftarrow z$ 
1135   end
1136    $\text{MAXIMAL}_{\text{up}}$ 
1137   return  $\text{IMDOM}$ 
1138 end
1139 Procedure  $\text{MAXIMAL}_{\text{up}}$ 
1140    $\text{workqueue} \leftarrow \text{COND}_G$ 
1141    $\text{oldest} \leftarrow \perp$ 
1142   while  $\text{workqueue} \neq \emptyset$  do
1143      $x \leftarrow \text{removeFront}(\text{workqueue})$ 
1144     assert  $\text{IMDOM}[x] = \perp$ 
1145     if  $\text{oldest} = x$  then
1146       return
1147     end
1148     if  $\text{oldest} = \perp$  then
1149        $\text{oldest} \leftarrow x$ 
1150     end
1151      $a \leftarrow \text{lca}(\{y \mid x \rightarrow y\})$ 
1152      $z \leftarrow \begin{cases} \perp & \text{if } a = \perp \vee a = x \\ a & \text{otherwise} \end{cases}$ 
1153     if  $z \neq \perp$  then
1154        $\text{IMDOM}[x] \leftarrow z$ 
1155        $\text{oldest} \leftarrow \perp$ 
1156     else
1157        $\text{pushBack}(\text{workqueue}, x)$ 
1158     end
1159   end
1160 end

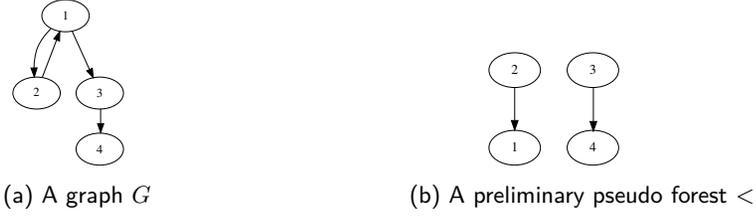
```

1161 **Algorithm 2:** An efficient algorithm for the computation of $<_{\text{MAX}}$. COND_G denotes the set
1162 of *conditional* nodes, i.e.: nodes with more than one successor. workqueue is ordered by
1163 any fixed ordering on nodes N .

1164
1165

1166 Here, “bound for S ” means that the path cannot escape sink S . To illustrate the iteration
1167 for $\sqsubseteq_{\text{SINK}}$, consider figure 12b. For node 3 we have already established $4 <^* 3$ for the sink
1168 node $4 \in S$, but we have not yet established $4 <^* 2$. This suggests that we must – whenever
1169 $\text{lca}(\{y \mid x \rightarrow y\}) = \perp$ – choose some successor node y of x such that already $s <^* y$ for
1170 some sink node s . We call such nodes y *processed*, and maintain a set PROCD of all such
1171 nodes. Algorithm 3 presents the computation of $<_{\text{SINK}}$, the additional procedures performing
1172 the iteration are given in Figure 13.

1173 Algorithm 3 first initializes ISDOM for sink nodes and nodes with one successor. Remember
1174 that any nontrivial sink S_i contains a $<_{\text{SINK}}$ -cycle. For each sink S_i , we therefore initialize
1175 ISDOM to be such a cycle in arbitrary order. We also choose a representative s_i for each sink
1176

Fig. 12. Computing an initial approximation \prec_0 .

S_i and mark all nodes in S_i as processed. For all nodes outside sinks with one successor, the initialization of ISDOM is identical to the one in algorithm 2. Once a successor is processed, we mark all nodes that reach this node through ISDOM-chains as processed.

Next, we construct in $SINK_{up}$ a preliminary ISDOM that fulfills $ISDOM^* \supseteq \sqsubseteq_{SINK}$ but might be too optimistic: For nodes x , $ISDOM[x]$ might exist even though it should be \perp ; or it might be a node that is too small to be a common ancestor of the successors of x (but the correct lca is an ancestor of $ISDOM[x]$). We choose such an ISDOM of x as soon as one of its successors is processed. When calculating the lca , we only consider the successors which have already been processed. If the resulting lca is \perp , we choose an arbitrary successor as lca . Now, we set $ISDOM[x]$ to be this lca . We also set x and all nodes that reach x through ISDOM-chains as processed. This succeeds for any x with distance k to a sink at attempt k at the latest (this can be shown by induction on k), so this algorithm terminates.

Then, these spurious postdominances are eliminated during $SINK_{down}$. For each conditional node x outside sinks the lca of its successors is calculated. If it is part of a sink S_i , its distinguished representative s_i is chosen instead. If it is different from the current ISDOM (either a different node or \perp), ISDOM is updated and all nodes possibly affected by this change are put back in the worklist: These are all conditional nodes n having a successor y that reaches x through ISDOM-chains. This is done until the worklist is empty.

As an example, consider Figure 14 (left). In the initial phase, we set $ISDOM[8]=9$ and $ISDOM[9]=8$ for the non-trivial sink. For its representative, let's assume we choose 8. We mark all sink nodes 6, 8 and 9 as processed. Then, we handle non-condition nodes. We set $ISDOM[4]=9$ and mark 4 as processed. After that, we set $ISDOM[5]=7$ (but cannot mark it as processed since 7 is not). Finally we set $ISDOM[7]=8$ and mark both 7 and 5 as processed.

In $SINK_{up}$, 1 has a single processed successor, namely 4. Thus $ISDOM[1]=4$, and 1 is processed. For 2, we have two processed successors, but $lca(6,7)=\perp$. Let's suppose we choose $ISDOM[2]=7$; 2 is also marked as processed. Finally, 3 has two processed successors and $lca(5,7)=7$, so we set $ISDOM[3]=7$ and mark 3 as processed. This finishes $SINK_{up}$.

In $SINK_{down}$, we first check $x=1$. Since we still have $ISDOM[2]=7$, $lca(\{2,3,4\})=8$. This is also the representative of this sink, so we set $ISDOM[1]=8$. For $x=2$, we now find that $ISDOM[2]=\perp$. This change puts 1 back into the worklist. For $x=3$, no change occurs, since $lca(5,7)=ISDOM[3]=7$. For $x=4$, we have $lca(9)=9$. The representative of this sink is 8, so we set $ISDOM[4]=8$. We would also have to put 1 back into the worklist if it wasn't there already. For $x=1$, the updated $ISDOM[2]$ now means we find $lca(\{2,3,4\})=\perp$, so we set $ISDOM[2]=\perp$. This finishes the calculation of ISDOM.

5.2 Postdominance Frontiers in Graphs Without Unique Exit

We will now derive algorithms for \rightarrow_{ntscd} and \rightarrow_{nticd} , based on \sqsubseteq_{MAX} and \sqsubseteq_{SINK} . In particular, we generalize Cytron's idea to split up the postdominance frontier into an "up" and a "local"

```

1226 Input : A CFG  $G$ 
1227 Data: A pseudo-forest  $\prec$  represented as a map  $\text{ISDOM} : N \hookrightarrow N$  s.t.  $\text{ISDOM}[n] = m$  iff
1228  $m \prec n$ 
1229 Output: A transitive reduction  $\prec_{\text{SINK}}$  of  $\sqsubseteq_{\text{SINK}}$ 
1230 begin
1231    $\{S_1, \dots, S_n\} \leftarrow \{S_i \mid S_i \in \text{sccs}(G), \neg \exists s \rightarrow n. s \in S_i \wedge n \notin S_i\}$ 
1232    $S \leftarrow \bigcup S_i$ 
1233   for  $1 \leq i \leq n$  do
1234      $s_i \leftarrow$  any node in  $S_i$ 
1235     for  $n_j \in S_i$  in any fixed ordering  $n_1, \dots, n_{k_i}$  of  $S_i$  do
1236        $\text{ISDOM}[n_j] \leftarrow n_{j+1 \bmod k_i}$  unless  $k_i = 1$ 
1237        $\text{processed}(n_j)$ 
1238     end
1239   end
1240   for  $x \in N, x \notin S, \{z \mid x \rightarrow z\} = \{z\}, z \neq x$  do
1241      $\text{ISDOM}[x] \leftarrow z$ 
1242     if  $z \in \text{PROCD}$  then  $\text{processed}(x)$ 
1243   end
1244    $\text{SINK}_{\text{up}}$ 
1245    $\text{SINK}_{\text{down}}$ 
1246   return  $\text{ISDOM}$ 
1247 end

```

1247 **Algorithm 3:** Computation of transitive reduction \prec_{SINK} of $\sqsubseteq_{\text{SINK}}$. Not shown is the
1248 procedure $\text{processed}(x)$ which updates PROCD given a node x s.t. $s \prec^* x$ for some sink
1249 node s , by following linear segments ending in x upwards.

1250
1251
1252 part, and to follow the tree structure (parent links) while iterating. The latter also works for
1253 pseudo-forests.

1254 To describe this idea in detail, first remember that in graphs with unique exit node n_x ,
1255 standard postdominance $\sqsubseteq_{\text{POST}}$ is always a partial order, while in arbitrary graphs, \sqsubseteq_{MAX}
1256 and $\sqsubseteq_{\text{SINK}}$ may lack anti-symmetry, and may thus contain cycles of nodes postdominating
1257 each other. In the following we therefore reconstruct Cytron's algorithm with our generalized
1258 definition for postdominance frontiers. In particular, the following definitions replace Cytron's
1259 definitions from [10]: instead of Cytron's original $\sqsubseteq_{\text{POST}}$ we use our new $1\text{-}\sqsubseteq$, and instead of
1260 Cytron's original $\text{ipdom}_{\sqsubseteq_{\text{POST}}}$ we use $\text{ipdom}_{\sqsubseteq}$. We will thus be able to define the generalized
1261 algorithm in a self-contained way.

1262
1263 **Definition 5.2** (Immediate \sqsubseteq -Postdominance). Given a binary relation \sqsubseteq on nodes, a node
1264 x is said to $1\text{-}\sqsubseteq$ -postdominate z if there exists some node $y \neq x$ such that $x \sqsubseteq y \sqsubseteq z$. The
1265 set $\text{ipdom}_{\sqsubseteq}(n)$ is defined by

$$1266 \quad \text{ipdom}_{\sqsubseteq}(n) = \left\{ m \mid \forall m' \in N. m' 1\text{-}\sqsubseteq n \implies m' \sqsubseteq m \right\}$$

1267
1268
1269 In contrast to strict postdominance, $x 1\text{-}\sqsubseteq x$ might hold, namely if there is a cycle
1270 $x \sqsubseteq y \sqsubseteq x$ for $x \neq y$. $\text{ipdom}_{\sqsubseteq}(x)$ is the set of *immediate* postdominators: it contains the
1271 postdominators of x that all (other) postdominators of x postdominate.

1272 As an example, consider the CFG in Figure 14 (left) with \sqsubseteq_{MAX} -postdominance. We have
1273 $\text{ipdom}_{\sqsubseteq}(5) = \{7\}$ since $7 1\text{-}\sqsubseteq 5$ and each $1\text{-}\sqsubseteq$ -postdominator of 5 also postdominates 7.

1274

```

1275 Procedure SINKup
1276 workqueue  $\leftarrow$  CONDG \ S in any order
1277 while workqueue  $\neq$   $\emptyset$  do
1278    $x \leftarrow$  removeFront(workqueue)
1279   assert ISDOM [x] =  $\perp$   $\wedge$   $x \notin$  PROCD
1280   SUCCS  $\leftarrow$  { y | x  $\rightarrow$  y, y  $\in$  PROCD }
1281   if SUCCS =  $\emptyset$  then
1282     | z  $\leftarrow$   $\perp$ 
1283   else
1284     | a  $\leftarrow$  lca (SUCCS)
1285     | z  $\leftarrow$  { any y  $\in$  SUCCS if a =  $\perp$ 
1286                 a otherwise
1287   end
1288   if z  $\neq$   $\perp$  then
1289     | ISDOM [x]  $\leftarrow$  z
1290     | processed (x)
1291   else
1292     | pushBack(workqueue, x)
1293   end
1294 end

```

```

Procedure SINKdown
workset  $\leftarrow$  { n | n  $\in$ 
CONDG \ S, ISDOM[n]  $\neq$   $\perp$  }
while workset  $\neq$   $\emptyset$  do
  x  $\leftarrow$  removeMin(workset)
  a  $\leftarrow$  lca ({ y | x  $\rightarrow_G$  y })
  z  $\leftarrow$  {  $\perp$  if a =  $\perp$ 
            si if a  $\in$  Si
            a otherwise
  assert ISDOM[x] =  $\perp$   $\Rightarrow$  z =  $\perp$ 
  if z  $\neq$  ISDOM[x] then
    | workset  $\leftarrow$  workset  $\cup$ 
    | { n  $\in$  CONDG \ S |  $\exists n \rightarrow y$ . x  $<^*$  y }
    | ISDOM [x]  $\leftarrow$  z
  end
end

```

Fig. 13. Upward and downward iteration for algorithm 3

```

1295
1296
1297
1298
1299 Input : A transitive reduction  $<$  of  $\sqsubseteq$ 
1300 Input : A map SCC from nodes  $x$  to the strongly connected component  $c$  of  $<$  s.t  $x \in c$ 
1301 Input : A topological sorting sccs of all strongly connected components of  $<$ .
1302 Output: pdf $_{\sqsubseteq}$ 
1303 for scc  $\in$  sccs do
1304   | local  $\leftarrow$  { y | x  $\in$  scc, y  $\rightarrow$  x,  $\neg \exists x' \in$  scc.  $x' <$  y }
1305   | up  $\leftarrow$  { y |  $\underbrace{x \in \text{scc}, x < z, y \in \text{PDF}[z]}_{z \in \text{scc}_{<}}$ ,  $\neg \exists x' \in$  scc.  $x' <$  y }
1306   | for x  $\in$  scc do PDF[x]  $\leftarrow$  local  $\cup$  up
1307   |  $\underbrace{\neg \exists x' \in \text{scc}. x' < y}_{y \in \text{scc}_{<}}$ 
1308   |  $\underbrace{y \in \text{scc}_{<}}_{y \in \text{scc}_{<}}$ 
1309 end

```

Algorithm 4: Computation of pdf $_{\sqsubseteq}$

1310
1311
1312
1313 8 \notin ipdom $_{\sqsubseteq}$ (5) because 7 1- \sqsubseteq 5 but not 7 \sqsubseteq 8. For the cycle of 8 and 9, each of those
1314 1- \sqsubseteq -postdominates itself and the other one, so we have ipdom $_{\sqsubseteq}$ (8) = ipdom $_{\sqsubseteq}$ (9) = {8, 9}.

1315 Next we need a generalized notion of Cytron's postdominance frontiers. Intuitively, the
1316 postdominance frontier contains all nodes that are one step away from having x as a
1317 postdominator.
1318

1319 **Definition 5.3** (\sqsubseteq -Postdominance Frontiers).

$$1320 \text{pdf}_{\sqsubseteq}(x) = \left\{ y \mid \text{for some } s \text{ s.t. } y \rightarrow s : \begin{array}{l} \neg x \text{ 1-}\sqsubseteq y \\ x \sqsubseteq s \end{array} \right\}$$

1323

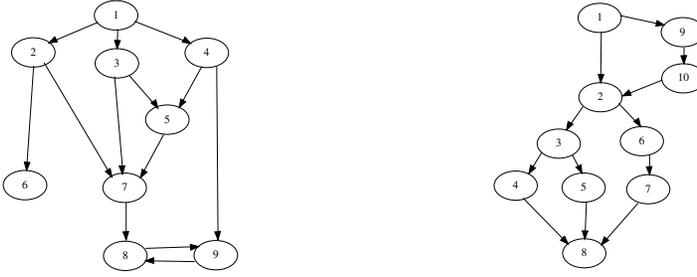


Fig. 14. Two example CFGs

Consider again Figure 14 (left) with \sqsubseteq_{MAX} -postdominance. We have $\text{pdf}_{\sqsubseteq}(5) = \{3, 4\}$, since 5 neither postdominates 3 or 4, but postdominates a successor of those nodes (namely 5 itself). $1 \notin \text{pdf}_{\sqsubseteq}(5)$ since 5 postdominates no successor of 1. For the node 7 we have $\text{pdf}_{\sqsubseteq}(7) = \{1, 2, 4\}$. Note that $3 \notin \text{pdf}_{\sqsubseteq}(7)$ since 7 postdominates 3.

The following lemma generalizes Cytron’s insight that CD is essentially the same as postdominance frontiers:

Lemma 5.1. For $n \neq m$, we have

$$n \rightarrow_{\text{ntscd}} m \iff n \in \text{pdf}_{\sqsubseteq_{\text{MAX}}}(m)$$

and

$$n \rightarrow_{\text{nticd}} m \iff n \in \text{pdf}_{\sqsubseteq_{\text{SINK}}}(m)$$

Due to this lemma, we easily obtain $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{nticd}}$ once we have an algorithm for pdf_{\sqsubseteq} . For the latter, we – following Cytron – partition $\text{pdf}_{\sqsubseteq}(x)$ into two parts: those y contributed *locally*, and those y contributed by nodes z which are immediately \sqsubseteq -postdominated by x (implying $x \sqsubseteq z$). Informally, the local part $\text{pdf}_{\sqsubseteq}^{\text{local}}(x)$ of $\text{pdf}_{\sqsubseteq}(x)$ comprises all nodes from which one can get to x in one step, but which do not have x as a postdominator. On the other hand, if $y \in \text{pdf}_{\sqsubseteq}(z)$ and $\text{ipdom}_{\sqsubseteq}(z)$ is not the join point of all of y ’s branching, then y is in the “upper” part $\text{pdf}_{\sqsubseteq}^{\text{up}}(z)$. This is formalized in

Definition 5.4 (\sqsubseteq -Postdominance Frontiers: *local* and *up* part).

$$\text{pdf}_{\sqsubseteq}^{\text{local}}(x) = \{y \mid \neg x \sqsubseteq y, y \rightarrow x\}$$

$$\text{pdf}_{\sqsubseteq}^{\text{up}}(z) = \{y \in \text{pdf}_{\sqsubseteq}(z) \mid \forall x \in \text{ipdom}_{\sqsubseteq}(z). \neg x \sqsubseteq y\}$$

Under suitable conditions, $\text{pdf}_{\sqsubseteq}^{\text{up}}$ and $\text{pdf}_{\sqsubseteq}^{\text{local}}$ indeed partition pdf_{\sqsubseteq} . This is made precise in the following

Observation 5.3. Let \sqsubseteq be transitive and reflexive. Also, identify $\text{ipdom}_{\sqsubseteq}$ with the relation $\{(x, z) \mid x \in \text{ipdom}_{\sqsubseteq}(z)\}$, and assume $\text{ipdom}_{\sqsubseteq}^* = \sqsubseteq$. Then

$$\text{pdf}_{\sqsubseteq}(x) = \text{pdf}_{\sqsubseteq}^{\text{local}}(x) \cup \bigcup_{\{z \mid x \in \text{ipdom}_{\sqsubseteq}(z)\}} \text{pdf}_{\sqsubseteq}^{\text{up}}(z)$$

Fortunately, \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ are reflexive and transitive (but, as explained, not antisymmetric); thus the partitioning can be applied. For an example, consider again Figure 14

1373 (left) with \sqsubseteq_{MAX} -postdominance. We have $\text{pdf}_{\sqsubseteq}^{\text{local}}(5) = \{3, 4\}$. Since 5 only postdominates
 1374 itself trivially, we have $5 \in \text{ipdom}_{\sqsubseteq}(z)$ for no node z , and observation 5.3 indeed gives
 1375 $\text{pdf}_{\sqsubseteq}(5) = \{3, 4\}$. We have $\text{pdf}_{\sqsubseteq}^{\text{local}}(7) = \{2\}$. Since we have $\{z \mid 7 \in \text{ipdom}_{\sqsubseteq}(z)\} = \{3, 5\}$,
 1376 we need to calculate $\text{pdf}_{\sqsubseteq}^{\text{up}}(3)$ and $\text{pdf}_{\sqsubseteq}^{\text{up}}(5)$. For 5, we have already seen $\text{pdf}_{\sqsubseteq}(5) = \{3, 4\}$.
 1377 But $\text{pdf}_{\sqsubseteq}^{\text{up}}(5)$ contains only node 4, since 7 actually postdominates 3! Since $\text{pdf}_{\sqsubseteq}^{\text{up}}(3) = \{1\}$,
 1378 observation 5.3 results in $\text{pdf}_{\sqsubseteq}(7) = \{1, 2, 4\}$, as expected.

1380 The next definition provides properties which will enable a fixpoint computation of
 1381 $\text{pdf}_{\sqsubseteq}^{\text{local}}(x)$ and $\text{pdf}_{\sqsubseteq}^{\text{up}}(z)$.

1382 **Definition 5.5.** \sqsubseteq is closed under \rightarrow , if it admits the rules

$$1383 \frac{y \rightarrow x \quad x' \sqsubseteq y \quad x' \neq y}{x' \sqsubseteq x} \text{CL} \rightarrow$$

1386 \sqsubseteq lacks joins if it admits the rules

$$1387 \frac{\begin{array}{l} x \in \text{ipdom}_{\sqsubseteq}(v) \quad v \sqsubseteq s \\ x \in \text{ipdom}_{\sqsubseteq}(z) \quad z \sqsubseteq s \end{array} \quad z \neq v}{v \in \text{ipdom}_{\sqsubseteq}(z) \vee z \in \text{ipdom}_{\sqsubseteq}(v)} \text{NoJoin}$$

1389 Informally, the premise of the last rule is “split” at s (into v and z), and joined at x . The
 1390 conclusion demands that this cannot happen unless v and z are immediate neighbours.

1394 **Lemma 5.2.** 🎲 Both \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ are closed under \rightarrow , and lack joins.

1396 As promised, the following theorems provide, under the “lacks join” assumption for \sqsubseteq ,
 1397 simplified formulae for $\text{pdf}_{\sqsubseteq}^{\text{local}}(x)$ and $\text{pdf}_{\sqsubseteq}^{\text{up}}(z)$.

1399 **Observation 5.4.** Let \sqsubseteq be transitive, and closed under \rightarrow . Then

$$1400 \text{pdf}_{\sqsubseteq}^{\text{local}}(x) = \{y \mid \neg x \in \text{ipdom}_{\sqsubseteq}(y), y \rightarrow x\}$$

1402 **Observation 5.5.** Let \sqsubseteq be transitive, reflexive, lacking joins, and closed under \rightarrow . Also
 1403 assume $\text{ipdom}_{\sqsubseteq}^* = \sqsubseteq$. Then, given some z with $x \in \text{ipdom}_{\sqsubseteq}(z)$

$$1404 \text{pdf}_{\sqsubseteq}^{\text{up}}(z) = \{y \in \text{pdf}_{\sqsubseteq}(z) \mid \neg x \in \text{ipdom}_{\sqsubseteq}(y)\}$$

1406 As both \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ satisfy the assumptions of the last theorems, these theorems
 1407 immediately lead to an efficient rule system for computing $\text{pdf}_{\sqsubseteq}(x)$. The first rule initializes
 1408 $\text{pdf}_{\sqsubseteq}(x)$ to its “local” part; the second rule applies the formula for the “upper” part, until
 1409 a fixpoint is reached. Of course, $\text{ipdom}_{\sqsubseteq}$ must be computed beforehand.

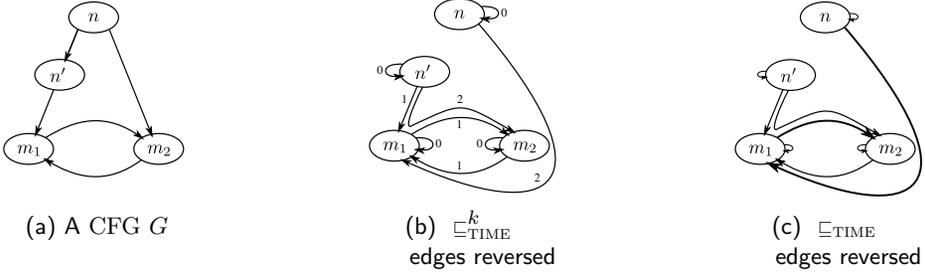
1411 **Definition 5.6.** The monotone rule system for computing $\text{pdf}_{\sqsubseteq}(x)$ is given by

$$1412 \frac{x \notin \text{ipdom}_{\sqsubseteq}(y) \quad y \rightarrow x}{y \in \text{pdf}_{\sqsubseteq}(x)} \quad \frac{x \notin \text{ipdom}_{\sqsubseteq}(y) \quad x \in \text{ipdom}_{\sqsubseteq}(z) \quad y \in \text{pdf}_{\sqsubseteq}(z)}{y \in \text{pdf}_{\sqsubseteq}(x)}$$

1416 The smallest fixpoint of this rule system can be computed by a standard worklist algorithm.
 1417 Additionally, we can exploit transitive reductions. Given any transitive reduction $<$ of \sqsubseteq ,

1418 (1) compute the strongly connected components sccs of the graph $(N, <)$, in a correspond-
 1419 ing topological order. These can either be provided by the algorithm computing $<$, or
 1420 by Tarjan’s algorithm [32].

1421

Fig. 15. An irreducible graph with *intransitive* \subseteq_{TIME}

(2) compute pdf_{\subseteq} by traversing the condensed graph in that order *once*.

This concludes the algorithm for generalized postdominance frontiers $\text{pdf}_{\subseteq}(x)$; and thus for $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{nticd}}$. For the actual computation, we propose the following optimization: By *precomputing* the set $\text{scc}_{<} = \{y \mid \exists x' \in \text{scc}. x' < y\}$ for each scc, we can use this for both the tests on y , and for enumerating z .

To illustrate the fixpoint iteration for $\text{pdf}_{\subseteq}(x)$, consider once more the CFG in Figure 14 (left). The “local” rule gives us e.g. $1 \in \text{pdf}_{\subseteq}(3)$, $3 \in \text{pdf}_{\subseteq}(5)$, $4 \in \text{pdf}_{\subseteq}(5)$ and $2 \in \text{pdf}_{\subseteq}(7)$. With the “up” rule we can now get $4 \in \text{pdf}_{\subseteq}(7)$ by instantiating the rule with $x = 7$, $y = 4$ and $z = 5$. Note that we indeed have shown $4 \in \text{pdf}_{\subseteq}(5)$ earlier and we have $7 \notin \text{ipdom}_{\subseteq}(4)$ as well as $7 \in \text{ipdom}_{\subseteq}(5)$. In contrast, if we try to use $3 \in \text{pdf}_{\subseteq}(5)$ to show $3 \in \text{pdf}_{\subseteq}(7)$ (which is false), $7 \notin \text{ipdom}_{\subseteq}(3)$ would have to hold. But $7 \in \text{ipdom}_{\subseteq}(3)$, so the right rule is not applicable, and we are prevented from showing $3 \in \text{pdf}_{\subseteq}(7)$.

5.3 Timing Sensitive Postdominance Frontiers

In order to develop efficient algorithms for the computation of timing sensitive postdominance \subseteq_{TIME} and timing sensitive control-dependence $\rightarrow_{\text{tscd}}$, let us first recall that our algorithms for \subseteq_{MAX} and $\rightarrow_{\text{ntscd}}$ rely on the fact that \subseteq_{MAX} is *transitive*:

- (1) Transitivity of \subseteq_{MAX} allows us to efficiently compute and represent \subseteq_{MAX} in form of its transitive reduction $<_{\text{MAX}}$. Here, $<_{\text{MAX}}$ turned out to be a pseudo-forest.
- (2) Transitivity of \subseteq_{MAX} , and the fact that

$$\text{ipdom}_{\subseteq_{\text{MAX}}}^* = \subseteq_{\text{MAX}}$$

allows us to use algorithm 4 to efficiently compute $\rightarrow_{\text{ntscd}}$ via $\text{pdf}_{\subseteq_{\text{MAX}}}$.

Disregarding for now that $\rightarrow_{\text{tscd}}$ is defined in terms of the *ternary* relation $n \subseteq_{\text{TIME}}^k m$, and not in terms of its binary “ $\exists k$. -closure” $n \subseteq_{\text{TIME}} m$, let us investigate first if $n \subseteq_{\text{TIME}} m$ is – in general – transitive. Consider the (irreducible) CFG in Figure 15a. Here, every maximal path starting in n first reaches m_1 after two steps, hence $m_1 \subseteq_{\text{TIME}} n$. Also, every maximal path starting in m_1 first reaches m_2 after one step, hence $m_2 \subseteq_{\text{TIME}} m_1$. But it is for *no* number k of steps the case that $m_2 \subseteq_{\text{TIME}}^k n$, hence: $\neg m_1 \subseteq_{\text{TIME}} n$. In summary, \subseteq_{TIME} is *not* transitive.

Fortunately, situations as in Figure 15 are the only ones in which \subseteq_{TIME} is not transitive:

Theorem 5.1. 🌟 Let G be any *reducible* CFG. Then \subseteq_{TIME} is transitive.

Theorem 5.2. 🌟 Let G be any CFG with unique exit node n_x . Then \subseteq_{TIME} is transitive.

1471 In practice, many programs have reducible CFGs *or* a unique exit; then $\sqsubseteq_{\text{TIME}}$ is transitive
 1472 by the above two theorems. Whenever $\sqsubseteq_{\text{TIME}}$ is transitive, we can use algorithm 4 to compute
 1473 $\rightarrow_{\text{tscd}}$. And if not, in [16] we present an algorithm for $\rightarrow_{\text{tscd}}$ which works even if $\sqsubseteq_{\text{TIME}}$ is
 1474 not transitive. But it is much more complex and thus not described in this article. Note that
 1475 even our transitive “restriction” is more general than the restriction to *structured* CFGs
 1476 which is often required in literature on timing leaks, such as in e.g. [1, 17, 28].

1477 Still, even under the $\sqsubseteq_{\text{TIME}}$ transitivity assumption, we are not done. Compared to the
 1478 above \sqsubseteq_{MAX} algorithm, we must deal with the ternary $n \sqsubseteq_{\text{TIME}}^k m$ instead of the binary \sqsubseteq_{MAX} .
 1479 To this end, remember that for $m \neq n$,

$$1480 \quad n \in \text{pdf}_{\sqsubseteq_{\text{MAX}}}(m) \Leftrightarrow n \rightarrow_{\text{ntscd}} m$$

1482 To obtain the analogous result for $\rightarrow_{\text{tscd}}$, we first need to “conservatively” redefine the
 1483 notion pdf_{\sqsubseteq} of \sqsubseteq -postdominance in order to obtain a notion appropriate for non-transitive
 1484 relations \sqsubseteq . Remember that in definition 5.3, we defined for any binary relation \sqsubseteq :

$$1485 \quad \text{pdf}_{\sqsubseteq}(m) = \left\{ n \mid \begin{array}{l} \neg m \sqsubseteq n \\ \text{for some } n' \text{ s.t. } n \rightarrow_G n' : m \sqsubseteq n' \end{array} \right\}$$

1488 Syntactically, we will stick with this definition, but will modify the notion of $1\text{-}\sqsubseteq$ -
 1489 postdominance. The new definition is

1490 **Definition 5.7** ($1\text{-}\sqsubseteq$ -Postdominance, redefinition). Given a relation $\sqsubseteq \subseteq N \times N$, a node
 1491 $x \in N$ is said to $1\text{-}\sqsubseteq$ -postdominate z if $x \sqsubseteq z$ and there exists some node $y \neq x$ such that

$$1493 \quad x \sqsubseteq y \sqsubseteq z$$

1494 The only change is the new requirement $x \sqsubseteq z$, which of course was redundant up to
 1495 this section, since any relation \sqsubseteq we considered (i.e.: $\sqsubseteq_{\text{POST}}$, \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$) was transitive.
 1496 Implicitly, this change also affects immediate \sqsubseteq -postdominance $\text{ipdom}_{\sqsubseteq}$ – see definition 5.2.

1498 **Theorem 5.3.** 🌟 Let $n \neq m \in N$. Then

$$1500 \quad n \in \text{pdf}_{\sqsubseteq_{\text{TIME}}}(m) \Leftrightarrow n \rightarrow_{\text{tscd}} m$$

1502 Theorem 5.3 holds for *arbitrary* graphs, and establishes that indeed, timing sensitive
 1503 postdominance frontiers are essentially timing sensitive control dependence.

1504 But in order to use the generalized postdominance frontiers algorithm from subsection 5.2
 1505 at least for transitive $\sqsubseteq_{\text{TIME}}$, we also need the two other two requirements of that algorithm.
 1506 These two do, indeed, hold even for *arbitrary* graphs:

1507 **Observation 5.6.** Let $\sqsubseteq = \sqsubseteq_{\text{TIME}}$. Then \sqsubseteq is closed under \rightarrow_G , and

$$1510 \quad \text{pdf}_{\sqsubseteq}^{\text{local}}(x) = \left\{ y \mid \begin{array}{l} \neg x \in \text{ipdom}_{\sqsubseteq}(y) \\ y \rightarrow x \end{array} \right\}$$

1512 **Observation 5.7.** Let $\sqsubseteq = \sqsubseteq_{\text{TIME}}$. Then \sqsubseteq lacks joins and is closed under \rightarrow_G , and given
 1513 some z with $x \in \text{ipdom}_{\sqsubseteq}(z)$:

$$1515 \quad \text{pdf}_{\sqsubseteq}^{\text{up}}(z, x) = \left\{ y \in \text{pdf}_{\sqsubseteq}(z) \mid \neg x \in \text{ipdom}_{\sqsubseteq}(y) \right\}$$

1517 All that is required now is an algorithm to compute $\sqsubseteq_{\text{TIME}}$. For graphs that are reducible,
 1518 or have a unique exit node, this can be done by modifying algorithm 3 to work on \mathbb{N} -labeled
 1519

```

1520 Input: A  $\mathbb{N}$  labeled pseudo-forest  $<$ , represented as a map  $\text{IDOM} : N \hookrightarrow N \times \mathbb{N}$  s.t.
1521  $\text{IDOM}[n] = (m, k)$  iff  $m <^k n$ 
1522 Input: Numbers  $k_0^n, k_0^m \in \mathbb{N}$  and nodes  $n_0, m_0$ 
1523 Output:  $\text{lca}_<((n_0, k_0^n), (m_0, k_0^m))$  if it exists, or  $\perp$  otherwise.
1524 return  $\text{lca}((n_0, k_0^n, [n_0 \mapsto k_0^n]), (m_0, k_0^m, [m_0 \mapsto k_0^m]))$ 
1525 Function  $\text{lca}(\pi_n, \pi_m)$ 
1526   Input: A cycle free  $<$ -path  $\pi_n = n_0, \dots, n$  ending in  $n$ , represented by a tuple  $(n, k^n, \text{KS}_n)$ 
1527     where  $\text{KS}_n$  is a map on the nodes  $n$  appearing in  $\pi_n$  s.t.  $k^n = \text{KS}_n[n]$  and for any
1528     such  $n$ :  $\text{KS}_n[n] = k_0^n + \sum_i k_i$  where  $n <^{k_c} \dots <^{k_1} n_0$  in  $\pi_n$ 
1529   Input: A  $<$ -path  $\pi_m = m_0, \dots, m$  likewise
1530   if  $k^n > k^m$  then return  $\text{lca}(\pi_m, \pi_n)$ 
1531   if  $n \in \pi_m \wedge k^n = \text{KS}_m[n]$  then return  $(n, k^n)$ 
1532   if  $n \in \pi_m \wedge k^n \neq \text{KS}_m[n]$  then return  $\perp$ 
1533   switch  $\text{IDOM}[n]$  do
1534     case  $\perp$  do return  $\perp$ 
1535     case  $(n', k^{n'})$  do
1536       if  $n' \in \pi_n$  then return  $\perp$ 
1537        $\text{KS}_n[n'] \leftarrow k^n + k^{n'}$ 
1538       return  $\text{lca}((n', k^n + k^{n'}, \text{KS}_n), \pi_m)$ 
1539     end
1540   end
1541 end

```

Algorithm 5: A timing sensitive *least common ancestor* algorithm for graphs with transitive $\sqsubseteq_{\text{TIME}}$.

pseudo-forests, i.e.: pseudo forests $<$ with edges $n <^k m$ indicating that m must first be reached from n after $k \in \mathbb{N}$ steps. The result is a \mathbb{N} -labeled pseudo-forest $<$ with

$$m \sqsubseteq_{\text{TIME}} n \iff \exists k_1, \dots, k_c. m <^{k_c} \dots <^{k_1} n$$

for some number $c \geq 0$ of edges in $<$. One possible implementation of the required least common ancestor computation in \mathbb{N} -labeled pseudo-forests is shown in algorithm 5.

For an example, consider Figure 14 (right). Before the first call to lca , IDOM contains only trivial relations for nodes with exactly one successor, e.g. $\text{IDOM}[4] = (8, 1)$. To calculate $\text{IDOM}[2]$, we need to call lca with its successors, namely $\text{lca}((3, 1), (6, 1))$. But there, we find that $\text{IDOM}[3]$ is still empty, so the call returns \perp .

When calculating $\text{IDOM}[3]$, we call $\text{lca}((4, 1), (5, 1))$. We find that $\text{IDOM}[4] = (8, 1)$, so we extend this $<$ -path and call $\text{lca}([(4, 1), (8, 2)], (5, 1))$. There, since the left path is now longer, we swap the arguments and call $\text{lca}((5, 1), [(4, 1), (8, 2)])$. Now, we find that $\text{IDOM}[5] = (8, 1)$, so we extend this path and call $\text{lca}([(5, 1), (8, 2)], [(4, 1), (8, 2)])$. Now, since the final element of the left path, namely 8, is also contained in the right one with the same distance of 2, we finally can return $(8, 2)$ as the lca and update $\text{IDOM}[3] = (8, 2)$.

Now we can analyse $\text{IDOM}[2]$ again. Since $\text{IDOM}[3]$ has now an entry, we can extend the path $(3, 1)$ to $[(3, 1), (8, 3)]$. After extending $(6, 1)$ to $[(6, 1), (7, 2)]$ and then $[(6, 1), (7, 2), (8, 3)]$, both paths contain 8 with the same distance 3, so we update $\text{IDOM}[2]$ to $(8, 3)$.

On the contrary, if we try to calculate $\text{IDOM}[1]$ and call $\text{lca}((2, 1), (9, 1))$, the left path get extended to $[(2, 1), (8, 4)]$ and the right path to $[(9, 1), (10, 2), (2, 3)]$. Now, both paths contain the same node 2, but with different distances 1 and 3. Therefore, the lca is \perp .

1569 6 MEASUREMENTS

1570 We evaluated the performance of our algorithms on a) control flow graphs of Java methods,
 1571 as generated by the JOANA system for various third party Java programs; b) randomly
 1572 generated graphs $G = (N, E)$ usually with $|E| = 2|N|$, as generated by the standard generator
 1573 from the JGraphT [26] library. In some cases, we additionally use ladder graphs¹⁵, which
 1574 are used to represent bad case behaviour.

1575 All benchmarks in this section were made on a desktop computer with an Intel i7-6700
 1576 CPU at 3.40GHz, and 64 GB RAM. We implemented the algorithms in Java, using OpenJDK
 1577 Java 9 VM. All benchmarks were run using the Java Microbenchmark Harness JMH [9].

1578 Unless explicitly stated otherwise, all data points represent the average over $n + 1$ runs
 1579 of the benchmark, where n is at least the number of runs which can be finished within 1
 1580 second. For example, the data point at $|N| = 21076$, time = 18ms in Figure 16a stands for
 1581 the average of at least ≈ 50 runs of the benchmark that finished within 1 second. On the
 1582 other hand, the data point in at $|N| = 65000$, time = 88s in Figure 16c results from only one
 1583 run of the benchmark.

1584 The purpose of these benchmarks is to give a general idea of the scalability of the
 1585 algorithms. For example, the benchmarks in the upper left and upper right of Figure 18
 1586 suggest that our new algorithm for the computation of nontermination sensitive control
 1587 dependence $\rightarrow_{\text{ntscd}}$ appears to scale almost linearly for “average” CFGs, while Ranganath’s
 1588 original algorithm [29] clearly grows super-linearly for such graphs. The benchmarks can be
 1589 summarized as follows:

- 1590 (1) For “average” CFGs, our algorithms for $\rightarrow_{\text{ntscd}}$, $\rightarrow_{\text{nticd}}$, and $\rightarrow_{\text{tscd}}$ offer performance
 1591 “almost linear” in the size of the graph.
- 1592 (2) But for “bad case” CFGs, some algorithms perform decidedly super-linear, and become
 1593 impractical for very large such graphs.
 1594

1595 6.1 Nontermination Sensitive Postdominance

1596 Algorithm 2 computes maximal path postdominance \sqsubseteq_{MAX} , represented as a pseudo-forest
 1597 $<_{\text{MAX}}$. This algorithm requires the computation of least common ancestors $\text{lca}_{<}$ in pseudo-
 1598 forests $<$, for which we use algorithm 1.

1600 Algorithm 2 repeatedly iterates in a fixed node order. Alternatively, one can implement a
 1601 chaotic iteration, by reinserting into a workset those nodes affected by modification to the
 1602 pseudo-forest. Both these variants do not specify an iteration order (e.g.: Algorithm 2 does
 1603 not specify the initial order of nodes in the workqueue). By default, the implementation
 1604 orders the nodes reversed-topologically (as computed by an implementation of Kosaraju’s
 1605 Algorithm for strongly connected components, with nodes in the same strongly connected
 1606 component ordered arbitrarily).

1607 For Java CFG and randomly generated graphs (neither necessarily with unique exit
 1608 node), the chaotic iteration (\oplus) and Algorithm 2 (\blacktriangledown) behave similarly (Figure 16a and
 1609 Figure 16b). Ladder graphs expose non-linear *bad-case* behavior (Figure 16c). This is even
 1610 more pronounced when we deliberately choose a bad iteration order (Figure 16d).

1611 6.2 Nontermination Insensitive Postdominance

1612 Algorithm 3 computes sink path postdominance $\sqsubseteq_{\text{SINK}}$, represented as a pseudo-forest $<_{\text{SINK}}$.
 1613 Just as before, it uses algorithm 1 for the computation of least common ancestors $\text{lca}_{<}$.

1614
 1615
 1616 ¹⁵Ladder graphs consist of two rising chains, one-to-one connected at every node. Just like a ladder.
 1617

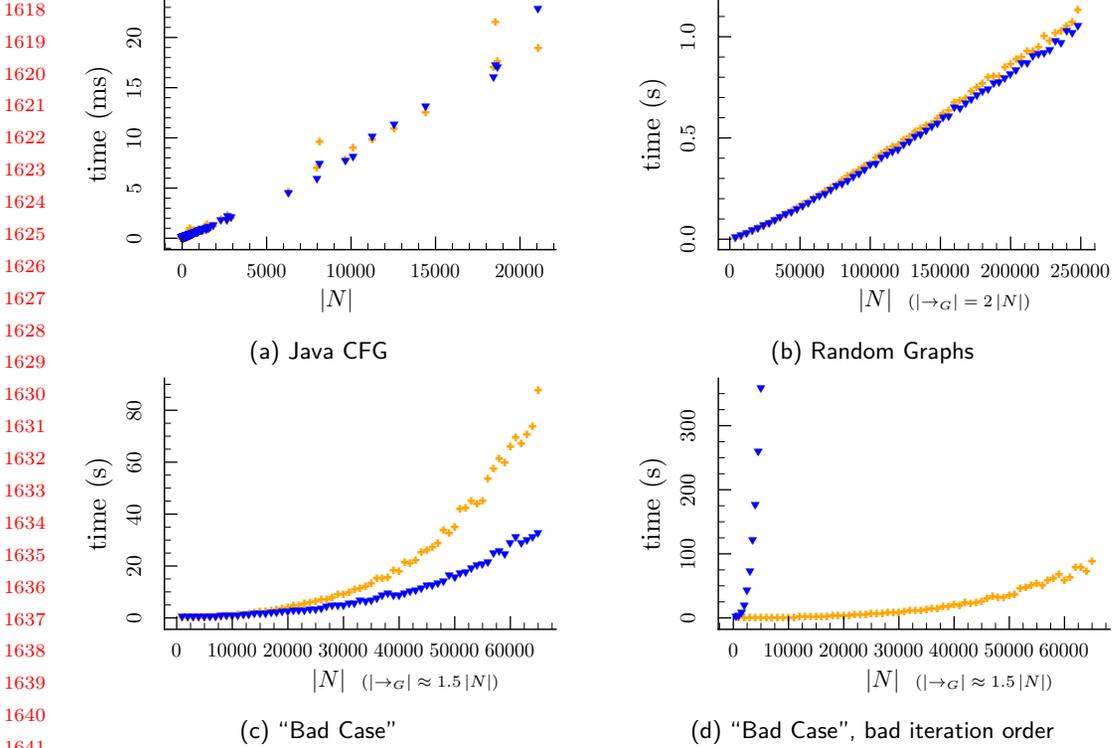


Fig. 16. Computation of \langle_{MAX} . The orange line shows chaotic iteration performance, the blue line shows algorithm 2.

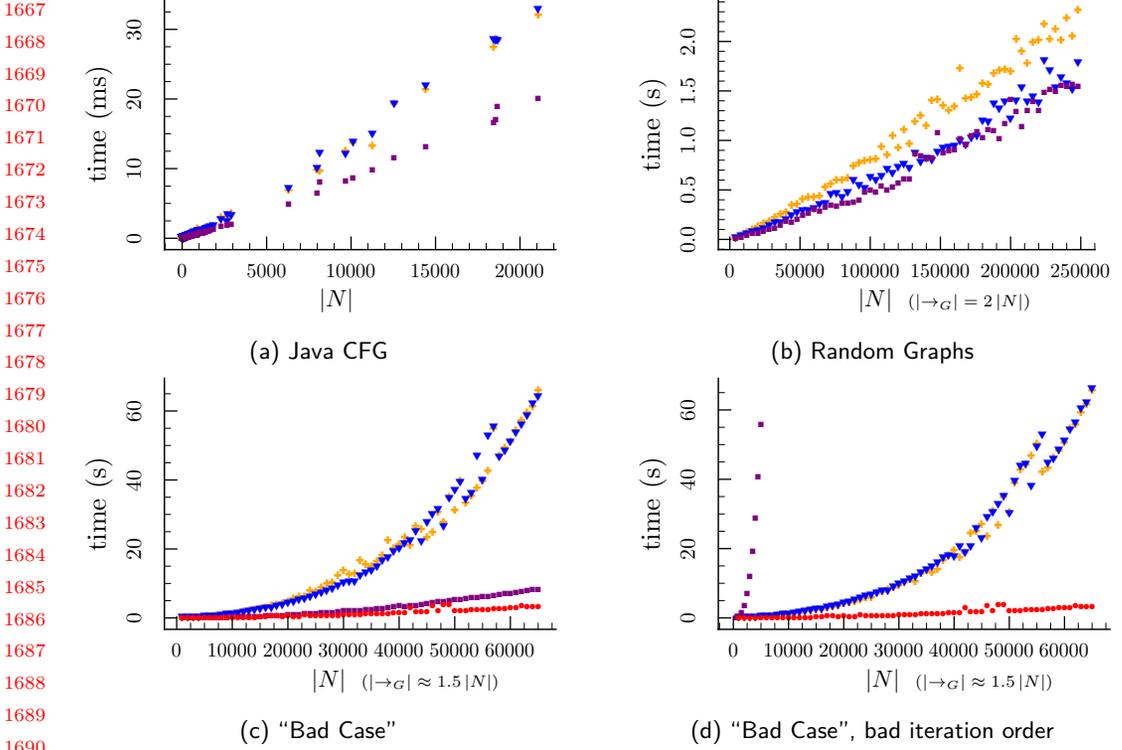
Algorithm 3 implements chaotic iteration. We also implemented a variant of Algorithm 3 in which the downward fixed point phase repeatedly iterates a workqueue of nodes in a fixed node order. The implementations order the nodes reversed-topologically. Unlike before, this ordering does not require an additional step, since the strongly connected component computation it can be obtained from is necessary anyway, in order to find *control sinks*.

Instead of computing least common ancestors lca_{\langle} by chasing (pseudo-tree) pointers, it can also be computed by comparison of postorder numbers, as in [8].

For Java CFGs (Figure 17a) the fixed-iteration order variant of Algorithm 3 (\blacktriangledown) performs on par with the Algorithm 3 as stated (\oplus). For randomly generated graphs (Figure 17b) the variant (\blacktriangledown) appears to perform a bit better than the original (\oplus) for very large graphs, roughly on-par with the implementation based on postorder numbers (\blacksquare).

Using reversed-topological iteration order, ladder graphs (Figure 17c) expose non-linear *bad-case* behavior only for Algorithm 3 (\oplus) and its variant (\blacktriangledown). Even with a bad iteration order, performance for these two algorithm is not much worse (Figure 17d). On the other hand, the postorder number based implementation (\blacksquare) is affected heavily by iteration order.

The ladder graphs we use are unique-exit-node ladder graphs. This also allows us to directly compare with an implementation of the algorithm by Lengauer and Tarjan [24] (\bullet).

Fig. 17. Computation of \langle_{SINK} .

6.3 Generalized Postdominance Frontiers

When algorithm 4 is instantiated with \langle_{MAX} , this yields an algorithm for $\rightarrow_{\text{ntscd}}$. The benchmarks for $\rightarrow_{\text{ntscd}}$ include the computation time of both algorithm 4 and \langle_{MAX} (\blacktriangledown). We compare with an implementation of Ranganath’s algorithm [29] ($+$). For Java CFG and randomly generated graphs, the latter becomes impractical for moderately sized graphs, while algorithm 4 performs well even for very large graphs (Figure 18, upper left and right). Ladder graphs expose non-linear *bad-case* behavior even for algorithm 4 (Figure 18c). This cannot be circumvented, since in these ladder graphs, the size of the relation $\rightarrow_{\text{ntscd}}$ is quadratic in the number of nodes.

6.4 Timing Sensitive CD

Whenever $\sqsubseteq_{\text{TIME}}$ is transitive, we can use algorithm 4 to compute timing sensitive control dependence $\rightarrow_{\text{tscd}}$. We thus measure the computation time for $\rightarrow_{\text{tscd}}$ on graphs for which $\sqsubseteq_{\text{TIME}}$ is transitive. These are control flow graphs from Java programs in subfigures (a), randomly generated graphs (b), and ladder graphs (c). We use algorithm 4, and obtain a transitive reduction \langle_{TIME} of $\sqsubseteq_{\text{TIME}}$ via the modification of algorithm 3 that uses the upwards iteration of algorithm 5. The benchmarks for $\rightarrow_{\text{tscd}}$ in Figure 19 include the computation time of all sub-algorithms ($+$). Ladder graphs expose non-linear *bad-case* behavior.

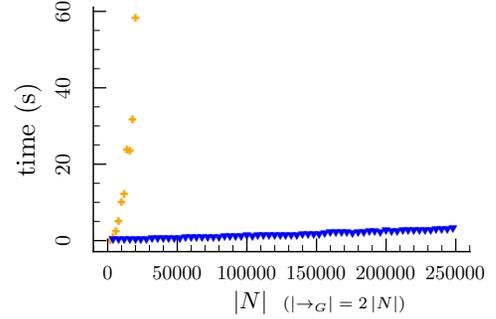
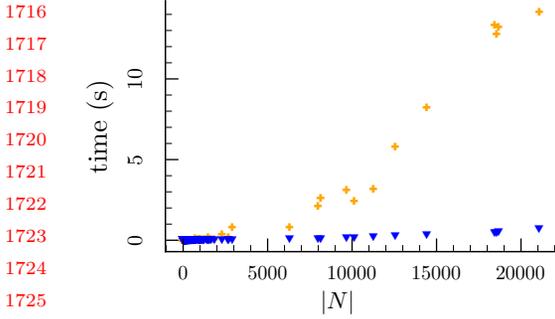
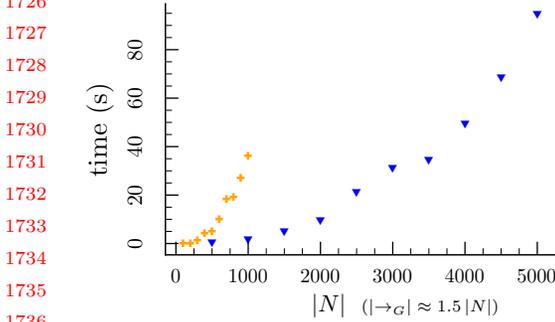


Fig. 18. Computation of \rightarrow_{ntsced} .



(c) "Bad Case"

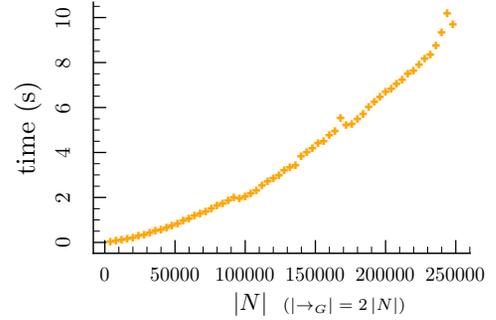
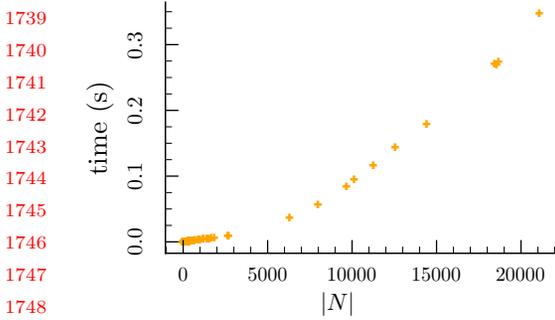
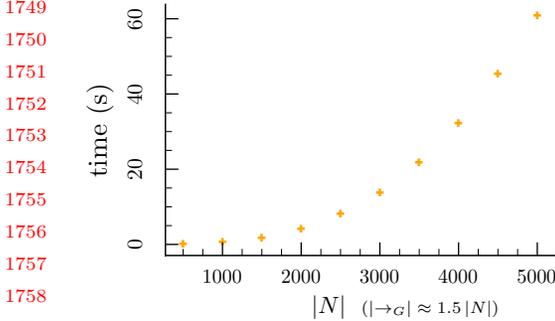


Fig. 19. Computation of \rightarrow_{tsced} .



(c) "Bad Case"

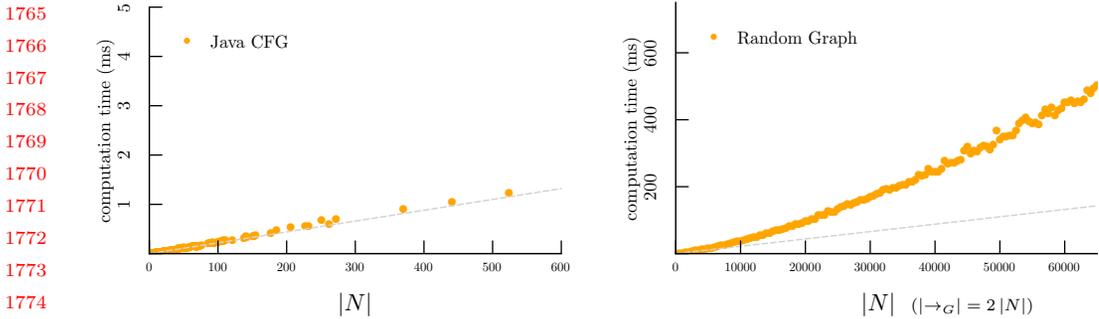


Fig. 20. Computation of $\rightarrow_{\text{nticd}}$ via algorithm 4 based on algorithm 3.

7 FUTURE WORK

This article concentrated on the definition of $\rightarrow_{\text{tscd}}$, and on efficient algorithms. Ongoing work includes

- provide Isabelle proofs for the last 7 “observations” in section 5.
- provide formal correctness proofs for the algorithms in section 5.
- implement and evaluate the $\rightarrow_{\text{tscd}}$ algorithm which can handle nontransitive $\sqsubseteq_{\text{TIME}}$, which was mentioned in section 5.3.
- provide a theoretical complexity analysis of the algorithms, and more measurements.
- transform out timing leaks as in [1], but for arbitrary CFGs (based on $\rightarrow_{\text{tscd}}$).
- apply $\rightarrow_{\text{tscd}}$ to improve IFC and probabilistic noninterference; in particular improve precision of the so-called “RLSOD” algorithm [6, 7, 12] which is used in JOANA.

Initial work on some of these topics can be found in the first author’s dissertation [16]. A long-time goal is an interprocedural, context-sensitive extension of $\rightarrow_{\text{tscd}}$.

8 CONCLUSION

Ranganath and Amtoft opened the door to control dependencies in nonterminating programs. Inspired by this work, we presented 1. new, efficient algorithms for Ranganath’s nontermination-(in)sensitive control dependencies; 2. definitions and algorithms for time-sensitive control dependencies; 3. application of the latter to timing leaks in software security. Our algorithms are based on systematic generalizations of Cytron’s postdominance frontier algorithm. Important properties of the new algorithms have been proven using the Isabelle machine prover; and their performance has been studied. We believe that *time-sensitive control dependencies* will prove useful for many applications in program analysis, code optimization, and software security.

Acknowledgements. The work described in this paper was funded by Deutsche Forschungsgemeinschaft Grant Sn11-12/3 in the scope of SPP 1496 “Reliably Secure Software”, and by BMBF Grant 01BY1172 in the scope of the Security Competence Center KASTEL.

Preliminary versions of parts of this paper have been published in the first author’s dissertation [16].

REFERENCES

- [1] Johan Agat. Transforming out timing leaks. In *POPL ’00: Proceedings of the 27th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 40–53, New York, NY, USA, 2000. ACM.

- 1814 [2] A.V. Aho, M.R. Garey, and J.D. Ullman. The transitive reduction of a directed graph. *SIAM Journal*
1815 *on Computing*, 1 (2):131–137, 1972.
- 1816 [3] Torben Amtoft. Slicing for modern program structures: a theory for eliminating irrelevant loops.
1817 *Information Processing Letters*, 106(2):45 – 51, 2008.
- 1818 [4] Daniel J. Bernstein. Cache-timing attacks on AES. Technical report, 2005.
- 1819 [5] David W. Binkley, Mariano Ceccato, Mark Harman, Filippo Ricca, and Paolo Tonella. Tool-supported
1820 refactoring of existing object-oriented code into aspects. *IEEE Trans. Software Eng.*, 32(9):698–717,
1821 2006.
- 1822 [6] Simon Bischof, Joachim Breitner, Jürgen Graf, Martin Hecker, Martin Mohr, and Gregor Snelting. Low-
1823 deterministic security for low-nondeterministic programs. *Journal of Computer Security*, 26:335–366,
1824 2018.
- 1825 [7] Joachim Breitner, Jürgen Graf, Martin Hecker, Martin Mohr, and Gregor Snelting. On improvements of
1826 low-deterministic security. In *Proc. Principles of Security and Trust (POST)*, volume 9635 of *Lecture*
1827 *Notes in Computer Science*, pages 68–88. Springer Berlin Heidelberg, 2016.
- 1828 [8] Keith D. Cooper, Timothy J. Harvey, and Ken Kennedy. A simple, fast dominance algorithm. *Software*
1829 *Practice & Experience*, 2001.
- 1830 [9] Oracle Corporation. Code tools: jmh, 2020.
- 1831 [10] Ron Cytron, Jeanne Ferrante, Barry K. Rosen, Mark N. Wegman, and F. Kenneth Zadeck. Efficiently
1832 computing static single assignment form and the control dependence graph. *ACM Trans. Program.*
1833 *Lang. Syst.*, 13(4):451–490, October 1991.
- 1834 [11] Jeanne Ferrante, Karl J. Ottenstein, and Joe D. Warren. The program dependence graph and its use in
1835 optimization. *ACM Trans. Program. Lang. Syst.*, 9(3):319–349, July 1987.
- 1836 [12] Dennis Giffhorn and Gregor Snelting. A new algorithm for low-deterministic security. *International*
1837 *Journal of Information Security*, 14(3):263–287, April 2015.
- 1838 [13] Jürgen Graf, Martin Hecker, Martin Mohr, and Gregor Snelting. Tool demonstration: JOANA. In *Proc.*
1839 *Principles of Security and Trust (POST)*, volume 9635 of *Lecture Notes in Computer Science*, pages
1840 89–93. Springer Berlin Heidelberg, 2016.
- 1841 [14] Christian Hammer and Gregor Snelting. Flow-sensitive, context-sensitive, and object-sensitive informa-
1842 tion flow control based on program dependence graphs. *International Journal of Information Security*,
1843 8(6):399–422, Dec 2009.
- 1844 [15] Matthew S. Hecht and Jeffrey D. Ullman. Analysis of a simple algorithm for global data flow problems.
1845 In *Proceedings of the 1st Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming*
1846 *Languages*, POPL '73, pages 207–217, New York, NY, USA, 1973. ACM.
- 1847 [16] Martin Hecker. *Timing Sensitive Dependency Analysis and its Application to Software Security*. PhD
1848 thesis, Karlsruher Institut für Technologie, Fakultät für Informatik, 2020.
- 1849 [17] Daniel Hedin and David Sands. Timing aware information flow security for a javacard-like bytecode.
1850 *Electron. Notes Theor. Comput. Sci.*, 141(1):163–182, 2005.
- 1851 [18] Susan Horwitz, Jan Prins, and Thomas W. Reps. On the adequacy of program dependence graphs for
1852 representing programs. In *POPL*, pages 146–157, 1988.
- 1853 [19] Susan Horwitz, Thomas Reps, and David Binkley. Interprocedural slicing using dependence graphs.
1854 *ACM Trans. Program. Lang. Syst.*, 12(1):26–60, January 1990.
- 1855 [20] Neil Jones, Carsten Gomard, and Peter Sestoft. *Partial Evaluation and Automatic Program Generation*.
1856 Prentice Hall, 1993.
- 1857 [21] Vineeth Kashyap, Ben Wiedermann, and Ben Hardekopf. Timing- and termination-sensitive secure
1858 information flow: Exploring a new approach. In *32nd IEEE Symposium on Security and Privacy*, pages
1859 413–428, 2011.
- 1860 [22] Paul Kocher, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard,
1861 Thomas Prescher, Michael Schwarz, and Yuval Yarom. Spectre attacks: Exploiting speculative execution.
1862 *CoRR*, abs/1801.01203, 2018.
- [23] Jens Krinke. Context-sensitive slicing of concurrent programs. In *Proc. FSE*, pages 178–187. ACM,
2003.
- [24] Thomas Lengauer and Robert Endre Tarjan. A fast algorithm for finding dominators in a flowgraph.
ACM Trans. Program. Lang. Syst., 1(1):121–141, January 1979.
- [25] Steven Muchnik. *Advanced Compiler Design and Implementation*. Morgan Kaufmann.
- [26] Barak Naveh and Stephane Popinet. JGraphT: A java library of graph theory data structures and
algorithms. <https://jgrapht.org/>, 2003–2019.

- 1863 [27] A. Podgurski and L. A. Clarke. A formal model of program dependences and its implications for software
1864 testing, debugging, and maintenance. *IEEE Transactions on Software Engineering*, 16(9):965–979, Sep
1865 1990.
- 1866 [28] Willard Rafnsson, Limin Jia, and Lujo Bauer. Timing-sensitive noninterference through composition.
1867 In Matteo Maffei and Mark Ryan, editors, *Principles of Security and Trust (POST)*, volume 10204 of
1868 *Lecture Notes in Computer Science*, pages 3–25. Springer, 2017.
- 1869 [29] Venkatesh Prasad Ranganath, Torben Amtoft, Anindya Banerjee, John Hatcliff, and Matthew B. Dwyer.
1870 A new foundation for control dependence and slicing for modern program structures. *ACM Trans.*
1871 *Program. Lang. Syst.*, 29(5), August 2007.
- 1872 [30] Thomas Reps, Susan Horwitz, Mooly Sagiv, and Genevieve Rosay. Speeding up slicing. In *Proc. FSE*
1873 *'94*, pages 11–20, New York, NY, USA, 1994. ACM.
- 1874 [31] Andrei Sabelfeld and David Sands. Probabilistic noninterference for multi-threaded programs. In
1875 *Proceedings of the 13th IEEE Computer Security Foundations Workshop, CSFW '00, Cambridge,*
1876 *England, UK, July 3-5, 2000*, pages 200–214, 2000.
- 1877 [32] Robert Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–
1878 160, 1972.
- 1879 [33] Michael Joseph Wolfe. *High Performance Compilers for Parallel Computing*. Addison-Wesley Longman
1880 Publishing Co., Inc., Boston, MA, USA, 1995.
- 1881
- 1882
- 1883
- 1884
- 1885
- 1886
- 1887
- 1888
- 1889
- 1890
- 1891
- 1892
- 1893
- 1894
- 1895
- 1896
- 1897
- 1898
- 1899
- 1900
- 1901
- 1902
- 1903
- 1904
- 1905
- 1906
- 1907
- 1908
- 1909
- 1910
- 1911