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 , then x decides not only whether y is executed but also how many timesteps 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 us 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 work not only 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.* 44, 1, Article 2 (December 2021), 37 pages.
<https://doi.org/10.1145/3486003>

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 **control dependencies (CDs)** to

The work described in this article 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. Authors' address: M. Hecker, S. Bischof, and G. Snelting, Karlsruhe Institute of Technology, Fakultät für Informatik, Am Fasanengarten 5, 76131 Karlsruhe, Germany; email: 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.

0164-0925/2021/12-ART2 \$15.00

<https://doi.org/10.1145/3486003>

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

CDs, originally introduced in References [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., Reference [25]) or for program transformations such as partial evaluation (e.g., Reference [20]) or refactoring (e.g., Reference [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

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

on node (1); node (3) is not control dependent on (1) but on (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 Reference [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 , and so on; 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 Reference [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, then 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 $\rightarrow_{\text{cd}}$ but also works for graphs without unique exit. $\rightarrow_{\text{ntscd}}$ is identical to $\rightarrow_{\text{wcd}}$ (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 runtime between x and y varies: $x \rightarrow_{\text{tscd}} y$. If this variation depends on secret data, and can be measured by an attacker, then 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 Reference [21].

Now, let us illustrate the differences between the different kinds of control dependencies. 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, then we might avoid reaching (10) by staying in the loop forever, so (3) decides if (10) will be executed.

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

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

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

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

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

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

2 CONTROL DEPENDENCE IN GRAPHS WITHOUT UNIQUE EXIT

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

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 post-domination 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 that 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 that never pass through m : $\forall k \exists \pi : \text{len}(\pi) = k \wedge m \notin \pi$. If this happens, then $\sqsubseteq_{\text{SPOST}}$ 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 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 Reference [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

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, Ranganath et al. [29] proposed control dependence for arbitrary graphs based on the notions of *maximal* and *sink* paths.

A maximal path is a path that cannot be extended (i.e., is infinite, or ends in some node n without successor). However, 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 that infinitely loops in S .

Definition 2.1 (Implicit in Reference [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 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.

²All examples and measurements in this article are based on CFGs that 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.

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

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

Note that this definition is identical to the original CD definition, with $\sqsubseteq_{MAX}$, respectively, $\sqsubseteq_{SINK}$ instead of $\sqsubseteq_{POST}$. In fact, for graphs with unique exit node, we have $\rightarrow_{ntscd} = \rightarrow_{wcd}$ and $\rightarrow_{nticd} = \rightarrow_{cd}$ [29]. Figure 3 shows $\rightarrow_{ntscd}$ and $\rightarrow_{nticd}$ 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 $\sqsubseteq_{SINK}$ can be characterized as a greatest, respectively, 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 that transforms a set of dominance relationships $\{x \sqsubseteq y\}$ into a new set $D(\{(x \sqsubseteq y)\})$. If such a functional is monotone, then it has a least as well as a greatest fixpoint.

THEOREM 2.1. 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 $\mu D = \sqsubseteq_{MAX}$, and $\nu D = \sqsubseteq_{SINK}$.

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

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 $\rightarrow_{ntscd}$ and $\rightarrow_{nticd}$ are needed, and indeed Ranganath et al. proposed such algorithms. The algorithm for $\rightarrow_{ntscd}$ from Reference [29] can in principle be thought of as a simple least fixed point computation of the set of nodes m such that $m \sqsubseteq_{MAX} n$ but only for nodes n that are successors of *branching* nodes.

We, however, discovered a more general and systematic algorithmic approach, which exploits the above fix-point theorem. It is based on the insight that Cytron’s efficient algorithm for dominance frontiers can be generalized to an abstract notion of “dominance” and thus can be used

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 Reference [8] is based. We will not describe P in detail but note that $\sqsubseteq_{POST} = \nu P$ [15].

Fig. 4. The canonical irreducible graph, where neither $n \rightarrow_{\text{ntscd}} m_1$ nor $n \rightarrow_{\text{ntscd}} m_2$.

for $\sqsubseteq_{\text{POST}}$, $\sqsubseteq_{\text{SPOST}}$, $\sqsubseteq_{\text{MAX}}$, $\sqsubseteq_{\text{SINK}}$, $\rightarrow_{\text{ntscd}}$, $\rightarrow_{\text{ntcd}}$, and in particular for $\rightarrow_{\text{tscd}}$. We will present all algorithms in a separate section (Section 5), as they demand a rather heavy technical machinery.

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

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

We omit the formal definition, but provide an intuition: In Reference [29], the necessity of $\rightarrow_{\text{dod}}$ was motivated by an irreducible⁶ graph, such as the graph shown in Figure 4. Here, neither m_1 nor m_2 is nontermination sensitively control dependent on n : $\neg(n \rightarrow_{\text{ntscd}} m_1) \wedge \neg(n \rightarrow_{\text{ntscd}} m_2)$. But the decision at n determines which node is executed next: Leaving n via the left branch will execute m_1 before m_2 but leaving n via the right branch will execute m_2 before m_1 . Thus, $n \rightarrow_{\text{dod}} (m_1, m_2)$ holds. Ranganath and Amtoft used $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{dod}}$ to define a sound notion of *nontermination-sensitive backward slicing*. This is consistent with the fact that CDs are fundamental for slicing and PDGs.

3 TIMING-SENSITIVE CONTROL DEPENDENCE

3.1 Why Time Sensitivity Matters

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

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

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

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

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

Fig. 5. Control Flow Graphs for AES Sbox substitution.

sbox is a constant array that typically spans multiple memory blocks, while $r1$, $r2$, $r3$ are registers. Thus the value of $r2$ in one iteration may influence whether the read of $r3$ in a later iteration is served from cache (namely if, earlier, the corresponding memory block was already loaded into the cache), or from main memory. This makes a difference in execution time and can be observed by an attacker; who may thus be able to infer the value of $r2$. 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 5 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 $r3 := \text{sbox}[r2]$ 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 $r2 := \text{state}[r1]$, 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):

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.

Fig. 6. Dependence of execution time of m_x on n .

(secret input) $\rightarrow_{dd}$ state $\rightarrow_{dd}$ r2 $\rightarrow_{dd}$ (3) $\rightarrow_{tscd}$ (5). This means that the time until (5) is reached depends on secret input. Thus time-sensitive CDs reveal that sbx access is not constant time (in contrast to the AES specification); opening a door to cache-leak based attacks.

If a $\rightarrow_{tscd}$ dependency is not (indirectly) data dependent on secret data, then it does *not* generate a timing leak. For example, in the AES CFG, (2) $\rightarrow_{tscd}$ (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 6(a): m_x is guaranteed to be executed, no matter which branch is taken at n , so we have $\neg(n \rightarrow_{ntscd} m_x)$. 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 $n \rightarrow_{tscd} m_x$. In Figure 6(b), however, m_x will always be executed 4 steps after n , so there we have $\neg(n \rightarrow_{tscd} m_x)$.

We will now formally define $\rightarrow_{tscd}$. 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 $\rightarrow_{tscd}$ of *timing-sensitive* control dependence. It will be based on $\sqsubseteq_{\text{TIME}}$ the same way that $\rightarrow_{ntscd}$ is based on $\sqsubseteq_{\text{MAX}}$.
- (4) Prove soundness and minimality of $\rightarrow_{tscd}$.

To start with, remember that $\sqsubseteq_{\text{MAX}}$ 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 6(a) 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 6(a) m_x can be reached via two different paths, *with varying execution time*. 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 $i < k$, then 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 $m \sqsubseteq_{\text{MAX}} n$, then 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 $m \sqsubseteq_{\text{TIME}} n$, then m must appear in all maximal paths from n at the same position. Therefore in Figure 6(a), $m_x \sqsubseteq_{\text{TIME}} n$ does not hold, while in Figure 6(b), $m_x \sqsubseteq_{\text{TIME}} n$ 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 $\rightarrow_{\text{ntcd}}$, 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 $\rightarrow_{\text{ntscd}}$, respectively, $\rightarrow_{\text{cd}}$; with $\sqsubseteq_{\text{TIME}}^k$ instead of $\sqsubseteq_{\text{MAX}}$, respectively, $\sqsubseteq_{\text{POST}}$. Thus $\rightarrow_{\text{tscd}}$ has the same formal structure as classical CD and its later extensions.

In Figure 6(a), we have $m_x \sqsubseteq_{\text{TIME}}^3 n'$ but $\neg(m_x \sqsubseteq_{\text{TIME}}^3 n'')$, and thus we have $n \rightarrow_{\text{tscd}} m_x$; while in Figure 6(b) 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 7(b). Figure 7(c) and Figure 7(d) 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: $\neg(7 \xrightarrow{*}_{\text{ntscd}} 8)$, 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 $2 \xrightarrow{*}_{\text{tscd}} 8$ 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 $\sqsubseteq_{\text{MAX}}$ is the least fixed point of the rule system D,

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

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

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

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

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 $\rightarrow_{\text{dod}}$: 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 $\rightarrow_{\text{dod}}$. Now remember that Ranganath and Amtoft introduced $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{dod}}$ 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 $\rightarrow_{\text{dod}}$. As seen above, the necessity of $\rightarrow_{\text{dod}}$ was motivated by an irreducible graph, such as the graph in Figure 4. But while in Figure 4 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 equivalence. We imagine an attacker who tries to infer secret values (such as r_2 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, respectively, 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 (unclocked) 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.

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

¹⁰This kind of security analysis is called 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.

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

Definition 3.6. A clocked trace is a trace where every step is additionally annotated with a time stamp. We write $t \text{ @ } [i]$ 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

$$t^{\text{@}} = (n_e, n_1) \text{ @ } [0], (n_1, n_2) \text{ @ } [1], \dots$$

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

¹¹The assumption of a fixed, static S , and batchlike 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.

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 , then 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, then 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 dependencies* 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 unclocked traces are S -equivalent, their clocked versions may be different. This is the essence of time-sensitivity. For illustration consider Figure 6(a), 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, then 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:

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

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. However, the program in Figure 6(b) has no timing leak: Even if we annotate each edge in the observable trace with its execution time, all inputs i, i' starting in m have the same observable *clocked* trace

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

This discussion motivates the following definition of timing leaks:

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

$$i \sim_S i' \implies t_i \Big|_S = t_{i'} \Big|_S.$$

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

To prevent a timing leak, it is necessary that *all nodes that influence the timing of observable nodes $\in S$ are observable itself*. Otherwise, a secret node might influence the timing of an observable node. For example, Figure 6(a) contains—as described above—a timing leak if we assume $S = \{m, m_x\}$.

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

Indeed $n \rightarrow_{\text{tscd}} m_x$, but *not* $n \in S$. With $S' = S \cup \{n\} = \{m, n, m_x\}$, the timing leak disappears: While the timing of m_x still differs, i and i' are now distinguishable for the attacker, so this timing difference does not give additional information.

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

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

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

$$t_i^\circ|_B = t_{i'}^\circ|_B.$$

COROLLARY 3.1. If $BS(S) \subseteq S$, then 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 $x \in BS(S)$: $x \notin S$, then 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

$$t_i^\circ|_{B'} \neq t_{i'}^\circ|_{B'}.$$

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 dependencies in addition to control dependencies. Data dependencies 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, and so on, 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., References [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 that may influence M ; other nodes that 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, then 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

In the abstract, we mentioned the infamous AES cache timing leaks that 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, respectively, 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 $a, b, \dots$, and uncacheable registers $r1, r2, \dots$.

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 9(b), 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 $(9, [x, d, c, b])$ represents all those concrete configurations at control node 9 in which the concrete micro-architectural cache contains cached values for the variables $[x, d, c, b]$, in this order (with arbitrary concrete macro-architectural state).

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

In the example, executions can reach the control node $m = 15$ at cache states represented by either $[b, y, c, x]$, or by $[b, y, d, x]$. 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 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 y , which is the only cacheable variable accessed by the edge $15 \xrightarrow{r_2:=y} 16$ at m .

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

Now consider the read of 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 c : $(21, [b, y, c, x])$ and $(21, [a, y, b, c])$. The other two do *not* contain c : $(21, [a, y, b, d])$ and $(21, [b, y, d, x])$. The cache states containing c are reachable from configurations at control node $n' = 16$. At the same time, cache states *not* containing c are *also* reachable from configurations at control node $n' = 16$. But in fact, whether 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 $(16, [y, b, c, x])$. But whenever $m = 21$ is reached from this abstract configuration, c is in the cache (either $(21, [b, y, c, x])$ or $(21, [a, y, b, c])$). The second abstract configuration at which $n' = 16$ can be reached is $(16, [y, b, d, x])$. But whenever $m = 21$ is reached from that configuration, c is *not* in the cache ($(21, [a, y, b, d])$ or $(21, [b, y, d, x])$).

However, the cache status of 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 9(b).

Note that through a small modification of the program, the cache status of 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., e, f) right before $m = 21$, then *all* cache states at m would *not* have contained c . This is because these two reads would have evicted c even from $[b, y, c, x]$ (and $[a, y, b, c]$).

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

Figure 10 shows the micro-architectural-aware CFG G' for Figure 9; together with an explicit timing cost model C' . A cache-miss is assumed to take 10 units of time, while a cache-hit takes 2 units.¹² At node 14, the read from b takes either 2 or 10 units of time, since b there might either be in the cache, or not.¹³ Hence in G' , node 14 has *two* artificial successors: The read from b takes either 2 or 10 units of time, since b there might either be in the cache, or not. However, node 15 still has only one successor, reached with timing cost $3 = 2 + 1$ (cache access plus register access), since we found that there the variable y is always in cache.

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

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

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

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

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

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

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

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

5 ALGORITHMS

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

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

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

Both $\sqsubseteq_{\text{MAX}}$ and $\sqsubseteq_{\text{SINK}}$ will always be represented by their transitive reductions; allowing efficient construction algorithms. A transitive reduction $<$ of a transitive relation $\sqsubseteq$ is 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 $\sqsubseteq_{\text{MAX}}$ 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 $\sqsubseteq_{\text{MAX}}$ 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 $\sqsubseteq_{\text{MAX}}$ and $\sqsubseteq_{\text{SINK}}$ are reflexive and transitive but not necessarily antisymmetric.

2. Any transitive, reflexive reduction $<_{\text{MAX}}$ of $\sqsubseteq_{\text{MAX}}$ 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 $\sqsubseteq_{\text{MAX}}$ 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* $\sqsubseteq_{\text{MAX}}$ -postdominate node 3, because the loop at 3 may not terminate. However, 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.

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

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

```

Input      :A pseudo-forest  $\prec$ , represented as a map  $\text{IMDOM} : N \hookrightarrow N$  s.t.  $\text{IMDOM}[n] = m$  iff
                $m \prec n$ .
Input      :Nodes  $m_0, n_0$ 
Output     :A least common ancestor of  $n_0, m_0$ , or  $\perp$  if there is none.
begin
  | return  $\text{lca}(n_0, m_0)$ 
end

Function  $\text{lca}(\pi_n, \pi_m)$ 
  Input      :A  $\prec$ -path  $\pi_n = n_0, \dots, n$  ending in  $n$ 
  Input      :A  $\prec$ -path  $\pi_m = m_0, \dots, m$  ending in  $m$ 
  if  $m \in \pi_n$  then return  $m$ 
  switch  $\text{IMDOM}[n]$  do
    | case  $\perp$  do return  $\text{lin}[\pi_n](\pi_m)$ 
    | case  $n'$  do
      | if  $n' \in \pi_n$  then
      | | return  $\text{lin}[\pi_n](\pi_m)$ 
      | end
      | return  $\text{lca}(\pi_m, \pi_n n')$ 
    | end
  end
end

Function  $\text{lin}[\pi_n](\pi_m)$ 
  Input      :A  $\prec$ -path  $\pi_m = m_0, \dots, m$  ending in  $m$ 
  Implicit    :A  $\prec$ -path  $\pi_n = n_0, \dots, n$  ending in  $n$ 
  switch  $\text{IMDOM}[m]$  do
    | case  $\perp$  do return  $\perp$ 
    | case  $m'$  do
      | if  $m' \in \pi_n$  then return  $m'$ 
      | if  $m' \in \pi_m$  then return  $\perp$ 
      | return  $\text{lin}[\pi_n](\pi_m m')$ 
    | end
  end
end

```

We will now present new algorithms to compute $\sqsubseteq_{\text{MAX}}$ and $\sqsubseteq_{\text{SINK}}$. The representation of both $\sqsubseteq_{\text{MAX}}$ 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 $\sqsubseteq_{\text{MAX}}$ is a standard fixpoint iteration. Beginning with the empty pseudo-forest, new edges are added to $\prec_{\text{MAX}}$ according to Theorem 2.1 until a fixpoint is reached. Since $\sqsubseteq_{\text{MAX}}$ is efficiently represented by a pseudo-forest $\prec_{\text{MAX}}$, it is straightforward to derive an efficient algorithm for the computation of $\sqsubseteq_{\text{MAX}}$, see Algorithm 2. In addition, we need an efficient implementation of set-intersection in the representation $\prec$, i.e., a *least common ancestor* algorithm $\text{lca}_\prec$ for pseudo-forests; see Algorithm 1.

Algorithm 1 calculates the least common ancestor of n_0 and m_0 by alternately extending $\prec$ -paths from n_0 and m_0 one by one. If the newly added node is already contained in the other path, then 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), then only the other path is extended (procedure

Fig. 11. Nontermination-sensitive Postdominance.

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 12 (left), these would be $\text{IMDOM}[5] = 7$, $\text{IMDOM}[7] = 8$, $\text{IMDOM}[8] = 9$, and $\text{IMDOM}[9] = 8$.

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

In the case of the graph in Figure 12 (left), assume we iterate in order $\text{COND}_G = [1, 2, 3, 4]$, which becomes our first workqueue. For $x = 1$, we calculate $\text{lca}(\{2, 3, 4\})$. Let us suppose we try to calculate $\text{lca}(2, 3) = \text{lca}([2], [3])$ first. Since $\text{IMDOM}[2] = \perp$, we immediately call $\text{lin}[[2]]([3])$, but since $\text{IMDOM}[3] = \perp$ as well, we return $\perp$ as $\text{lca}(2, 3)$. But then $\text{lca}(\{2, 3, 4\}) = \perp$ as well. 1 is therefore put back into the queue, so we now have $\text{workqueue} = [2, 3, 4, 1]$ and $\text{oldest} = 1$. For $x = 2$, when calculating $\text{lca}(6, 7)$, we have $\text{IMDOM}[6] = \perp$, so we immediately call $\text{lin}[[6]]([7])$. During the recursion in lin , we extend $[7]$ to $[7, 8]$ and $[7, 8, 9]$ (since no new node is in $[6]$). The next step would be $\text{IMDOM}[9] = 8$. Since $8 \in [7, 8, 9]$ (which would lead to a loop), we return $\perp$. 2 is therefore put back into the queue, so we now have $\text{workqueue} = [3, 4, 1, 2]$ and $\text{oldest} = 1$. For $x = 3$, when calculating $\text{lca}(5, 7)$, we have $\text{IMDOM}[5] = 7$ and $7 \in [7]$, so we return 7 as our lca. Since we have an lca, we update $\text{IMDOM}[3] = 7$, keep 3 out of the workqueue (so $\text{workqueue} = [4, 1, 2]$) and set $\text{oldest} = \perp$. For $x = 4$, when calculating $\text{lca}(9, 5)$, we extend both paths alternately until the path starting in 9 would enter a loop, then only the path starting in 5 is extended. Then we will find that 8 is our lca, since it is in both paths. In detail, $\text{lca}([9], [5]) = \text{lca}([5], [9, 8]) = \text{lca}([9, 8], [5, 7]) = \text{lin}[[9, 8]]([5, 7]) = \text{lin}[[9, 8]]([5, 7, 8]) = 8$. We update $\text{IMDOM}[4] = 8$, keep 4 out of the workqueue (so $\text{workqueue} = [1, 2]$) and keep $\text{oldest} = \perp$. Now we are back to $x = 1$. When calculating $\text{lca}(2, 3)$, we now have $\text{IMDOM}[3] = 7$, so we can extend $[3]$ until we get $[3, 7, 8, 9]$. But still, no node is contained in $[2]$, so we still have $\perp$ as our lca. We put 1 back into the workqueue (so $\text{workqueue} = [2, 1]$) and set $\text{oldest} = 1$. After finding for $x = 2$ that $\text{lca}(6, 7) = \perp$, we put 2 back into the workqueue

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

```

Input : A CFG  $G$ 
Data: A pseudo-forest  $<$  represented as a map  $\text{IMDOM} : N \hookrightarrow N$  s.t.  $\text{IMDOM}[n] = m$  iff  $m < n$ 
Output: A transitive reduction  $<_{\text{MAX}}$  of  $\sqsubseteq_{\text{MAX}}$ 
begin
  for  $x \in N$ ,  $\{z \mid x \rightarrow z\} = \{z\}$ ,  $z \neq x$  do
     $\text{IMDOM}[x] \leftarrow z$ 
  end
   $\text{MAXIMAL}_{\text{up}}$ 
  return  $\text{IMDOM}$ 
end

Procedure  $\text{MAXIMAL}_{\text{up}}$ 
   $\text{workqueue} \leftarrow \text{COND}_G$ 
   $\text{oldest} \leftarrow \perp$ 
  while  $\text{workqueue} \neq \emptyset$  do
     $x \leftarrow \text{removeFront}(\text{workqueue})$ 
    assert  $\text{IMDOM}[x] = \perp$ 
    if  $\text{oldest} = x$  then
      return
    end
    if  $\text{oldest} = \perp$  then
       $\text{oldest} \leftarrow x$ 
    end
     $a \leftarrow \text{lca}(\{y \mid x \rightarrow y\})$ 
     $z \leftarrow \begin{cases} \perp & \text{if } a = \perp \vee a = x \\ a & \text{otherwise} \end{cases}$ 
    if  $z \neq \perp$  then
       $\text{IMDOM}[x] \leftarrow z$ 
       $\text{oldest} \leftarrow \perp$ 
    else
       $\text{pushBack}(\text{workqueue}, x)$ 
    end
  end
end

```

(so $\text{workqueue} = [1, 2]$) and keep $\text{oldest} = 1$. But now our next element in the queue is our oldest, so we are done.

The computation of $\sqsubseteq_{\text{SINK}}$ is slightly more complicated. As it is a greatest fixpoint, in principle we must start with $N \times N$ and reduce it according to the rules; until the greatest fixpoint is reached. But $N \times N$ cannot be represented by a pseudo-forest. Hence we need to initialize the fixed point iteration with an approximation $\sqsubseteq_0$ of $\sqsubseteq_{\text{SINK}}$ (i.e., $\sqsubseteq_0 \supseteq \sqsubseteq_{\text{SINK}}$) that is representable by a pseudo-forest $<_0$. We can build $<_0$ by interleaving a traversal of a preliminary pseudo forest $<$ with $\text{lca}_<$ computations. Consider the preliminary $<$ in Figure 13(b). We need to establish $3 < 1$, but find that $\text{lca}_<(\{2, 3\}) = \perp$ for the successors of 1. We would like to assume *both* $3 < 1$ and $2 < 1$, the latter of which would then be invalidated in the (downward) fixed point iteration. But then $<$ no longer would be a pseudo forest. If we assumed just $2 < 1$, then we would obtain a $<_0$ such that *not*:

Fig. 12. Two example CFGs.

Fig. 13. Computing an initial approximation $<_0$.

$<_0^* \supseteq \sqsubseteq_{\text{SINK}}$, so we need to make the assumption $3 < 1$. This example illustrates how the fixpoint iteration must proceed. It is based on the following:

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

Here, “bound for S ” means that the path cannot escape sink S . To illustrate the iteration for $\sqsubseteq_{\text{SINK}}$, consider Figure 13(b). For node 3 we have already established $4 <^* 3$ for the sink node $4 \in S$, but we have not yet established $4 <^* 2$. This suggests that we must—whenever $\text{lca}_{<}(\{y \mid x \rightarrow y\}) = \perp$ —choose some successor node y of x such that already $s <^* y$ for some sink node s . We call such nodes y *processed*, and maintain a set PROCD of all such nodes. Algorithm 3 presents the computation of $<_{\text{SINK}}$, the additional procedures performing the iteration are given in Figure 14.

Algorithm 3 first initializes ISDOM for sink nodes and nodes with one successor. Remember that any nontrivial sink S_i contains a $<_{\text{SINK}}$ -cycle. For each sink S_i , we therefore initialize ISDOM to be such a cycle in arbitrary order. We also choose a representative s_i for each sink 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 $\text{ISDOM}^* \supseteq \sqsubseteq_{\text{SINK}}$ but might be too optimistic: For nodes x , $\text{ISDOM}[x]$ 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 $\text{ISDOM}[x]$). 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 that have already been processed. If the resulting lca is $\perp$, then we choose an arbitrary successor as lca . Now, we set $\text{ISDOM}[x]$ 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.

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

```

Input : A CFG  $G$ 
Data: A pseudo-forest  $<$  represented as a map  $\text{ISDOM} : N \hookrightarrow N$  s.t.  $\text{ISDOM}[n] = m$  iff  $m < n$ 
Output: A transitive reduction  $<_{\text{SINK}}$  of  $\sqsubseteq_{\text{SINK}}$ 
begin
   $\{S_1, \dots, S_n\} \leftarrow \{S_i \mid S_i \in \text{sccs}(G), \neg \exists s \rightarrow n. s \in S_i \wedge n \notin S_i\}$ 
   $S \leftarrow \bigcup S_i$ 
  for  $1 \leq i \leq n$  do
     $s_i \leftarrow \text{any node in } S_i$ 
    for  $n_j \in S_i$  in any fixed ordering  $n_1, \dots, n_{k_i}$  of  $S_i$  do
       $\text{ISDOM}[n_j] \leftarrow n_{j+1 \bmod k_i}$  unless  $k_i = 1$ 
       $\text{processed}(n_j)$ 
    end
  end
  for  $x \in N, x \notin S, \{z \mid x \rightarrow z\} = \{z\}, z \neq x$  do
     $\text{ISDOM}[x] \leftarrow z$ 
    if  $z \in \text{PROCD}$  then  $\text{processed}(x)$ 
  end
   $\text{SINK}_{\text{up}}$ 
   $\text{SINK}_{\text{down}}$ 
  return  $\text{ISDOM}$ 
end

```

Then, these spurious postdominances are eliminated during $\text{SINK}_{\text{down}}$. For each conditional node x outside sinks the lca of its successors is calculated. If it is part of a sink S_i , then its distinguished representative s_i is chosen instead. If it is different from the current ISDOM (either a different node or $\perp$), then 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 12 (left). In the initial phase, we set $\text{ISDOM}[8] = 9$ and $\text{ISDOM}[9] = 8$ for the non-trivial sink. For its representative, let us assume we choose 8. We mark all sink nodes 6, 8, and 9 as processed. Then, we handle non-condition nodes. We set $\text{ISDOM}[4] = 9$ and mark 4 as processed. After that, we set $\text{ISDOM}[5] = 7$ (but cannot mark it as processed, since 7 is not). Finally, we set $\text{ISDOM}[7] = 8$ and mark both 7 and 5 as processed.

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

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

Procedure SINK_{up}

```

workqueue  $\leftarrow$  CONDG \ S in any order
while workqueue  $\neq \emptyset$  do
   $x \leftarrow$  removeFront(workqueue)
  assert ISDOM[ $x$ ] =  $\perp \wedge x \notin$  PROCD
  SUCCS  $\leftarrow \{y \mid x \rightarrow y, y \in \text{PROCD}\}$ 
  if SUCCS =  $\emptyset$  then
     $z \leftarrow \perp$ 
  else
     $a \leftarrow \text{lca}(\text{SUCCS})$ 
     $z \leftarrow \begin{cases} \text{any } y \in \text{SUCCS} & \text{if } a = \perp \\ a & \text{otherwise} \end{cases}$ 
  end
  if  $z \neq \perp$  then
    ISDOM[ $x$ ]  $\leftarrow z$ 
    processed( $x$ )
  else
    pushBack(workqueue,  $x$ )
  end
end

```

Procedure SINK_{down}

```

workset  $\leftarrow \{n \mid n \in \text{COND}_G \setminus S, \text{ISDOM}[n] \neq \perp\}$ 
while workset  $\neq \emptyset$  do
   $x \leftarrow$  removeMin(workset)
   $a \leftarrow \text{lca}(\{y \mid x \rightarrow_G y\})$ 
   $z \leftarrow \begin{cases} \perp & \text{if } a = \perp \\ s_i & \text{if } a \in S_i \\ a & \text{otherwise} \end{cases}$ 
  assert ISDOM[ $x$ ] =  $\perp \Rightarrow z = \perp$ 
  if  $z \neq \text{ISDOM}[x]$  then
    workset  $\leftarrow$  workset  $\cup \{n \in \text{COND}_G \setminus S \mid \exists n \rightarrow y. x <^* y\}$ 
    ISDOM[ $x$ ]  $\leftarrow z$ 
  end
end

```

Fig. 14. Upward and downward iteration for Algorithm 3.

5.2 Postdominance Frontiers in Graphs without Unique Exit

We will now derive algorithms for $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{nticd}}$, based on $\sqsubseteq_{\text{MAX}}$ and $\sqsubseteq_{\text{SINK}}$. In particular, we generalize Cytron’s idea to split up the postdominance frontier into an “up” and a “local” part, and to follow the tree structure (parent links) while iterating. The latter also works for pseudo-forests.

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

Definition 5.2 (Immediate $\sqsubseteq$ -Postdominance). Given a binary relation $\sqsubseteq$ on nodes, a node 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 set $\text{ipdom}_{\sqsubseteq}(n)$ is defined by

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

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

As an example, consider the CFG in Figure 12 (left) with $\sqsubseteq_{\text{MAX}}$ -postdominance. We have $\text{ipdom}_{\sqsubseteq}(5) = \{7\}$, since $7 1\text{-}\sqsubseteq 5$ and each $1\text{-}\sqsubseteq$ -postdominator of 5 also postdominates 7. $8 \notin \text{ipdom}_{\sqsubseteq}(5)$, because $7 1\text{-}\sqsubseteq 5$ but not $7 \sqsubseteq 8$. For the cycle of 8 and 9, each of those $1\text{-}\sqsubseteq$ -postdominates itself and the other one, so we have $\text{ipdom}_{\sqsubseteq}(8) = \text{ipdom}_{\sqsubseteq}(9) = \{8, 9\}$.

ALGORITHM 4: Computation of $\text{pdf}_{\sqsubseteq}$

Input : A transitive reduction $<$ of $\sqsubseteq$
Input : A map SCC from nodes x to the strongly connected component c of $<$ s.t $x \in c$
Input : A topological sorting sccs of all strongly connected components of $<$.
Output: $\text{pdf}_{\sqsubseteq}$

for $\text{scc} \in \text{sccs}$ **do**
 $\text{local} \leftarrow \{y \mid x \in \text{scc}, y \rightarrow x, \underbrace{\neg \exists x' \in \text{scc}. x' < y}_{y \in \text{scc}_{<}}\}$
 $\text{up} \leftarrow \{y \mid \underbrace{x < z, y \in \text{PDF}[z]}_{z \in \text{scc}_{<}}, \underbrace{\neg \exists x' \in \text{scc}. x' < y}_{y \in \text{scc}_{<}}\}$
 for $x \in \text{scc}$ **do** $\text{PDF}[x] \leftarrow \text{local} \cup \text{up}$
end

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

Definition 5.3 ($\sqsubseteq$ -Postdominance Frontiers).

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

Consider again Figure 12 (left) with $\sqsubseteq_{\text{MAX}}$ -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 $\text{pdf}_{\sqsubseteq}$. 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. However, 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).

$$\begin{aligned} \text{pdf}_{\sqsubseteq}^{\text{local}}(x) &= \{y \mid \neg x \sqsubseteq y, y \rightarrow x\} \\ \text{pdf}_{\sqsubseteq}^{\text{up}}(z) &= \{y \in \text{pdf}_{\sqsubseteq}(z) \mid \forall x \in \text{ipdom}_{\sqsubseteq}(z). \neg x \sqsubseteq y\}. \end{aligned}$$

Under suitable conditions, $\text{pdf}_{\sqsubseteq}^{\text{up}}$ and $\text{pdf}_{\sqsubseteq}^{\text{local}}$ indeed partition $\text{pdf}_{\sqsubseteq}$. 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, $\sqsubseteq_{\text{MAX}}$ 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 12 (left) with $\sqsubseteq_{\text{MAX}}$ -postdominance. We have $\text{pdf}_{\sqsubseteq}^{\text{local}}(5) = \{3, 4\}$. Since 5 only postdominates itself trivially, we have $5 \in \text{ipdom}_{\sqsubseteq}(z)$ for no node z , and Observation 5.3 indeed gives $\text{pdf}_{\sqsubseteq}(5) = \{3, 4\}$. We have $\text{pdf}_{\sqsubseteq}^{\text{local}}(7) = \{2\}$. Since we have $\{z \mid 7 \in \text{ipdom}_{\sqsubseteq}(z)\} = \{3, 5\}$, we need to calculate $\text{pdf}_{\sqsubseteq}^{\text{up}}(3)$ and $\text{pdf}_{\sqsubseteq}^{\text{up}}(5)$. For 5, we have already seen $\text{pdf}_{\sqsubseteq}(5) = \{3, 4\}$. But $\text{pdf}_{\sqsubseteq}^{\text{up}}(5)$ contains only node 4, since 7 actually postdominates 3! Since $\text{pdf}_{\sqsubseteq}^{\text{up}}(3) = \{1\}$, Observation 5.3 results in $\text{pdf}_{\sqsubseteq}(7) = \{1, 2, 4\}$, as expected.

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

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

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

$\sqsubseteq$ lacks joins if it admits the rules

$$\frac{\begin{array}{ccc} x \in \text{ipdom}_{\sqsubseteq}(v) & v \sqsubseteq s & \\ x \in \text{ipdom}_{\sqsubseteq}(z) & z \sqsubseteq s & z \neq v \end{array}}{v \in \text{ipdom}_{\sqsubseteq}(z) \vee z \in \text{ipdom}_{\sqsubseteq}(v)} \text{NoJoin}.$$

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

LEMMA 5.2. Both $\sqsubseteq_{\text{MAX}}$ and $\sqsubseteq_{\text{SINK}}$ are closed under $\rightarrow$, and lack joins.

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

OBSERVATION 5.4. Let $\sqsubseteq$ be transitive, and closed under $\rightarrow$. Then

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

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

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

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

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

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

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

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

- (1) compute the strongly connected components sccs of the graph $(N, <)$, in a corresponding topological order. These can either be provided by the algorithm computing $<$, or by Tarjan's algorithm [32].
- (2) compute $\text{pdf}_{\sqsubseteq}$ by traversing the condensed graph in that order *once*.

This concludes the algorithm for generalized postdominance frontiers $\text{pdf}_{\sqsubseteq}(x)$; 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 $\text{pdf}_{\sqsubseteq}(x)$, consider once more the CFG in Figure 12 (left). The “local” rule gives us, e.g., $1 \in \text{pdf}_{\sqsubseteq}(3)$, $3 \in \text{pdf}_{\sqsubseteq}(5)$, $4 \in \text{pdf}_{\sqsubseteq}(5)$ and $2 \in \text{pdf}_{\sqsubseteq}(7)$. With the “up” rule we can now get $4 \in \text{pdf}_{\sqsubseteq}(7)$ by instantiating the rule with $x = 7$, $y = 4$ and $z = 5$. Note that we indeed have shown $4 \in \text{pdf}_{\sqsubseteq}(5)$ earlier and we have $7 \notin \text{ipdom}_{\sqsubseteq}(4)$ as well as $7 \in \text{ipdom}_{\sqsubseteq}(5)$. In contrast, if we try to use $3 \in \text{pdf}_{\sqsubseteq}(5)$ to show $3 \in \text{pdf}_{\sqsubseteq}(7)$ (which is false), then $7 \notin \text{ipdom}_{\sqsubseteq}(3)$ would have to hold. But $7 \in \text{ipdom}_{\sqsubseteq}(3)$, so the right rule is not applicable, and we are prevented from showing $3 \in \text{pdf}_{\sqsubseteq}(7)$.

5.3 Timing-sensitive Postdominance Frontiers

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

- (1) Transitivity of $\subseteq_{\text{MAX}}$ allows us to efficiently compute and represent $\subseteq_{\text{MAX}}$ in form of its transitive reduction $<_{\text{MAX}}$. Here, $<_{\text{MAX}}$ turned out to be a pseudo-forest.
- (2) Transitivity of $\subseteq_{\text{MAX}}$, 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 $n \sqsubseteq_{\text{TIME}}^k m$, and not in terms of its binary “ $\exists k$. -closure” $n \sqsubseteq_{\text{TIME}} m$, let us investigate first if $n \sqsubseteq_{\text{TIME}} m$ is—in general—transitive. Consider the (irreducible) CFG in Figure 15(a). Here, every maximal path starting in n first reaches m_1 after two steps, hence $m_1 \sqsubseteq_{\text{TIME}} n$. Also, every maximal path starting in m_1 first reaches m_2 after one step, hence $m_2 \sqsubseteq_{\text{TIME}} m_1$. But it is for *no* number k of steps the case that $m_2 \sqsubseteq_{\text{TIME}}^k n$; hence, $\neg m_1 \sqsubseteq_{\text{TIME}} n$. In summary, $\subseteq_{\text{TIME}}$ is *not* transitive.

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

THEOREM 5.1. Let G be any reducible CFG. Then $\subseteq_{\text{TIME}}$ is transitive.

THEOREM 5.2. Let G be any CFG with unique exit node n_x . Then $\sqsubseteq_{\text{TIME}}$ is transitive.

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

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

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

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

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

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

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

$$x \sqsubseteq y \sqsubseteq z.$$

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

THEOREM 5.3. Let $n \neq m \in N$. Then

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

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

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

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

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

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

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

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

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

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

Input: A $\mathbb{N}$ labeled pseudo-forest $<$, represented as a map $\text{IDOM} : N \hookrightarrow N \times \mathbb{N}$ s.t. $\text{IDOM}[n] = (m, k)$ iff $m <^k n$

Input: Numbers $k_0^n, k_0^m \in \mathbb{N}$ and nodes n_0, m_0

Output: $\text{lca}_{<}((n_0, k_0^n), (m_0, k_0^m))$ if it exists, or $\perp$ otherwise.

return $\text{lca}((n_0, k_0^n, [n_0 \mapsto k_0^n]), (m_0, k_0^m, [m_0 \mapsto k_0^m]))$

Function $\text{lca}(\pi_n, \pi_m)$

Input: A cycle free $<$ -path $\pi_n = n_0, \dots, n$ ending in n , represented by a tuple (n, k^n, KS_n) where KS_n is a map on the nodes n appearing in π_n s.t. $k^n = \text{KS}_n[n]$ and for any such n : $\text{KS}_n[n] = k_0^n + \sum_i k_i$ where $n <^{k_c} \dots <^{k_1} n_0$ in π_n

Input: A $<$ -path $\pi_m = m_0, \dots, m$ likewise

if $k^n > k^m$ **then return** $\text{lca}(\pi_m, \pi_n)$

if $n \in \pi_m \wedge k^n = \text{KS}_m[n]$ **then return** (n, k^n)

if $n \in \pi_m \wedge k^n \neq \text{KS}_m[n]$ **then return** $\perp$

switch $\text{IDOM}[n]$ **do**

case $\perp$ **do return** $\perp$

case $(n', k^{n'})$ **do**

if $n' \in \pi_n$ **then return** $\perp$

$\text{KS}_n[n'] \leftarrow k^n + k^{n'}$

return $\text{lca}((n', k^n + k^{n'}, \text{KS}_n), \pi_m)$

end

end

end

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 12 (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 $\text{lca}((2, 1), (9, 1))$, then 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 $\perp$.

6 MEASUREMENTS

We evaluated the performance of our algorithms on (a) control flow graphs of Java methods, as generated by the JOANA system for various third party Java programs; (b) randomly generated graphs $G = (N, E)$ usually with $|E| = 2|N|$, as generated by the standard generator from the JGraphT [26]

Fig. 16. Computation of $\langle \text{MAX} \rangle$. The orange line shows chaotic iteration performance, the blue line shows Algorithm 2.

library. In some cases, we additionally use ladder graphs,¹⁵ which are used to represent bad case behaviour.

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

Unless explicitly stated otherwise, all data points represent the average over $n + 1$ runs of the benchmark, where n is at least the number of runs that can be finished within 1 s. For example, the data point at $|N| = 21,076$, time = 18 ms in Figure 16(a) stands for the average of at least ≈ 50 runs of the benchmark that finished within 1 s. However, the data point in at $|N| = 65,000$, time = 88s in Figure 16(c) results from only one run of the benchmark.

The purpose of these benchmarks is to give a general idea of the scalability of the algorithms. For example, the benchmarks in the upper left and upper right of Figure 17 suggest that our new algorithm for the computation of nontermination-sensitive control dependence $\rightarrow_{\text{ntscd}}$ appears to scale almost linearly for “average” CFGs, while Ranganath’s original algorithm [29] clearly grows super-linearly for such graphs. The benchmarks can be summarized as follows:

- (1) For “average” CFGs, our algorithms for $\rightarrow_{\text{ntscd}}$, $\rightarrow_{\text{nticd}}$, and $\rightarrow_{\text{tscd}}$ offer performance “almost linear” in the size of the graph.

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

Fig. 17. Computation of $\rightarrow_{ntscd}$.

- (2) But for “bad case” CFGs, some algorithms perform decidedly super-linear, and become impractical for very large such graphs.

6.1 Nontermination Sensitive Postdominance

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

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

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

6.2 Nontermination Insensitive Postdominance

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

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

Fig. 18. Computation of $<_{\text{SINK}}$.

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 $\text{lca}_<$ by chasing (pseudo-tree) pointers, it can also be computed by comparison of postorder numbers, as in Reference [8].

For Java CFGs (Figure 18(a)) the fixed-iteration order variant of Algorithm 3 (▼) performs on par with the Algorithm 3 as stated (+). For randomly generated graphs (Figure 18(b)) the variant (▼) appears to perform a bit better than the original (+) for very large graphs, roughly on-par with the implementation based on postorder numbers (■).

Using reversed-topological iteration order, ladder graphs (Figure 18(c)) expose non-linear *bad-case* behavior only for Algorithm 3 (+) and its variant (▼). Even with a bad iteration order, performance for these two algorithm is not much worse (Figure 18(d)). However, the postorder number based implementation (■) 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] (●).

6.3 Generalized Postdominance Frontiers

When Algorithm 4 is instantiated with $<_{\text{MAX}}$, this yields an algorithm for $\rightarrow_{\text{ntscd}}$. The benchmarks for $\rightarrow_{\text{ntscd}}$ include the computation time of both Algorithm 4 and $<_{\text{MAX}}$ (▼). 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 17, upper left and right). Ladder graphs expose non-linear *bad-case*

Fig. 19. Computation of $\rightarrow_{nticd}$ via Algorithm 4 based on Algorithm 3. Fig. 20. Computation of $\rightarrow_{tscd}$.

behavior even for Algorithm 4 (Figure 17(c)). This cannot be circumvented, since in these ladder graphs, the size of the relation $\rightarrow_{ntscd}$ is quadratic in the number of nodes. Likewise, Algorithm 4 can be instantiated with $<_{SINK}$. Performance of the resulting $\rightarrow_{nticd}$ is shown in Figure 19.

6.4 Timing-sensitive CD

Whenever $\sqsubseteq_{TIME}$ is transitive, we can use Algorithm 4 to compute timing-sensitive control dependence $\rightarrow_{tscd}$. We thus measure the computation time for $\rightarrow_{tscd}$ on graphs for which $\sqsubseteq_{TIME}$ is transitive. These are control flow graphs from Java programs in subfigures 20(a), randomly generated graphs 20(b), and ladder graphs (c). We use Algorithm 4, and obtain a transitive reduction $<_{TIME}$ of $\sqsubseteq_{TIME}$ via the modification of Algorithm 3 that uses the upwards iteration of Algorithm 5.

The benchmarks for $\rightarrow_{\text{tscd}}$ in Figure 20 include the computation time of all sub-algorithms (+). Ladder graphs expose non-linear *bad-case* behavior.

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 that 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] that 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; and (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.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] Johan Agat. 2000. Transforming out timing leaks. In *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL’00)*. ACM, New York, NY, 40–53. <https://doi.org/10.1145/325694.325702>
- [2] A. V. Aho, M. R. Garey, and J. D. Ullman. 1972. The transitive reduction of a directed graph. *SIAM J. Comput.* 1 (2) (1972), 131–137.
- [3] Torben Amtoft. 2008. Slicing for modern program structures: A theory for eliminating irrelevant loops. *Inf. Process. Lett.* 106, 2 (2008), 45–51. <https://doi.org/10.1016/j.ipl.2007.10.002>
- [4] Daniel J. Bernstein. 2005. *Cache-Timing Attacks on AES*. Technical Report.
- [5] David W. Binkley, Mariano Ceccato, Mark Harman, Filippo Ricca, and Paolo Tonella. 2006. Tool-Supported Refactoring of Existing Object-Oriented Code into Aspects. *IEEE Trans. Softw. Eng.* 32, 9 (2006), 698–717.
- [6] Simon Bischof, Joachim Breitner, Jürgen Graf, Martin Hecker, Martin Mohr, and Gregor Snelting. 2018. Low-deterministic security for low-nondeterministic programs. *J. Comput. Secur.* 26, 3 (2018), 335–366. <https://doi.org/10.3233/JCS-17984>
- [7] Joachim Breitner, Jürgen Graf, Martin Hecker, Martin Mohr, and Gregor Snelting. 2016. On improvements of low-deterministic security. In *Principles of Security and Trust, Lecture Notes in Computer Science*, Vol. 9635. Springer, Berlin, 68–88. https://doi.org/10.1007/978-3-662-49635-0_4
- [8] Keith D. Cooper, Timothy J. Harvey, and Ken Kennedy. 2001. *A Simple, Fast Dominance Algorithm*. Technical Report. Rice University.
- [9] Oracle Corporation. 2020. *Code Tools: jmh*. Retrieved from <https://github.com/AIDanial/cloc>.

- [10] Ron Cytron, Jeanne Ferrante, Barry K. Rosen, Mark N. Wegman, and F. Kenneth Zadeck. 1991. Efficiently computing static single assignment form and the control dependence graph. *ACM Trans. Program. Lang. Syst.* 13, 4 (Oct. 1991), 451–490. <https://doi.org/10.1145/115372.115320>
- [11] Jeanne Ferrante, Karl J. Ottenstein, and Joe D. Warren. 1987. The program dependence graph and its use in optimization. *ACM Trans. Program. Lang. Syst.* 9, 3 (Jul. 1987), 319–349. <https://doi.org/10.1145/24039.24041>
- [12] Dennis Giffhorn and Gregor Snelting. 2015. A new algorithm for low-deterministic security. *Int. J. Inf. Secur.* 14, 3 (Apr. 2015), 263–287.
- [13] Jürgen Graf, Martin Hecker, Martin Mohr, and Gregor Snelting. 2016. Tool demonstration: JOANA. In *Principles of Security and Trust*, Lecture Notes in Computer Science, Vol. 9635. Springer Berlin Heidelberg, 89–93.
- [14] Christian Hammer and Gregor Snelting. 2009. Flow-sensitive, context-sensitive, and object-sensitive information flow control based on program dependence graphs. *Int. J. Inf. Secur.* 8, 6 (01 Dec. 2009), 399–422. <https://doi.org/10.1007/s10207-009-0086-1>
- [15] Matthew S. Hecht and Jeffrey D. Ullman. 1973. Analysis of a simple algorithm for global data flow problems. In *Proceedings of the 1st Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL'73)*. ACM, New York, NY, 207–217. <https://doi.org/10.1145/512927.512946>
- [16] Martin Hecker. 2020. *Timing Sensitive Dependency Analysis and its Application to Software Security*. Ph.D. Dissertation. Karlsruher Institut für Technologie, Fakultät für Informatik.
- [17] Daniel Hedin and David Sands. 2005. Timing aware information flow security for a JavaCard-like bytecode. *Electron. Notes Theor. Comput. Sci.* 141, 1 (2005), 163–182.
- [18] Susan Horwitz, Jan Prins, and Thomas W. Reps. 1988. On the Adequacy of Program Dependence Graphs for Representing Programs. In *Proceedings of the Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL'88)*. 146–157.
- [19] Susan Horwitz, Thomas Reps, and David Binkley. 1990. Interprocedural slicing using dependence graphs. *ACM Trans. Program. Lang. Syst.* 12, 1 (Jan. 1990), 26–60. <https://doi.org/10.1145/77606.77608>
- [20] Neil Jones, Carsten Gomard, and Peter Sestoft. 1993. *Partial Evaluation and Automatic Program Generation*. Prentice Hall.
- [21] Vineeth Kashyap, Ben Wiedermann, and Ben Hardekopf. 2011. Timing- and termination-sensitive secure information flow: Exploring a new approach. In *Proceedings of the 32nd IEEE Symposium on Security and Privacy*. 413–428.
- [22] Paul Kocher, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. 2018. Spectre attacks: Exploiting speculative execution. arXiv:1801.01203. Retrieved from <http://arxiv.org/abs/1801.01203>.
- [23] Jens Krinke. 2003. Context-sensitive slicing of concurrent programs. In *Proceedings of the Symposium on the Foundations of Software Engineering (FSE'03)*. ACM, 178–187.
- [24] Thomas Lengauer and Robert Endre Tarjan. 1979. A fast algorithm for finding dominators in a flowgraph. *ACM Trans. Program. Lang. Syst.* 1, 1 (Jan. 1979), 121–141.
- [25] Steven Muchnik. 1997. *Advanced Compiler Design and Implementation*. Morgan Kaufmann.
- [26] Barak Naveh and Stephane Popinet. 2003–2019. JGraphT: A Java Library of Graph Theory Data Structures and Algorithms. Retrieved from <https://jgraph.org/>.
- [27] A. Podgurski and L. A. Clarke. 1990. A formal model of program dependences and its implications for software testing, debugging, and maintenance. *IEEE Trans. Softw. Eng.* 16, 9 (Sep. 1990), 965–979.
- [28] Willard Rafnsson, Limin Jia, and Lujo Bauer. 2017. Timing-sensitive noninterference through composition. In *Principles of Security and Trust*, Lecture Notes in Computer Science, Matteo Maffei and Mark Ryan (Eds.), Vol. 10204. Springer, 3–25.
- [29] Venkatesh Prasad Ranganath, Torben Amtoft, Anindya Banerjee, John Hatcliff, and Matthew B. Dwyer. 2007. A new foundation for control dependence and slicing for modern program structures. *ACM Trans. Program. Lang. Syst.* 29, 5, Article 27 (Aug. 2007), 27–es. <https://doi.org/10.1145/1275497.1275502>
- [30] Thomas Reps, Susan Horwitz, Mooly Sagiv, and Genevieve Rosay. 1994. Speeding up slicing. In *Proceedings of the Symposium on the Foundations of Software Engineering (FSE'94)*. ACM, New York, NY, 11–20.
- [31] Andrei Sabelfeld and David Sands. 2000. Probabilistic Noninterference for Multi-Threaded Programs. In *Proceedings of the 13th IEEE Computer Security Foundations Workshop (CSFW'00)*. 200–214.
- [32] Robert Tarjan. 1972. Depth-First Search and Linear Graph Algorithms. *SIAM J. Comput.* 1, 2 (1972), 146–160.
- [33] Michael Joseph Wolfe. 1995. *High Performance Compilers for Parallel Computing*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA.

Received July 2020; revised August 2021; accepted August 2021