

**Timing Sensitive Dependency Analysis and its Application to Software
Security**

Zur Erlangung des akademischen Grades eines

DOKTORS DER INGENIEURWISSENSCHAFTEN

von der KIT-Fakultät für Informatik des
Karlsruher Instituts für Technologie (KIT)
genehmigte

DISSERTATION

von

Martin Hecker
aus Hannover

Tag der mündlichen Prüfung: 20.07.2020

1. Referent: Prof. Dr.-Ing. Gregor Snelting

2. Referent: Prof. Dr. Markus Müller-Olm

Zusammenfassung

Ich präsentiere neue Verfahren zur statischen Analyse von Ausführungszeit-sensitiver Informationsflusskontrolle in Softwaresystemen. Ich wende diese Verfahren an zur Analyse nebenläufiger Java Programme, sowie zur Analyse von Ausführungszeit-Seitenkanälen in Implementierungen kryptographischer Primitive.

Methoden der Informationsflusskontrolle zielen darauf ab, Fluss von Informationen (z.B.: zwischen verschiedenen externen Schnittstellen einer Software-Komponente) anhand expliziter Richtlinien einzuschränken. Solche Methoden können daher zur Einhaltung sowohl von Vertraulichkeit als auch Integrität eingesetzt werden. Der Ziel *korrekter* statischer Programmanalysen in diesem Umfeld ist der Nachweis, dass in allen Ausführungen eines gegebenen Programms die zugehörigen Richtlinien eingehalten werden. Ein solcher Nachweis erfordert ein Sicherheitskriterium, welches formalisiert, unter welchen Bedingungen dies der Fall ist.

Jedem formalen Sicherheitskriterium entspricht implizit ein Programm- und Angreifermodell. Einfachste *Nichtinterferenz*-Kriterien beschreiben beispielsweise nur nicht-interaktive Programme. Dies sind Programme die nur bei Beginn und Ende der Ausführung Ein- und Ausgaben erlauben. Im zugehörigen Angreifer-Modell kennt der Angreifer das Programm, aber beobachtet nur bestimmte (*öffentliche*) Aus- und Eingaben oder stellt diese bereit. Ein Programm ist nicht-interferent, wenn der Angreifer aus seinen Beobachtungen keinerlei Rückschlüsse auf *geheime* Aus- und Eingaben terminierender Ausführungen machen kann. Aus nicht-terminierenden Ausführungen hingegen sind dem Angreifer in diesem Modell Schlussfolgerungen auf geheime Eingaben erlaubt.

Seitenkanäle entstehen, wenn einem Angreifer aus Beobachtungen *realer* Systeme Rückschlüsse auf vertrauliche Informationen ziehen kann, welche im formalen Modell unmöglich sind. Typische Seitenkanäle (also: in vielen formalen Sicherheitskriterien unmodelliert) sind neben Nichttermination beispielsweise auch Energieverbrauch und die Ausführungszeit von Programmen. Hängt diese von geheimen Eingaben ab, so kann ein Angreifer aus der beobachteten Ausführungszeit auf die Eingabe (z.B.: auf den Wert einzelner geheimer Parameter) schließen.

In meiner Dissertation präsentiere ich neue Abhängigkeitsanalysen, die auch Nichtterminations- und Ausführungszeitkanäle berücksichtigen. In Hinblick auf Nichtterminationskanäle stelle ich neue Verfahren zur Berechnung von Programm-Abhängigkeiten vor. Hierzu entwickle ich ein vereinheitlichendes Rahmenwerk, in welchem sowohl Nichttermination-sensitive als auch Nichttermination-insensitive Abhängigkeiten aus zueinander dualen *Postdominanz*-Begriffen resultieren. Für Ausführungszeitkanäle entwickle ich neue Abhängigkeitsbegriffe und dazugehörige Verfahren zu deren Berechnung. In zwei Anwendungen untermauere ich die These:

Ausführungszeit-sensitive Abhängigkeiten ermöglichen korrekte statische Informationsfluss-Analyse unter Berücksichtigung von Ausführungszeitkanälen.

Basierend auf Ausführungszeit-sensitiven Abhängigkeiten entwerfe ich hierfür neue Analysen für nebenläufige Programme. Ausführungszeit-sensitive Abhängigkeiten sind dort selbst für Ausführungszeit-insensitive Angreifermodelle relevant, da dort *interne* Ausführungszeitkanäle zwischen unterschiedlichen Ausführungsfäden extern beobachtbar sein können. Meine Implementierung für nebenläufige Java Programme basiert auf dem Programmanalyse-System JOANA.

Außerdem präsentiere ich neue Analysen für Ausführungszeitkanäle aufgrund *mikro-architektureller* Abhängigkeiten. Exemplarisch untersuche ich Implementierungen von AES256 Blockverschlüsselung. Bei

einigen Implementierungen führen Daten-Caches dazu, dass die Ausführungszeit abhängt von Schlüssel und Geheimtext, wodurch diese aus der Ausführungszeit inferierbar sind. Für andere Implementierungen weist meine automatische statische Analyse (unter Annahme einer einfachen konkreten Cache-Mikroarchitektur) die Abwesenheit solcher Kanäle nach.

Abstract

I present new methods for the static analysis of timing sensitive information flow control in software systems. I apply these methods in the analysis of concurrent Java programs, as well as the analysis of timing side-channels in implementations of cryptographic primitives.

Methods for information flow control aim to control the flow of information (e.g.: between different external interfaces of a software component) with respect to explicit flow policies. Such methods can protect confidentiality as well as integrity of information. In this setting, the goal of *sound* static program analysis is to prove that in all executions of a given program, a corresponding flow policy is respected. Such a proof requires a security criterion which formalizes under which condition this is indeed the case.

Every formal security criterion implicitly corresponds to a program and attacker model. Simple *non-interference* criteria, for example, apply only to non-interactive programs. These are programs which allow input and output only at the beginning and the end of the execution. In the corresponding attacker model, the attacker does know the program, but observes only certain (*public*) output and input, or provides those. A program is non-interferent if from these observations the attacker cannot infer any properties of *secret* input and output of terminating executions. From non-terminating executions in this model, the attacker *is* allowed to infer properties of secret input.

Side channels occur if an attacker can infer properties of secret information from observations of *real* systems which are impossible to infer in the formal model. Typical side channels include (i.e., typically unmodeled are): nontermination as well as energy consumption and

timing channels: By observing the program's execution time, the attacker (partially or completely) infers the program's secret input on which the execution time depends.

In my dissertation I present new dependency analysis sensitive to both nontermination and timing channels. With respect to nontermination, I introduce new methods for the computation of control dependencies. For this purpose, I develop a generalized framework in which nontermination-sensitive as well as nontermination-insensitive control dependencies result from (mutually dual) *postdominance* notions. For timing channels, I develop new notions of control dependencies, and corresponding methods for their computation. In two applications, I substantiate my thesis:

Timing sensitive control dependencies facilitate sound static information flow control in the presence of timing channels.

For this purpose, I develop new analysis for concurrent programs, based on timing sensitive control dependencies. In this setting, timing sensitive control dependencies are relevant even for timing insensitive attacker models, since there *internal* timing channels between concurrent threads of execution may become observable externally. My implementation for concurrent Java programs is based on the JOANA program analysis framework.

I also present new analysis for timing channels caused by *micro-architectural* dependencies. I illustrate these by a study of AES256 block cipher implementations. For some implementations, *data caches* cause the execution time to depend on key and plaintext, making them inferrable from the execution time. For other implementations, my automatic static analysis proves (assuming a simple concrete cache micro-architecture) the absence of such timing channels.

Contents

| | |
|---|-------------|
| Zusammenfassung | i |
| Abstract | v |
| Symbols | xiii |
| 1 Introduction | 1 |
| 1.1 Contributions | 4 |
| 2 Methodology | 7 |
| I Dependency Analysis in Arbitrary Graphs | 11 |
| 3 Control Dependence in Arbitrary Graphs | 13 |
| 3.1 Generalized Control Dependence | 15 |
| 3.2 An Algorithm for Generalized Control Dependence . . | 18 |
| 3.3 Related Work | 29 |
| 4 Nontermination (In-)Sensitive Control Dependence | 31 |
| 4.1 Nontermination (In-)Sensitive Control Dependence in Arbitrary Graphs | 33 |
| 5 Postdominator Pseudoforests | 39 |
| 5.1 Fixed-Point Characterizations for Postdominance . . . | 41 |
| 5.2 Nontermination Sensitive Pseudoforests | 44 |
| 5.3 Nontermination Insensitive Pseudoforests | 51 |
| 5.4 Reduction to Postdominance Trees | 59 |

| | | |
|-----------|---|------------|
| 6 | Order Dependence | 63 |
| 6.1 | Decisive Order Dependence | 66 |
| 6.2 | Nontermination Sensitive Order Dependence | 72 |
| 6.3 | Soundness of Nontermination Sensitive Slices | 78 |
| 6.4 | Weak Order Dependence | 85 |
| 6.5 | Soundness of Nontermination Insensitive Slices | 89 |
| 6.6 | A Trace-Based Notion of Infinite Delay | 92 |
| 6.7 | Nontermination Insensitive Order Dependence | 97 |
| 6.8 | Soundness of Nontermination Insensitive Order Dependence | 113 |
| 7 | Slicing | 117 |
| 7.1 | Nontermination Insensitive Slicing | 121 |
| 7.2 | Nontermination Sensitive Slicing | 136 |
| 7.3 | Weak Order Control Slices | 138 |
| 7.4 | Weak Control Closures | 139 |
| 7.5 | The Role of \rightarrow_{nticd} for Nontermination Insensitive Slices | 141 |
| 8 | Performance Benchmarks | 143 |
| 8.1 | Nontermination Sensitive Postdominance | 146 |
| 8.2 | Nontermination Insensitive Postdominance | 148 |
| 8.3 | Generalized Postdominance Frontiers | 150 |
| 8.4 | Control Slices | 152 |
| 8.5 | Nontermination Insensitive Order Dependence | 154 |
| 8.6 | Nontermination Insensitive Slices | 156 |
| 8.7 | Timing Sensitive Algorithms | 158 |
| II | Timing Sensitive Dependency Analysis | 161 |
| 9 | Timing Sensitive Control Dependence | 163 |
| 9.1 | Timing Sensitive Control Dependence | 167 |
| 9.2 | Timing Sensitive Post Postdominance Frontiers | 172 |
| 9.3 | Transitive Timing Sensitive Postdominance | 179 |
| 9.4 | Algorithms for Timing Sensitive Control Dependence | 188 |

| | | |
|------------|---|------------|
| 9.5 | Soundness and Minimality of Timing Sensitive Control Dependence | 197 |
| 9.6 | Timing Sensitive Control Dependence in Graphs with Unique Exit Node | 198 |
| 10 | Timing Dependence | 201 |
| 10.1 | Timing Dependence | 204 |
| 10.2 | Computation of Timing Dependence | 206 |
| 11 | Timing Stratification | 207 |
| 11.1 | Timing Sensitive Control Dependence for Arbitrary CFG with Cost Model | 210 |
| 11.2 | Timing-Stratification | 213 |
| 11.3 | An Algorithm for Timing Stratification | 217 |
| III | Timing Sensitive Software Security | 221 |
| 12 | Transforming Out Timing Leaks in Arbitrary CFG | 223 |
| 12.1 | An Naive Algorithm | 227 |
| 12.2 | A More Precise Algorithm | 229 |
| 13 | Micro-Architectural Dependencies | 233 |
| 13.1 | Introduction | 235 |
| 13.2 | Control Flow Graphs | 238 |
| 13.3 | Micro-Architectural Dependencies | 244 |
| 13.4 | Limitations of Micro-Architectural Dependencies | 256 |
| 13.5 | Timing Dependence for Micro-Architectural Dependencies | 258 |
| 13.6 | Arrays | 263 |
| 14 | Cache Timing Attacks on AES256 | 267 |
| 14.1 | AES256 Encryption | 268 |
| 14.2 | Pre-Caching | 272 |
| 14.3 | Constant Time S-Box Substitution | 274 |
| 14.4 | Validation | 276 |

| | | |
|-----------|--|------------|
| 15 | Approximate Cache Dependencies | 281 |
| 15.1 | Data Dependence | 282 |
| 15.2 | Local Cache-Cache Dependencies | 283 |
| 15.3 | Local State-Cache Dependence | 288 |
| 15.4 | Transitive Cache Dependencies | 291 |
| 15.5 | Approximate Cache Dependencies | 293 |
| 15.6 | Improving the Precision | 295 |
| 15.7 | Approximation in AES256 Implementations | 296 |
| 15.8 | Related Work | 300 |
| 16 | Timing Sensitivity in Concurrent Programs | 303 |
| 16.1 | Probabilistic Noninterference | 305 |
| 16.2 | Observability of Internal Timing Leaks | 306 |
| 16.3 | The RLSOD Criterion | 308 |
| 16.4 | A Statistical Test for Probabilistic Noninterference | 317 |
| 16.5 | Imprecision of the RLSOD criterion | 325 |
| 16.6 | Timing Sensitivity for Probabilistic Noninterference | 331 |
| 17 | Timing Sensitivity with JOANA | 337 |
| 17.1 | Precision of the TIMING criterion for Java | 339 |
| 17.2 | Scalability of the TIMING criterion for Java | 341 |
| 18 | Summary and Future Work | 343 |
| A | Proofs | 349 |
| A.1 | Nontermination (In-)Sensitive Control Dependence in Arbitrary Graphs | 350 |
| A.2 | Postdominator Pseudoforests | 354 |
| A.3 | Order Dependence | 358 |
| B | Nontermination (In-)Sensitive Control Dependence | 361 |
| B.1 | Analysis of previous Algorithms | 362 |
| B.2 | Duality of Nontermination (In-)Sensitivity | 373 |
| B.3 | New Algorithms | 379 |
| C | A Slicing Algorithm using C-Edges | 385 |

| | | |
|----------|---|------------|
| D | Algorithm Variants | 395 |
| | D.1 Another Algorithm for \rightarrow_{nticd} | 400 |
| | D.2 Efficient $lca_{<}$ via Postorder Numbers | 402 |
| E | Generalizations for CFG with Timing Cost Model | 409 |

Symbols

What's all this? Looks like Darth Vader's bathroom.

(Michal Knight — Knight Rider, Knight of the Phoenix)

New contributions of this thesis are **highlighted**.

Part I

| nontermination sensitive | | p. | p. | nontermination insensitive | | p. |
|--|--|-----|-----|---|--|-----|
| postdominance | \sqsupset_{MAX} | 36 | 36 | \sqsupset_{SINK} | postdominance | |
| ..pseudo-forest | $<_{\text{MAX}}$ | | | $<_{\text{SINK}}$ | ..pseudo-forest | |
| control dependence | $\rightarrow_{\text{ntscd}}$ | 34 | 34 | $\rightarrow_{\text{nticd}}$ | control dependence | |
| order dependence | $\rightarrow_{\text{ntsod}}$ | 72 | 99 | $\rightarrow_{\text{ntiod}}$ | order dependence | |
| "nearest dominator" (generalized) | $\rightarrow_{\text{ntind}}$ | 122 | 136 | $\rightarrow_{\text{ntsnd}}$ | "nearest dominator" (generalized) | |
| | | | | | | p. |
| | standard postdominance | | | \sqsupset_{POST} | | 17 |
| | standard control dependence | | | \rightarrow_{cd} | | 17 |
| | decisive order dependence | | | \rightarrow_{dod} | | 66 |
| | weak order dependence | | | \rightarrow_{wod} | | 85 |
| | generalized 1-postdominance | | | $1\text{-}\sqsupset$ | | 19 |
| | generalized immediate postdominators | | | $\text{ipdom}_{\sqsupset}(n)$ | | 19 |
| | generalized postdominance frontier | | | $\text{PDF}_{\sqsupset}(n)$ | | 21 |
| | least common ancestor | | | $\text{lca}_{<}((, n), m)$ | | 46 |
| | backward slice | | | $(\rightarrow a \cup \rightarrow b)^*$ | | 74 |
| | generic graph transformations | | | $G_{M \not\rightarrow}, G \rightarrow^* M, \dots$ | | 75 |
| | input equivalence (unlabeled CFG) | | | $i \sim_S i'$ | | 80 |
| | trace equivalence (unlabeled CFG) | | | $t \sim_S t'$ | | 80 |
| | next-observable | | | $\text{obs}_S(n)$ | | 89 |
| | observation equivalence (unlabeled CFG) | | | $i \sim_{\mathcal{T}_M}^{\omega} i'$ | | 94 |
| | weakly deciding nodes | | | $\text{WD}_G(M)$ | | 139 |
| | weak control closure | | | $\text{WCC}_G(M)$ | | 139 |

Part II

| | | p. |
|--|---|-----|
| timing sensitive postdominance | $\sqsupset_{\text{TIME}[\text{FIRST}]}$ | 167 |
| timing sensitive postdominance (transitive) | \sqsupset_{TIME} | 179 |
| ..pseudo-forest | $<_{\text{TIME}}$ | 179 |
| timing sensitive control dependence | $\rightarrow_{\text{tscd}}$ | 168 |
| “fuel available” | F_n | 184 |
| timing dependence | \rightarrow_{td} | 205 |

Part III

| | | p. |
|---|-------------------------------|-----|
| timing dependence (in concurrent CFG) | $\rightarrow_{\text{timing}}$ | 331 |
| macro-architectural state | σ_M | 238 |
| micro-architectural state | σ_μ | 239 |
| micro-architectural dependence | $\rightarrow_{\mu d}$ | 254 |
| approximate cache micro-architectural dependence | $\rightarrow_{\mu d}^\#$ | 293 |

Notation

For any binary relation R , I write R^+ for its transitive closure, and R^* for its transitive reflexive closure. Whenever I denote a binary relation by a symbol similar to \sqsubseteq , I write \sqsupset for its inverse, and \sqsubset for the relation $\sqsubseteq \setminus \{(x, x)\}$. Similarly for the symbol \leq .

```

static int runs = 2000000;
static void main() {
    h = input();
    for (int i=31; i>0; i--) {
        int b = h & (1 << 30);
        A a = new A();
        a.start();
        if (b != 0) {
            delay(runs);
        }
        x = 1;
        h = h << 1;
        a.join();
        print(x);
    }
}

class A extends Thread {
    public void run() {
        delay(runs/2);
        x = 0;
    }
}

```

```

static int delay(int t) {
    int n = 1;
    for (int k=1; k < 100; k++) {
        for (int j=1; j < t; j++) {
            n = n * k;
        }
    }
    return n;
}

```

(a) Internal Timing Channel

```

...
for (int i=0; i<16; i++) {
    state[i] = sbox[state[i]]
}
...

```

(b) External Timing Channel

1 Introduction

Consider the first program code example above. It reliably prints¹ the input value obtained from `h = input()`. It is propagated due to the relative *execution time* of the main thread, and the thread A started in the main loop. The execution time of each iteration of the main loop *depends* on the input, and the program exhibits an *internal* timing channel. The second example is part of an implementation of a AES256 block cipher. If run on modern CPUs with *data caches*, and given enough observations of the *execution time* of encrypt operations, an attacker can infer the AES256 key (e.g., [Ber05; BM06]). The implementation exhibits an *external* timing channel.

¹ for example, using OpenJDK11 JVM, on a Intel Xeon Gold 6230 CPU

Timing channel attacks on software systems exploit that the execution time of some program part depends on *secret* input. They occur in numerous attack scenarios. For example:

- The attacker passively observes the execution time of a user-provided program on the users computer. The attacker is either a legitimate client of the program, or monitors the programs network traffic. The attacker may be able to provoke multiple computation with the same secret input, which allows him to take multiple samples of the computation time. This scenario includes several attacks on cryptographic protocols and implementations (e.g., [BB05; Ber05]), such as that shown in the second example.
- The attacker provides a program to the user, who tries to automatically verify that the program leaks no secret information back to the attacker. The attacker may have included code similar to the first example. Alternatively, the user wants to verify that a program provided by himself does not by leak secret information. The verification can be done either by the user himself, or a trusted third party. This scenario includes, for example, every World Wide Web site that includes JavaScript or WebAssembly code, or *app stores* for mobile devices (e.g., [Lor+14]).
- The attackers and the users program run on shared hardware. The attacker is able to observe the users computation through side channels arising from the use of shared resources such as data caches. This scenario includes the “Meltdown” and “Spectre” attacks ([Lip+18; Koc+19]).

In the first two scenarios, timing attacks can in principle be detected by analysis of a single program, while in the third scenario, the interaction of two programs on the shared hardware needs to be taken into consideration. There, the attacker measures the execution time of *his* program.

Information flow control aims to detect and prevent illegitimate flow of information, as defined by a formal security property. Numerous different security properties exist, for example: standard non-interference (for *batch* execution of sequential programs), strategy non-interference (for interactive programs), possibilistic and probabilistic non-interference (for concurrent programs).

Information flow control methods can either be purely static or dynamic. Depending on the program model (abstract syntax tree or control flow graph), static methods are typically based on security type-systems, or program dependency graphs [FOW87; HS09]. Program dependency graphs include some form of *data dependencies* and *control dependencies*. Data dependencies capture the dependencies of the data-state on previous data-state, while control dependencies capture the dependence of the control state on (choices made at) previous control state. Hence data and control dependencies account for *explicit* and *implicit* channels, but not for *timing* channels.

All useful information flow control security properties are proper *hyperproperties*[CS10], which means that their violation cannot be detected purely dynamically (i.e.: by monitoring of one execution of the problem). Therefore, even dynamic information flow control must include *some* form of static analysis. Control dependence can be applied here, as well [Jus+11; XZ07].

In this thesis, I propose a new notion of timing sensitive control dependence. I demonstrate by that this new notion can be used for static timing sensitive information flow control in the first two attack scenarios, and programs like those in the two examples. For the third attack scenario, I suspect that timing sensitive control dependence could in principle be relevant as well, but make no attempt to demonstrate this.

1.1 Contributions

My central contributions (in Part II of this thesis) is the new notion of *timing sensitive control dependence* (Chapter 9), and the related notion of *timing dependence* (Chapter 10). Just as control and data dependencies represent dependencies for data and control state, timing dependencies represent dependencies in the timing state of a configuration, i.e.: the part of an configuration that models how much time has passed during the programs execution. In Chapter 11, I introduce the notion of timing cost model that is implicit in Chapter 9 and Chapter 10. I also introduce the technical notion of *timing stratification*, which sheds light on the relation between timing sensitive control dependence and *nontermination sensitive control dependence*.

In Part III of this thesis I demonstrate that timing dependence and timing sensitive control dependence enable sound static information flow analysis for *internal* timing channels (for concurrent programs, in Chapter 16) and (external) timing channels (due to the micro-architecture of modern CPUs, in Chapter 13). The two central contributions of this part are a new notion of *micro-architectural* dependencies (Section 13.6), and a new timing dependence based criterion for probabilistic noninterference (Section 16.6). I apply *micro-architectural* dependencies in a case study on cache-based timing attacks in implementations of the AES256 block cipher (Chapter 14). At the expense of some precision, I develop an efficient approximation to *micro-architectural* dependencies for cache micro-architectures in Chapter 15. I also explain in Chapter 12 how under simple timing cost models, timing leaks can sometimes be automatically “transformed out”.

I derive algorithms for the computation of timing sensitive control dependence and timing dependence from algorithms for *nontermination sensitive control dependence* in arbitrary graphs. My algorithms for micro-architectural dependence on the other hand rely on algorithms for nontermination *insensitive* control dependence and slices. In that light, the central contributions of Part I of thesis then simply are new algorithms (Chapter 3, Chapter 5) for nontermination sensi-

tive and insensitive control dependence (Chapter 4), as well as new algorithms for corresponding slices (Chapter 7). In fact if one is interested only in timing sensitive information flow control, one may want to skip Part I on first reading of this thesis, and trust that indeed: Algorithms for nontermination sensitive and insensitive control dependence in *arbitrary* graphs can be obtained by suitable generalizations of algorithms for standard control dependence in graphs with unique exit node (Section 3.1). But I also want to highlight the following other contributions from Part I, which I find to be of interest independent from their application to timing sensitive analysis: 1. A new fixed-point characterizations of nontermination sensitive and insensitive postdominance (Section 5.1). 2. Two new notions for nontermination sensitive and insensitive *order* dependence (Chapter 6). 3. Two new notions of soundness for nontermination insensitive slicing (Section 6.5, Section 6.6). 4. Reductions of several different notions of nontermination insensitive slicing to slices from nontermination insensitive control dependence (Section 7.5).

2 Methodology

The first principle is that you must not fool yourself — and you are the easiest person to fool. So you have to be very careful about that. And after you've not fooled yourself, it's easy not to fool other scientists. You just have to be honest in a conventional way after that.

(Richard P. Feynman — Cargo Cult Science)

In this thesis, I claim numerous new results (e.g., soundness and sometimes minimality) for various forms of program dependencies. The conventional method to support such claims is by mathematical proof, which can range from offhand sketches to fully formalize and mechanically checked proof scripts for systems like Isabelle/HOL ([NWP02]), Coq ([Tea17]) or Lean ([Mou+15]). In this thesis, I provide semi-formal proofs for some, but by no means for all new results. For most other results, I gathered evidence by performing extensive *random testing* of the claimed result. Specifically, for such results I

- implemented all involved predicates and algorithms in the pure functional programming language Haskell,
- used a randomized test data generator to generate *inputs*, and
- verified the claimed result for a large number of random inputs.

For example, Observation 5.3.7 involves one algorithm and a relation \exists_{SINK} , and reads:

Let G be any CFG. Then Algorithm 6 terminates with a result \langle_{SINK} s.t. \rangle_{SINK} is a transitive reduction of $\sqsubseteq_{\text{SINK}}$.

In order to randomly test this observation, I use a rather inefficient “reference” implementation of the relation $\sqsubseteq_{\text{SINK}}$ to check the more efficient to Algorithm 6. The “input” graphs G are obtained from the randomized graph generator for the Haskell Functional Graph Library ([EM17; Mil17; Erw01]). I then test the observation on random graphs of size up to 100 nodes, by running the random tests for several days. Here, one million of such graphs can be tested in approximately one hour on a standard desktop PC.

In this example I cannot use a naive “reference” implementation of the relation $\sqsubseteq_{\text{SINK}}$. In fact, a direct implementation of its definition would require me to check infinitely many paths in input graphs G . Instead I use a fixed point characterization of $\sqsubseteq_{\text{SINK}}$ (Theorem 5.1.2) which *can* be naively implemented in Haskell. While this reference implementation is effective, it is not very efficient for larger graphs (e.g.: larger than 100 nodes). For example, it would take me at least 16 hours to check this observation for one million random graphs of size up to 500 nodes.

For every **Observation** in this thesis, I provide random test properties in the Haskell QuickCheck framework ([CH02]). For example, the property for Observation 5.3.7 reads¹

```
observation_5_3_7 = testProperty "isinkdomOf^* == sinkdomOfGfp" $
  \ (ARBITRARY(g)) ->
    toSuccMap $ trc $ (fromSuccMap $ PDOM.isinkdomOf g)
    == PDOM.sinkdomOfGfp g
```

All properties are available in a ready-to-run virtual machine image[Hec20]. They can be found in the Haskell module `Program.Properties.DissObservations`.

¹ up to renaming of variables, and type annotations

Ideally, of course, every observation would be supported not only by evidence from randomized tests, but by formal proofs. In fact for some key results of Part I, proofs in the Isabelle/HOL proof assistant are available due to ongoing work by Simon Bischof[Bis19].

Part I

Dependency Analysis in Arbitrary Graphs

Choice. The problem is: choice.

(Neo, *The Matrix Reloaded* (2003),
The Wachowskis)

3 Control Dependence in Arbitrary Graphs

I am never forget the day I first meet the great Lobachevsky. In *one* word he told me secret of success in mathematics: Plagiarize!

Let no one else's work evade your eyes
Remember why the good Lord made your eyes
So don't shade your eyes
But plagiarize, plagiarize, plagiarize

Only be sure always to call it please "research"

(Tom Lehrer — Lobachevsky)

My overall strategy to obtain definitions and algorithms for time sensitive control dependence is as follows:

1. Generalize existing algorithms for standard (nontermination insensitive) control dependence such that they become applicable also to nontermination sensitive control dependence.
2. Define timing sensitive control dependence as modification of nontermination sensitive control dependence. Then modify the generalized algorithms for nontermination sensitive control dependence such that they become applicable also to timing sensitive control dependence.

In principle, I could have tried to follow this strategy under the customary assumption that (control flow) graphs always come with a *unique* exit node. Instead I chose to drop this assumption, and develop generalizations that apply to *arbitrary* graphs. I do this not only out of theoretical curiosity, since this investment pays off at the latest in Chapter 13, in which I will utilize control dependence for certain graphs (*derived* from control flow graphs) that do *not* have unique exit nodes, even if the control flow graphs they derive from do have a unique exit node.

In Section 3.1 of this chapter, I define the framework of generalized control dependence for arbitrary graphs, of which standard control dependence is an instantiation. Then in Section 3.2, I develop a new algorithm for the computation of generalized control dependence. I do so by generalizing the standard approach from [Cyt+91].

3.1 Generalized Control Dependence

For the purpose of the first two parts of this thesis, a *control flow graph* is any directed, unlabeled graph. Standard definitions and algorithm for control dependence require the existence of *unique* exit-nodes, but the new generalized definitions introduced in this chapter do not.

Definition 3.1.1 (Control Flow Graph). A pair $G = (N, E)$ of a finite set N of nodes, and a set $E \subseteq N \times N$ of directed edges is called a *control flow graph*. I write

- $n \rightarrow_G m$ whenever $(n, m) \in E$, and sometimes omit the subscript G when it is clear from the context: $n \rightarrow m$.

A node $n_e \in N$ is called the *unique entry node* of G whenever it is the only in N such that

- it is an entry node, i.e.: n_e has no incoming edge in G , and
- every node is reachable from n_e , i.e.: $\forall m \in N. n_e \rightarrow^* m$

Similarly, a node $n_x \in N$ is called the *unique exit node* of G whenever it is the only in N such that

- it is an exit node, i.e.: n_x has no outgoing edge in G , and
- n_x is reachable from every node, i.e.: $\forall m \in N. m \rightarrow^* n_x$

A path $\pi = n_1, \dots, n_k, \dots$ is a (possibly empty or infinite) sequence of nodes $n_i \in N$ such that $(n_i, n_{i+1}) \in E$ for all $n_i, n_{i+1} \in \pi$. Given a finite path $\pi_1 = n_1, \dots, n_k$ and a path $\pi_2 = m_1, \dots$, I write π_1, π_2 for

- the path $n_1, \dots, n_k, m_1, \dots$ if neither π_1 nor π_2 is empty, and $(n_k, m_1) \in E$
- the path π_2 if π_1 is empty
- the path π_1 if π_2 is empty

A finite path $\pi = n_0, \dots, n_k$ is said to have length k . A finite path π of length k is said to be between n and m if $n = n_0$ and $m = n_k$. In this case, I write $n \xrightarrow{\pi^k} m$, and sometimes omit k . With ${}_n\Pi_m$ I denote the set of all (finite) paths between n and m .

Intuitively, a node n (immediately) controls whether a node m is executed, if a (control flow) *choice* made at n — i.e.: the choice at which successor n_l of some choice node n to continue — can force control flow to “eventually” reach m , while a different choice (to continue at some other successor n_r of n) does *not* force control flow to eventually reach m .

This intuitive characterization is ambiguous: what exactly does it mean for control flow at n_l to be forced to “eventually” reach m ? Different realization of this concept will lead to different control dependence notion. Formally, control dependence is parameterized by a binary relation \sqsubseteq on nodes that specifies a notion of *eventuality*.

Definition 3.1.2 (\sqsubseteq -Control Dependence). Let $G = (N, E)$ be a CFG, and $\sqsubseteq \subseteq N \times N$ any binary relation on nodes¹ (where $m \sqsubseteq n$ is meant to be read: any path from n “eventually” reaches m , or: m “postdominates” n). Then m is said to be \sqsubseteq -control dependent on n in G iff there exist nodes n_l and n_r such that

- $n \rightarrow n_l$ and $n \rightarrow n_r$,
- $m \sqsubseteq n_l$, and
- $\neg m \sqsubseteq n_r$

One contribution of this thesis is a new *timing sensitive* instantiation of \sqsubseteq , which I will introduce in Section 9.1. For now, I will review the instantiation of \sqsubseteq that leads to *standard* control dependence.

Given a graph G with unique exit node n_x , it is natural to say that *any* path “eventually” reaches n_x . For arbitrary nodes m , it is then equally

¹ more formally: a function that maps a CFG G to such a relation \sqsubseteq

natural to say that any path from n reaches m if m *postdominates* n (w.r.t the unique exit node n_x), i.e.: if m occurs on every (necessarily finite) path between n and n_x :

Definition 3.1.3 (Standard Postdominance). Let G be a CFG with unique exit node n_x . Then m (standard) postdominates n , and I write $m \sqsubseteq_{\text{POST}} n$, iff all paths in G from n to n_x contain m , i.e.: iff

$$\forall \pi \in {}_n\Pi_{n_x}. m \in \pi$$

Also, for *standard* control dependence (i.e.: $\sqsubseteq_{\text{POST-control}}$ dependence) I just write:

$$n \rightarrow_{\text{cd}} m$$

Standard control dependence $n \rightarrow_{\text{cd}} m$ then holds iff there exist successors n_l and n_r of n such that

- $m \sqsubseteq_{\text{POST}} n_l$ (i.e.: m standard postdominates n_l), and
- $\neg m \sqsubseteq_{\text{POST}} n_r$

Remark 3.1.1. This definition of standard control dependence $n \rightarrow_{\text{cd}} m$ appears to be due to [Wol95]. A more classical definition is found in [FOW87], and reads:

Definition 3.1.4 (Standard Control Dependence, Alternative Definition). m is standard control dependent on n , iff

1. $\neg m \sqsubseteq_{\text{POST}} n$
2. there exists a path $n \xrightarrow{n, \pi, m} m$ such that
 - $m \notin \pi$, and
 - $\forall n' \in \pi. m \sqsubseteq_{\text{POST}} n'$

A proof of equivalence of these two definition can be found in [Was10].

3.2 An Algorithm for Generalized Control Dependence

Efficient computation of standard control dependence for CFGs with unique exit node n_x is possible due to the following facts:

1. Every node $n \neq n_x$ has a unique *immediate* post-dominator $\text{ipdom}(n) \in N$, i.e. a node m such that $m \sqsupset_{\text{POST}} n$, and $m' \sqsupset_{\text{POST}} m$ for all nodes m' s.t. $m' \sqsupset_{\text{POST}} n$.
2. Hence, the mapping ipdom induces a tree rooted in n_x .
3. This *postdominator tree* can be computed efficiently (e.g., [LT79]).
4. Given the postdominator tree, the *postdominance frontier* $\text{PDF}(m) \subseteq N$ for nodes $m \in N$ can be computed efficiently ([Cyt+91]).

Here, $\text{PDF}(m)$ is the set of all nodes n such that

- m does not strictly post-dominate n
 - m post-dominates some direct successor n' of n ,
5. $n \rightarrow_{\text{cd}} m$ iff $n \in \text{PDF}(m)$ for $m \neq n$.

In order to understand item 5, observe that item 2 of Definition 3.1.4 is equivalent to m postdominating some direct G -successor n' of n ([Cyt+91], Lemma 11).

In this section, I develop suitable generalizations of the notions *post-dominance* and *postdominance frontier*, and present a new generalized algorithm for the *computation* of postdominance frontiers. This algorithm makes use of a generalization of postdominator trees. Later in Section 4.1, I will provide efficient algorithms for their computation in *arbitrary* CFGs (specifically: CFGs which lack a unique exit node n_x), and hence obtain an efficient algorithm for the computation of control dependence in arbitrary CFGs.

Generalized Immediate Postdominance

The algorithm from [Cyt+91] computes the postdominance frontiers for the standard postdominance relation $\sqsupseteq_{\text{POST}}$ ². In this subsection, I identify a set of sufficient properties of otherwise arbitrary postdominance relations \sqsupseteq that are enough to allow efficient postdominance frontiers algorithms for such arbitrary postdominance relations. Specifically, these postdominance relations will not need to be anti-symmetric (as $\sqsupseteq_{\text{POST}}$ is). In a first step, I provide a substitute for the notion of *strict* postdominance.

Definition 3.2.1 (Immediate \sqsupseteq -Postdominance). Given a relation $\sqsupseteq \subseteq N \times N$, a node $x \in N$ is said to 1- \sqsupseteq -postdominate z if there exists some node y such that

$$x \sqsupseteq y \sqsupseteq z$$

The set $\text{ipdom}_{\sqsupseteq}(n)$ of *immediate \sqsupseteq -postdominators* of n is defined by

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

Remark 3.2.1. 1- \sqsupseteq differs from strict postdominance \sqsupseteq whenever \sqsupseteq contains “cycles”. If, for example,

$$x \sqsupseteq y \sqsupseteq x$$

then x 1- $\sqsupseteq x$, but not: $x \sqsupseteq x$. If, on the other hand, \sqsupseteq is a partial order, then

$$m \sqsupseteq n \iff m \text{ 1-}\sqsupseteq n$$

Corollary 3.2.1. Immediate \sqsupseteq -postdominance generalizes immediate postdominance for CFGs with unique exit node n_x :

- $\text{ipdom}_{\sqsupseteq_{\text{POST}}}(n_x) = \emptyset$

² or the dominance frontiers for the fixed dominance relation \sqsupseteq_{DOM} , which is exactly the postdominance relation in the graph G^{-1} with inverted edges

- $\text{ipdom}_{\sqsupseteq_{\text{POST}}}(n) = \{\text{ipdom}(n)\}$ for $n \neq n_x$

Remark 3.2.2. Essentially, $1\text{-}\sqsupseteq$ -postdominance behaves like standard strict postdominance, but allows sets of “equivalent” (w.r.t \sqsupseteq) nodes, i.e. nodes m, m' such that $m \sqsupseteq m'$ and $m' \sqsupseteq m$.

Informally:

$$\begin{aligned} 1\text{-}\sqsupseteq &\triangleq \sqsupseteq_{\text{POST}} \text{ for relations possibly lacking anti-symmetry} \\ \text{ipdom}_{\sqsupseteq} &\triangleq \text{ipdom}_{\sqsupseteq_{\text{POST}}} \text{ lifted to sets of } \sqsupseteq \text{-equivalent nodes} \end{aligned}$$

Specifically, if \sqsupseteq is reflexive and transitive, it admits the following rules:

Lemma 3.2.1.

$$\begin{aligned} &\frac{x \sqsupseteq y \quad y \sqsupseteq x \quad x \neq y}{x \in \text{ipdom}_{\sqsupseteq}(y) \wedge y \in \text{ipdom}_{\sqsupseteq}(x)} \text{EQ}_1^{\sqsupseteq} \\ &\frac{x \sqsupseteq y \quad y \sqsupseteq x}{\text{ipdom}_{\sqsupseteq}(x) = \text{ipdom}_{\sqsupseteq}(y)} \text{EQ}_2^{\sqsupseteq} \\ &\frac{x \sqsupseteq y \quad y \sqsupseteq x \quad x \in \text{ipdom}_{\sqsupseteq}(z)}{y \in \text{ipdom}_{\sqsupseteq}(z)} \text{EQ}_3^{\sqsupseteq} \end{aligned}$$

Rule $\text{EQ}_1^{\sqsupseteq}$ says that if x and y are in a \sqsupseteq -“cycle”, then they are immediate \sqsupseteq -postdominators of each other. By $\text{EQ}_2^{\sqsupseteq}$, they then also have the same set of immediate \sqsupseteq -postdominators. $\text{EQ}_3^{\sqsupseteq}$ says that if x is an immediate \sqsupseteq -postdominator of some other node y , then also is y .

Proof: On page 350 in the appendix. □

Postdominance Frontiers

Intuitively, the postdominance frontier of a node x is the set of nodes y such that y is at the frontier of regions of nodes postdominating x , i.e.: only *after* making a CFG choice $y \rightarrow_G s$ at y am I guaranteed to “eventually” meet x . I propose the following general definition:

Definition 3.2.2 (\sqsupseteq -Postdominance Frontiers). Given a CFG $G = (N, E)$, a relation $\sqsupseteq \subseteq N \times N$ and a node $x \in N$, the \sqsupseteq -postdominance frontier $\text{PDF}_{\sqsupseteq}(x)$ is defined by

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

The key to efficient computation of PDF_{\sqsupseteq} is to partition $\text{PDF}_{\sqsupseteq}(x)$ into two parts: those nodes y contributed to $\text{PDF}_{\sqsupseteq}(x)$ *locally*, and those nodes y contributed to $\text{PDF}_{\sqsupseteq}(x)$ by nodes z which are immediately \sqsupseteq -postdominated by x (implying $x \sqsupseteq z$).

Definition 3.2.3 (\sqsupseteq -Postdominance Frontiers: *local* part). Given a CFG $G = (N, E)$, a relation $\sqsupseteq \subseteq N \times N$ and a node $x \in N$, the \sqsupseteq -postdominance frontiers *local* part $\text{PDF}_{\sqsupseteq}^{\text{local}}(x)$ is defined by

$$\text{PDF}_{\sqsupseteq}^{\text{local}}(x) = \left\{ y \mid \begin{array}{l} \neg x \sqsupseteq y \\ y \rightarrow_G x \end{array} \right\}$$

Definition 3.2.4 (\sqsupseteq -Postdominance Frontiers: *up* part). Given a CFG $G = (N, E)$, a relation $\sqsupseteq \subseteq N \times N$ and a node $z \in N$, the part $\text{PDF}_{\sqsupseteq}^{\text{up}}(z)$ of nodes $y \in \text{PDF}_{\sqsupseteq}(z)$ that z contributes *upwards* (i.e.: to immediate \sqsupseteq -postdominators x of z) is defined by

$$\text{PDF}_{\sqsupseteq}^{\text{up}}(z) = \left\{ y \in \text{PDF}_{\sqsupseteq}(z) \mid \forall x \in \text{ipdom}_{\sqsupseteq}(z). \neg x \sqsupseteq y \right\}$$

Remark 3.2.3. Definition 3.2.2, 3.2.3 and 3.2.4 generalize definitions from ([Cyt+91], Section 4.2), by parameterizing w.r.t the underlying

postdominance relation \sqsupseteq , and replacing \sqsupseteq by the generalized notion of $1\text{-}\sqsupseteq$ -postdominance.

Specifically, I have:

$$\text{PDF}_{\sqsupseteq\text{POST}}(x) = \left\{ y \mid \text{for some } s \text{ s.t. } y \rightarrow_G s : \neg x \sqsupseteq\text{POST } y \right\}$$

$$\text{PDF}_{\sqsupseteq\text{POST}}^{\text{local}}(x) = \left\{ y \mid \begin{array}{l} \neg x \sqsupseteq\text{POST } y \\ y \rightarrow_G x \end{array} \right\}$$

$$\text{PDF}_{\sqsupseteq\text{POST}}^{\text{up}}(z) = \left\{ y \in \text{PDF}_{\sqsupseteq\text{POST}}(z) \mid \neg \text{ipdom}(z) \sqsupseteq\text{POST } y \right\}$$

Note that in [Cyt+91], dominance frontiers are defined for *predominance*, and CFGs with unique entry node n_e . In contrast, I only consider (generalizations of) *postdominance*. Consequently, my definitions differ from those in [Cyt+91] also w.r.t the direction of CFG edges.

I now state the conditions under which $\text{PDF}_{\sqsupseteq}^{\text{up}}$ and $\text{PDF}_{\sqsupseteq}^{\text{local}}$ are indeed suitable partitions of PDF_{\sqsupseteq} .

Lemma 3.2.2. Assume a CFG $G = (N, E)$, a node $z \in N$ and a relation $\sqsupseteq \subseteq N \times N$ such that \sqsupseteq is transitive and reflexive. Also, identify $\text{ipdom}_{\sqsupseteq}$ with the relation $\{(x, z) \mid x \in \text{ipdom}_{\sqsupseteq}(z)\}$, and assume that its transitive and reflexive closure is \sqsupseteq , i.e.:

$$\text{ipdom}_{\sqsupseteq}^* = \sqsupseteq$$

Then:

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

Proof: On page 350 in the appendix. □

Remark 3.2.4. For CFG with unique exit node n_x , the requirement

$$\text{ipdom}_{\sqsupseteq}^* = \sqsupseteq$$

for $\sqsupseteq = \sqsupseteq_{\text{POST}}$ is trivially true because there, the unique transitive reduction $>$ of $\sqsupseteq_{\text{POST}}$ (i.e.: the $\sqsupseteq_{\text{POST}}$ -dominance tree) is exactly $\text{ipdom}_{\sqsupseteq_{\text{POST}}}$:

$$m > n \iff m \in \text{ipdom}(n)$$

This equivalence between $\text{ipdom}_{\sqsupseteq}$ and transitive reductions of \sqsupseteq will *not* hold in general.

Algorithmically, the definitions of $\text{PDF}_{\sqsupseteq}^{\text{UP}}$ and $\text{PDF}_{\sqsupseteq}^{\text{local}}$ are not satisfactory, for they imply reachability checks

$$x \text{ 1-}\sqsupseteq y$$

Now I identify the conditions under which these checks can be replaced by cheaper checks that only involve *immediate* \sqsupseteq -postdominators.

Definition 3.2.5 (Closed under \rightarrow_G). A relation $\sqsupseteq \subseteq N \times N$ is said to be closed under \rightarrow_G if it admits the rule:

$$\frac{y \rightarrow_G x \quad x' \sqsupseteq y \quad x' \neq y}{x' \sqsupseteq x} \text{CL}_{\rightarrow_G}$$

Remark 3.2.5. Of course, $\sqsupseteq_{\text{POST}}$ is closed under \rightarrow_G .

Lemma 3.2.3. Let $\sqsupseteq \subseteq N \times N$ be transitive, and closed under \rightarrow_G . Then

$$\text{PDF}_{\sqsupseteq}^{\text{local}}(x) = \left\{ y \mid \begin{array}{l} \neg x \in \text{ipdom}_{\sqsupseteq}(y) \\ y \rightarrow_G x \end{array} \right\}$$

Proof: On page 352 in the appendix. □

Informally, the next technical observation says that (unless x and y are \sqsubseteq -equivalent) when starting at y and going \sqsubseteq -upwards towards x , there exist some nodes w, v such that I am not \sqsubseteq -equivalent to x upon reaching v , but *become* \sqsubseteq -equivalent to x by taking just one more \geq step towards x (from v to w).

Observation 3.2.1. Let \geq be any relation such that $\geq^* = \sqsubseteq$, and $x \sqsubseteq y$ but $\neg y \sqsubseteq x$. Then there exists some w, v such that

$$\begin{aligned} x \sqsubseteq w \geq v \sqsubseteq y \\ \text{and } w \sqsubseteq x \\ \text{but: } \neg v \sqsubseteq x \end{aligned}$$

The requirement of \sqsubseteq being closed under \rightarrow_G allowed me to simplify $\text{PDF}_{\sqsubseteq}^{\text{local}}$. The next additional requirement will allow me to simplify $\text{PDF}_{\sqsubseteq}^{\text{up}}$.

Definition 3.2.6 (Lack of Joins). A relation $\sqsubseteq \subseteq N \times N$ is said to *lack joins*, if it admits the rule:

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

In order to justify the phrase *lacking joins*, consider Figure 3.1, where arrows $n \rightarrow m$ signify³ $m \in \text{ipdom}_{\sqsubseteq}(n)$. Rule NoJoin demands that in Figure 3.1a, there are no joins at x for “paths” starting in s , unless

³ i.e.: arrows are *downwards* with respect to \sqsubseteq , that is, they are in accord with the inverse \sqsupseteq of \sqsubseteq . The justification of this choice of presentation is that it presents postdominance relations (or their transitive reductions) in accord with the “general direction” of control flow. Specifically, if in a CFG I have an edge $n \rightarrow_G m$ such that m is the *only* G -successor of n , I will have $m \in \text{ipdom}_{\sqsubseteq}(n)$ (and usually: $m > n$ in a transitive reduction of \sqsubseteq). By drawing an edge from n to m , this then is in accord with the G -edge $n \rightarrow_G m$.

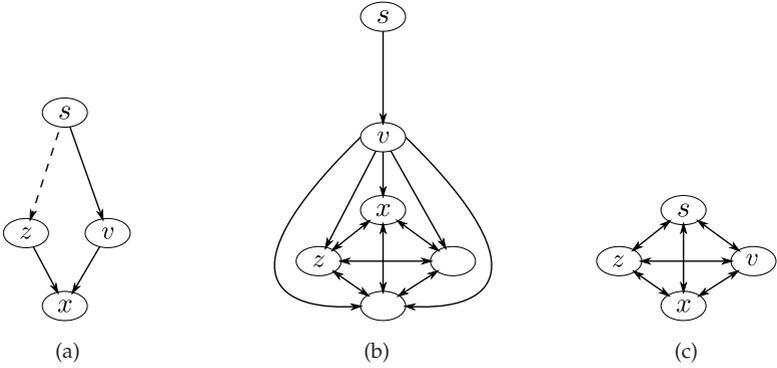


Figure 3.1: Illustration of Definition 3.2.6.

either $z \rightarrow v$ or $v \rightarrow z$, i.e.: I cannot have $s \rightarrow z$ in Figure 3.1a. On the other hand, the situations in Figure 3.1b and Figure 3.1c are allowed.

Remark 3.2.6. Standard postdominance $\sqsubseteq_{\text{POST}}$ lacks joins, because (given a unique exit node n_x) $\text{ipdom}_{\sqsubseteq_{\text{POST}}}$ forms a tree rooted in n_x , and $\sqsubseteq_{\text{POST}} = \text{ipdom}^*$.

Lemma 3.2.4. Let $\sqsubseteq \subseteq N \times N$ be transitive, reflexive, lacking joins, and closed under \rightarrow_G . Also assume

$$\text{ipdom}_{\sqsubseteq}^* = \sqsubseteq$$

Then, given some $x \in \text{ipdom}_{\sqsubseteq}(z)$:

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

Proof: On page 352 in the appendix. □

Definition 3.2.7 (Efficient PDF Partitioning). In summary, I say that a relation \sqsubseteq admits an efficient PDF partitioning if

1. it is reflexive and transitive (but not necessarily anti-symmetric)

2. $\text{ipdom}_{\sqsupseteq}^* = \sqsupseteq$

3. it admits the rules

$$\frac{y \rightarrow_G x \quad x' \sqsupseteq y \quad x' \neq y}{x' \sqsupseteq x} \text{CL} \rightarrow_G$$

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

In that case, I have

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

$$\text{PDF}_{\sqsupseteq}^{\text{local}}(x) = \left\{ y \mid \begin{array}{l} \neg x \in \text{ipdom}_{\sqsupseteq}(y) \\ y \rightarrow_G x \end{array} \right\}$$

$$\text{PDF}_{\sqsupseteq}^{\text{up}}(z) = \left\{ y \in \text{PDF}_{\sqsupseteq}(z) \mid \neg x \in \text{ipdom}_{\sqsupseteq}(y) \right\} \quad \text{for } x \in \text{ipdom}_{\sqsupseteq}(z)$$

If \sqsupseteq admits an efficient PDF partitioning, an algorithm to compute PDF is immediately available by

1. computing $\text{ipdom}_{\sqsupseteq}$
2. computing — using, e.g., any suitable work-list algorithm — the least fixed point of the monotone functional defined by the rules

$$\frac{\neg x \in \text{ipdom}_{\sqsupseteq}(y) \quad y \rightarrow_G x}{y \in \text{PDF}_{\sqsupseteq}(x)}$$

$$\frac{\neg x \in \text{ipdom}_{\sqsupseteq}(y) \quad x \in \text{ipdom}_{\sqsupseteq}(z) \quad y \in \text{PDF}_{\sqsupseteq}(z)}{y \in \text{PDF}_{\sqsupseteq}(x)}$$

It is natural to choose a *topological* iteration order during fixed-point iteration, and also to choose an efficient representation of $\text{ipdom}_{\sqsubseteq}$. A naive representation of $\text{ipdom}_{\sqsubseteq}$ (e.g.: using a generic data structure that maps nodes to set of nodes) is inefficient in general: recall from rule $\text{EQ}_3^{\sqsubseteq}$ and $\text{EQ}_2^{\sqsubseteq}$ that $\text{ipdom}_{\sqsubseteq}(z)$ may consists of many $x \neq x'$ such that $x \sqsubseteq x'$ and $x' \sqsubseteq x$, and also that then: $\text{ipdom}_{\sqsubseteq}(x) = \text{ipdom}_{\sqsubseteq}(x')$.

Given *any* transitive reduction $>$ of \sqsubseteq , I can

1. compute the strongly connected components of the graph $(N, <)$, and a corresponding topological sorting, both of which can either be provided implicitly by the algorithm computing $<$, or can be done simultaneously by, e.g., Tarjan's algorithm [Tar72].
2. compute PDF by traversing the condensed graph in that order *once*.

For the computation of *nontermination insensitive* control dependence and *nontermination sensitive* control dependence (Section 4.1), I will use algorithm 1. Note that the test

$$\neg x \in \text{ipdom}_{\sqsubseteq}(y)$$

becomes

$$\neg \exists x' \in \text{scc. } x' > y$$

due to the characterizations of $\text{ipdom}_{\sqsubseteq_{\text{SINK}}}$ and $\text{ipdom}_{\sqsubseteq_{\text{MAX}}}$ in Observation 5.3.2 and Observation 5.2.2. Also note that by computing the set $\text{scc}_{>} = \{y \mid \exists x' \in \text{scc. } x' > y\}$ *once* per scc, I can use this for both the tests on y , and for enumerating z .

Of course, for $\sqsubseteq = \sqsubseteq_{\text{POST}}$, algorithm 1 becomes the well known algorithm 2 from [Cyt+91].

Input : A transitive reduction $<$ of \sqsubseteq

Input : A topological sorting sccs of all strongly connected components of $<$.

Output: PDF_{\sqsubseteq}

for $\text{scc} \in \text{sccs}$ **do**

$\text{local} \leftarrow \{y \mid x \in \text{scc}, y \rightarrow_G x, \underbrace{\neg \exists x' \in \text{scc}. x' > y}_{y \in \text{scc}_{>}}\}$

$\text{up} \leftarrow \{y \mid \underbrace{x \in \text{scc}, x > z, y \in \text{DF}[z]}_{z \in \text{scc}_{>}}, \underbrace{\neg \exists x' \in \text{scc}. x' > y}_{y \in \text{scc}_{>}}\}$

for $x \in \text{scc}$ **do**

$\text{DF}[x] \leftarrow \text{local} \cup \text{up}$

end

end

Algorithm 1: Computation of PDF

Input : The postdominance tree $>$

Output: $\text{PDF}_{\sqsubseteq_{\text{POST}}}$

for each node x **in a bottom up traversal of** $>$ **do**

$\text{local} \leftarrow \{y \mid y \rightarrow_G x, \neg x > y\}$

$\text{up} \leftarrow \{y \mid x > z, y \in \text{DF}[z], \neg x > y\}$

$\text{DF}[x] \leftarrow \text{local} \cup \text{up}$

end

Algorithm 2: Computation of $\text{PDF}_{\sqsubseteq_{\text{POST}}}$, [Cyt+91]

3.3 Related Work

I generalized from standard postdominance $\sqsubseteq_{\text{POST}}$ to relations \sqsubseteq that admit an efficient PDF partitioning. Specifically, they need not be anti-symmetric, and CFGs are not required to have a unique exit node n_x . Other generalization exists, but still require a unique exit node n_x . For example, the authors of [BP96] generalize to P -dominance relations, and identify conditions that guarantee that the dominance relation remains anti-symmetric, and that its transitive reduction is unique and forest-structured. They apply their framework to obtain algorithms to compute “weak” control dependence. The framework from [CR06] also only applies to CFG with unique exit node n_x .

Summary

- The standard algorithm by Cytron et al for standard control dependence can be generalized to a class of postdominance relations that do not necessarily reduce to a tree.

4 Nontermination (In-)Sensitive Control Dependence

But remember not a game new under the sun
Everything you did has already been done
I know all the tricks from bricks to kingston

(Lauryn Hill — Lost Ones)

In [Ran+07], the authors introduce both **non-termination sensitive** and **-insensitive** notions of **control dependence** ($\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{nticd}}$) suitable for *modern program structures*. Specifically, they develop these notions for control flow graphs that do not have a unique exit node n_x . They also propose suitable notions of correctness for nontermination sensitive slicing, and give proof of correspondence for slices involving $\rightarrow_{\text{ntscd}}$. They do not, however, provide a notion of slicing correctness for nontermination insensitive slicing, or for $\rightarrow_{\text{nticd}}$. Instead in [Amt08], the authors need a new kind of nontermination insensitive dependence. In general, neither $\rightarrow_{\text{ntscd}}$ nor $\rightarrow_{\text{nticd}}$ slices are sound:

- Sound nontermination sensitive slices must be closed not only under $\rightarrow_{\text{ntscd}}$ and data dependence, but also under *decisive order dependence* ([Ran+07]).
- Sound nontermination *insensitive* slices must be closed not under $\rightarrow_{\text{nticd}}$ and data dependence, but instead under *weak order dependence* and data dependence ([Amt08]).

The authors claim algorithms for the computation of nontermination sensitive and insensitive control dependence, and decisive order dependence, all with running time polynomial in the size of the control flow graph.

In this chapter, I merely repeat definitions and basic results from [Ran+07]. However for those interested, a detailed treatment of the algorithms from [Ran+07] can be found in Appendix B. There in Section B.1, I explain why their algorithms for $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{nticd}}$ are

flawed. In particular, there is no immediately obvious way to repair the algorithm for $\rightarrow_{\text{nticd}}$. In order to find a correct algorithm for $\rightarrow_{\text{nticd}}$, in Section B.2 I try to explain how nontermination-sensitive control dependence naturally corresponds to *least* fixed points, while nontermination-insensitive control dependence corresponds to *greatest* fixed points. I do this by an analogy with *liveness* and *safety* properties. Indeed in Section B.3, I propose a fixed-point based characterization for $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{nticd}}$ that immediately yields correct algorithms both for $\rightarrow_{\text{ntscd}}$ and $\rightarrow_{\text{nticd}}$. There, $\rightarrow_{\text{nticd}}$ and $\rightarrow_{\text{ntscd}}$ are obtained from least and greatest fixed point, respectively, of the *same* functional, yielding the informal slogan:

nontermination sensitivity = least fixed point
nontermination insensitivity = greatest fixed point

4.1 Nontermination (In-)Sensitive Control Dependence in Arbitrary Graphs

Standard postdominance $\sqsubseteq_{\text{POST}}$ requires a unique exit node n_x to formalize the notion of a node m always “eventually” executing after an execution of n . In order to formalize this notions for arbitrary graphs (possibly lacking a unique exit node n_x), in [Ran+07] the authors introduce the notion of *control sinks*.

Definition 4.1.1. Mostly following [Ran+07], I define:

- A path π is called *maximal* if it is infinite, or its last node is an exit node (i.e.: it has no outgoing edges). With

$${}_n\Pi_{\text{MAX}}^G = \{\pi \mid \pi = n, \dots \text{ and } \pi \text{ is maximal}\}$$

I denote the set of maximal paths starting in n .

- A set $S \subseteq N$ is called a *control sink* if it is a strongly connected component of G , and successors y of nodes $x \in S$ remain in S :

$$\frac{x \rightarrow_G y \quad x \in S}{y \in S}$$

- A control sink is called *trivial* if it consists of one node. This is then either a (not necessarily unique) exit node, or a node with a self-edge (and no other outgoing edges).
- A maximal path π is called *sink-bound* (for some control-sink S) if
 - there is some node $n_s \in \pi \cap S$
 - if S is non-trivial, then all nodes in S appear in π infinitely often.

I also call such paths simply *sink paths*. With

$${}_n\Pi_{\text{SINK}}^G = \{\pi \mid \pi = n, \dots \text{ and } \pi \text{ is sink-bound}\}$$

I denote the set of sink paths starting in n .

- The set of *conditional nodes* of G is

$$\text{COND}_G = \{n \mid \exists n_l \neq n_r. n \rightarrow_G n_l \wedge n \rightarrow_G n_r\}$$

Definition 4.1.2. Let $m, n \in N$ be nodes in a graph G . Then m is said to be *non-termination sensitively/insensitively control-dependent on n* , written $n \rightarrow_{\text{ntscd}} m$ / $n \rightarrow_{\text{nticd}} m$, if there exists edges

$$n \rightarrow_G n_l, \quad n \rightarrow_G n_r$$

such that all maximal/sink- paths from n_l contain m , but not all maximal/sink- paths from n_r do:

$n \rightarrow_{\text{ntscd}} m$ iff

- $\forall \pi \in {}_{n_l}\Pi_{\text{MAX}}. m \in \pi$
- $\neg \forall \pi \in {}_{n_r}\Pi_{\text{MAX}}. m \in \pi$

$n \rightarrow_{\text{nticd}} m$ iff

- $\forall \pi \in {}_{n_l}\Pi_{\text{SINK}}. m \in \pi$
- $\neg \forall \pi \in {}_{n_r}\Pi_{\text{SINK}}. m \in \pi$

In Figure 4.1b, I show nontermination insensitive control dependence for the graph G in Figure 4.1a. It is nontermination insensitive because, for example, node 9 is *not* dependent on node 3, which can in principle delay the execution of node 9 infinitely.

The graph G has one non-trivial control sink $\{6, 7, 8, 11, 13\}$. Note that every node in this sink is control-dependent on node 2, but within this sink, there are *no* control dependencies. Intuitively, one might have expected, for example, node 11 to be dependent on node 7. Although node 7 can prevent node 11 from being executed in principle, it can only do so by delaying it infinitely (i.e.: by always choosing the successor 8). But nontermination-insensitivity here causes $\rightarrow_{\text{nticd}}$ to ignore such executions. The fact that there are no such $\rightarrow_{\text{nticd}}$ -dependencies within control-sinks means that $\rightarrow_{\text{nticd}}$ alone cannot be used for slicing in information-flow applications.

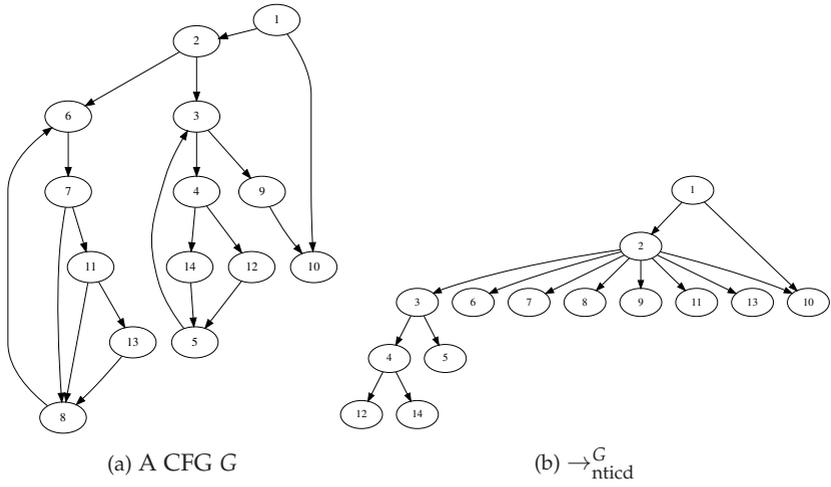


Figure 4.1: Nontermination Insensitive Control Dependence

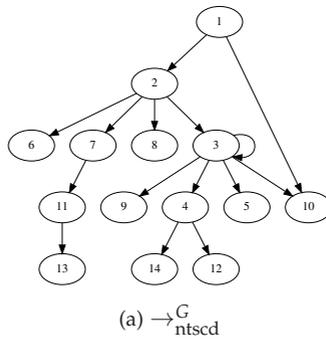


Figure 4.2: Nontermination Insensitive Control Dependence

In Figure 4.2a, I show nontermination sensitive control dependence for the same graph CFG from Figure 4.1a. It is sensitive because, for example, node 9 is dependent on node 3, which can delay the execution of node 9 infinitely. Also, node 11 is control dependent on node 7. In this graph G , $\rightarrow_{\text{ntscd}}$ induces sound slices, but this is not so for every graph G .

Definition 4.1.3. In the words of Definition 3.1.2 from Section 3.2, $\rightarrow_{\text{nticd}}$ is just $\sqsupseteq_{\text{SINK}}$ -control dependence, and $\rightarrow_{\text{ntscd}}$ is just \sqsupseteq_{MAX} -control dependence, where

$$\begin{aligned} m \sqsupseteq_{\text{SINK}}^G n &\Leftrightarrow \forall \pi \in {}_n\Pi_{\text{SINK}}^G. m \in \pi \\ m \sqsupseteq_{\text{MAX}}^G n &\Leftrightarrow \forall \pi \in {}_n\Pi_{\text{MAX}}^G. m \in \pi \end{aligned}$$

I say that $\sqsupseteq_{\text{SINK}}^G$ is a nontermination *insensitive* form of postdominance, while $\sqsupseteq_{\text{MAX}}^G$ is a nontermination *sensitive*.

Nontermination insensitive control dependence $\rightarrow_{\text{nticd}}$ generalizes standard control dependence to graphs without unique exit node n_x . Nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$ generalizes *weak* control-dependence [PC90] to graphs without unique exit node. Intuitively, the difference between $\rightarrow_{\text{nticd}}$ and $\rightarrow_{\text{ntscd}}$ is that $\rightarrow_{\text{nticd}}$ assumes “every loop to terminate”. Thus I have $n \rightarrow_{\text{ntscd}} m$ even if the only possibility of conditional node n to prevent the execution of m is by infinitely often choosing the same successor n_l . On the other hand, $\rightarrow_{\text{nticd}}$ assumes such loops to terminate, i.e.: it expects n to eventually choose a *different* successor n_r .

Indeed, Ranganath et al. give proofs for the following properties:

Theorem 4.1.1 (Ranganath et al., Theorem 1). Let G have a unique exit node n_x . Then

$$n \rightarrow_{\text{nticd}} m \iff n \rightarrow_{\text{cd}} m$$

Theorem 4.1.2 (Ranganath et al., Theorem 4). Let G be *any* CFG. Then

$$n \rightarrow_{\text{nticd}} m \implies n \rightarrow_{\text{ntscd}} m$$

Summary

- Nontermination sensitive and insensitive control dependence generalize standard control dependence notions to graphs without unique exit node.
- But these notions alone do not lead to sound slices.
- They can be defined uniformly, by reference to nontermination sensitive and insensitive postdominance \exists_{MAX} and \exists_{SINK} .

5 Postdominator Pseudoforests

You come to me for advice, but you can't cope with anything you don't recognize. Hmm. So we'll have to tell you something you already know but make it sound like news, eh. Well, business as usual, I suppose.

(Douglas Adams — The Ultimate Hitchhiker's Guide to the Galaxy)

Neither the (originally flawed) Algorithm 14 for $\rightarrow_{\text{ntscd}}$ due to [Ran+07], nor my correct algorithm Algorithm 16 for $\rightarrow_{\text{ntcd}}$ from Appendix B offer performance comparable to the efficient-in-practice algorithm for standard control dependence \rightarrow_{cd} .

But with the algorithm for generalized control dependence from Section 3.2, I can immediately derive efficient-in-practice algorithm for both $\rightarrow_{\text{ntcd}}$ and $\rightarrow_{\text{ntscd}}$ if

1. $\sqsubseteq_{\text{SINK}}$ and \sqsubseteq_{MAX} admit efficient PDF partitionings, and
2. I can provide efficient-in-practice algorithm for the computation of some transitive reduction $<$ of $\sqsubseteq_{\text{SINK}}$ and \sqsubseteq_{MAX} .

In order to do this, I first give in Section 5.1 fixed-point characterizations of $\sqsubseteq_{\text{SINK}}$ and \sqsubseteq_{MAX} . Then in Section 5.2 and Section 5.3 I generalize the notion of standard postdominator trees to obtain the new notion of postdominator *pseudo-forests*. These will materialize as transitive reductions of nontermination sensitive and insensitive postdominance \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$, and together with the fixed-point characterizations of $\sqsubseteq_{\text{SINK}}$ and \sqsubseteq_{MAX} immediately lead to efficient-in-practice algorithms for their computation.

For the nontermination sensitive case $\rightarrow_{\text{ntscd}}$, I am not aware of previously existing efficient algorithms even for graphs with unique exit

node n_x .¹ For the nontermination insensitive case, on the other hand, such algorithms exist ([LT79; Cyt+91]). Indeed if I were only interested in (“offline”) algorithms for $\rightarrow_{\text{nticd}}$, I would not need the development in Section 5.3, because there exists a reduction $G \mapsto G_S$ such that G_S is a graph with unique-exit node, and nontermination insensitive control dependence $\rightarrow_{\text{nticd}}^G$ in G can easily be obtained from standard control dependence $\rightarrow_{\text{cd}}^{G_S}$ in G_S . I explain this reduction in Section 5.4. Later in Section 6.7 however, I *will* make use of the algorithms developed in Section 5.3 in order to support an *incremental* recomputation of a series $(\sqsupset_{\text{SINK}}^m)_{m \in M}$ of postdominance pseudo-forests $\sqsupset_{\text{SINK}}^m$ for a sequence $(G_m)_{m \in M}$ of related graphs.

¹ In [BP96], the authors provide an efficient algorithm for the computation of the *Augmented Loop-Postdominator-Tree*, which for CFG with unique exit nodes is derived by augmenting a nontermination-sensitive postdominance forest. Using this tree, they give an algorithm to answer weak (i.e.: nontermination-sensitive) control dependence *queries*, but do not give an explicit algorithm to efficiently compute the whole weak control relation.

5.1 Fixed-Point Characterizations for Postdominance

It is implicit already in, e.g., [HU73], that $\sqsupseteq_{\text{POST}}$ (i.e.: the postdominance relation underlying standard control dependence for graphs with unique end node n_x) has a *greatest* fixed point characterization:

Theorem 5.1.1. Let G be a CFG with unique end node n_x and P be the rule-system

$$\frac{}{n \sqsupseteq n} P^{\text{self}} \quad \frac{\forall p \rightarrow_G x. m \sqsupseteq x \quad p \neq n_x}{m \sqsupseteq p} P^{\text{suc}}$$

Then $\sqsupseteq_{\text{POST}} = \nu P$ over the lattice $(N \times N, \subseteq)$.

Proof: On page 354 in the appendix. □

Note that when computing $\sqsupseteq_{\text{POST}}$ as the greatest fixed point of P , fixed point iteration starts with the initial value $\top = N \times N$, i.e. with the assumption

$$m \sqsupseteq n \text{ for all } m, n$$

regardless whether m is even reachable in G from n . This is correct because an assumption $m \sqsupseteq n$ for $\neg n \rightarrow_G^* m$ must be validated by $m \sqsupseteq x$ for all G -successors x of n . But eventually I must reach n_x *without* ever reaching m before, and at n_x i cannot validate $m \sqsupseteq n_x$.

Usually, P is presented as a system of equations:

$$\begin{aligned} \text{ipdom}_{\text{POST}}(n_x) &= \{n_x\} \\ \text{ipdom}_{\text{POST}}(n) &= \{n\} \cup \bigcap_{n \rightarrow_G x} \text{ipdom}_{\text{POST}}(x) \quad n \neq n_x \end{aligned}$$

An algorithm loosely based on P is given in [CHK01]. Here, the authors — while computing *immediate* postdominators — do *not* explic-

itly (or implicitly!) start with the assumption that any node m may post-dominate any node n . Instead, they make crucial use of the fact that any preliminary *guess* $\text{ipdom}(n) = x$ for any G -successor x of n is always O.K. in the sense that x will at least always *reach* the *true* $\text{ipdom}(n)$ (by following ipdom upwards the dominator tree).

Following the idea that $\rightarrow_{\text{ntcd}}$ is a generalization of standard control dependence (i.e.: $\sqsubseteq_{\text{POST}}$ control dependence), and $\sqsubseteq_{\text{SINK}}$ a generalization of $\sqsubseteq_{\text{POST}}$, I must suspect that the greatest fixed point of *some* generalization of P describes $\sqsubseteq_{\text{SINK}}$ for CFGs without unique end nodes. Moreover, the slogan

$$\begin{aligned} \text{non-termination insensitivity} &= \text{greatest fixed point} \\ \text{non-termination sensitivity} &= \text{least fixed point} \end{aligned}$$

then also suggests that it's *least* fixed point characterizes \sqsubseteq_{MAX} . In fact:

Theorem 5.1.2. Let G be any CFG and D be the rule-system

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

Then

$$\begin{aligned} \nu D &= \sqsubseteq_{\text{SINK}} \\ \mu D &= \sqsubseteq_{\text{MAX}} \end{aligned}$$

Proof: On page 354 in the appendix. □

Remark 5.1.1. A formal proof of Theorem 5.1.2 checked by the Isabelle/HOL proof assistant is available from Simon Bischof [Bis19].

When generalizing to D from P , I had to add the reachability constraint $p \rightarrow_G^* m$, since otherwise I could — for example: given some control sink S , and nodes $n \in S$ but $m \notin S$ — validate $m \sqsubseteq_{\text{SINK}} n$

merely by mutual application of rule D^{succ} (to all nodes in S), even if m is not even reachable from n . Given this reachability constraint, I do not need a generalization of the constraint $p \neq n_x$, since it's main function was to prevent application of rule P^{succ} to nodes p that do not have *any* successors (of which n_x is the only one in CFG with unique exit node). But in D , this is already accomplished by the reachability constraint $p \rightarrow_G^* m$.

Remark 5.1.2. Rule system D can be read as a simplified version of the following rule system D_3 from the appendix.

$$\frac{m \in \overline{\text{next}_{\text{COND}}[x]} \quad D_3^{\text{lin}}}{m \sqsubseteq x} \quad \frac{n = \text{next}_{\text{COND}}[x] \quad \forall n \rightarrow_G x. m \sqsubseteq x \quad x \rightarrow_G^* m}{m \sqsubseteq x} D_3^{\text{cond}}$$

But D_3 , then, can be read as just a variant of the rules S_3 (Definition B.2.5 in the appendix).

Both D fixed-points can be computed naively, using an explicit representations of the relation \sqsubseteq . For more efficient algorithms, I need instead a sparse representations of \sqsubseteq , which is available both for \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ in form of *pseudo forests*.

5.2 Nontermination Sensitive Pseudoforests

Before I devise an efficient algorithm for nontermination sensitive postdominance \sqsupseteq_{MAX} , let me first affirm that, indeed, \sqsupseteq_{MAX} admits efficient PDF partitionings, as required if I want to use my generalized postdominance frontiers algorithm from Section 3.2. It is easy to see that the first requirement is met:

Lemma 5.2.1. \sqsupseteq_{MAX} is reflexive and transitive.

Proof (Sketch): By definition.

In order to establish the remaining requirements, I investigate the structure of \sqsupseteq_{MAX} .

Observation 5.2.1. Let G be any CFG, and $>_{\text{MAX}}$ any transitive reduction of \sqsupseteq_{MAX} . Then the graph

$$(N, <_{\text{MAX}}) = (N, \{ (m, n) \mid n >_{\text{MAX}} m \})$$

is a *pseudo forest*, i.e.: a graph with at most one successor at each node n .

A glance at Figure 5.1b will justify the name *pseudo forest*: Visible are five independent pseudo-trees (where roots sometimes consist of *multiple* nodes $n \in N$) with roots: 1, 2, 3, {6, 7, 8}, 10. The roots are shown near the bottom, and arrows \rightarrow follow $<_{\text{MAX}}$.

Observation 5.2.2. Let $>_{\text{MAX}}$ be any transitive reduction of \sqsupseteq_{MAX} . Then

$$\begin{aligned} x \perp_{\sqsupseteq_{\text{MAX}}} z &\iff x >_{\text{MAX}}^+ z \\ x \in \text{ipdom}_{\sqsupseteq_{\text{MAX}}}(z) &\iff x' >_{\text{MAX}}^* x >_{\text{MAX}}^* x' \\ &\text{for some } x' \text{ s.t. } x' >_{\text{MAX}} z \end{aligned}$$

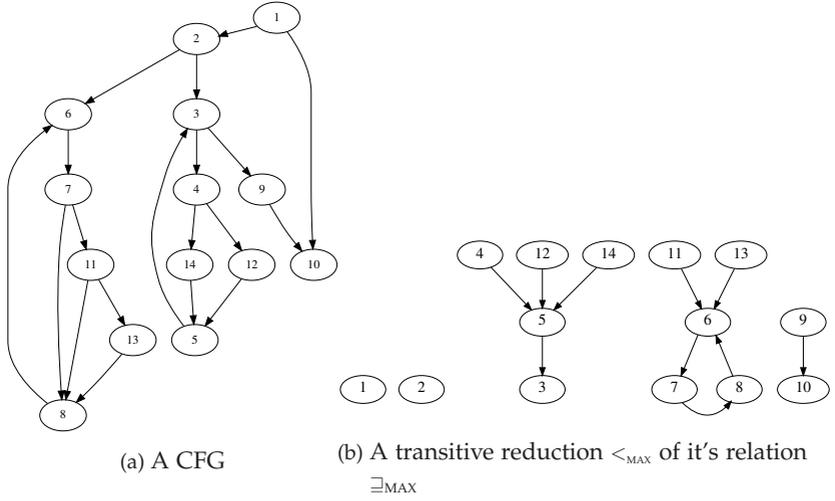


Figure 5.1: A Dominator Pseudo-Forest

and consequently:

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

Lemma 5.2.2. \sqsupseteq_{MAX} lacks joins.

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

Proof: On page 355 in the appendix. □

Lemma 5.2.3. \sqsupseteq_{MAX} is closed under \rightarrow_G .

$$\frac{y \rightarrow_G x \quad x' \sqsupseteq_{\text{MAX}} y \quad x' \neq y}{x' \sqsupseteq_{\text{MAX}} x} \text{CL}_{\rightarrow_G}$$

Proof: By definition. □

The Algorithm for computing a pseudo-forest $<_{\text{MAX}}$ will require the computation of *least common ancestors* for (preliminary) pseudo-forests:

Definition 5.2.1. Given a pseudo-forest $<$ and two nodes n, m

$$\text{lca}_<(n, m) = \left\{ a \mid \begin{array}{l} n <^* a \wedge m <^* a \\ \forall a'. (n <^* a' \wedge m <^* a') \implies a <^* a' \end{array} \right\}$$

More generally: for any set S of nodes:

$$\text{lca}_<(S) = \left\{ a \mid \begin{array}{l} \forall n \in S. n <^* a \\ \forall a'. (\forall n \in S. n <^* a') \implies a <^* a' \end{array} \right\}$$

A least common ancestor $\text{lca}_<(S)$ of a set S can be computed by iterating over nodes $n \in S$:

Proposition 5.2.1. Given a pseudo-forest $<$ and nodes x, y, z , if $a_{xy} \in \text{lca}_<(x, y)$ and $a_{xyz} \in \text{lca}_<(a_{xy}, z)$, then $a_{xyz} \in \text{lca}_<(\{x, y, z\})$.

The least common ancestor can naively be computed by, e.g., Algorithm 3, and be extended to sets via Algorithm 4. See Algorithm 18 in Appendix D for the slightly more performant variant used in all benchmarks.

Input : A pseudo-forest \prec , represented as a map
 $\text{IMDOM} : N \leftrightarrow N$ s.t. $\text{IMDOM}[n] = m$ iff $n < m$.

Input : Nodes m_0, n_0

Output : A least common ancestor of n_0, m_0 , or \perp if there is none.

```

begin
  | return lca ( n_0, m_0 )
end
Function 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  $m$ 
  | end
  | if  $n \in \pi_m$  then
  |   | return  $n$ 
  | end

  | switch {  $\text{IMDOM}[n]$  } \  $\pi_n$  do
  |   | case  $\emptyset$  do
  |     | switch {  $\text{IMDOM}[m]$  } \  $\pi_m$  do
  |       | case  $\emptyset$  do
  |         | return  $\perp$ 
  |         end
  |       | case {  $m'$  } do
  |         | return lca (  $\pi_m m', \pi_n$  )
  |         end
  |       end
  |     end
  |   | case {  $n'$  } do
  |     | return lca (  $\pi_m, \pi_n n'$  )
  |     end
  |   end
end
end

```

Algorithm 3: A least common ancestor algorithm.

Input : A pseudo-forest \prec , represented as a map $\text{IMDOM} : N \leftrightarrow N$
 s.t. $\text{IMDOM}[n] = m$ iff $n < m$.

Input : A set $\{n\} \cup S$ of at least one Node

Output: A least common ancestor of $\{n\} \cup S$, or \perp if there is none.

begin

$a \leftarrow n$

for $s \in S$ **do**

$a \leftarrow \text{lca}(a, s)$

end

end

Algorithm 4: A least common ancestor algorithm for sets.

5.2.1 An Algorithm for Nontermination Sensitive Postdominance

I can now present Algorithm 5 for the computation of some transitive reduction $>_{\text{MAX}}$ of \exists_{MAX} . It can be understood as a least fixed point computation (of μD_3) using an *efficient representation* $<_{\text{MAX}} \in \alpha(\exists_{\text{MAX}})$ of \exists_{MAX} . Informally:

$$\begin{aligned} \gamma(\prec) &= \prec^* \\ \alpha(\exists) &= \{\prec \mid \prec^* = \sqsubseteq\} \end{aligned}$$

with (abstract) rule system $D_3^\#$

$$\frac{m \rightarrow_G x \quad \neg m \in \text{COND}}{m <_{\text{MAX}} x}$$

$$\frac{a = \epsilon a. a \in \text{lca}_{<_{\text{MAX}}}(\{x \mid m \rightarrow_G x\}) \quad m \in \text{COND} \quad m \rightarrow_G^* a}{m <_{\text{MAX}} a}$$

In order to guarantee termination, the algorithm will choose

$$\epsilon a. a \in \text{lca}_{<_{\text{MAX}}} (\dots)$$

consistently (i.e.: always choose the same a for a given input.) The Algorithm 5 is a straight-forward implementation of a chaotic least fixed point iteration. The update of workset can be implemented reasonably efficiently if the set $\{y \mid y <^* x\}$ can be enumerated efficiently, which can, for example, be done by maintaining a second data structure $\text{IMDOM} : N \rightarrow 2^N$ representing $>$.

Observation 5.2.3. Let G be any CFG. Then Algorithm 5 terminates with a result $<_{\text{MAX}}$ s.t. $>_{\text{MAX}}$ is a transitive reduction of \sqsupseteq_{MAX} .

Also see Subsection A.2.1 in the appendix for a sketch a proof of Observation 5.2.3.

This chaotic iteration can be improved slightly by making use of the fact that for each node $x \in \text{COND}$, the value $\text{IMDOM}[x]$ is never changed from some z_0 back to \perp , and changed from some z_0 to some z only for $z_0, z \in \text{ipdom}_{\sqsupseteq_{\text{MAX}}}(x)$: never re-insert nodes n s.t. $\text{IMDOM}[n] \neq \perp$.

For the very same reason, a sequential (non-chaotic) variant of this least fixed point computation can be implemented reasonably efficiently: The resulting algorithm (Algorithm 19 in Appendix D) maintains a FIFO queue which at all times contains those nodes $x \in \text{COND}$ for which $\text{IMDOM}[x] = \perp$, re-inserting nodes x only if no $\text{lca} \neq \perp$ was found. It terminates as soon as the whole queue was traversed once without it becoming smaller.

Input : A CFG G

Data: A pseudo-forest $<$ represented as a map $\text{IMDOM} : N \leftrightarrow N$ s.t.
 $\text{IMDOM}[n] = m$ iff $n < m$

Output: A transitive reduction $<_{\text{MAX}}$ of \exists_{MAX}

begin

for $x \in N$, $\{z \mid x \rightarrow_G z\} = \{z\}$, $z \neq x$ **do**

$\text{IMDOM}[x] \leftarrow z$

end

$\text{MAXIMAL}_{\text{up}}$

return IMDOM

end

Procedure $\text{MAXIMAL}_{\text{up}}$

$\text{workset} \leftarrow \text{COND}_G$

while $\text{workset} \neq \emptyset$ **do**

$x \leftarrow \text{remove}(\text{workset})$

$a \leftarrow \text{lca}(\{y \mid x \rightarrow_G y\})$

$z \leftarrow \begin{cases} \perp & \text{if } a = \perp \vee a = x \\ a & \text{otherwise} \end{cases}$

assert $z \neq \text{IMDOM}[x] \Rightarrow z \neq \perp$

assert $z \neq \text{IMDOM}[x] \wedge \text{IMDOM}[x] = z_0 \Rightarrow z_0 <^* z <^* z_0$

if $z \neq \text{IMDOM}[x]$ **then**

$\text{workset} \leftarrow \text{workset} \cup \{n \in \text{COND} \mid n \neq x, \exists n \rightarrow_G y. y <^* x\}$

$\text{IMDOM}[x] \leftarrow z$

end

end

end

Algorithm 5: An algorithm for the computation of some $<_{\text{MAX}}$.

5.3 Nontermination Insensitive Pseudoforests

Just as Algorithm 5 for a transitive reduction $>_{\text{MAX}}$ of \sqsubseteq_{MAX} corresponds to a computation of the least fixed point μD , the algorithm for a transitive reduction $>_{\text{SINK}}$ of $\sqsubseteq_{\text{SINK}}$ will resemble a *greatest* fixed point computation (of the functional D). The key of obtaining an efficient algorithm will be

1. To avoid explicit reachability checks $m \rightarrow_G^* a$. These were unnecessary in the least fixed point computation because there, reachability was established during the iteration starting from $\perp = \emptyset \subseteq N \times N$. In the greatest fixed-point computation, however, these would become necessary if I were to start a greatest fixed point computation at $\top = N \times N$.
2. To instead find a *pseudo forest* $<$ such that

$$>^* \supseteq \sqsubseteq_{\text{MAX}}$$

to start the fixed point iteration from.

I first affirm that $\sqsubseteq_{\text{SINK}}$ does, indeed, admit efficient PDF partitionings. Again, it is easy to see that the first requirement is met:

Lemma 5.3.1. $\sqsubseteq_{\text{SINK}}$ is reflexive and transitive.

Proof (Sketch): By definition.

Also, the structure of transitive reductions $>_{\text{SINK}}$ is not different from before:

Observation 5.3.1. Let G be any CFG, and $>_{\text{SINK}}$ any transitive reduction of $\sqsubseteq_{\text{SINK}}$. Then the graph

$$(N, <_{\text{SINK}}) = (N, \{ (m, n) \mid n >_{\text{SINK}} m \})$$

is a *pseudo forest*.

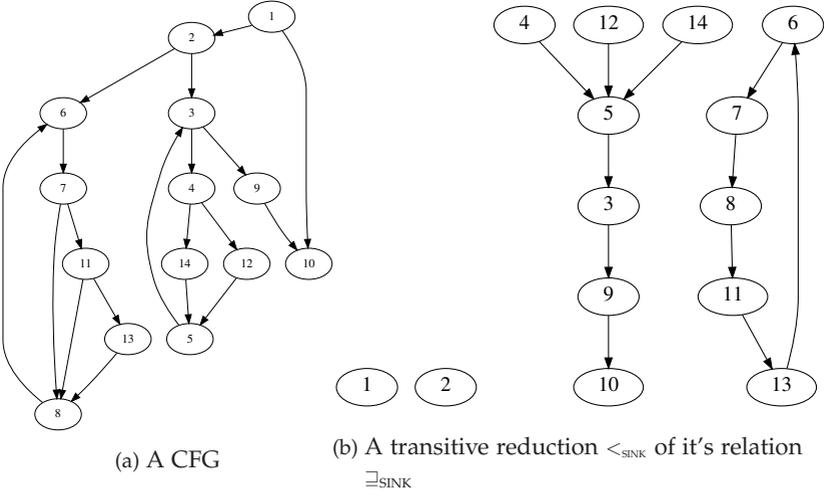


Figure 5.2: A Dominator Pseudo Forest

In Figure 5.2b such a pseudo forest \prec_{SINK} is shown for the previously used example CFG. Visible are four independent pseudo-trees with roots: 1, 2, 10, {6, 7, 8, 11, 13}.

Observation 5.3.2. Let \prec_{SINK} be any transitive reduction as above. Then

$$\begin{aligned}
 x \perp_{\text{SINK}} z &\iff x \succ_{\text{SINK}}^+ z \\
 x \in \text{ipdom}_{\text{SINK}}(z) &\iff x \succ_{\text{SINK}}^* x' \succ_{\text{SINK}}^* z \\
 &\quad \text{for some } x' \text{ s.t. } x' \succ_{\text{SINK}} z
 \end{aligned}$$

and consequently:

$$\text{ipdom}_{\text{SINK}}^* = \text{SINK}$$

Just as before, SINK lacks joins, and is closed under \rightarrow_G :

Observation 5.3.3.

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

Proof: Just as for Lemma 5.2.2. □

Observation 5.3.4.

$$\frac{y \rightarrow_G x \quad x' \exists_{\text{SINK}} y \quad x' \neq y}{x' \exists_{\text{SINK}} x} \text{CL} \rightarrow_G$$

Additionally, the sinks of the CFG determine $<_{\text{SINK}}$ to some degree:

Observation 5.3.5. Let $>_{\text{SINK}}$ be a transitive reduction of \exists_{SINK} . Then the multi-node roots of the pseudo-forest $<_{\text{SINK}}$ are exactly the multi-node sinks of G . Also, any single-node sink of G is a single-node root in $<_{\text{SINK}}$.

The Algorithm 6 for the computation of $<_{\text{SINK}}$ then works in two phases. In the first phase SINK_{up} , it computes some approximation $<_0$ “just above” $<_{\text{SINK}}$, i.e.: a pseudo-forest $<_0$ such that

$$>_0^* \supseteq \exists_{\text{SINK}}$$

The phase $\text{SINK}_{\text{down}}$ then computes $<_{\text{SINK}}$ by (greatest) fixed-point iteration from above.

In SINK_{up} , I necessarily have to be more lenient than I was in $\text{MAXIMAL}_{\text{up}}$. There, in order to establish $x < a$, I needed positive evidence that $a \in \text{lca}_{<}(\{y \mid x \rightarrow_G y\})$. Now, even if $\text{lca}_{<}(\{y \mid x \rightarrow_G y\}) = \emptyset$, I will sometimes have to chose an *assumption* $x < a$ for some reasonable a , which I later either validate, weaken or abandon in $\text{SINK}_{\text{down}}$.

Consider the preliminary pseudo forest $<$ in Figure 5.3b. I need to establish $1 < 3$ (i.e.: that 3 is an “immediate postdominator” of 1), but I find that $\text{lca}_<(\{2, 3\}) = \emptyset$. Now I would *like* to assume *both*, $1 < 3$ and $1 < 2$, the latter of which would then be invalidated in phase $\text{SINK}_{\text{down}}$. But then $<$ no longer would be a pseudo forest! If I were to assume just $1 < 2$, I would obtain a $>_0$ to start the phase $\text{SINK}_{\text{down}}$ with such that *not*: $>_0^* \supseteq \sqsupset_{\text{MAX}}$, so I need to make the assumption $1 < 3$. But what is a generally applicable criterion to decide which assumption to make?

Observation 5.3.6. Let $>_{\text{SINK}}$ be a transitive reduction of \sqsupset_{SINK} . Then whenever $x <_{\text{SINK}} y$, there is at most one sink S such that any path starting in x is bound for S , and any such S is the same sink that y is bound for.

Note that in this example, for node 3 I have already established $3 <^* 4$ for the sink node $4 \in S$, but I have not yet established $2 <^* 4$ (nor: $2 <^* s$ for *any* sink node s). This suggest that I must — whenever $\text{lca}_<(\{y \mid x \rightarrow_G y\}) = \emptyset$ — choose some G -successor node y of x such that already $y <^* s$ for some sink node s . I will call such nodes y *processed*, and maintain in Algorithm 6 a set PROCD of all such nodes. Not shown is the procedure $\text{newProcessed}(x)$ which updates PROCD given a node x s.t. $x <^* s$ for some sink node s , by following linear segments ending in x upwards.

The second phase $\text{SINK}_{\text{down}}$ then corresponds to a (downward) fixed-point iteration that computes of νD starting from $>_0^*$ (instead of: from \top). The second phase is essentially the same as the first and only phase $\text{MAXIMAL}_{\text{up}}$ of Algorithm 5. Remember that there, I computed μD from \perp . Yet, here are some differences between $\text{SINK}_{\text{down}}$ and $\text{MAXIMAL}_{\text{up}}$. In $\text{MAXIMAL}_{\text{up}}$, I can make use of the fact that for each n , I need to set $\text{IMDOM}[n]$ at most once, while in $\text{SINK}_{\text{down}}$, I may need to update $\text{ISDOM}[n]$ several times. Also, $\text{SINK}_{\text{down}}$ only updates non-sink nodes – sink-nodes were already dealt with during the initialization. Finally, whenever the algorithm would normally set $\text{IMDOM}[x]$ to s for some node s in sink S_i , the algorithm instead sets

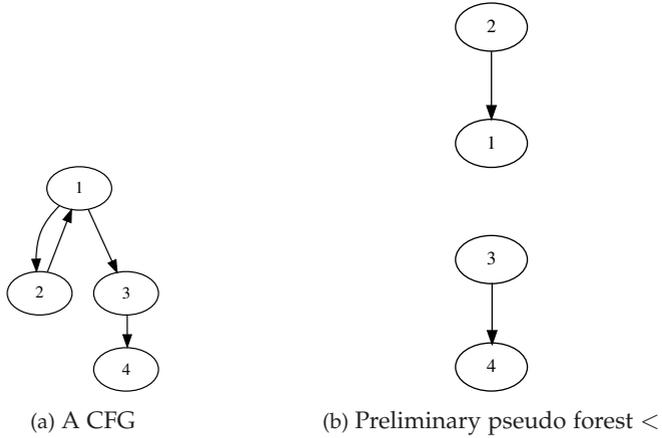


Figure 5.3

$\text{IMDOM}[x]$ to s_i for some canonical sink node $s_i \in S_i$. This is required to ensure termination, since otherwise, the algorithm might oscillate between choosing $x < s$ and $x < s'$ for two nodes $s, s' \in S_i^2$. Also note that during the workset update with n , here necessarily $n \neq x$.

Dually to $\text{MAXIMAL}_{\text{up}}$, the algorithm $\text{SINK}_{\text{down}}$ can be made slightly more efficient by making use of the fact that once $\text{ISDOM}[x] = \perp$ has been established, it must remain so: never insert such nodes into the workset. Again, chaotic iteration can be replaced with sequential iteration, making use of this same fact.

Observation 5.3.7. Let G be any CFG. Then Algorithm 6 terminates with a result $<_{\text{SINK}}$ s.t. $>_{\text{SINK}}$ is a transitive reduction of $\sqsubseteq_{\text{SINK}}$.

This result concludes Chapter 5 satisfactorily: I have provided algorithms for the computation of both \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ for *arbitrary* CFG, both of which

² In phase $\text{MAXIMAL}_{\text{up}}$ of Algorithm 5, the fact that $<$ approximated \sqsubseteq_{MAX} was enough to guarantee termination, but here, $<$ approximating $\sqsubseteq_{\text{SINK}}$ is not.

Input : A CFG G

Output: A transitive reduction $<_{\text{SINK}}$ of \exists_{SINK}

begin

$\{S_1, \dots, S_n\} \leftarrow \{S_i \mid S_i \in \text{scc}(G), \neg \exists s \rightarrow_G n. s \in S_i \wedge n \notin S_i\}$
using any efficient scc algorithm.

$S \leftarrow \bigcup S_i$

for $1 \leq i \leq n$ **do**

$s_i \leftarrow$ any node in S_i

for $n_j \in S$ 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** $k_i = 1$

processed(n_j)

end

end

for $x \in N, x \notin S, \{z \mid x \rightarrow_G z\} = \{z\}, z \neq x$ **do**

$\text{ISDOM}[x] \leftarrow z$

if $z \in \text{PROCD}$ **then** processed(x)

end

SINK_{up}

$\text{SINK}_{\text{down}}$

return ISDOM

end

Algorithm 6: An algorithm for the computation of some $<_{\text{SINK}}$.

- use efficient data structures (pseudo forests) to compute and represent $\exists_{\text{MAX}}, \exists_{\text{SINK}}$
- do not require explicit reachability checks $m \rightarrow_G^* n$

I derived these algorithm from simple fixed-points characterizations of $\exists_{\text{MAX}}, \exists_{\text{SINK}}$ via D , following the least/greatest fixed point duality.

Later in Chapter 8 validate that these algorithms are efficient in practice.

```

Procedure SINKup
  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  $\leftarrow$  {  $y \mid x \rightarrow_G y, y \in$  PROCD }
    if SUCCS =  $\emptyset$  then
       $z \leftarrow \perp$ 
    else
       $a \leftarrow$  lca(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)
    end
    else
      pushBack(workqueue, x)
    end
  end
end

```

Upward Iteration for Algorithm 6.

Procedure SINK_{down}

$$\text{workset} \leftarrow \{ n \mid n \in \text{COND}_G \setminus S, \text{ISDOM}[n] \neq \perp \}$$
while workset $\neq \emptyset$ **do** $x \leftarrow \text{removeMin}(\text{workset})$ $a \leftarrow \text{lca}(\{ y \mid x \rightarrow_G y \})$

$$z \leftarrow \begin{cases} \perp & \text{if } a = \perp \\ s_i & \text{if } a \in S_i \\ a & \text{otherwise} \end{cases}$$
assert $\text{ISDOM}[x] = \perp \Rightarrow z = \perp$ **if** $z \neq \text{ISDOM}[x]$ **then**

$$\text{workset} \leftarrow \text{workset} \cup \{ n \in \text{COND} \setminus S \mid \exists n \rightarrow_G y. y <^* x \}$$

$$\text{ISDOM}[x] \leftarrow z$$
end**end****end**

Downward Iteration for Algorithm 6. Set workset is ordered by any fixed ordering of nodes N .

5.4 Reduction to Postdominance Trees

In this section, I present another approach to the computation of $\sqsupseteq_{\text{SINK}}$.

By Observation 5.3.5, the roots (whether trivial or non-trivial) of pseudo-forests $<_{\text{SINK}}$ (i.e.: the $<_{\text{SINK}}$ -cycles) are exactly the control-sinks in G , which are exactly the strongly connected components of G that have no outgoing edge. Within such a root (i.e.: within such a set $M \subseteq N$ of nodes), any two nodes $m_1, m_2 \in M$ do $\sqsupseteq_{\text{SINK}}$ -postdominate each other, hence no node is ever nontermination insensitively control dependent on any node $m \in M$. At the same time, a node n (possibly not in M) is $\sqsupseteq_{\text{SINK}}$ -postdominated by some node $m \in M$ if and only if *every* node $m \in M$ is. In other words: nodes $m \in M$ are equivalent with regard to $\sqsupseteq_{\text{SINK}}$ -postdominance. This suggests that it is possible to reduce the computation of $<_{\text{SINK}}$ to a computation a CFG obtained by *condensing* sinks in G .

The Idea behind the following Construction was suggested by Maximilian Wagner:

Observation 5.4.1. Let $G = (N, E)$ be any CFG, and S_1, \dots, S_n its sinks. For $S = \bigcup S_i$, a fresh node $n_x \notin N$, and given for each sink S_i a representative node $s_i \in S_i$, let

$$\begin{aligned} G_S &= (N_S, E_S) \\ N_S &= \{n_x, s_1, \dots, s_n\} \cup (N \setminus S) \\ E_S &= \{(n, m) \mid n \rightarrow_G m, n, m \notin S\} \\ &\cup \{(n, s_i) \mid n \rightarrow_G s \text{ for some } s \in S_i\} \\ &\cup \{(s_i, n_x) \mid 1 \leq i \leq n\} \end{aligned}$$

Then for nodes $x, m \in N \setminus S$:

$$m \sqsupseteq_{\text{SINK}}^G x \iff m \sqsupseteq_{\text{POST}}^{G_S} x$$

and for $x \in N \setminus S$ and $m \in S_i$:

$$m \sqsupseteq_{\text{SINK}}^G x \iff x \in S_i \vee s_i \sqsupseteq_{\text{POST}}^{G_S} x$$

This immediately suggests Algorithm 7 for the computation of some transitive reduction \succ_{SINK} of $\sqsupseteq_{\text{SINK}}$

Input : A CFG G

Output: A transitive reduction \succ_{SINK} of $\sqsupseteq_{\text{SINK}}$ represented as a map
 $\text{ISINKDOM} : N \leftrightarrow N$ s.t. $\text{ISINKDOM}[n] = m$ iff $n <_{\text{SINK}} m$

begin

$\{S_1, \dots, S_n\} \leftarrow \{S_k \mid S_k \in \text{scc}(G), \neg \exists s \rightarrow_G n. s \in S_k, n \notin S_k\}$
 using any efficient scc algorithm.

$S \leftarrow \bigcup S_i$

for $1 \leq i \leq n$ **do**

$s_i \leftarrow$ any node in S_i

end

$\text{IPDOM} \leftarrow$ ipdom(G_S) by any efficient postdominance algorithm
 for graphs with unique end-node

for $m \in N \setminus S$ **do**

$\text{ISINKDOM}[n] \leftarrow \text{IPDOM}[n]$

end

for $1 \leq i \leq n$ **do**

$n_1, \dots, n_k \leftarrow$ any ordering of S_i

for $1 \leq j \leq k$ **do**

$\text{ISINKDOM}[n_j] \leftarrow n_{j+1 \bmod k}$

end

end

return ISINKDOM

end

Algorithm 7: An algorithm for the computation of some $<_{\text{SINK}}$.

Observation 5.4.2. Let G be any CFG. Then Algorithm 7 terminates with a result $<_{\text{SINK}}$ s.t. \succ_{SINK} is a transitive reduction of $\sqsupseteq_{\text{SINK}}$.

Remark 5.4.1. Wagner originally proposed to use the construction in Observation 5.4.1 to compute $\rightarrow_{\text{nticd}}$ directly, using any $\sqsubseteq_{\text{POST-control}}$ dependence algorithm on G_S . This is possible, as well.

Summary

- Nontermination sensitive and insensitive postdominance \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ can be defined uniformly, as the least and greatest fixed point of a single functional.
- \sqsubseteq_{MAX} and $\sqsubseteq_{\text{SINK}}$ transitively reduce not to trees, but to pseudoforests.
- These can be used to obtain least- and greatest fixed point algorithms.

6 Order Dependence

When you want to know how things really work,
study them when they're coming apart.

(William Gibson — Zero History)

In the context of this thesis, the ultimate purpose of control- and other dependencies is to statically establish information flow security. In fact it is well known that — given a CFG with unique exit node n_x — non-termination insensitive, batch-style *non-interference* is soundly approximated by *slicing* w.r.t standard control dependence \rightarrow_{cd} and data dependence $\rightarrow_{\text{data}}$ (see, e.g., [Was10]). Informally, in this setting

$$\text{IFC} = (\rightarrow_{\text{cd}} \cup \rightarrow_{\text{data}})^*$$

Nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$ and nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$ were explicitly designed for CFG without unique exit node n_x , in the extended setting of interactive, non-terminating programs. Ranganath et al. found that here, it is *not* enough to move from standard control dependence \rightarrow_{cd} to $\rightarrow_{\text{ntscd}}$ or $\rightarrow_{\text{nticd}}$. In this setting, the attacker does not only see the values of some program configurations at the termination of the program (i.e.: when control flow reaches n_x), but also observes events during execution of possibly never-terminating sections in the CFG (e.g., during the execution of control *sinks*). Specifically, the attacker observes the *order* of execution of some such nodes.

With regard to non-termination sensitive notions of IFC, in [Ran+07] the authors define the ternary notion of *decisive order dependence*, written $n \rightarrow_{\text{dod}} (m_1, m_2)$, which formalizes the notion that n controls the *order* in which two nodes m_1, m_2 are executed. Backward-Slicing w.r.t \rightarrow_{dod} then proceeds along the rule

$$\frac{m_1 \in S \quad m_2 \in S \quad n \rightarrow_{\text{dod}} (m_1, m_2)}{n \in S}$$

and in the nontermination sensitive setting, the authors show that (informally):

$$\text{IFC} = \text{BISIM} = (\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{dod}} \cup \rightarrow_{\text{data}})^*$$

by proving that any *sliced* program (w.r.t $\rightarrow_{\text{ntscd}}, \rightarrow_{\text{dod}}, \rightarrow_{\text{data}}$) is weakly *bi*-similar to its unsliced original program.

As for the nontermination insensitive setting, in [Amt08] the authors define a ternary notion of *weak order dependence*, written $n \rightarrow_{\text{wod}} (m_1, m_2)$, and then establish in this setting:

$$\text{SIM} = (\rightarrow_{\text{wod}} \cup \rightarrow_{\text{data}})^*$$

i.e.: they show that the sliced program (w.r.t $\rightarrow_{\text{wod}}, \rightarrow_{\text{data}}$) can weakly simulate the unsliced original program (but *not* the other way around, because the original program may get stuck in an infinite loop without observable actions, while the sliced program can proceed because the loop was sliced away).

In Section 6.1, I review \rightarrow_{dod} , and then observe how those regions of a graph in which \rightarrow_{dod} is non-empty can be characterized in terms of nontermination-sensitive postdominance $>_{\text{MAX}}$. This leads to practical algorithm for the computation of \rightarrow_{dod} .

Then in Section 6.2, I propose the new *nontermination sensitive order dependence* relation $\rightarrow_{\text{ntsod}}$ which can substitute for \rightarrow_{dod} during slicing, but is often smaller and more efficient to compute than \rightarrow_{dod} .

In Section 6.3, I quickly review the bisimulation soundness criterion for nontermination sensitive slicing in *labeled* control flow graphs from [Ran+07], and then propose a new *trace* based soundness criterion for nontermination sensitive slicing in *unlabeled* graphs. With regard to this criterion, the slices obtained from $\rightarrow_{\text{ntsod}}$ and \rightarrow_{dod} (together with $\rightarrow_{\text{ntscd}}$) will not only be sound but also *minimal*.

Later in Section 6.4, I will argue that \rightarrow_{wod} is not appropriate for the purpose of information flow control, both for principal and for practical reasons. In particular:

- There exist always-terminating programs for which the $\rightarrow_{\text{wod}}, \rightarrow_{\text{data}}$ backward slice of publicly “observable” nodes contain no private nodes, yet still, intuitively, private information is leaked.
- The relation \rightarrow_{wod} is very large even for CFG with unique exit node n_x .

Because it is difficult to formulate a trace-based soundness criterion for nontermination insensitive slicing, I first suggest in Section 6.5 a simple criterion based on the notion of *next observable* nodes.

Only then in Section 6.6 I offer my attempt to actually define a trace based criterion for nontermination *insensitive* slicing, using a notion of *infinite delay*.

In Section 6.7 I define *nontermination insensitive order dependence* ($\rightarrow_{\text{ntiod}}$) such that in the non-termination insensitive setting

$$\text{IFC} = \text{INFDEL} = \text{NEXTOBS} = (\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}} \cup \rightarrow_{\text{data}})^*$$

The relation $\rightarrow_{\text{ntiod}}$ will in practice be much smaller than \rightarrow_{wod} . In particular, and unlike \rightarrow_{wod} , the relation $\rightarrow_{\text{ntiod}}$ is empty for CFG with unique exit node n_x . The relation $\rightarrow_{\text{ntiod}}$ may still become large for CFG with large control-sinks, but at least I will provide an algorithm for the computation of $\rightarrow_{\text{ntiod}}$ that in practice scales no worse than this size. This algorithm will be based on $\text{SINK}_{\text{down}}$ of Algorithm 6 from Section 5.2.

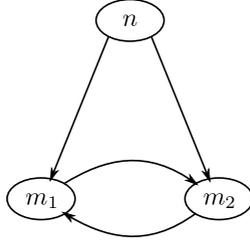


Figure 6.1: The canonical irreducible graph, where neither $n \rightarrow_{\text{ntscd}} m_1$ nor $n \rightarrow_{\text{ntscd}} m_2$.

6.1 Decisive Order Dependence

In [Ran+07], the authors note that for graphs such the canonical irreducible graph (Figure 6.1), nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$ is not enough to guarantee observational equivalence: there, neither $n \rightarrow_{\text{ntscd}} m_1$ nor $n \rightarrow_{\text{ntscd}} m_2$, but — assuming that both m_1 and m_2 are observable — the decision made at n will decide whether the sequence m_1, m_2, m_1, \dots or the sequence m_2, m_1, m_2, \dots is observed. In other words: An observer making either of these observations will learn what decision was made at n .

This is not merely a technical problem with $\rightarrow_{\text{ntscd}}$. In fact, *no* binary (dependence) relation can precisely capture observable equivalence for such graphs (also see: Observation 6.7.1 on page 99, or section 3.2 in [Ran+07]).

The authors remedy this by introducing the notion of *decisive order dependence*, a ternary relation \rightarrow_{dod} where $n \rightarrow_{\text{dod}} (m_1, m_2)$ iff n decisively decides in which order m_1 and m_2 may be observed. In order to state their formal definition, I need the following notation:

Definition 6.1.1. Given a path π , I write $m_1 \sqsupseteq_{m_2} \pi$ if π contains m_1 before any occurrence of m_2 , i.e.:

- $m_1 \in \pi$
- for the shortest prefix π_0 of π such that $m_1 \in \pi_0$: $m_2 \notin \pi_0$

Definition 6.1.2 (\rightarrow_{dod} , ([Ran+07])). Let G be any CFG, and n, m_1, m_2 be distinct nodes. Then m_1, m_2 are decisively order dependent on n , written: $n \rightarrow_{\text{dod}} (m_1, m_2)$, iff

- (a) All maximal paths from n contain both m_1 and m_2 , i.e.:

$$m_1 \sqsupset_{\text{MAX}} n \quad \text{and} \quad m_2 \sqsupset_{\text{MAX}} n$$

- (b) There exists some successor n_l of n such that all maximal paths π_l starting in n_l contain m_1 before any occurrence of m_2 , i.e.:

$$\exists n \rightarrow_G n_l. \forall \pi \in {}_{n_l}\Pi_{\text{MAX}}. m_1 \sqsupset_{m_2} \pi$$

- (c) There exists some successor n_r of n such that all maximal paths π_r starting in n_r contain m_2 before any occurrence of m_1 , i.e.:

$$\exists n \rightarrow_G n_r. \forall \pi \in {}_{n_r}\Pi_{\text{MAX}}. m_2 \sqsupset_{m_1} \pi$$

The authors proof that \rightarrow_{dod} is no larger than it needs to be, in the following sense:

Lemma 6.1.1 ([Ran+07]), Lemma 3). Let G be reducible. Then \rightarrow_{dod} is empty.

Unfortunately, this does not characterize the “regions” where triples $n \rightarrow_{\text{dod}} (m_1, m_2)$ may be found if G is *not* reducible. But such a characterization, presumably, is necessary for any efficient algorithm computing \rightarrow_{dod} . I now give such a characterization.

Lemma 6.1.2. Let G be any CFG.

- (i) Whenever $n \rightarrow_{\text{dod}} (m_1, m_2)$, then

$$m_1 \sqsupset_{\text{MAX}} m_2 \quad \text{and} \quad m_2 \sqsupset_{\text{MAX}} m_1$$

(ii) Whenever $n \rightarrow_{\text{dod}} (m_1, m_2)$, and $m \sqsupseteq_{\text{MAX}} n$ for $m \neq n$, then

$$m_1 \sqsupseteq_{\text{MAX}} m \quad \text{and} \quad m \sqsupseteq_{\text{MAX}} m_1$$

as well as

$$m_2 \sqsupseteq_{\text{MAX}} m \quad \text{and} \quad m \sqsupseteq_{\text{MAX}} m_2$$

(iii) Whenever $n \rightarrow_{\text{dod}} (m_1, m_2)$, then

$$\text{neither } n \sqsupseteq_{\text{MAX}} m_1 \text{ nor } n \sqsupseteq_{\text{MAX}} m_2$$

Proof: On page 358 in the appendix. □

As shown before in Section 5.2, \sqsupseteq_{MAX} can be efficiently computed and represented as transitive reduction $>_{\text{MAX}}$, which turned out to be a pseudo-forest. The situation in Lemma 6.1.2, rephrased in terms of $>_{\text{MAX}}$ — is depicted in Figure 6.2, and formally reads:

Corollary 6.1.1. Let G be any CFG, and $>_{\text{MAX}}$ any transitive reduction of \sqsupseteq_{MAX} .

For $n \rightarrow_{\text{dod}} (m_1, m_2)$, let $M = \{ m \mid m_1 <_{\text{MAX}}^* m \} = \{ m \mid m_2 <_{\text{MAX}}^* m \}$ be the $<_{\text{MAX}}$ -cycle both m_1, m_2 are part of. Then

- $n \notin M$
- $n <_{\text{MAX}} m$ for some $m \in M$.

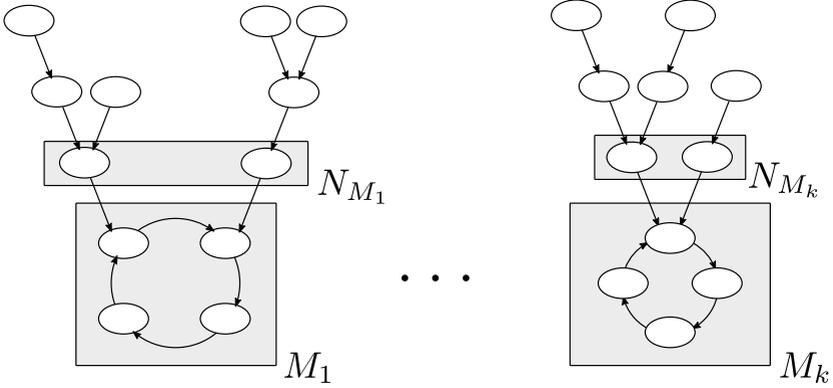


Figure 6.2: The situation in Corollary 6.1.1, exemplified. Given $m_1, m_2 \in M_i$, the set of n such that $n \rightarrow_{\text{dod}} (m_1, m_2)$ is contained in N_{M_i} .

6.1.1 Algorithms

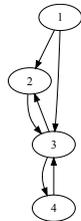
In [Ran+07], the authors propose a semi-naive algorithm (Figure 7, page 38) for the computation of \rightarrow_{dod} . As a subroutine, it implements a check $\text{DEP}(n, m_1, m_2)$ ¹ such that

$$\text{DEP}(n, m_1, m_2) \iff \text{clauses (b), (c) in Definition 6.1.2 hold}$$

The author then propose to compute² \rightarrow_{dod} as the set of all triples (n, m_1, m_2) such that

$$n \in \text{COND}, m_1 \rightarrow_G^* m_2, m_2 \neq m_1, m_2 \rightarrow_G^* m_1, \text{DEP}(n, m_1, m_2) \quad (6.1)$$

Let me first point out that this is incorrect in general. An example is shown on the right. Here, $\neg 3 \rightarrow_{\text{dod}} (2, 4)$, since while clauses (b) and (c) hold, clause (a) does *not*: Neither $2 \sqsupset_{\text{MAX}} 3$ nor $3 \sqsupset_{\text{MAX}} 2$.



¹ originally: DEPENDENCE

² via “generate-and-test” in the implied order

The obvious fix is to replace the checks

$$m_1 \rightarrow_G^* m_2, m_2 \rightarrow_G^* m_1 \quad \text{by} \quad m_1 \sqsupset_{\text{MAX}} n, m_2 \sqsupset_{\text{MAX}} n$$

and obtain

$$n \in \text{COND}, m_1, m_2 \sqsupset_{\text{MAX}} n, m_2 \neq m_1, \text{DEP}(n, m_1, m_2) \quad (6.2)$$

but the authors do not provide an algorithm to check (let alone: enumerate such m_i given n) this. Given my Algorithm 5 for the computation of some $<_{\text{MAX}}$, however, this is both trivial.

Furthermore, Corollary 6.1.1 allows me to significantly reduce the number of queries $\text{DEP}(n, m_1, m_2)$. Let $>$ be some transitive reduction of \sqsupset_{MAX} , and \mathbb{M} be the set of $<$ -cycles M , which are easily enumerated given $<$. Also easily enumerated for each M is the set $N_M = \{n \in \text{COND} \mid n \notin M, \exists m \in M. n < m\}$. The resulting scheme — which can be further optimized by noting that \rightarrow_{dod} is symmetric w.r.t m_1, m_2 — then is:

$$M \in \mathbb{M}, m_1 \neq m_2 \in M, n \in N_M, \text{DEP}(n, m_1, m_2) \quad (6.3)$$

Using my own implementation of the algorithm DEP (which I will not repeat, here), the empirical comparison of (6.2) and (6.3) confirms Lemma 6.1.2:

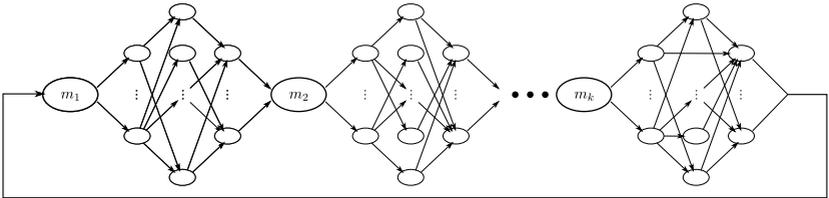
Observation 6.1.1. Given any CFG, the relation computed via 6.2 equals that computed via 6.3.

Given a CFG $G = (N, E)$, the complexity of DEP is, in [Ran+07], given as $\mathcal{O}(|E| \times |N| \times \log(|N|))$.

The original (flawed) scheme (6.1) requires $\Theta(|N|^3)$ queries to DEP , while for my scheme (6.3), the number of queries is highest if the maximal size

$$\max_{M \in \mathbb{M}} |M|$$

among $<_{\max}$ -cycles M is maximized (since each M and N_M are disjoint). Hence, the scheme (6.3) requires $\mathcal{O}(|N|^2)$ queries to DEP, but I claim that in CFG for virtually any programming language and any but the most untypical program, $\max_{M \in \mathbb{M}} |M|$ is much smaller than $|N|$. This is because any $<_{\max}$ -cycle $M = \{m_1, \dots, m_k\}$ is just the spine of a sub-cycle free control-sink, i.e., a G -subgraph of the form



6.2 Nontermination Sensitive Order Dependence

I do not attempt to further improve upon the computation scheme (6.3) for \rightarrow_{dod} I gave in Subsection 6.1.1.

Instead, in this section, I propose a new (indecisively) **nontermination sensitive order dependence** relation $\rightarrow_{\text{ntsod}}$ which, in the application of information flow control, can be substituted for \rightarrow_{dod} . Informally:

$$(\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{dod}})^* = (\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{ntsod}})^*$$

For this new relation $\rightarrow_{\text{ntsod}}$, I will give an efficient-in-practice algorithm. This will be possible because I can rephrase $\rightarrow_{\text{ntsod}}$ in terms of $\rightarrow_{\text{ntscd}}$ for certain subgraphs of G . Hence, I can once again use Algorithm 1 for the computation of suitable generalized postdominance frontiers, based on Algorithm 5 (or Algorithm 19).

Definition 6.2.1 ($\rightarrow_{\text{ntsod}}$). Let G be any CFG, and n, m_1, m_2 be distinct nodes. Then m_1 is **nontermination sensitively order dependent** on n with respect to m_2 , written: $n \rightarrow_{\text{ntsod}} (m_1, m_2)$, iff

- (a) All maximal paths from n contain both m_1 and m_2 , i.e.:

$$m_1 \sqsupseteq_{\text{MAX}} n \quad \text{and} \quad m_2 \sqsupseteq_{\text{MAX}} n$$

- (b) There exists some successor n_l of n such that all maximal paths π_l starting in n_l contain m_1 before any occurrence of m_2 , i.e.:

$$\exists n \rightarrow_G n_l. \quad \forall \pi \in {}_{n_l}\Pi_{\text{MAX}}. m_1 \sqsupseteq_{m_2} \pi$$

- (c) There exists some successor n_r of n such that *not* all maximal paths π_r starting in n_r contain m_1 before any occurrence of m_2 , i.e.:

$$\exists n \rightarrow_G n_r. \quad \neg \forall \pi \in {}_{n_r}\Pi_{\text{MAX}}. m_1 \sqsupseteq_{m_2} \pi$$

Note that (as opposed to \rightarrow_{dod}) this definition is *not* symmetric in m_1, m_2 . Also note that after holding fast m_2 , and subject to the constraint (a), the binary relation $\cdot \rightarrow_{\text{ntsod}} (\cdot, m_2)$ follows the scheme of \sqsubseteq control dependence (established in Definition 3.1.2), using $\sqsubseteq = \sqsubseteq_{\text{MAX}[m_2]}$ and the following notation:

Definition 6.2.2 ($m_1 \sqsubseteq_{\text{MAX}[m_2]} n$). Let G be some CFG.

$$m_1 \sqsubseteq_{\text{MAX}[m_2]}^G n \iff \forall \pi \in {}_n\Pi_{\text{MAX}}. m_1 \sqsubseteq_{m_2} \pi$$

A minor modification to the rule system D from Theorem 5.1.2 yields the following rule system:

Proposition 6.2.1. Let G be a CFG, m_2 any of it's nodes, and D_{m_2} be the rule-system

$$\frac{n \neq m_2 \quad D_{m_2}^{\text{self}}}{n \sqsubseteq n} \quad \frac{\forall p \rightarrow_G x. m \sqsubseteq x \quad p \rightarrow_G^* m \quad p \neq m_2 \quad D_{m_2}^{\text{suc}}}{m \sqsubseteq p}$$

Then $\sqsubseteq_{\text{MAX}[m_2]} = \mu D_{m_2}$.

6.2.1 Comparison with Decisive Order Dependence

The relation $\rightarrow_{\text{ntsod}}$ is meant to replace \rightarrow_{dod} . In fact, its symmetric core is just \rightarrow_{dod} :

Observation 6.2.1. Let G be any CFG. Then

$$n \rightarrow_{\text{dod}} (m_1, m_2) \iff n \rightarrow_{\text{ntsod}} (m_1, m_2) \wedge n \rightarrow_{\text{ntsod}} (m_2, m_1)$$

At the same time, $\rightarrow_{\text{ntsod}}$ is not too large for application to information flow control:

Observation 6.2.2. Let G be any CFG, and assume $n \rightarrow_{\text{ntsod}} (m_1, m_2)$, but $\neg n \rightarrow_{\text{ntsod}} (m_2, m_1)$. Then there exists some n' with both $n' \rightarrow_{\text{ntsod}} (m_1, m_2)$ and $n' \rightarrow_{\text{ntsod}} (m_2, m_1)$ such that

$$n \xrightarrow{*}_{\text{ntscd}} n'$$

This observation immediately implies that slicing backwards from any set M along $\rightarrow_{\text{nticd}}$, \rightarrow_{dod} is equivalent to slicing backwards along $\rightarrow_{\text{nticd}}$, $\rightarrow_{\text{ntsod}}$. I use the following notation to state this as Proposition 6.2.2:

Definition 6.2.3. Given a binary relation $\cdot \rightarrow \cdot$ and a ternary relation $\cdot \Rightarrow (\cdot, \cdot)$ on nodes N , their backward slice — written $(\rightarrow \cup \Rightarrow)^*$ — is the function mapping any set $M \subseteq N$ to the smallest set $S \supseteq M$ satisfying

$$\frac{n \rightarrow m \quad m \in S}{n \in S} \quad \frac{n \Rightarrow (m_1, m_2) \quad m_1 \in S \quad m_2 \in S}{n \in S}$$

I write both

$$\begin{aligned} & (\rightarrow)^* \text{ for } (\rightarrow \cup \emptyset)^* \\ \text{and } & (\Rightarrow)^* \text{ for } (\emptyset \cup \Rightarrow)^* \end{aligned}$$

The following result directly follows from Observation 6.2.1 and Observation 6.2.2.

Proposition 6.2.2.

$$(\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{dod}})^* = (\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{ntsod}})^*$$

Structural Properties

Structurally, the relation $\rightarrow_{\text{ntsod}}$ is constrained to the very same regions within $>_{\text{MAX}}$ that \rightarrow_{dod} is:

Observation 6.2.3. Lemma 6.1.2 still holds after replacing all occurrences of $n \rightarrow_{\text{dod}} (m_1, m_2)$ with $n \rightarrow_{\text{ntsod}} (m_1, m_2)$.

6.2.2 An Algorithm Based on Nontermination Sensitive Postdominance

Owing to Observation 6.2.3, I can limit the search for triples $n \rightarrow_{\text{ntsod}} (m_1, m_2)$ the same way I previously (in Subsection 6.1.1) did for \rightarrow_{dod} . But I can do more: for fixed m_2 , the relation $\cdot \rightarrow_{\text{ntsod}} (\cdot, m_2)$ is essentially $\sqsupseteq_{\text{MAX}[m_2]}$ -control dependence, with $\sqsupseteq_{\text{MAX}[m_2]}$ almost equal to \sqsupseteq_{MAX} :

$$\begin{aligned} \forall \pi \in n_s \Pi_{\text{MAX}}. m_1 \sqsupseteq_{m_2} \pi &\Leftrightarrow m_1 \sqsupseteq_{\text{MAX}[m_2]} n_s \\ \text{instead of } \forall \pi \in n_s \Pi_{\text{MAX}}. m_1 \in \pi &\Leftrightarrow m_1 \sqsupseteq_{\text{MAX}} n_s \end{aligned}$$

By choosing — given m_2 — a suitable subgraph of G , it is not hard to reduce $\sqsupseteq_{\text{MAX}[m_2]}$ in G to \sqsupseteq_{MAX} in the subgraph. To do this, I use the following notation:

Definition 6.2.4. Let $G = (N, E)$ be some CFG, $n, m \in N$, $M \subseteq N$. Then I define

$$\begin{aligned} G_{M \nrightarrow} &= (N, E \setminus \{ (m, n) \mid n \in N, m \in M \}) \\ G_{m \nrightarrow} &= G_{\{m\} \nrightarrow} \end{aligned}$$

In other words: $G_{M \nrightarrow}$ is the subgraph obtained from G by deleting all outgoing edges of nodes in M .

$$\begin{aligned} N^{\rightarrow^* M} &= \{ n \mid \exists m \in M. n \rightarrow_G^* m \} & N^{M \rightarrow^*} &= \{ n \mid \exists m \in M. m \rightarrow_G^* n \} \\ G^{\rightarrow^* M} &= (N^{\rightarrow^* M}, E|_{N^{\rightarrow^* M}}) & G^{M \rightarrow^*} &= (N^{M \rightarrow^*}, E|_{N^{M \rightarrow^*}}) \end{aligned}$$

$$\begin{aligned} N^{\rightarrow^* m} &= N^{\rightarrow^* \{m\}} & N^{m \rightarrow^*} &= N^{\{m\} \rightarrow^*} \\ G^{\rightarrow^* m} &= G^{\rightarrow^* \{m\}} & G^{m \rightarrow^*} &= G^{\{m\} \rightarrow^*} \end{aligned}$$

G^{\rightarrow^*M} is the subgraph of G consisting of those nodes that can reach some node in M , and $G^{M\rightarrow^*}$ is the subgraph of G consisting of those nodes that are reachable from some node in M .

$$\begin{aligned} G_{m\nrightarrow}^{\rightarrow^*M} &= \left(G^{\rightarrow^*M}\right)_{m\nrightarrow} & G^{M_1\rightarrow^*M_2} &= \left(G^{M_1\rightarrow^*}\right)^{\rightarrow^*M_2} \\ G_{m\nrightarrow}^{M\rightarrow^*} &= \left(G^{M\rightarrow^*}\right)_{m\nrightarrow} & G_{M'\nrightarrow}^{\rightarrow^*M} &= \left(G^{\rightarrow^*M}\right)_{M'\nrightarrow} \end{aligned}$$

$G^{M_1\rightarrow^*M_2}$ is the ‘‘chop’’ subgraph of G between M_1 and M_2 , and $G_{m\nrightarrow}^{M\rightarrow^*}$ is the graph obtained from G by retaining only those nodes that are reachable from some node in M , and deleting all outgoing edges of node m .

Observation 6.2.4. Let $G = (N, E)$ be any CFG, and $n_s, m_1, m_2 \in N$, $m_1 \neq m_2$. Then any maximal G -path starting in n_s contains m_1 before any occurrence of m_2 iff in the graph $G_{m_2\nrightarrow}$ obtained by removing all outgoing edges of m_2 , any maximal path starting in n_s contains m_1 , i.e.:

$$m_1 \sqsupset_{\text{MAX}[m_2]}^G n_s \iff m_1 \sqsupset_{\text{MAX}}^{G_{m_2\nrightarrow}} n_s$$

Combining Observation 6.2.3 and Observation 6.2.4 in a form suggesting an algorithm, I make the following observation:

Observation 6.2.5. Let $G = (N, E)$ be any CFG, $> = >_{\text{MAX}}$ any transitive reduction of \sqsupset_{MAX}^G , \mathbb{M} the set of $<$ -cycles, and let --- as before in Figure 6.2 — for $M \in \mathbb{M}$ the set of conditional nodes $n \notin M$ s.t. $n < m$ for some $m \in M$ be denoted by N_M .

Then, $n \rightarrow_{\text{ntsod}} (m_1, m_2)$ if and only if in the CFG

$$G_{m_2} := G_{m_2\nrightarrow}^{N_M\rightarrow^*M}$$

m_1 is $\rightarrow_{\text{ntsod}}$ -control dependent on n , where $M \in \mathbb{M}$ is the $<$ -cycle containing both m_1 and m_2 , i.e.:

$$n \rightarrow_{\text{ntsod}} (m_1, m_2) \iff \exists M \in \mathbb{M}. m_1, m_2 \in M \wedge n \in N_M \wedge n \rightarrow_{\text{ntsod}}^{G_{m_2}} m_1$$

Informally:

$$\rightarrow\text{ntsod} \approx \left(\sum_{M \in \mathbf{M}} |M| \right) \times \rightarrow\text{ntscd}$$

6.3 Soundness of Nontermination Sensitive Slices

In [Ran+07], the soundness criterion for nontermination sensitive slicing requires a weak bisimulation between the original and sliced program. The authors label CFG-nodes with statements reading and modifying *states* $\sigma \in \Sigma$, and assume a corresponding CFG-execution semantics $G \vdash s \rightarrow s'$ for $s, s' \in N \times \Sigma$. Given a set S of (observable) nodes, they differentiate between S -visible, and silent (τ) transitions:

Definition 6.3.1 ([Ran+07], Definition 20).

- for $s = (n, \sigma)$, $G \vdash s \xrightarrow{n} s'$ iff $G \vdash s \rightarrow s' \wedge n \in S$
- for $s = (n, \sigma)$, $G \vdash s \xrightarrow{\tau} s'$ iff $G \vdash s \rightarrow s' \wedge n \notin S$
- $G \vdash \xrightarrow{\tau} = G \vdash \xrightarrow{\tau}^*$
- $G \vdash s \xrightarrow{n} s''$ iff $G \vdash s \xrightarrow{\tau} s'$, $G \vdash s' \xrightarrow{n} s''$ for some s'

Furthermore, given a set S of nodes and statement-labeled CFG G , the authors define the corresponding *sliced* (labeled) CFG G_S such that labels remain unchanged for $n \notin S$, and become *no-ops*³ for $n \in S$.

Then, the authors show that, for S closed under $\rightarrow_{\text{ntscd}}$, $\rightarrow_{\text{data}}$, \rightarrow_{dod} , there exists some weak bisimulation R between G and G_S , i.e. a binary relation R on configurations s such that:

- If $s_1 R s_2$ and $G \vdash s_1 \xrightarrow{\tau} s'_1$, then $G_S \vdash s_2 \xrightarrow{\tau} s'_2$ for some s'_2 such that $s'_1 R s'_2$, and
if $s_1 R s_2$ and $G_S \vdash s_2 \xrightarrow{\tau} s'_2$, then $G \vdash s_1 \xrightarrow{\tau} s'_1$ for some s'_1 such that $s'_1 R s'_2$
- If $s_1 R s_2$ and $G \vdash s_1 \xrightarrow{n} s'_1$, then $G_S \vdash s_2 \xrightarrow{n} s'_2$ for some s'_2 such that $s'_1 R s'_2$, and

³ i.e.: operations that leave the state component σ unchanged

if $s_1 R s_2$ and $G_S \vdash s_2 \xrightarrow{n} s'_2$, then $G \vdash s_1 \xrightarrow{n} s'_1$ for some s'_1 such that $s'_1 R s'_2$

Thanks to Proposition 6.2.2, the same holds for S closed under $\rightarrow_{\text{ntscd}}$, $\rightarrow_{\text{data}}$, $\rightarrow_{\text{ntsod}}$, so I am almost ready to conclude my treatment of non-termination sensitive order dependencies.

Before I do that, however, I will observe that in the simplified setting of unlabeled CFG (i.e.: those without statements), slicing w.r.t $\rightarrow_{\text{ntscd}}$, $\rightarrow_{\text{ntsod}}$ (but not: $\rightarrow_{\text{data}}$, since data dependencies no longer make sense) satisfies a different *trace based* criterion.

The advantages of this criterion will be that

1. At least for very small CFG⁴ ($|N| \leq 25$), it can be exhaustively tested.
2. It will turn out that — with regards to this criterion — slices will not only be sound, but also *minimal*.⁵

Definition 6.3.2. Let G be any (unlabeled) CFG, and $N_x = \{n_x \mid \neg \exists m. n \rightarrow_G m\}$ be the set of exit nodes. An *input* i to G is a pair $i = (n_e, \epsilon)$ of a node n_e (the entry node) and a map $\epsilon : \text{COND}_G \rightarrow N$ (the choice made by i at n) such that $n \rightarrow_G \epsilon(n)$ for all conditional nodes n .

Given input i , the deterministic transition relation

- $G, i \vdash n \rightarrow n'$ if $n' = \epsilon(n)$ for $n \in \text{COND}_G$
- $G, i \vdash n \rightarrow n'$ if $n \rightarrow_G n'$ for $n \notin \text{COND}_G$

⁴ and sets of nodes S arising as backward slides from small slicing criteria $M \subseteq S$

⁵ presumably, $\rightarrow_{\text{ntscd}}$, $\rightarrow_{\text{data}}$, $\rightarrow_{\text{ntsod}}$ -slices are also minimal for statement-labeled CFG up to the *def-use abstraction*, i.e.: whenever $n \in S$, there exists *some* labeling for the same CFG with the same def- and use sets such that the newly-labeled CFG behaves observably differently for two observably-equivalent inputs. A corresponding result for syntax-tree based slicing is given [Dan+05], but I am unaware of any prior result with regard to the slicing of arbitrary CFG.

determines a unique (possibly infinite) path in G

$$\pi_i^G = n_e, \dots$$

such that $n \in \text{COND}_G$ implies $n' = \epsilon(n)$ for any two consecutive nodes n, n' .

Similarly, it determines a unique sequence t_i^G of *partial edges* $(n, n') \in E \cup (N_x \times \{\perp\})$ that is either finite with — for some exit node n_x —

$$t_i^G = (n_e, n_1), (n_1, n_2), \dots, (n_k, n_x), (n_x, \perp)$$

or infinite with

$$t_i^G = (n_e, n_1), (n_1, n_2), \dots$$

such that $n \in \text{COND}_G$ implies $n' = \epsilon(n)$ for any edge (n, n') . I call such sequences *traces*.

Remark 6.3.1. Given such an input i , the trace t_i in G is uniquely determined. This is not so different from models of labeled CFG with deterministic state transitions at non-conditional nodes, and mutually exclusive (and: usually exhaustive) choices (depending on state σ) made at conditional nodes. Usually in those models, the entry node n_e is fixed, and the input consists of an initial state σ_0 .

Definition 6.3.3. Given an input $i = (n_e, \epsilon)$ to G and a set S of (“observable”) nodes, the S -observable input $i|_S$ of i is the restriction $(n_e, \epsilon|_S)$ of the choices ϵ to nodes S . Furthermore, two inputs i, i' in G are called S -equivalent, and I write $i \sim_S i'$, iff

$$i|_S = i'|_S$$

Definition 6.3.4. Given a (possibly infinite) sequence t of partial edges, and a set S of (observable) nodes, the S -observable subsequence $t|_S$ of t is obtained from t by removing any occurrences

of (partial) edges $(n, _)$ s.t. $n \notin S$ from t (possibly transforming an infinite sequence into a finite).

Furthermore, two traces t, t' in G are called nontermination sensitively S -equivalent, and I write $t \sim_S t'$, iff

$$t|_S = t'|_S$$

Because in this “stateless” notion of input and execution of CFG, choices $\epsilon(n)$ made at $n \in \text{COND}$ are deterministic (depending only on n , and not on some state σ), any path π_i is either

- a finite path of the form

$$\pi_{\text{fin}} = n_e, \dots, n_x$$

for some n_x such that $\neg n_x \rightarrow_G$, or

- an infinite path of the form

$$\pi_{\text{inf}} = \underbrace{n_e, \dots, n}_{\pi_{\text{fin},0}}, \pi, \pi, \dots$$

with a finite prefix $\pi_{\text{fin},0}$ and an infinitely-repeating, finite and minimal cycle-segment π , i.e.: a path

$$\pi = n', \dots, m$$

with no node occurring twice, and such that $n \rightarrow_G n'$ and $m \rightarrow_G n'$

This fact allows me, given any input i , to finitely⁶ represent its execution path π_i via $\gamma(\pi_{\text{fin}}) = \pi_{\text{fin}}$ and $\gamma(\pi_{\text{fin},0}, \pi) = \pi_{\text{inf}}$. I can represent any execution *trace* t_i similarly.

⁶ and: uniquely

It is self-evident that, given such a finite representation of some t , the finite representation of $t|_S$ can be efficiently computed. Also, as will be required later, it can be efficiently checked whether $t \sqsubseteq_{\mathcal{T}} t'$, i.e.: whether t is a prefix of t' , and the finite representation of the concatenation $t t'$ can be computed efficiently, with $tt' = t$ if t is infinite.

In summary, I can confirm the following observation empirically:

Observation 6.3.1 (Soundness of $\rightarrow_{\text{ntsod}}, \rightarrow_{\text{ntscd}}$). Let G be any CFG, and $M \subseteq N$ a set of nodes (the slicing criterion). Let $S = (\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{ntsod}})^*(M)$ be the backward slice w.r.t M . Then, for any inputs i, i' such that

$$i \sim_S i'$$

I have

$$t_i \sim_S t_{i'}$$

Observation 6.3.2 (Minimality of $\rightarrow_{\text{ntsod}}, \rightarrow_{\text{ntscd}}$). Let G be any CFG, and $M \subseteq N$ a set of nodes (the slicing criterion). Let $S = (\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{ntsod}})^*(M)$ be the backward slice w.r.t M . Then, for any $n \in S, n \notin M$ and $S' = S \setminus \{n\}$, there exist inputs i, i' such that

$$i \sim_{S'} i'$$

but:

$$\neg t_i \sim_{S'} t_{i'}$$

The fact that $\rightarrow_{\text{ntscd}}, \rightarrow_{\text{ntsod}}$ slices are minimal also retrospectively establishes that in fact, *some* kind of ternary dependence relation is required, and there is not merely a problem with the binary relation $\rightarrow_{\text{ntscd}}$. To show why, I repeat in Figure 6.3 the canonical irreducible graph from Figure 6.1. Assume that there exists some binary relation \rightarrow_{xcd} such that for all slicing criteria M , the slice $(\rightarrow_{\text{xcd}})^*(M)$ is sound and minimal. For $M = \{m_1, m_2\}$, I then have $(\rightarrow_{\text{xcd}})^*(M) = \{n, m_1, m_2\}$. Since \rightarrow_{xcd} is binary, then necessarily $n \rightarrow_{\text{xcd}} m_1 \vee n \rightarrow_{\text{xcd}} m_2$. But if $n \rightarrow_{\text{xcd}} m_1$, then for the alternative

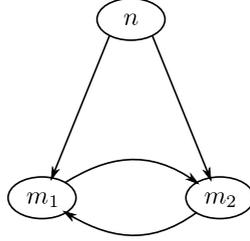


Figure 6.3: Neither $n \rightarrow_{\text{ntscd}} m_1$ nor $n \rightarrow_{\text{ntscd}} m_2$, but $n \rightarrow_{\text{ntsod}} (m_1, m_2)$ and $n \rightarrow_{\text{ntsod}} (m_2, m_1)$.

slicing criterion $M' = \{m_1\}$, I have $(\rightarrow_{\text{xccd}})^*(M') \supseteq \{n, m_1\}$, which is *not* minimal, since $(\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{ntsod}})^*(M') = \{m_1\}$.

Observation 6.3.3. Let $G = (N, E)$ be the canonical irreducible graph from Figure 6.1 on page 66. Then there exists no binary relation \rightarrow such that for all $M \subseteq N$

$$(\rightarrow)^*(M) = (\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{ntsod}})^*(M)$$

Remark 6.3.2. I must stress that when applying slicing (be it nontermination sensitive or not) for the purpose of information flow control, the set S of nodes deemed observable will usually *not* be implicitly defined to be the slice $S = (\dots)^*(M)$ of some criterion M . Instead, the set of observable nodes will remain fixed and fully determined by a user-provided specification, and *not* depend on the chosen slicing technique. For example, the user might provide a set L of nodes deemed public (or: *low* observable), and a disjoint set H of secret (or: *high*-observable) nodes. Slicing will then determine whether

$$h \in (\dots)^*(L)$$

for some $h \in H$, in which case I must expect some inputs i, i' with different choices $\epsilon(h) \neq \epsilon'(h)$ made at h made such that

$$i \sim_L i'$$

but:

$$\neg t_i \sim_L t_{i'}$$

If, on the other hand, no such h exists, I will conclude that any (L -observable) difference between $t_i|_L$ and $t_{i'}|_L$ is purely due to different choices $\epsilon(x) \neq \epsilon'(x)$ made at nodes $x \notin H$ (i.e.: that are *not* considered *high*), and hence that any L -observer cannot learn choices made at $h \in H$.

6.4 Weak Order Dependence

In this and the following sections, I turn from nontermination sensitive notions of order dependency to nontermination *insensitive* notions. The first such notion for arbitrary CFG was proposed in [Amt08]:

Definition 6.4.1 (\rightarrow_{wod}). Let G be any CFG, and n, m_1, m_2 be nodes. Then m_1, m_2 are **weakly order dependent** on n , written: $n \rightarrow_{\text{wod}} (m_1, m_2)$, iff

- (a) There exists some path

$$\pi_1 = n, \dots, m_1 \quad \text{with} \quad m_2 \notin \pi_1$$

- (b) There exists some path

$$\pi_2 = n, \dots, m_2 \quad \text{with} \quad m_1 \notin \pi_2$$

- (c) There exists some successor n' of n such that one of the following holds:

- m_1 is reachable from n' , and $m_1 \in \pi'_1$ for all paths

$$\pi'_1 = n', \dots, m_2$$

- m_2 is reachable from n' , and $m_2 \in \pi'_2$ for all paths

$$\pi'_2 = n', \dots, m_1$$

Note that $n \rightarrow_{\text{wod}} (m_1, m_2)$ implies that n, m_1, m_2 are distinct. Also note that — unlike \rightarrow_{dod} and $\rightarrow_{\text{ntsod}}$ —, nodes such that $n \rightarrow_{\text{wod}} (m_1, m_2)$ are *not* constraint to any particular regions of the CFG (e.g., $n \rightarrow_{\text{wod}} (m_1, m_2)$ does *not*, in general, imply that $m_1 \exists_{\text{SINK}} m_2$ and $m_2 \exists_{\text{SINK}} m_1$, or similar). In fact, \rightarrow_{wod} is designed to “cover” the whole CFG, in the sense that slicing is meant *not* to additionally re-

quire closure w.r.t $\rightarrow_{\text{nticd}}$ (unlike slicing with \rightarrow_{dod} , which additionally required closure under $\rightarrow_{\text{ntscd}}$).

Assuming a statement-labeled CFG (see: Definition 6.3.1 etc. pp.), the authors proceed to show that, for S closed under $\rightarrow_{\text{data}}$, \rightarrow_{wod} , there exists some weak simulation R between G and G_S (i.e.: the labeled graph obtained from G by replacing statements of nodes $n \notin S$ with no-ops):

- If $s_1 R s_2$ and $G \vdash s_1 \xRightarrow{n} s'_1$, then $G_S \vdash s_2 \xRightarrow{n} s'_2$ for some s'_2 such that $s'_1 R s'_2$.

The key difference between this result and the corresponding result for nontermination sensitive slicing is that there, the original and the sliced graph were (weakly) *bi*-similar, i.e.: any observable step in the original graph could be matched in the sliced graph *and vice versa*, while weakly similarity for nontermination insensitive slicing as just stated does *not* demand observable steps in the sliced graph to be matched in the original graph, because non-termination may have been “sliced away”.

While this simulation-based correctness-notion is appropriate for *batch-style* non-interference for CFG with unique exit node n_x (see, e.g., [Was10] for formal proof), and while some notions of non-interference that allow observations before a programs termination are similar, (e.g.: TINI in [Ask+08]), I find it unsatisfactory because under this notion, even always-terminating programs may leak (even: “more than a bit”). To see this, consider Figure 6.4a, and imagine this to be the CFG of the program that at each conditional node n_i tests $h == i$ for some secret input variable h . Also, let $M = \{m_1, \dots, m_k\}$ be the set of observable node. Then the observable trace of this CFG (assuming entry node $n_e = n_1$) is

$$(m_1, n_2), \dots, (m_i, n_{i+1}) \text{ iff the input to } h \text{ is } i,$$

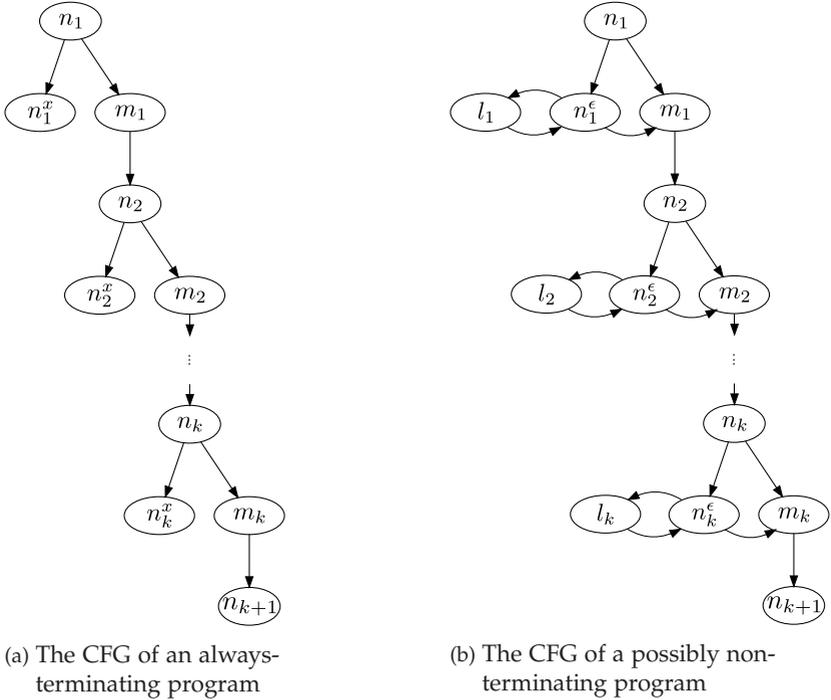


Figure 6.4: The CFG of a program that leaks more than a bit

but it is easy to see that the simulation criterion holds for $S = M$. In fact, for any pair $m_i, m_{i'} \in M$, there exist *no* $n \in N$ such that $n \rightarrow_{\text{wod}} (m_i, m_{i'})$, so the set

$$S = M = (\rightarrow_{\text{data}} \cup \rightarrow_{\text{wod}})^*(M)$$

is closed under $\rightarrow_{\text{data}}, \rightarrow_{\text{wod}}$.

Remark 6.4.1. The example in Figure 6.4a can be understood to paraphrase the example *Program 2* from [Ask+08], show in Figure 6.5. I essentially unrolled the `for` loop, and replaced the `while` loop with a

```
1 for i = 0 to k (  
2   output i on public_channel  
3   if (i = secret) then (while true do skip)  
4 )
```

Figure 6.5: Example from [Ask+08]

transition to an exit node n_i^x . In the light of this, I concede that my misgivings of the simulation criterion are possibly not fundamental, since it appears that any always-terminating, leaking program satisfying the simulation criterion that involves multiple exit nodes n_i^x can be transformed back into a program which any reasonable nontermination insensitive criterion must accept, simply by replacing n_i^x with a loop, exiting into m_i (see, e.g., Figure 6.4b).

6.5 Soundness of Nontermination Insensitive Slices

So how can I characterize what nodes – in my opinion – are missing in \rightarrow_{wod} slices? To do this, let me first recall (from [Amt08]) the notion and property of *next observable nodes* crucial for establishing the simulation property of \rightarrow_{wod} , $\rightarrow_{\text{data}}$ slices. That notion captures the following intuition: once an “unobservable region” is reached during execution, any decision made in this region (e.g., at nodes n within that region) shall have no influence on which other node is observed next when leaving this region — since otherwise, observing some node m instead of another (a priori possible) possible node m' will tell an observer which decision was made at n .

Definition 6.5.1 ([Amt08], Definition 3). Given a CFG G , a node n , and a set S of nodes, the set $\text{obs}_S^G(n)$ of next-observable nodes is the set of nodes $m \in S$ with the property that in G there exists a non-empty path $\pi = n_1 \dots, n_k$ with $n_1 = n$ and $n_k = m$ such that $n_i \notin S$ for $1 \leq i < k$.

Note that for $n \in S$, $\text{obs}_S^G(n) = \{n\}$.

Lemma 6.5.1 ([Amt08], Lemma 5). Let G be any CFG, and assume that S is closed under \rightarrow_{wod} . Then for all n in G , $\text{obs}_S^G(n)$ is at most singleton.

This property is so crucial to the correctness of slicing that other authors, in effect⁷, take it to be the *definition* of correctness.

Definition 6.5.2 (adapting [Dan+11], Definition 34 and Definition 9). A node $n \notin S$ is S -weakly committing in G if $\text{obs}_S^G(n)$ is at most a singleton.

Definition 6.5.3 (adapting [Dan+11], Definition 35). A set $S \subseteq N$ is weakly control closed in $G = (N, E)$ iff all vertices $n \notin S$ reachable from S are S -weakly committing in G .

⁷ see: [Dan+11], Theorem 45

The additional new requirement I want to pose — which, in general, does *not* hold for sets S closed under \rightarrow_{wod} (or, for that matter, about weakly control closed sets S) — is the following: Once execution reaches an S -unobservable region, no decision made in this region can influence whether the region can eventually be left, or not.

Formally — since obviously $\text{obs}_S^G(n') \subseteq \text{obs}_S^G(n)$ whenever $n \rightarrow_G n'$ for $n \notin S$ — I propose:

Definition 6.5.4. Given a CFG $G = (N, E)$, and a set of (“observable”) nodes S , I say that a node $n \notin S$ *retains all possible next observations* iff

$$\text{obs}_S^G(n) = \text{obs}_S^G(n')$$

for all nodes n' such that $n \rightarrow_G n'$.

I say that *all possible next observations are retained outside S* iff this is the case for all $n \notin S$.

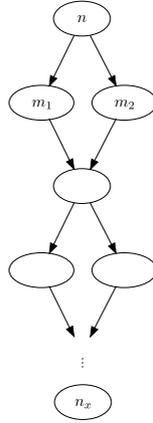
Recalling Figure 6.4a for the set $S = M = \{m_1, \dots, m_k\}$ closed under \rightarrow_{wod} , you will find that S is weakly control closed, but that *not* all possible next observations are retained outside S . For example, $\text{obs}_S^G(n_1) = \{m_1\}$, but after the step $n_1 \rightarrow_G n_1^x$, I have $\text{obs}_S^G(n_1) = \emptyset$.

On the other hand in Figure 6.4a for the same set $S = M = \{m_1, \dots, m_k\}$, not only is S weakly control closed, but also all possible next observations are retained outside S , since $\text{obs}_S(n_i^\epsilon) = \text{obs}_S(l) = \text{obs}_S(m_i) = \{m_i\}$.

In general, the following holds:

Observation 6.5.1. If all possible next observations are retained outside S , then S is weakly control closed.

I propose that a slice S is to be deemed nontermination insensitively sound if all possible next observations are retained outside S . This criterion has a somewhat *syntactic* flavor: it does not explicitly reference some notion of execution of the graph G . Instead, it is directly stated

Figure 6.6: A CFG with infeasibly large \rightarrow_{wod} .

in terms of reachability in the graph. Contrast this with Section 6.3. There, in the nontermination sensitive case, I gave a notion of input and executions (“traces”), and used a notion of trace-equivalence to give a soundness (and minimality) criterion for control slices. In other words, I gave a *semantic* criterion.

In an attempt to obtain a *semantic* criterion for control slices also in the nontermination *insensitive* case, in Section 6.6 I propose a similar, trace-based notion of execution in arbitrary graphs.

Then in Section 6.7, I will propose the new dependency notion $\rightarrow_{\text{ntiod}}$. Slices S w.r.t. $\rightarrow_{\text{nticd}}$, $\rightarrow_{\text{ntiod}}$ will not only be weakly control closed, but all possible next observations outside S will be retained, as well. They will also be minimal along all such slices. More importantly, though, $\rightarrow_{\text{ntiod}}$ will address the following criticism of \rightarrow_{wod} as a foundation for practical slicing: Since it “covers” the whole CFG, and since it is a ternary relation, explicit representations of \rightarrow_{wod} become infeasibly large even for graphs with unique exit node n_x . To get an Idea, consider Figure 6.6. Not only do I have $n \rightarrow_{\text{wod}} (m_1, m_2)$, but in fact $n \rightarrow_{\text{wod}} (m_1, m'_2)$ for all $m'_2 \notin \{n, m_1\}$.

6.6 A Trace-Based Notion of Infinite Delay

What is an appropriate “semantic” notion of nontermination insensitive slicing correctness? In this chapter, I will give a notion which I find satisfying for the “stateless” notion⁸ of input and execution for CFG from Section 6.3. I consider this result preliminary, if only for the fact that it is not obvious to me how this notion can be generalized to a more traditional stateful notion of input and execution.

Reconsider Figure 6.4b. In the stateless notion of execution, consider inputs $i = (n_1, \epsilon)$ with entry node n_1 . Then, if $\epsilon(n_j) = m_j$ for all j , the execution trace t_i is

$$(n_1, m_1), (m_1, n_2), \dots, (n_k, m_k), (m_k, n_{k+1}), (n_{k+1}, \perp)$$

If, on the other hand, $\epsilon(n_j) = n_j^\epsilon$ for some j , the choice $\epsilon(n_j^\epsilon)$ made at n_j^ϵ determines whether an execution reaching n_j can continue towards m_j , or if it *infinitely delays* this continuation by choosing $\epsilon(n_j^\epsilon) = l_j$. Suppose, again, that $M = \{m_1, \dots, m_k\}$ is the set of observable nodes. Then — as was the case for Figure 6.4a — any observable trace is of the form

$$t_i|_M = (m_1, n_2), \dots, (m_j, n_{j+1})$$

with

$$\begin{array}{ll} j = k \text{ and} & t_i = \dots, (m_j, n_{j+1}), \left(n_{k+1}, n_{k+1}^\epsilon \right) \\ \text{or} & t_i = \dots, (m_j, n_{j+1}), \left(n_{j+1}, n_{j+1}^\epsilon \right), t, t, t, \dots \end{array}$$

and

$$t = \left(n_{j+1}^\epsilon, l_{j+1} \right), \left(l_{j+1}, n_{j+1}^\epsilon \right)$$

⁸ i.e.: with execution steps between configurations consisting only of nodes n instead of pairs (n, σ) with variable state σ .

Also: the only (“up to infinite delay”) possible observable continuations of $t_i|_M$ are

$$\begin{aligned}
 t'|_M &= t_i|_M, (m_{j+1}, n_{j+2}) \\
 t''|_M &= t_i|_M, (m_{j+1}, n_{j+2}), (m_{j+2}, n_{j+3}) \\
 &\vdots \\
 &= t_i|_M, (m_{j+1}, n_{j+2}), (m_{j+2}, n_{j+3}), \dots, (m_k, n_{k+1})
 \end{aligned}$$

and form a prefix-chain

$$t_i|_M \sqsubseteq_{\mathcal{T}} t'|_M \sqsubseteq_{\mathcal{T}} t''|_M \sqsubseteq_{\mathcal{T}} \dots$$

Let me — by example — explain how these infinitely delayed observations can be explained. Consider the observation $t'|_M$. It can be obtained as the concatenation of the finite observation $t_i|_M$ with an observation $t_{i'}|_M$ determined by an input $i' = (n'_e, \epsilon')$ which “breaks out of the loop”: i' starts the execution at some node n'_e in the loop L which infinitely delays the observation (m_{j+1}, n_{j+2}) , i.e.:

$$n'_e \in L = \{n_{j+1}^\epsilon, l_{j+1}\}$$

and then makes choices ϵ' *observably-consistent* with those made by i , i.e.

$$\epsilon'|_M = \epsilon|_M$$

For example:

$$\epsilon'(n) = \begin{cases} m_{j+1} & \text{for } n = n_{j+1}^\epsilon \\ n_{j+2}^\epsilon & \text{for } n = n_{j+2} \\ l_{j+2} & \text{for } n = n_{j+2}^\epsilon \\ \epsilon(n) & \text{otherwise} \end{cases}$$

where the first choice breaks out of the loop L , while the next two choices force the execution to remain in the *following* loop $L' = \{n_{j+2}^\epsilon, l_{j+2}\}$.

Definition 6.6.1. Let G be any CFG, M be any set (of observable nodes), and $i = (n_e, \epsilon)$ any input to G . Then the set $\mathcal{T}_i^{\omega, G}|_M$ (or just $\mathcal{T}_i^\omega|_M$) of observable behavior up to infinite delay is defined to be

$$\left\{ t_i|_M \right\} \quad \text{if } t_i \text{ is finite}$$

$$\left\{ t_i|_M t_{i'}|_M \mid i' = (n'_e, \epsilon'), \epsilon'|_M = \epsilon|_{M'}, n'_e \in t \right\} \quad \text{if } t_i \text{ is infinite}$$

where for infinite t_i :

$$t_i = (n_e, n'), \dots, t, t, t, \dots$$

for some finite prefix followed by the infinitely repeating cycle t .

Obviously, $t_i|_M \in \mathcal{T}_i^\omega|_M$ and $t_i|_M \sqsubseteq_{\mathcal{T}} t$ for all $t \in \mathcal{T}_i^\omega|_M$.

Definition 6.6.2. Let G be any CFG, M be any set (of observable nodes), and i, i' two inputs G . Then input i, i' are said to have M -equivalent observable behavior *up to infinite delay* iff the observable behavior up to infinite delay of i and i' form non-disjoint $\sqsubseteq_{\mathcal{T}}$ -ascending chains.

Formally: $i \sim_{\mathcal{T}_M}^\omega i'$ iff

- (a) $\mathcal{T}_i^\omega|_M = \{t_1, t_2, t_3, \dots\}$ with $t_1 \sqsubseteq_{\mathcal{T}} t_2 \sqsubseteq_{\mathcal{T}} t_3 \sqsubseteq_{\mathcal{T}} \dots$
- (b) $\mathcal{T}_{i'}^\omega|_M = \{t'_1, t'_2, t'_3, \dots\}$ with $t'_1 \sqsubseteq_{\mathcal{T}} t'_2 \sqsubseteq_{\mathcal{T}} t'_3 \sqsubseteq_{\mathcal{T}} \dots$ and
- (c) $\mathcal{T}_i^\omega|_M \cap \mathcal{T}_{i'}^\omega|_M \neq \emptyset$

The requirement item (c) appears to be remarkably permissive. Why can I not demand $\mathcal{T}_i^\omega|_M = \mathcal{T}_{i'}^\omega|_M$? Or, for example,

$$\mathcal{T}_i^\omega|_M \subseteq \mathcal{T}_{i'}^\omega|_M \quad \vee \quad \mathcal{T}_{i'}^\omega|_M \subseteq \mathcal{T}_i^\omega|_M$$

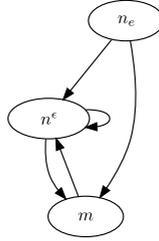


Figure 6.7: Infinite observable behavior up to infinite delay

To see why, consider Figure 6.7, for $M = \{m\}$, $i = (n_e, \epsilon)$, $i' = (n_e, \epsilon')$ and

$$\begin{aligned} \epsilon(n_e) &= n^\epsilon & \epsilon(n^\epsilon) &= n^\epsilon \\ \epsilon'(n_e) &= m & \epsilon'(n^\epsilon) &= n^\epsilon \end{aligned}$$

and thus

$$t_i|_M = \sqcup \quad \text{and} \quad t_{i'}|_M = (m, n^\epsilon)$$

Intuitively, both inputs have just the same observable behavior up to infinite delay: an infinite sequence $t = (m, n^\epsilon), (m, n^\epsilon), \dots$. But formally

$$\mathcal{T}_i^\omega|_M = \{t, \sqcup\} \quad \text{and} \quad \mathcal{T}_{i'}^\omega|_M = \{t, (m, n^\epsilon)\}$$

Neither are these sets of observable traces the same, nor is one a subset of the other.

6.6.1 Other Criteria for Nontermination Insensitive Slicing

Aside from my two new notions (Definition 6.5.4 in Section 6.4 and Definition 6.6.2 in this section) and the simulation criterion (Section 6.4), several other approaches defining correctness for nontermination insensitive slicing are possible.

In his thesis [Gif12], Dennis Giffhorn, too, proposes a notion of infinite delay in order to define low-equivalence between traces. Unlike the notions I developed in this section, Giffhorn's definitions resort to the notion of dynamic control dependence, which, in turn resorts to (standard) control-dependence for CFG with unique end node. Hence, his notions *presuppose* that (dynamic) control dependence are appropriate to specify the meaning of *infinite delay*. I, on the other hand, developed the notions in this chapter in order to establish the adequacy of the notion of $\rightarrow_{\text{ntiod}}$ (to be defined in the following section) for the application of slicing of arbitrary CFG, and hence needed to give a criterion only based on observable traces $t_i|_M$.

Other notions are based on *transfinite* semantics of programs. In such semantics, for example, the program

```

1 x = 0;
2 while(true) {
3   x = 1;
4 }
5 x = 2;
```

might be assigned the transfinite state-sequence

$$\underbrace{[x \mapsto 0], [x \mapsto 1], [x \mapsto 1], \dots, [x \mapsto 2]}_{\omega}$$

corresponding to the ordinal $\omega + 1$. Correctness of nontermination insensitive slicing then demands some correspondence between the transfinite semantics of the original and the transfinite semantics of the sliced program. In the example, slicing w.r.t x at the program's end and obtaining the program $x = 2$ with the semantics $[x \mapsto 2]$ is considered correct. For details, see, e.g., [GM03], and [Bar+10] for a critique.

In fact, my notion $\mathcal{T}_i^\omega|_M$ of infinite delay can be understood as an ad-hoc “finitization” of a poor-man's transfinite semantic: instead of considering actual transfinite sequences, I consider all possibilities of short-cutting loops in and behind an infinite trace t_i with finite observation $t_i|_M$.

6.7 Nontermination Insensitive Order Dependence

In this section, I develop a nontermination *insensitive* notion of order dependence. At first glance, this section is just a variation on Section 6.2, in which I developed a new notion of nontermination sensitive order dependence. Following the least-/greatest fixed point duality (Section B.2), I will mostly replace

| | | |
|---------------------------------------|------|--|
| nontermination sensitivity | with | nontermination sensitivity |
| μ | with | ν |
| \sqsupseteq_{MAX} | with | $\sqsupseteq_{\text{SINK}}$ |
| $>_{\text{MAX}}$ | with | $>_{\text{SINK}}$ |
| $\rightarrow_{\text{ntscd}}$ | with | $\rightarrow_{\text{nticd}}$ |
| $m_1 \sqsupseteq_{\text{MAX}[m_2]} n$ | with | $m_1 \sqsupseteq_{\text{SINK}[m_2]} n$ |
| $\rightarrow_{\text{ntsod}}$ | with | $\rightarrow_{\text{ntiod}}$ |

However, there will also be the following differences w.r.t. Section 6.2:

- The nodes n s.t. $n \rightarrow_{\text{ntiod}}(m_1, m_2)$ will *not* be restricted to nodes N_M “near the roots M of $<_{\text{SINK}}$ ”, but also appear *within* those roots M s.t. $m_1, m_2 \in M$.
- It is *not* plausible to argue that such roots M (i.e.: such $<_{\text{SINK}}$ cycles) are, in practice, small. In fact, given any CFG G_0 with unique entry and exit nodes n_e, n_x , the graph G with edges $E = E_0 \cup (n_e, n_x)$ will yield a $>_{\text{SINK}}$ -cycle $M = N$.
- Since with regards to the computation of $\rightarrow_{\text{ntiod}}$, I will (“again”) have

$$\rightarrow_{\text{ntiod}} \approx \left(\sum_{M \in \mathbb{M}} |M| \right) \times \rightarrow_{\text{nticd}}$$

I devise schemes that allow me — given $m_2 \in M$ and the postdominance-tree $>_{\text{SINK}}^{G_{m_2}}$ for one subgraph G_{m_2} — to compute

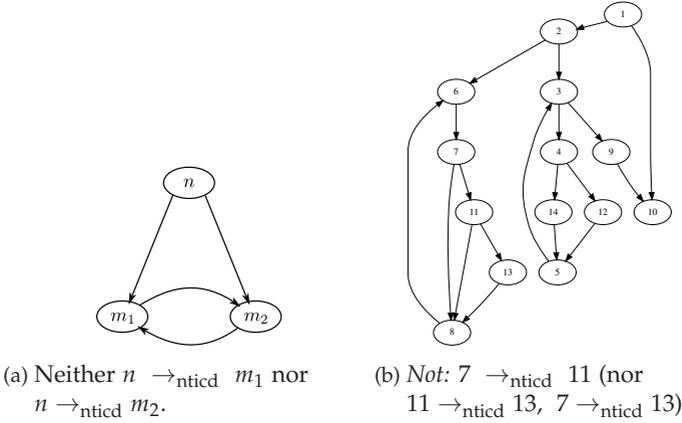


Figure 6.8: $\rightarrow_{\text{nticd}}$ is too small with respect to nontermination insensitive slicing

the postdominance-tree $\succ_{\text{SINK}}^{G_{m'_2}}$ for some other subgraph $G_{m'_2}$ incrementally, by (minimally) modifying $\succ_{\text{SINK}}^{G_{m_2}}$.

As argued in Section 6.4, I reject slicing w.r.t. \rightarrow_{wod} on the grounds that such slices are *too small*, since outside such slices, in general, possible next observations are *not* retained. My goal can only be then to devise a new notion $\rightarrow_{\text{ntiod}}$ such that

$$(\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}})^* \supseteq (\rightarrow_{\text{wod}})^*$$

but not *much* larger.

The fact that *some* ternary relation is required can (as earlier for non-termination sensitive slicing) be demonstrated via the canonical irreducible CFG (Figure 6.8a), but also in the reducible CFG from Figure 6.8b. Specifically, $\rightarrow_{\text{nticd}}$ is too small, but even more: *no* binary (dependence) relation \rightarrow_{xcd} precisely characterizes nontermination sensitive slicing in those CFG. Consider Figure 6.8b. Intuitively, I would want $7 \rightarrow_{\text{xcd}} 11$. But if 11 is the *only* observable node, then *all* observable traces reaching 6 are of the form $11, 11, \dots$, and hence a slice

including 7 would be imprecise. On the other hand I would *need* something like $7 \rightarrow_{\text{xcd}} 11$ as soon as, for example, node 8 is also observable.

Observation 6.7.1. Let $G = (N, E)$ be either of the CFG from Figure 6.8 (or: the subgraph $G|_{\{6,7,8,11,13\}}$ of Figure 6.8b). Then there exists no binary relation \rightarrow_{xcd} such that for all $M \subseteq N$

$$(\rightarrow_{\text{xcd}})^*(M) = (\rightarrow_{\text{wod}})^*(M)$$

The same observation will hold for \rightarrow_{wod} replaced with $\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}}$.

Definition 6.7.1 ($\rightarrow_{\text{ntiod}}$, corresponding to Definition 6.2.1). Let G be any CFG, and n, m_1, m_2 be distinct nodes. Then m_1 is **nontermination insensitively order dependent** on n with respect to m_2 , written: $n \rightarrow_{\text{ntiod}}(m_1, m_2)$, iff

- (a) All sink paths from n contain both m_1 and m_2 , i.e.:

$$m_1 \sqsupset_{\text{SINK}} n \quad \text{and} \quad m_2 \sqsupset_{\text{SINK}} n$$

- (b) There exists some successor n_l of n such that all sink paths π_l starting in n_l contain m_1 before any occurrence of m_2 , i.e.:

$$\exists n \rightarrow_G n_l. \quad \forall \pi \in n_l \Pi_{\text{SINK}}. m_1 \sqsupset_{m_2} \pi$$

- (c) There exists some successor n_r of n such that *not* all sink paths π_r starting in n_r contain m_1 before any occurrence of m_2 , i.e.:

$$\exists n \rightarrow_G n_r. \quad \neg \forall \pi \in n_r \Pi_{\text{SINK}}. m_1 \sqsupset_{m_2} \pi$$

Again (and as opposed to \rightarrow_{wod}), this definition is *not* symmetric in m_1, m_2 . Also note again that after holding fast m_2 , and subject to the constraint (a), the relation $\cdot \rightarrow_{\text{ntiod}}(\cdot, m_2)$ follows the scheme of \sqsupset control dependence (see Definition 3.1.2).

Definition 6.7.2 ($m_1 \sqsupseteq_{\text{SINK}[m_2]} n$, corresponding to Definition 6.2.2). Let G be some CFG.

$$m_1 \sqsupseteq_{\text{SINK}[m_2]}^G n \Leftrightarrow \forall \pi \in {}_n\Pi_{\text{SINK}}. m_1 \sqsupseteq_{m_2} \pi$$

As expected, $\sqsupseteq_{\text{MAX}[m_2]}$ is not the least but the greatest fixed point of a suitable rule system.

Proposition 6.7.1. Let G be a CFG, m_2 any of its nodes, and D_{m_2} be the rule-system from Proposition 6.2.1 on page 73. Then $\sqsupseteq_{\text{SINK}[m_2]} = \nu D_{m_2}$.

Comparison with Weak Order Dependence

The relation $\rightarrow_{\text{ntiod}}$ (together with $\rightarrow_{\text{nticd}}$) is meant to replace \rightarrow_{wod} . In fact, it is a subset:

Observation 6.7.2. Let G be any CFG. Then

$$n \rightarrow_{\text{ntiod}} (m_1, m_2) \Rightarrow n \rightarrow_{\text{wod}} (m_1, m_2)$$

Remember that \rightarrow_{wod} “covers” the whole CFG, while $\rightarrow_{\text{ntiod}}$ is meant to “cover” only those regions not covered by $\rightarrow_{\text{nticd}}$, i.e.: the control sinks. Also remember that \rightarrow_{wod} is symmetric, while $\rightarrow_{\text{ntiod}}$ is not.

In the other direction the situation is as follows:

Observation 6.7.3. Let G be any CFG, and assume

$$n \rightarrow_{\text{wod}} (m_1, m_2)$$

Then

$$\begin{aligned} n \rightarrow_{\text{ntiod}} (m_1, m_2) &\vee n \rightarrow_{\text{nticd}}^* m_1 \\ \vee n \rightarrow_{\text{ntiod}} (m_2, m_1) &\vee n \rightarrow_{\text{nticd}}^* m_2 \end{aligned}$$

Corollary 6.7.1.

$$(\rightarrow_{\text{wod}})^* \subseteq (\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}})^*$$

For example in Figure 6.4a on page 87 and with $M = \{m_1, \dots\}$, the slice $(\rightarrow_{\text{wod}})^*(M)$ contains only the nodes M , while the slice $(\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}})^*(M)$ additionally contains the nodes $\{n_1, \dots\}$. This is required by Definition 6.5.4, since if these nodes were missing, the slice would not retain next-observable nodes. Similarly in Figure 6.8b and with $M = \{8, 11\}$, the slice $(\rightarrow_{\text{wod}})^*(M)$ contains M and node 7, while $(\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}})^*(M)$ additionally contains the nodes $\{1, 2\}$.

Structural Properties

In support of algorithms for nontermination insensitive order dependence $\rightarrow_{\text{ntiod}}$, I now characterize the regions in which $\rightarrow_{\text{ntiod}}$ is non-empty.

Observation 6.7.4 (corresponding to a weakening of Lemma 6.1.2). Let G be any CFG.

- (i) Whenever $n \rightarrow_{\text{ntiod}} (m_1, m_2)$, then

$$m_1 \sqsupseteq_{\text{SINK}} m_2 \quad \text{and} \quad m_2 \sqsupseteq_{\text{SINK}} m_1$$

- (ii) Whenever $n \rightarrow_{\text{ntiod}} (m_1, m_2)$, and $m \sqsupseteq_{\text{SINK}} n$ for $m \neq n$. Then

$$m_1 \sqsupseteq_{\text{SINK}} m \quad \text{and} \quad m \sqsupseteq_{\text{SINK}} m_1$$

as well as

$$m_2 \sqsupseteq_{\text{SINK}} m \quad \text{and} \quad m \sqsupseteq_{\text{SINK}} m_2$$

As shown before, $\sqsupseteq_{\text{SINK}}$ can be efficiently computed and represented as transitive reduction $>_{\text{SINK}}$, which turned out to be a pseudo-forest.

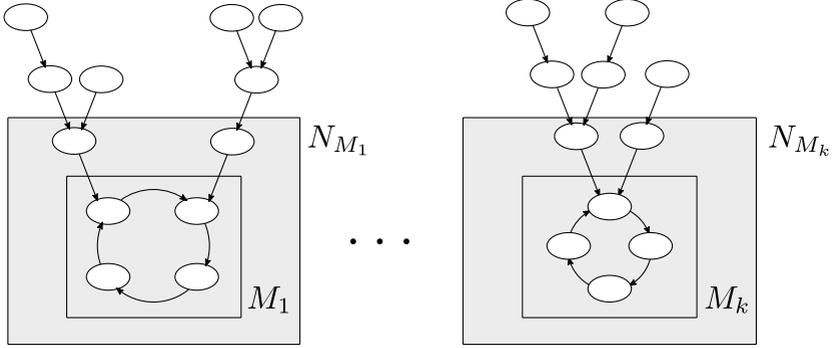


Figure 6.9: The situation in Corollary 6.7.2, exemplified. Given $m_1, m_2 \in M_i$, the set of n such that $n \rightarrow_{\text{ntiod}} (m_1, m_2)$ is contained in N_{M_i} .

The situation in Observation 6.7.4, rephrased in terms of $>_{\text{SINK}}$ — is depicted in Figure 6.9, and formally reads:

Corollary 6.7.2 (corresponding to a weakening of Corollary 6.1.1). Let G be any CFG, and $>_{\text{SINK}}$ any transitive reduction of $\sqsupseteq_{\text{SINK}}$.

For $n \rightarrow_{\text{ntiod}} (m_1, m_2)$, let $M = \{ m \mid m_1 <_{\text{SINK}}^* m \} = \{ m \mid m_2 <_{\text{SINK}}^* m \}$ be the $<_{\text{SINK}}$ -cycle both m_1, m_2 are part of. Then

- $n <_{\text{SINK}} m$ for some $m \in M$.

6.7.1 An Incremental Algorithm Based on Nontermination Insensitive Postdominance

Owing to Observation 6.7.4, I can limit the search for triples $n \rightarrow_{\text{ntiod}} (m_1, m_2)$ the same way I previously did for \rightarrow_{dod} . But, just as was the case for $\rightarrow_{\text{ntsod}}$, I can do more: for fixed m_2 , the rela-

tion $\cdot \rightarrow_{\text{ntiod}}(\cdot, m_2)$ is essentially $\sqsupset_{\text{SINK}[m_2]}$ -control dependence, with $\sqsupset_{\text{SINK}[m_2]}$ *almost equal* to \sqsupset_{SINK} : I have

$$\begin{aligned} \forall \pi \in {}_{n_s} \Pi_{\text{SINK}}. m_1 \sqsupset_{m_2} \pi &\Leftrightarrow m_1 \sqsupset_{\text{SINK}[m_2]} n_s \\ \text{instead of } \forall \pi \in {}_{n_s} \Pi_{\text{SINK}}. m_1 \in \pi &\Leftrightarrow m_1 \sqsupset_{\text{SINK}} n_s \end{aligned}$$

Again, by choosing — given m_2 — a suitable subgraph G , it is not hard to reduce $\sqsupset_{\text{SINK}[m_2]}$ in G to \sqsupset_{SINK} in the subgraph.

Observation 6.7.5 (corresponding to Observation 6.2.4). Let $G = (N, E)$ be any CFG, and $n_s, m_1, m_2 \in N$, $m_1 \neq m_2$. Then any G -sink-path starting in n_s contains m_1 before any occurrence of m_2 iff in the graph $G_{m_2 \not\rightarrow}$ obtained by removing all outgoing edges of m_2 , any sink-path starting in n_s contains m_1 , i.e.:

$$m_1 \sqsupset_{\text{SINK}[m_2]}^G n_s \quad \Leftrightarrow \quad m_1 \sqsupset_{\text{SINK}}^{G_{m_2 \not\rightarrow}} n_s$$

Again, combining Observation 6.7.4 and Observation 6.7.5 in a form suggesting an algorithm, I make the following observation:

Observation 6.7.6 (corresponding to Observation 6.2.5). Let $G = (N, E)$ be any CFG, $> = >_{\text{SINK}}$ any transitive reduction of \sqsupset_{SINK} , \mathbb{M} the set of $<$ -cycles (i.e.: the set of control sinks), and let for $M \in \mathbb{M}$ the set of conditional nodes n s.t. $n < m$ for some $m \in M$ be denoted by N_M .

Then, $n \rightarrow_{\text{ntiod}}(m_1, m_2)$ if and only if in the CFG

$$G_{m_2} := G_{m_2 \not\rightarrow}^{N_M \rightarrow^* M}$$

$m_1 \neq n$ is $\rightarrow_{\text{ntiod}}$ -control dependent on n , where $M \in \mathbb{M}$ is the $<$ -cycle containing both m_1 and m_2 , i.e.:

$$\begin{aligned} n \rightarrow_{\text{ntiod}}(m_1, m_2) &\Leftrightarrow \exists M \in \mathbb{M}. m_1, m_2 \in M \wedge n \in N_M \wedge n \rightarrow_{\text{ntiod}}^{G_{m_2}} m_1 \\ &\wedge m_1 \neq n \end{aligned}$$

Informally:

$$\rightarrow\text{ntiod} \approx \left(\sum_{M \in \mathbb{M}} |M| \right) \times \rightarrow\text{nticd}$$

Note that in Observation 6.2.5, $m_1 \neq n$ was implied by $m_1 \in M$, but this is not the case here, since here N_M , does not exclude M .

As argued in the opening words of this section, it is not unreasonable to expect large $M \in \mathbb{M}$, e.g.: $M = N$. A naive algorithm will simply compute $\rightarrow\text{nticd}$ in G_m for each sink-node m , by computing $<_{\text{SINK}}$ and then $\rightarrow\text{nticd}$ from scratch for each such graph. In order to achieve a reasonable performance for graphs with big $M \in \mathbb{M}$, I will replace the from-scratch $<_{\text{SINK}}$ computations by incremental computations.

Specifically, for each $M \in \mathbb{M}$, I will compute the

$$<_m = <_{\text{SINK}}^{G_m}$$

such that $>_m$ is a transitive reduction of $\supseteq_{\text{SINK}}^{G_m}$, i.e.:

$$\supseteq_m = \supseteq_{\text{SINK}}^{G_m}$$

for

$$\begin{aligned} \supseteq_m &= (>_m)^* \\ G_m &= (G_M)_{m \nrightarrow} \quad \text{and} \\ G_M &= G^{N_M \rightarrow^* M} \end{aligned}$$

for each $m \in M$ incrementally.

By construction, the following holds:

Proposition 6.7.2.

1. G_m is a graph with unique exit node m .
2. $<_m$ is a tree with root m (i.e.: $\forall n \neq m. \exists! n'. n <_m n'$, and $\neg \exists n. m <_m n$).

Observation 6.7.7. Let M be any control-sink of a CFG G (and hence: a $>_{\text{SINK}}^G$ -cycle), and use the above notation for any node $m \in M$. Let $m_2, m'_2 \in M, m_2 \neq m'_2$ such that $m_2 \rightarrow_{G_M} m'_2$. Furthermore, let

$$\begin{aligned} <'_{m_2} &= \{ (n, m) \mid n <_{m_2} m, n \neq m'_2, \neg n \rightarrow_{G_M} m'_2 \} \\ &\cup \{ (n, m'_2) \mid n \neq m'_2, n \rightarrow_{G_M} m'_2 \} \end{aligned}$$

be the tree which is obtained from $<_{m_2}$ by making m'_2 the new root, and letting all G_M -predecessors $n \neq m'_2$ of m'_2 point to m'_2 , and write \sqsubseteq'_{m_2} for $(>'_{m_2})^*$. Then the following holds:

1. If m_2 is the only G_M -predecessor of m'_2 , then $>'_{m_2}$ is a transitive reduction of $\sqsupset_{\text{SINK}}^{G_{m'_2}}$, i.e.:

$$\sqsupset'_{m_2} = \sqsupset_{m'_2}$$

2. Otherwise, \sqsupset'_{m_2} approximates $\sqsupset_{\text{SINK}}^{G_{m'_2}}$ from above, i.e.:

$$\sqsupset'_{m_2} \supseteq \sqsupset_{m'_2}$$

Also, if I write $<_0$ for the two-tree forest $(<_{m_2})_{m'_2 \nrightarrow}$, then for any nodes n such that $n \rightarrow_{G_M} m'_2$ or $\neg n <_0^* m_2$, I have:

$$\{ x \mid x \sqsupset'_{m_2} n \} = \{ x \mid x \sqsupset_{m'_2} n \}$$

I attempt to visualize this process in figure Figure 6.10.

With regard to computation costs, I note the following:

- The computation of $<'_{m_2}$ from $<_{m_2}$ is trivially cheap for any but the most ridiculous graphs G .
- The fact that \sqsupset'_{m_2} is a superset of $\sqsupset_{m'_2}$ allows me to compute $\sqsupset_{m'_2}$ by immediately starting a fixed-point computation from

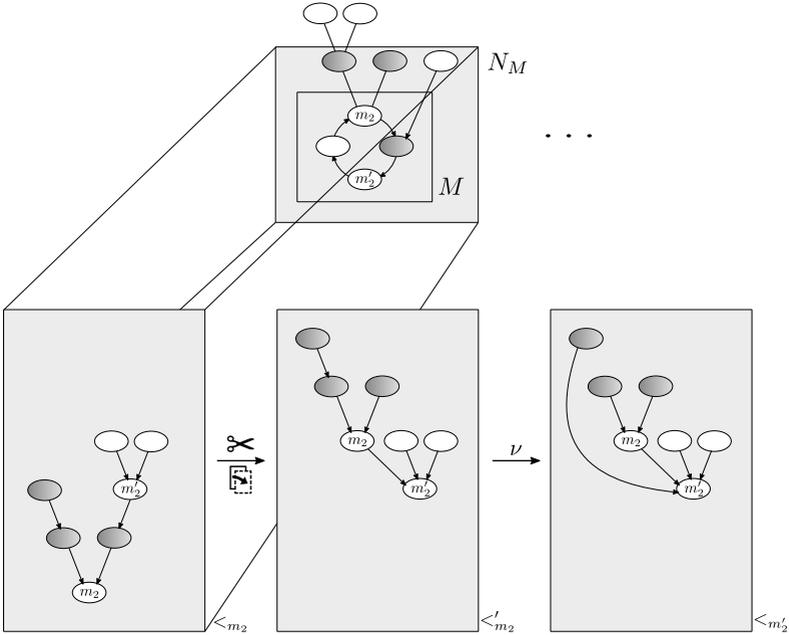


Figure 6.10: The Process in Observation 6.7.7. On top: \prec_{SINK}^G . The last step is only necessary if m_2 is not the only predecessor of m'_2 .

\sqsubseteq'_{m_2} (e.g., by using phase $\text{SINK}_{\text{down}}$ from Algorithm 6). I do not need to execute phase SINK_{up} .

- Even more, I am done for any node n such that $n \rightarrow_{G_M} m'_2$ or $\neg n <_0^* m_2$, and hence can modify workset-based algorithms such as $\text{SINK}_{\text{down}}$ to never put such nodes in the workset!

Now, Observation 6.7.7 only applies whenever $m_2 \rightarrow_{G_M} m'_2$, but this already allows me to compute \sqsubseteq_m for all $m \in M$ as follows:

1. Using any heuristic, enumerate a sequence π_1, \dots, π_k of (finite) G_M -paths in M such that

$$M = \bigcup_i \pi_i$$

2. For all

$$\pi_i = m, \dots$$

compute \sqsupseteq_m using any postdominator-tree algorithm (remember that G_m is a CFG with unique exit node m). Then, for all other nodes $m'_2 \in \pi_i$ with π_i -predecessor m_2 , compute $\sqsupseteq_{m'_2}$ using \sqsupseteq_{m_2} via Observation 6.7.7.

Any heuristic that chooses the π_i should then, presumably:

- Attempt to cover each node $m \in M$ only once.
- Try to minimize the number of π_i -neighbours m_2, m'_2 such that m_2 is not the only G_M -predecessor of m'_2 .
- Try to minimize the number k of π_i , since every $\pi_i = m, \dots$ requires an initial full computation of \sqsupseteq_m .

My implementation uses a greedy heuristic that finds paths starting in join nodes, and covers each node exactly once.

A Variant for Arbitrary Enumeration of Sink Nodes

Owing to the following observation, it turns out that the full computations of \sqsupseteq_m for $\pi_i = m, \dots$ is not even necessary. I attempt to visualize the corresponding process in figure Figure 6.11 on page 112.

Observation 6.7.8. Let M be any control-sink of a CFG G (and hence: a $>_{\text{SINK}}^G$ -cycle), and use the above notation for any node $m \in M$. Let $m_2 \neq m'_2$ be *any* nodes in M . Furthermore, let

$$\langle_0 = (\langle_{m_2})_{m'_2 \nrightarrow}$$

be the two-tree forest obtained by cutting the subtree with root m'_2 from the tree \langle_{m_2} rooted in m_2 , and write \exists_0 for $(>_0)^*$.

Also, let

$$P_{m'_2} = \{ n \mid m'_2 \exists_0 n \}$$

be the set of nodes n that are in the tree rooted in m'_2 ,

$$S_{m'_2} = \{ n \mid m_2 \rightarrow_G n, m'_2 \exists_0 n \}$$

be the set of G -successors n of m_2 that are in the tree rooted in m'_2 ,

$$C_{m_2} = \{ n \mid n \in \text{COND}_{G_{m'_2}}, m_2 \exists_0 n \}$$

be the set of conditional (in $G_{m'_2}$) nodes n that are in the tree rooted in m_2 .

Then

1. If $S_{m'_2}$ is non-empty, there is a unique least common ancestor $z \in \text{lca}_{\langle_0}(S_{m'_2})$ of all such $S_{m'_2}$, and the graph

$$\langle'_{m_2} = \langle_0 \cup \{(m_2, z)\}$$

is a tree (rooted in m'_2); specifically — in the words of *Algorithm 6* — every node is *processed*.

2. If $S_{m'_2}$ is empty, then for all $n \in P_{m'_2}$:

$$\{ x \mid x \exists_0 n \} \supseteq \{ x \mid x \exists_{m'_2} n \}$$

Also, let

$$\begin{aligned} <'_0 &= \{ (n, m) \mid n <_0 m, n \notin C_{m_2} \} \\ \cup &\begin{cases} \emptyset & \text{if } m_2 \in \text{COND}_{G_{m'_2}} \\ \{(m_2, m)\} & \text{if } m \text{ is the unique } G_{m'_2} \text{ - successor of } m_2 \end{cases} \end{aligned}$$

and let $<'_{m_2}$ be the tree obtained by “processing” the tree $<'_0$, i.e.: let $<'_{m_2}$ be the tree obtained by running phase SINK_{up} ⁹ from *Algorithm 6* with initial set $\text{PROCD} = P_{m'_2}$ of processed nodes, and initial workqueue = C_{m_2} .

In both cases, write \supseteq'_{m_2} for $(>'_{m_2})^*$. Then:

$$\supseteq'_{m_2} \supseteq \supseteq_{m'_2}$$

and also for $n \notin \{ n \mid m_2 \supseteq_0 n \} \supseteq C_{m_2}$:

$$\{ x \mid x \supseteq'_{m_2} n \} = \{ x \mid x \supseteq_{m'_2} n \}$$

With regard to computation costs, I note the following:

- The computation of $<'_{m_2}$ from $<_{m_2}$ is trivially cheap if $S_{m'_2} \neq \emptyset$, but requires a partial execution of phase SINK_{up} if it is empty.
- Again, the fact that \supseteq'_{m_2} is a superset of $\supseteq_{m'_2}$ allows me to compute $\supseteq_{m'_2}$ by immediately starting a fixed-point computation from \supseteq'_{m_2} (e.g., by using phase $\text{SINK}_{\text{down}}$ from *Algorithm 6*).
- Even more, I am done for any node $n \notin C_{m'_2}$, and hence can modify workset-based algorithms such as $\text{SINK}_{\text{down}}$ to never put such nodes in the workset!

Thanks to *Observation 6.7.8* — together with *Observation 6.7.7* — I can compute \supseteq_m for all $m \in M$ as follows:

⁹ for $G_{m'_2}$

1. Using any heuristic, enumerate any finite sequence¹⁰ π of nodes such that $M = \{ m \mid m \in \pi \}$
2. For

$$\pi = m, \dots$$

compute \sqsubseteq_m using any postdominator-tree algorithm. Then, for all other nodes $m'_2 \in \pi$ with π -predecessor m_2 , compute $\sqsubseteq_{m'_2}$ using \sqsubseteq_{m_2} via Observation 6.7.7 if $m_2 \rightarrow_{G_M} m'_2$, or via Observation 6.7.8 if not.

Any heuristic that chooses π should then, presumably:

- Attempt to cover each node $m \in M$ only once.
- Try to minimize the number of π -neighbours m_2, m'_2 such that m_2 is not a G_M -predecessor of m'_2 , and then
- try to minimize the number of π -neighbours m_2, m'_2 such that m_2 is not the only G_M -predecessor of m'_2 .

Using this scheme, and taking into account Observation 6.7.6, the computation of $\rightarrow_{\text{ntiod}}$ for a CFG G reduces to the computation of

1. \prec_{SINK}^G
2. the set \mathbb{M} of control-sinks M of G ¹¹
3. for each $M \in \mathbb{M}$ the subgraph G_M
4. for each $m_2 \in M$:
 - a) \prec_m as described
 - b) the set of $n \in N_M, m_1 \in M, m_1 \neq n$ such that $n \rightarrow_{\text{ntiod}}^{G_{m_2}} m_1$, obtaining $n \rightarrow_{\text{ntiod}} (m_1, m_2)$.

¹⁰ not necessarily a G_M -path!

¹¹ which are exactly the \prec_{SINK}^G -cycles, and also: exactly the strongly connected components of G that do not have edges leaving the component.

The computation item 4b can be done using any postdominance frontier algorithm.

Remark 6.7.1. I described a scheme to obtain

$$\prec_{m'_2} \text{ from } \prec_{m_2}$$

I did *not* investigate whether it is possible to directly obtain

$$\rightarrow_{\text{nticd}}^{G_{m'_2}} \text{ from } \rightarrow_{\text{nticd}}^{G_{m_2}}$$

Although this might very well be possible, it appears to be of limited practical value, since experiments suggest that in practice, the scheme I presented already yields an algorithm with execution time roughly linear in the size of the relation $\rightarrow_{\text{nticd}}$.

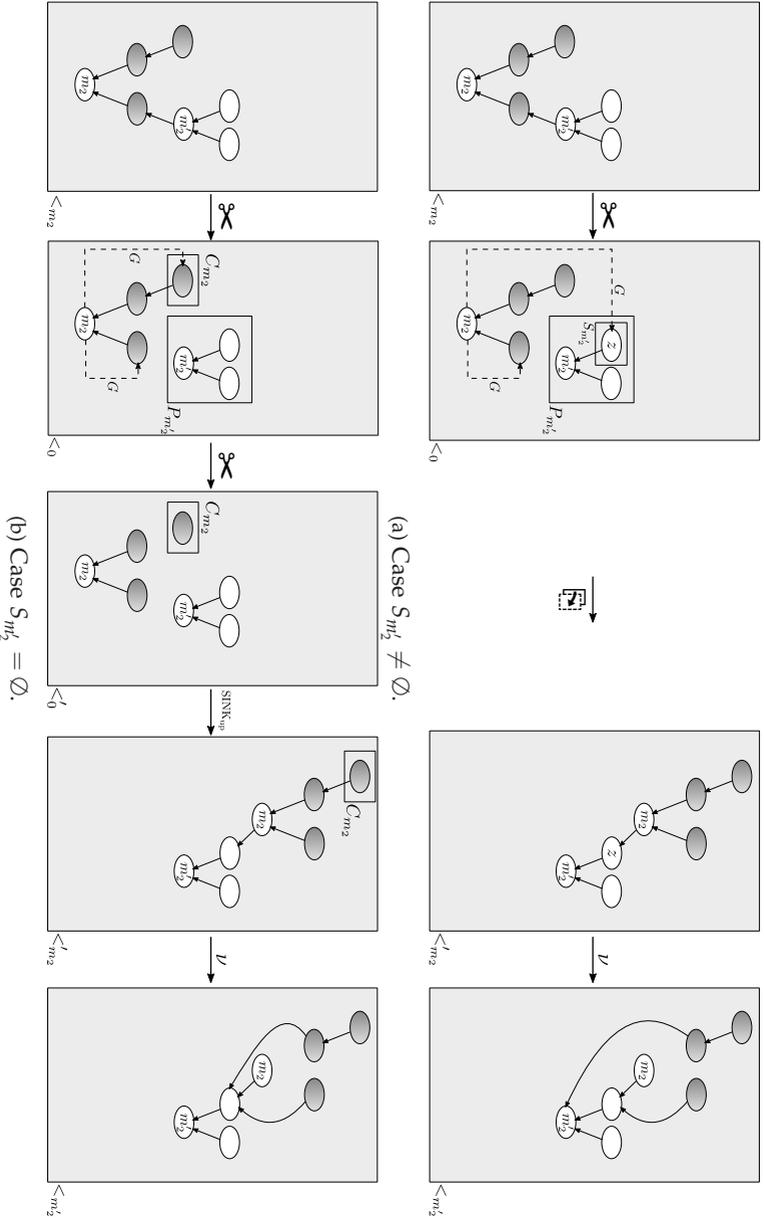


Figure 6.11: The Process in Observation 6.7.7. Dashed arrows are edges in G .

6.8 Soundness of Nontermination Insensitive Order Dependence

In Section 6.4, I recalled the simulation based correctness criterion for nontermination insensitive slicing which is satisfied by \rightarrow_{wod} ([Amt08]). I offered a critique of this criterion by example of an “always terminating” CFG that satisfies the simulation-criterion for some set $S = M$ of observable nodes, but intuitively leaks information.

The simulation criterion is based on “stateful” executions of CFG with configurations as pairs (n, σ) of nodes and variable assignments. It requires closure under $\rightarrow_{\text{data}}$. In [Amt08], it was shown that S being weakly control-closed is useful in proving the simulation property. In [Was10], for graphs with unique end nodes n_x , and assuming $n_x \in S$, it was shown that being weakly control-closed is outright the only property of control-dependency necessary to establish the simulation property, and in [Dan+11] this property is straight out taken to *define* correctness of CFG slicing.

While I did not propose a strengthened version of the simulation criterion, I *did* (in Section 6.5) propose a new requirement for (unlabeled) CFG (in addition to being weakly control-closed) that prevents information leak for always-terminating programs. I also proposed — in Section 6.6 — a new trace-based criterion with the same purpose.

In this section, I confirm that $\rightarrow_{\text{ntiod}}$ ¹² is, indeed, both sound and minimal with respect to these two new criteria. I also confirm that, for always-terminating CFG, slicing w.r.t $\rightarrow_{\text{ntiod}}$ is nontermination *sensitively* sound, i.e.: observable traces are uniquely defined by the observable input, which is *not* the case for slicing w.r.t \rightarrow_{wod} .

Observation 6.8.1 (Soundness of $\rightarrow_{\text{ntiod}}, \rightarrow_{\text{nticd}}$ w.r.t obs_S). Let G be any CFG, and $M \subseteq N$ a set of nodes (the slicing criterion). Let $S = (\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}})^*(M)$ be the backward slice w.r.t. M . Then all

¹² together with $\rightarrow_{\text{nticd}}$

possible next observations are retained outside S (see Definition 6.5.4 on page 90).

Observation 6.8.2 (Minimality of $\rightarrow_{\text{ntiod}}, \rightarrow_{\text{nticd}}$ w.r.t obs_S). Let G be any CFG, and $M \subseteq N$ a set of nodes (the slicing criterion). Let $S = (\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}})^*(M)$ be the backward slice w.r.t. M . Then, for any $n \in S, n \notin M$ and $S' = S \setminus \{n\}$, *not* all possible next observations are retained outside S'

Minimality of $\rightarrow_{\text{ntiod}}$ as stated in Observation 6.8.2 is directly amenable to empirical validation. For the following alternative notion of minimality, Simon Bischof prepared a machine checked proof in the Isabelle/HOL proof assistant:

Theorem 6.8.1 ([Bis19], Minimality of $\rightarrow_{\text{ntiod}}, \rightarrow_{\text{nticd}}$ w.r.t obs_S). Let G be any CFG, and $M \subseteq N$ a set of nodes (the slicing criterion). Let $S = (\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}})^*(M)$ be the backward slice w.r.t. M . Let $S' \supseteq M$ be any set of nodes that retains all possible next observations. Then $S' \supseteq S$.

Slicing via $\rightarrow_{\text{ntiod}}$ and $\rightarrow_{\text{nticd}}$ is also sound and minimal with regard to the nontermination insensitive trace based criterion of *infinite delay*.

Observation 6.8.3 (Soundness of $\rightarrow_{\text{ntiod}}, \rightarrow_{\text{nticd}}$ w.r.t $\mathcal{T}^\omega|_S$). Let G be any CFG, and $M \subseteq N$ a set of nodes (the slicing criterion). Let $S = (\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}})^*(M)$ be the backward slice w.r.t. M . Then, any S -equivalent inputs i, i' — i.e.: inputs such that

$$i \sim_S i'$$

— have S -equivalent observable behavior *up to infinite delay*¹³, i.e.:

$$i \sim_{\mathcal{T}_S}^\omega i'$$

¹³ see Definition 6.6.2 on page 94

Observation 6.8.4 (Minimality of $\rightarrow_{\text{ntiod}}, \rightarrow_{\text{nticd}}$ w.r.t $\mathcal{T}^\omega|_S$). Let G be any CFG, and $M \subseteq N$ a set of nodes (the slicing criterion). Let $S = (\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}})^*(M)$ be the backward slice w.r.t. M . Then, for any $n \in S, n \notin M$ and $S' = S \setminus \{n\}$, there exists S' -equivalent inputs i, i' — i.e.: inputs such that

$$i \sim_S i'$$

— that do *not* have S' -equivalent observable behavior *up to infinite delay*, i.e.:

$$\neg i \sim_{\mathcal{T}_S}^\omega i'$$

The two preceding observation are the nontermination *insensitive* analogues of observations 6.3.1 and 6.3.2 that concerned nontermination *sensitive* slicing. The following observation confirms that, for always-terminating CFG, nontermination insensitive slicing is as strict as non-termination sensitive slicing.

Observation 6.8.5 (Trace Equivalence for $\rightarrow_{\text{ntiod}}, \rightarrow_{\text{nticd}}$ in acyclic CFG). Let G be a CFG in which *all executions terminate*, i.e.: an acyclic graph. Let $M \subseteq N$ be a set of nodes (the slicing criterion), and let $S = (\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{ntiod}})^*(M)$ be the backward slice w.r.t. M . Then, any S -equivalent inputs i, i' — i.e.: inputs such that

$$i \sim_S i'$$

— have nontermination sensitively S -equivalent traces¹⁴, i.e.:

$$t_i \sim_S t_{i'}$$

Remember that this last observation, although relatively weak, does *not* hold for slicing under \rightarrow_{wod} , as was exemplified in Figure 6.4a on page 87.

¹⁴ see Definition 6.3.4 on page 80

Summary

- Together, decisive order dependence \rightarrow_{dod} and nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$ produce sound and minimal slices.
- The same holds for nontermination insensitive order dependence $\rightarrow_{\text{ntsod}}$.
- $\rightarrow_{\text{ntsod}}$ can naturally be reduced to $\rightarrow_{\text{ntscd}}$, which leads to an algorithm for $\rightarrow_{\text{ntsod}}$.
- Together, nontermination sensitive order dependence $\rightarrow_{\text{ntiod}}$ and nontermination insensitive control dependence $\rightarrow_{\text{nticd}}$ produce sound and minimal slices.
- Soundness and minimality of nontermination insensitive slicing can either be defined by a notion of *infinite delay*, or by a condition of *next observables*.
- $\rightarrow_{\text{ntiod}}$ can naturally be reduced to $\rightarrow_{\text{nticd}}$, which leads to an algorithm for $\rightarrow_{\text{ntiod}}$.

7 Slicing

If you do something too good, then, after a while, if you don't watch it, you start showing off.

(J.D. Salinger — The Catcher in the Rye)

In the preceding chapters, I devised new algorithms for the computation of (explicit representations of) the dependency relations

1. **Non-Termination Sensitive Control Dependence** $\rightarrow_{\text{ntscd}}$
 - a) indirectly via \langle_{MAX} and Algorithm 1
2. **Non-Termination Insensitive Control Dependence** $\rightarrow_{\text{nticd}}$
 - a) directly via Algorithm 16
 - b) indirectly via \langle_{SINK} and Algorithm 1
3. **Decisive¹ Order Dependence** \rightarrow_{dod}
 - a) via the generate-and-test scheme 6.3 on page 70
4. **Non-Termination Sensitive Order Dependence** $\rightarrow_{\text{ntsod}}$
 - a) via the scheme implied by Observation 6.2.3 and Observation 6.2.5
5. **Non-Termination Insensitive Order Dependence** $\rightarrow_{\text{ntiod}}$
 - a) via the scheme implied by Observation 6.7.4 and Observation 6.7.6,
 - b) improved upon via Observation 6.7.7 and Observation 6.7.8

Their natural application in information flow control is (backward)-slicing, i.e.: the computation of

¹ non-termination sensitive

$$((\rightarrow_{\text{nticd}} \cup \rightarrow_{\text{data}}) \cup \rightarrow_{\text{ntiod}})^*(M) \quad (7.1)$$

for sets of “observable” nodes M .

For graphs with unique end nodes n_x , this is equivalent to slice along standard control dependence \rightarrow_{cd} only, i.e.:

$$(\rightarrow_{\text{cd}} \cup \rightarrow_{\text{data}})^*(M)$$

Here, it usually is of little practical importance whether slicing is performed with regard to only one set M of nodes, or several such sets M, M', \dots : Computation of \rightarrow_{cd} is — in practice — cheap² so in both cases, one just does this, and then slices backwards from M, M', \dots in the obvious way.

On the other hand, for arbitrary graphs, I must slice as per (7.1). Here, in accordance with the informal slogan

$$\rightarrow_{\text{ntiod}} \approx \left(\sum_{M \in \mathcal{M}} |M| \right) \times \rightarrow_{\text{nticd}}$$

any algorithm computing $\rightarrow_{\text{ntiod}}$ in an explicit representation must be expected to run in $\simeq |N|^2$ steps for typical CFG, and $\mathcal{O}(|N|^3)$ worst case. To see this directly, consider the n -node looping ladder CFG in Figure 7.1 The variant of this graph which lacks node n and has no edge $(n-1) \rightarrow 0$ is the canonical example of a CFG for which \rightarrow_{cd} has size $\Theta(n^2)$. As is, this graph demonstrates that $\rightarrow_{\text{ntiod}}$ can be of size $\Theta(n^3)$. This is because for every odd node m_1 and every even node m_2 , I have

$$n \rightarrow_{\text{ntiod}} (m_1, m_2) \Leftrightarrow n < m_1 \wedge n \neq m_2$$

² In program slicing, \rightarrow_{cd} often is “mostly” a tree, and hence small, with $|\rightarrow_{\text{cd}}| \simeq |N|$. In fact, if one demands $n_e \rightarrow_G n_x$ and a reducible graphs G , it is a tree.

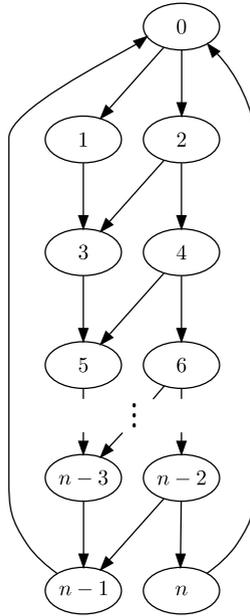


Figure 7.1: The n -node looping ladder CFG.

which means that for every odd m_1 , the set of such m_2, n is of size

$$s_{m_1} = \left\lfloor \frac{n}{2} \right\rfloor \cdot \frac{m_1 + 1}{2}$$

and hence:

$$|\rightarrow\text{ntiod}| \geq \sum_{m_1 \text{ odd}} s_{m_1} = \left\lfloor \frac{n}{2} \right\rfloor \left(\frac{n-1}{8} + \frac{1}{4} \right) (n-1) \in \Theta(n^3)$$

If I am interested in the slice (7.1) of only *one* set M of nodes, can I compute this without risking a size $\mathcal{O}(|N|^3)$ computation? In this chapter, I answer this question in the affirmative.

In Section 7.1 I obtain $\rightarrow_{\text{ntiod}}$, $\rightarrow_{\text{nticd}}$ backwards slice in G by a reduction to *one* computation of $\rightarrow_{\text{nticd}}^{G_M}$ for a suitable subgraph G_M of G . Then, I will show that it is even enough to compute nontermination sensitive postdominance $\sqsupseteq_{\text{SINK}}^{G_M}$ only, and then *directly* read off the $\rightarrow_{\text{nticd}}^{G_M}$ -backwards slice of M , without fully computing $\rightarrow_{\text{nticd}}^{G_M}$.

If only the $\rightarrow_{\text{nticd}}$ backwards slice in G is required (e.g.: if all control sinks in the graph G are trivial), I can generalize the technique used in [SG95] to compute the $\rightarrow_{\text{nticd}}$ backwards slice in G of any slicing criterion M . I present this generalization in Appendix C.

In Section 7.2 I present the nontermination sensitive analogue of the result from Section 7.1. In Section 7.3 I present the analogue of the result from Section 7.1 for \rightarrow_{wod} backward slices. In Section 7.4 I present the analogue of the result from Section 7.1 for *weak control closures*, and then summarize the results in Section 7.5.

7.1 Nontermination Insensitive Slicing

In this section, I will attack the simplified problem of computing the control slice only, i.e.: I disregard data dependencies $\rightarrow_{\text{data}}$. Given a single set M of “observable” nodes in an arbitrary CFG G , the appropriate reduction to *one* computation is remarkably simple. It will turn out that

$$\left(\rightarrow_{\text{nticd}}^G \cup \rightarrow_{\text{ntiod}}^G \right)^* (M) = \left(\rightarrow_{\text{nticd}}^{G_{M \not\rightarrow}} \right)^* (M) \quad (7.2)$$

Recall that $G_{M \not\rightarrow}$ is the graph obtained from G by removing all edges originating in M .

Intuitively, equation 7.2 holds because by Corollary 6.7.2 on page 102, $n \rightarrow_{\text{ntiod}}(m_1, m_2)$ can only hold for nodes m_1, m_2 that occur in a common cycle C in the pseudo-forest $<_{\text{sink}}$, and nodes n that either are in the same $<_{\text{sink}}$ cycle C , or not in any such cycle. By Observation 5.3.5, such cycles C are control sinks in G , i.e., they are strongly connected components of G that have no edge leaving the component. Then by deleting outgoing edges of the slicing criterion M , the component C becomes a region with exit nodes $C \cap M$, and any node n that formerly controlled the order of nodes m_1 and m_2 now controls which of the two nodes is executed *at all*.

Consider the example from Section 5.3, repeated in Figure 7.2a. The only non-trivial control sink of G is $C = \{6, 7, 8, 11, 13\}$. All nodes in C are nontermination insensitively control dependent on node 2 and 1, but these are the *only* nodes they are control dependent on. For example, node 11 is *not* nontermination insensitively control dependent on node 7. Instead, node 11, 8 together are nontermination insensitively *order* dependent on node 7: $7 \rightarrow_{\text{ntiod}}(11, 8)$. So consider the slicing criterion $M = \{7, 8\}$, for which the combined backward slice is

$$\left(\rightarrow_{\text{nticd}}^G \cup \rightarrow_{\text{ntiod}}^G \right)^* (M) = \{7, 8, 11, 2, 1\}$$

The resulting graph G_M is shown in Figure 7.2c, together with the corresponding nontermination insensitive control dependence in Fig-

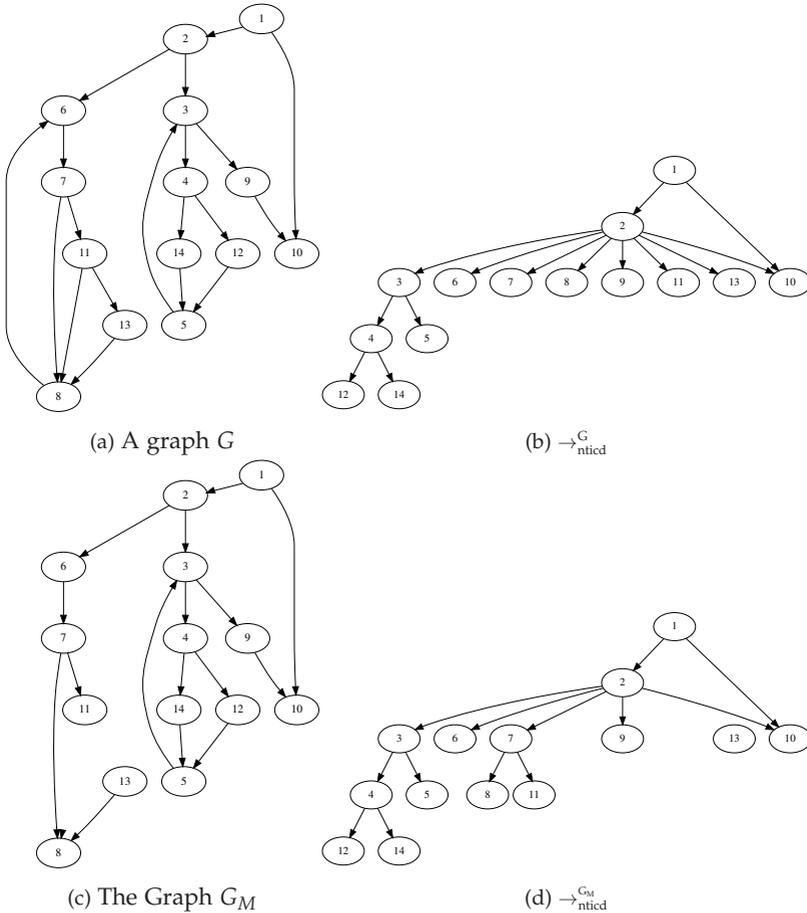


Figure 7.2: Equation 7.2 for $M = \{8, 11\}$

ure 7.2d. The choice at node $n = 7$ (which formerly decided the *order* of nodes 8, 11) now controls whether node 8 or node 11 is executed.

Before I can proof equation (7.2), I need a new characterization of the transitive closure of \rightarrow_{nticd} .

Definition 7.1.1. Let G be any CFG, and $m \neq n$ nodes in G . Then $m \notin \text{ipdom}_{\sqsupseteq_{\text{SINK}}}(n)$ is said to be on a path starting in n strictly before any immediate $\sqsupseteq_{\text{SINK}}$ -dominator of n — and I write $n \rightarrow_{\text{ntind}} m$ — iff there exists some G -successor x of n such that m is reachable in the graph obtained from G by removing all edges originating in $\text{ipdom}_{\sqsupseteq_{\text{SINK}}}(n)$, i.e.:

$$x \rightarrow_G^* \left(\text{ipdom}_{\sqsupseteq_{\text{SINK}}}(n) \right) \not\rightarrow m$$

This notion is most easily understood as a generalization of Weiser’s transitive notion \rightarrow_{nd} of control dependence ([Wei81]) for graphs with unique exit node n_x . Recall that (for $n \neq n_x$), $n \rightarrow_{\text{nd}} m$ iff $m \notin \{n, n'\}$ is on a G -path between n and the unique immediate post-dominator n' of n . The subscript $_{\text{nd}}$ is derived from the immediate postdominator, which Weiser calls *nearest (inverse) dominator*. It is easy to see that $\rightarrow_{\text{ntind}}$ is in fact a generalization of \rightarrow_{nd} , i.e.:

$$n \rightarrow_{\text{ntind}} m \iff n \rightarrow_{\text{nd}} m$$

for graphs with unique exit node n_x . More importantly, it relates to $\rightarrow_{\text{nticd}}$ the same way that \rightarrow_{nd} does to \rightarrow_{cd} :

Observation 7.1.1. Let G be any CFG, and $m \neq n$ nodes in G . Then

$$n \rightarrow_{\text{ntind}} m \iff n \rightarrow_{\text{nticd}}^* m$$

For the proof of (7.2), I also need the following properties of $\rightarrow_{\text{nticd}}$.

Observation 7.1.2. Let G be any CFG, and M any set of nodes in G . Then for all nodes n, m , and all nodes $m_0 \in M$ such that $m \neq m_0$:

$$m \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} n \implies m \sqsupseteq_{\text{SINK}[m_0]}^G n$$

In other words: if m must appear on every $G_{M \not\rightarrow}$ -sink-path starting in n , then m must also appear *before* m_0 on every G -sink-path starting in n .

Observation 7.1.3. Let G be any CFG, and M any set of nodes in G . Then for all nodes n, m :

$$m \sqsupseteq_{\text{SINK}}^{G_{M \neq \emptyset}} n \implies m \sqsupseteq_{\text{SINK}}^G n$$

The following fact is well-known, but hitherto only for the special case of graphs with unique exit-node n_x .

Observation 7.1.4. Let G be *any* CFG, and $m_1 \sqsupseteq_{\text{SINK}}^G n$ as well as $m_2 \sqsupseteq_{\text{SINK}}^G n$. Then, “up to $\sqsupseteq_{\text{SINK}}$ -equivalence”, the order in which paths starting in n visits m_1, m_2 is fixed, i.e.:

1. $m_1 \sqsupseteq_{\text{SINK}[m_2]}^G n$, or
2. $m_2 \sqsupseteq_{\text{SINK}[m_1]}^G n$, or
3. both $m_1 \sqsupseteq_{\text{SINK}}^G m_2$ and $m_2 \sqsupseteq_{\text{SINK}}^G m_1$

Immediate $\sqsupseteq_{\text{SINK}}$ -dominators are always visited before non-immediate postdominators:

Observation 7.1.5. Let G be *any* CFG and $n \neq m_2$. Also, let $m_2 \sqsupseteq_{\text{SINK}}^G n$ as well as not only $m_1 \sqsupseteq_{\text{SINK}}^G n$ but also

$$m_1 \in \text{ipdom}_{\sqsupseteq_{\text{SINK}}} (n)$$

Then, “up to $\sqsupseteq_{\text{SINK}}$ -equivalence”, any paths starting in n always visits m_1 before m_2 , i.e.:

1. $m_1 \sqsupseteq_{\text{SINK}[m_2]}^G n$, or
2. both $m_1 \sqsupseteq_{\text{SINK}}^G m_2$ and $m_2 \sqsupseteq_{\text{SINK}}^G m_1$, as well as

$$m_2 \in \text{ipdom}_{\sqsupseteq_{\text{SINK}}} (n)$$

I need the following three technical observations in the proof of (7.2).

Observation 7.1.6. Let G be any CFG, and M any set of nodes in G . Then whenever

$$m \sqsupseteq_{\text{SINK}}^G n \quad \text{but} \quad \neg m \sqsupseteq_{\text{SINK}}^{G_{M/\neq}} n$$

there exists some node $m_0 \in M$ and some $G_{M/\neq}$ -path π such that $m_0 \neq m$ and

$$\pi = n, \dots, m_0 \quad \text{such that} \quad m \notin \pi$$

Proof: From $\neg m \sqsupseteq_{\text{SINK}}^{G_{M/\neq}} n$ I obtain some $G_{M/\neq}$ sink-path π starting in n with $m \notin \pi$. Then $m_0 \in \pi$ for some $m_0 \in M$, since otherwise π is also a G sink-path with $m \notin \pi$, contradicting $m \sqsupseteq_{\text{SINK}}^G n$. With $m_0 \in \pi$, it must be of the form

$$\pi = n, \dots, m_0$$

□

Observation 7.1.7. Let G be any CFG, and M any set of nodes in G . Then whenever

$$m \sqsupseteq_{\text{SINK}}^G n \quad \text{but} \quad \neg m \sqsupseteq_{\text{SINK}}^{G_{M/\neq}} n$$

there exists some node $m_0 \in M$ such that

1. $m_0 \neq m$,
2. $m \sqsupseteq_{\text{SINK}}^G m_0$,
3. $n \rightarrow_{G_{M/\neq}}^* m_0 \rightarrow_G^+ m$, and
4. $\neg m \sqsupseteq_{\text{SINK}[m_0]}^G n$

Observation 7.1.8. Let G be any CFG, and M any set of nodes in G . Also, let

$$n \left(\rightarrow_{\text{nticd}}^{G_{M/\neq}} \right)^* m$$

for some $m \in M$. Then, for all $n' \neq n$:

$$\neg n' \sqsupseteq_{\text{SINK}}^{G_{M/\neq}} n$$

Proof (Sketch): By induction on $n \left(\rightarrow_{\text{nticd}}^{G_{M/\neq}} \right)^* m$.

Using these observations, I will proof (7.2) via two inductive arguments, starting with the reverse implication.

Theorem 7.1.1. Let G be any CFG, and M any set of nodes in G . Then for any n ,

$$n \in \underbrace{\left(\rightarrow_{\text{nticd}}^G \cup \rightarrow_{\text{ntiod}}^G \right)^* (M)}_{=:A} \iff n \in \underbrace{\left(\rightarrow_{\text{nticd}}^{G_{M \nrightarrow}} \cup \emptyset \right)^* (M)}_{=:B}$$

Proof: By induction on $n \in B$, via Definition 6.2.3. The case $n \in M$ is trivial. For the case $n \rightarrow_{\text{nticd}}^{G_{M \nrightarrow}} n_0$ for some $n_0 \in B$, i can assume $n_0 \in A$. Hence, i also can assume $n \neq n_0$, and also: $n \notin M$, since otherwise I am done. By definition, I have some $G_{M \nrightarrow}$ successors x, x' of n such that

$$\begin{aligned} n_0 \sqsupseteq_{\text{SINK}}^{G_{M \nrightarrow}} x' \quad \text{and hence, by Observation 7.1.3} \quad n_0 \sqsupseteq_{\text{SINK}}^G x' \\ \neg \quad n_0 \sqsupseteq_{\text{SINK}}^{G_{M \nrightarrow}} x \end{aligned}$$

Of course, the $G_{M \nrightarrow}$ -successors of n are exactly the G -successors of n .

If $n \rightarrow_{\text{nticd}}^G n_0$ I am done immediately, so I only have to deal with the opposite, i.e.: $\neg n \rightarrow_{\text{nticd}}^G n_0$. But then, from $n_0 \sqsupseteq_{\text{SINK}}^G x'$ i conclude

$$\forall y \text{ s.t. } n \rightarrow_G y. n_0 \sqsupseteq_{\text{SINK}}^G y \quad \text{and hence} \quad n_0 \sqsupseteq_{\text{SINK}}^G n$$

Specifically for x , I have

$$\begin{aligned} n_0 \sqsupseteq_{\text{SINK}}^G x \\ \text{but: } \neg \quad n_0 \sqsupseteq_{\text{SINK}}^{G_{M \nrightarrow}} x \end{aligned}$$

From Observation 7.1.7, I obtain some $m_0 \in M$ such that

$$\begin{aligned} m_0 \neq n_0 \qquad n_0 \sqsupseteq_{\text{SINK}}^G m_0 \\ x \rightarrow_{G_{M \nrightarrow}}^* m_0 \rightarrow_G^+ n_0 \quad \neg \quad n_0 \sqsupseteq_{\text{SINK}[m_0]}^G x \end{aligned}$$

Now, I discern the two cases $m_0 \sqsupseteq_{\text{SINK}}^G n_0$ and $\neg n_0 \sqsupseteq_{\text{SINK}}^G m_0$.

If $m_0 \sqsupseteq_{\text{SINK}}^G n_0$, I intend to show $n \rightarrow_{\text{ntiod}} (n_0, m_0)$ by directly confirming Definition 6.7.1. I already have $n_0 \sqsupseteq_{\text{SINK}}^G n$, and $m_0 \sqsupseteq_{\text{SINK}}^G n$ follows from $m_0 \sqsupseteq_{\text{SINK}}^G n_0$, so I have item (a) covered. For item (b), I want to show $n_0 \sqsupseteq_{\text{SINK}[m_0]}^G x'$, but this follows from $n_0 \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} x'$ and — since $m_0 \neq n_0$ and $m_0 \in M$ — Observation 7.1.2. For item (c), I already have $\neg n_0 \sqsupseteq_{\text{SINK}[m_0]}^G x$.

Now if on the other hand, $\neg m_0 \sqsupseteq_{\text{SINK}}^G n_0$, I intend to show $n \rightarrow_{\text{nticd}}^* m_0$ in G . Since m_0 in M and only needed to cover the case $n \notin M$, this is equivalent — by Observation 7.1.1 — to showing $n \rightarrow_{\text{ntind}} m_0$ in G . Writing

$$N' = \text{ipdom}_{\sqsupseteq_{\text{SINK}}^G} (n)$$

I will do this by exposing a path

$$\pi = x, \dots, m_0 \quad \text{such that} \quad \pi \cap N' = \emptyset$$

In order to do this, I differentiate the following cases:

1. $N' = \emptyset$
2. $N' = S$ for some control-sink S in G
3. $N' = \{n'\}$ for some node n'

It is indeed sufficient to cover these cases since $\sqsupseteq_{\text{SINK}}^G$ is the reflexive, transitive closure of some pseudo-forest $>$, and there either n has no $<$ -successor, in which case $N' = \emptyset$, or n has a unique $<$ -successor n' . Then, if n' lies in some $<$ -cycle S , S is a control-sink in G and $N' = S$. Otherwise, $N' = \{n'\}$ (c.f. Observation 5.3.2).

I already have $x \rightarrow_{G_{M \not\rightarrow}}^* m_0$ and hence $x \rightarrow_G^* m_0$, so case 1 is trivial.

For case 2, from $n_0 \sqsupseteq_{\text{SINK}}^G n$ I know that $n_0 \in S = N'$, and then from $\neg m_0 \sqsupseteq_{\text{SINK}}^G n_0$ I have $m_0 \notin S = N'$ (S is a “root” of $<$). But then for any G -path

$$\pi = x, \dots, m_0$$

I have $\pi \cap N' = \emptyset$, since i cannot leave a control-sink!

For case 3, I know that $n \neq n'$, since otherwise n is its only $<$ -successor, and hence $\neg n' \perp_{\text{SINK}} n$. If $n' = n_0$ then from $\neg n_0 \sqsupseteq_{\text{SINK}[m_0]}^G x$ and observing $n_0 \neq m_0$, I obtain a path

$$\pi = x, \dots, m_0 \quad \text{such that} \quad n_0 \notin \pi$$

which is exactly what I need.

I still need to cover the case $n' \neq n_0$, i.e.: $n_0 \notin N'$. Remembering that I still have $n \neq n_0$, $n_0 \sqsupseteq_{\text{SINK}}^G n$, $n' \in N'$, I use Observation 7.1.5 to conclude

$$n' \sqsupseteq_{\text{SINK}[n_0]}^G n$$

But since $\sqsupseteq_{\text{SINK}}^{G_{n_0 \not\rightarrow}}$ is closed under taking a step to $G_{n_0 \not\rightarrow}$ -successors (i.e.: closed under $\rightarrow_{G_{n_0 \not\rightarrow}}$) and both $n \neq n_0$ and $n \neq n'$, I have (via two applications of Observation 6.7.5):

$$n' \sqsupseteq_{\text{SINK}[n_0]}^G x'$$

from which I can obtain some path

$$\pi' = x', \dots, n' \quad \text{such that} \quad n_0 \notin \pi'$$

Again from $\neg n_0 \sqsupseteq_{\text{SINK}[m_0]}^G x$ (and still: observing that $m_0 \neq n_0$) I obtain a path

$$\pi = x, \dots, m_0 \quad \text{such that} \quad n_0 \notin \pi$$

But π cannot contain n' , since otherwise, I had a path $\pi'' = n', \dots, m_0$ with $n_0 \notin \pi''$, and together with π' a path

$$\pi' \pi'' = x', \dots, n', \dots, m_0 \quad \text{such that} \quad n_0 \notin \pi' \pi''$$

in contradiction with the fact that

$$n_0 \sqsupseteq_{\text{SINK}[m_0]}^G x'$$

which follows from $n_0 \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} x'$ via Observation 7.1.2 and $m_0 \neq n_0, m_0 \in M$. With $n' \notin \pi$, I have concluded case 3 and hence the proof. \square

What remains is the forward implication. The proof may — at first — appear to mirror the proof of Theorem 7.1.1. This is because the definition of $\sqsupseteq_{\text{nticd}}^G n_0$ requires $n_0 \sqsupseteq_{\text{SINK}} \sqsupseteq$ both in the positive and in the negative for some x, x' . But the actual situations are quite different from the former proof.

Theorem 7.1.2. Let G be any CFG, and M any set of nodes in G . Then for any n ,

$$n \in \underbrace{\left(\rightarrow_{\text{nticd}}^G \sqcup \rightarrow_{\text{ntiod}}^G \right)^* (M)}_{=:A} \Rightarrow n \in \underbrace{\left(\rightarrow_{\text{nticd}}^{G_{M \not\rightarrow}} \sqcup \emptyset \right)^* (M)}_{=:B}$$

Proof: By induction on $n \in A$, via Definition 6.2.3. The case $n \in M$ is trivial.

For the case $n \rightarrow_{\text{nticd}}^G n_0$ for some $n_0 \in A$, I can assume $n_0 \in B$. Hence, I also can assume $n \neq n_0$, and also: $n \notin M$, since otherwise I am done. By definition, I have some G successors x, x' of n such that

$$\begin{aligned} & n_0 \sqsupseteq_{\text{SINK}}^G x \\ \neg & n_0 \sqsupseteq_{\text{SINK}}^G x' \quad \text{and hence, by Observation 7.1.3} \quad \neg n_0 \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} x' \end{aligned}$$

Since $n \notin M$, the G -successors of n are exactly the $G_{M \not\rightarrow}$ -successors of n .

If $n \rightarrow_{\text{nticd}}^{G_{M \not\rightarrow}} n_0$ I am done immediately, so I only have to deal with the opposite, i.e.: $\neg n \rightarrow_{\text{nticd}}^{G_{M \not\rightarrow}} n_0$. But then, from $\neg n_0 \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} x'$ I conclude

$$\forall y \text{ s.t. } n \rightarrow_G y. \neg n_0 \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} y$$

Specifically for x , I have

$$\begin{aligned} n_0 &\sqsupseteq_{\text{SINK}}^G x \\ \text{but: } \neg n_0 &\sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} x \end{aligned}$$

From Observation 7.1.6, I obtain some $m_0 \in M$ with $m_0 \neq n_0$ and a $G_{M \not\rightarrow}$ -path

$$\pi = x, \dots, m_0 \quad \text{such that} \quad n_0 \notin \pi$$

I will show that there exists no $n' \neq n$ such that $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} n$. Once I have done that, I can conclude that

$$N' = \text{ipdom}_{\sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}}} (n) = \emptyset$$

and immediately get $n \rightarrow_{\text{nticd}}^* m_0$ in $G_{M \not\rightarrow}$ from $n \rightarrow_{\text{ntind}} m_0$ (in $G_{M \not\rightarrow}$), which then holds because of $x \rightarrow_{G_{M \not\rightarrow}}^* m_0$.

In order to show that there exists no $n' \neq n$ such that $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} n$, I assume the opposite. Then also $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} x$ and $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} x'$. I also obtain from $\neg n_0 \sqsupseteq_{\text{SINK}}^G x'$ some G -path

$$\pi' = x', \dots, z \quad \text{such that} \quad n_0 \notin \pi' \text{ and } \neg z \rightarrow_G^* n_0$$

I discern the two cases $\exists m' \in M. m' \in \pi'$ and its negation, both of which I have to lead to a contradiction. If there exists such a m' , I can choose m' as the one appearing first in π' . Then from $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} n$ I

conclude that n' must appear in π' before (or at) this occurrence. In summary, I have³

$$\pi' = x', \dots, \underbrace{n', \dots, m', \dots, z}_{=:\pi'_0} \quad \text{with} \quad n_0 \notin \pi' \quad \text{and} \quad \neg z \rightarrow_G^* n_0$$

Now, for the $G_{M \not\rightarrow}$ -path

$$\pi = x, \dots, m_0 \quad \text{such that} \quad n_0 \notin \pi$$

I infer from $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} n$ that $n' \in \pi$, and hence have a prefix $\pi_0 = x, \dots, n'$ with $n_0 \notin \pi$. But then I have

$$n_0 \notin \pi_0 \pi'_0 = x, \dots, n', \dots, m', \dots, z \quad \text{with} \quad \neg z \rightarrow_G^* n_0$$

in contradiction to $n_0 \sqsupseteq_{\text{SINK}}^G x$.

In order derive a contradiction from $\neg \exists m' \in M. m' \in \pi'$, I observe that then π' is not only a G -path but also a $G_{M \not\rightarrow}$ -path. This means that either $n' \in \pi'$, i.e.:

$$\pi' = x', \dots, \underbrace{n', \dots, z}_{=:\pi'_0} \quad \text{such that} \quad n_0 \notin \pi' \quad \text{and} \quad \neg z \rightarrow_G^* n_0$$

in which case I obtain a contradiction just as I did before, *or* there exists a continuation π'' of π' in $G_{M \not\rightarrow}$ such that $n' \in \pi''$, i.e.:

$$\pi' \pi'' = x', \dots, z, \dots, n'$$

But since $\neg z \rightarrow_G^* n_0$, I have $n_0 \notin \pi''$ and $\neg n' \rightarrow_G^* n_0$. Choosing $\pi_0 = x, \dots, n'$ with $n_0 \notin \pi_0$ as before, this contradicts $n_0 \sqsupseteq_{\text{SINK}}^G x$. This concludes the inductive case $n \rightarrow_{\text{nticd}}^G n_0$.

³ not demanding x', n', m', z to be mutually different

What is left is the inductive case $n \rightarrow_{\text{ntiod}}^G (m_1, m_2)$ for some $m_1, m_2 \in A$, in which case I can assume $m_1, m_2 \in B$. Hence, I also can assume $n \neq m_1$, $n \neq m_2$ and also: $n \notin M$, since otherwise I am done.

I have $m_1 \neq m_2$ by definition of $\rightarrow_{\text{ntiod}}$, so by Observation 7.1.8 and $m_2 \in B$ I conclude

$$\neg m_1 \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} m_2$$

Also by definition, I have some G successors x, x' of n such that

$$\begin{aligned} m_1 &\sqsupseteq_{\text{SINK}[m_2]}^G x \\ \neg m_1 &\sqsupseteq_{\text{SINK}[m_2]}^G x' \end{aligned}$$

as well as

$$\begin{aligned} m_1 &\sqsupseteq_{\text{SINK}}^G n \quad \text{and since } n \neq m_1 \quad m_1 \sqsupseteq_{\text{SINK}}^G x' \\ \neg m_2 &\sqsupseteq_{\text{SINK}}^G n \end{aligned}$$

Since $n \notin M$, the G -successors of n are exactly the $G_{M \not\rightarrow}$ -successors of n .

From $\neg m_1 \sqsupseteq_{\text{SINK}[m_2]}^G x'$ I obtain a G -path

$$\pi' = x', \dots, m_2 \quad \text{such that } m_1 \notin \pi'$$

If $n \rightarrow_{\text{nticd}}^{G_{M \not\rightarrow}} m_1$ I am done immediately, so I only have to deal with the opposite, i.e.: $\neg n \rightarrow_{\text{nticd}}^{G_{M \not\rightarrow}} m_1$. So either

$$\forall y \text{ s.t. } n \rightarrow_G y. \neg m_1 \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} y$$

or $\forall y \text{ s.t. } n \rightarrow_G y. m_1 \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} y$. But the latter cannot be, since if it did hold, I would have $m_1 \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} x'$ from which it would follow — since $m_1 \notin \pi'$ — that π' is not only a G -path, but also a $G_{M \not\rightarrow}$ -path, and also that any $G_{M \not\rightarrow}$ -sink-path starting in m_2 needed to contain m_1 , contradicting $\neg m_1 \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} m_2$.

I will show that there exists no $n' \neq n$ such that $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} n$. Once I have done that, I can conclude that

$$N' = \text{ipdom}_{\sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}}} (n) = \emptyset$$

and immediately get $n \rightarrow_{\text{ntind}}^* m_0$ in $G_{M \not\rightarrow}$ from $n \rightarrow_{\text{ntind}} m_0$ (in $G_{M \not\rightarrow}$) for some $m_0 \in M$, which must exist because of

$$\begin{aligned} m_1 \sqsupseteq_{\text{SINK}}^G x &\implies x \rightarrow_G^* m_1 \\ m_1 \in B &\implies \exists m'_0 \in M. m_1 \rightarrow_G^* m'_0 \end{aligned}$$

but any G -path x, \dots, m'_0 must either already be a $G_{M \not\rightarrow}$ -path, or interrupted by some $m_0 \in M$.

In order to show that there exists no $n' \neq n$ such that $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} n$, I assume the opposite. Then also $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} x$ and $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} x'$. I discern the two cases $\exists m' \in M. m' \in \pi'$ and its negation, both of which I have to lead to a contradiction. If there exists such a m' , I can choose m' as the one appearing first in π' . Then from $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} n$ I conclude that n' must appear in π' before (or at) this occurrence. In summary, I have⁴

$$\pi' = x', \dots, \underbrace{n', \dots, m', \dots, m_2}_{=: \pi'_0} \quad \text{with } m_1 \notin \pi'$$

Now, from $\neg m_1 \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} x$ I obtain some $m \in M$ and some $G_{M \not\rightarrow}$ -path

$$\pi = x, \dots, m \quad \text{such that } m_1 \notin \pi$$

and infer from $n' \sqsupseteq_{\text{SINK}}^{G_{M \not\rightarrow}} n$ that $n' \in \pi$, and hence have a prefix $\pi_0 = x, \dots, n'$ with $m_1 \notin \pi$. But then I have

$$m_1 \notin \pi_0 \pi'_0 = x, \dots, n', \dots, m', \dots, m_2$$

⁴ not demanding x', n', m', m_2 to be mutually different

in contradiction to $m_1 \sqsupset_{\text{SINK}[m_2]}^G x$.

In order derive a contradiction from $\neg \exists m' \in M. m' \in \pi'$, I observe that then π' is not only a G -path but also a $G_{M \not\rightarrow}$ -path, and $m_2 \notin M$. If $n' \in \pi'$, I derive a contradiction as before. Otherwise, I have $n' \neq m_2$, and because of $n' \sqsupset_{\text{SINK}}^{G_{M \not\rightarrow}} x'$, any $G_{M \not\rightarrow}$ -sink-path continuing π' ⁵ must contain n' , i.e.: $n' \sqsupset_{\text{SINK}}^{G_{M \not\rightarrow}} m_2$. But by Observation 7.1.8 and $m_2 \in B$ this contradicts $n' \neq m_2$. \square

This concludes the proof of Equation 7.2, but I am not quit finished: in many applications, I will need to slice w.r.t to some additional binary relation \rightarrow_d such as, e.g., data-dependence $\rightarrow_{\text{data}}$. I do not give formal proof but merely observe that under just one weak condition, this is possible as well:

Observation 7.1.9. Let G be any CFG, M any set of nodes, and \rightarrow_d be any binary relation on nodes compatible with \rightarrow_G , i.e.:

$$n \rightarrow_d m \implies n \rightarrow_G^* m$$

Then

$$\left(\rightarrow_{\text{nticd}}^G \cup \rightarrow_d \cup \rightarrow_{\text{ntiod}}^G \right)^* (M) = \left(\rightarrow_{\text{nticd}}^{G_{M \not\rightarrow}} \cup \rightarrow_d \right)^* (M)$$

7.1.1 A Direct Algorithm for Control Slicing

Given Equation 7.2, i.e.:

$$\left(\rightarrow_{\text{nticd}}^G \cup \rightarrow_{\text{ntiod}}^G \right)^* (M) = \left(\rightarrow_{\text{nticd}}^{G_{M \not\rightarrow}} \right)^* (M)$$

I can obviously compute the non-termination insensitive backward control slice of M using only $\rightarrow_{\text{nticd}}^{G_{M \not\rightarrow}}$, without the need to compute

⁵ i.e.: starting in some $G_{M \not\rightarrow}$ -successor of m_2

$\rightarrow_{\text{ntiod}}$. But I do not even need $\rightarrow_{\text{nticd}}^{G_{M \not\rightarrow}}$. Thanks to the fact that in $G_{M \not\rightarrow}$, nodes n (transitively) controlling nodes $m \in M$ have no successor in the corresponding pseudo-forest $\prec_{\text{SINK}}^{G_{M \not\rightarrow}}$, and thanks to the characterization of such n via $\rightarrow_{\text{ntind}}^{G_{M \not\rightarrow}}$, it is enough to compute $\prec_{\text{SINK}}^{G_{M \not\rightarrow}}$, and then enumerate those n following a reverse depth search in $G_{M \not\rightarrow}$ starting in M :

Lemma 7.1.1.

$$\left(\rightarrow_{\text{nticd}}^{G_{M \not\rightarrow}} \right)^* (M) = \{ n \mid n \rightarrow_{G_{M \not\rightarrow}}^* m, m \in M, \neg n \prec_{\text{SINK}}^{G_{M \not\rightarrow}} \sqcup \}$$

Proof: Directly from Observation 7.1.8 and Observation 7.1.1. \square

7.2 Nontermination Sensitive Slicing

Given the reduction to $(\rightarrow_{\text{nticd}})^*$ from equation (7.2), i.e.:

$$\left(\rightarrow_{\text{nticd}}^G \cup \rightarrow_{\text{ntiod}}^G \right)^* (M) = \left(\rightarrow_{\text{nticd}}^{G_{M/\not\rightarrow}} \right)^* (M)$$

it is natural to ask whether slices w.r.t. other ternary relations can be obtained similarly. For $\rightarrow_{\text{ntsod}}$ this is indeed the case:

Observation 7.2.1. Let G be any CFG, and M any set of nodes in G . Then

$$\left(\rightarrow_{\text{ntscd}}^G \cup \rightarrow_{\text{ntsod}}^G \right)^* (M) = \left(\rightarrow_{\text{ntscd}}^{G_{M/\not\rightarrow}} \right)^* (M)$$

I do not give proof, but merely note that the analogue of the characterization Observation 7.1.1 of the transitive closure of $\rightarrow_{\text{nticd}}$ does indeed hold, simply by replacing \sqsupset_{SINK} with \sqsupset_{MAX} .

Definition 7.2.1. Let G be any CFG, and $m \neq n$ nodes in G . Then $m \notin \text{ipdom}_{\sqsupset_{\text{MAX}}}(n)$ is said to be on a path starting in n strictly before any immediate \sqsupset_{MAX} -dominator of n — and I write $n \rightarrow_{\text{ntsnd}} m$ — iff

$$x \rightarrow_G^* \left(\text{ipdom}_{\sqsupset_{\text{MAX}}}(n) \right) \not\rightarrow m$$

Observation 7.2.2. Let G be any CFG, and $m \neq n$ nodes in G . Then

$$n \rightarrow_{\text{ntsnd}} m \iff n \rightarrow_{\text{ntscd}}^* m$$

Also, i do not need to compute $\rightarrow_{\text{ntscd}}^{G_{M/\not\rightarrow}}$, but it is enough to compute $\left\langle \rightarrow_{\text{MAX}}^{G_{M/\not\rightarrow}} \right\rangle$.

Observation 7.2.3. Let G be any CFG, and M any set of nodes in G . Also, let

$$n \left(\rightarrow_{\text{ntscd}}^{G_{M/\not\rightarrow}} \right)^* m$$

for some $m \in M$. Then, for all $n' \neq n$:

$$\neg n' \sqsupseteq_{\text{MAX}}^{G_{M \not\rightarrow}} n$$

Lemma 7.2.1.

$$\left(\rightarrow_{\text{ntscd}}^{G_{M \not\rightarrow}} \right)^* (M) = \{ n \mid n \rightarrow_{G_{M \not\rightarrow}}^* m, m \in M, \neg n <_{\text{MAX}}^{G_{M \not\rightarrow}} \perp \}$$

Proof: Directly from Observation 7.2.3 and Observation 7.2.2. □

Since ultimately, I am interested in slices respecting not only control- but also data-dependencies, I also need the following observation:

Observation 7.2.4. Let G be any CFG, M any set of nodes, and \rightarrow_d be any binary relation on nodes compatible with \rightarrow_G , i.e.:

$$n \rightarrow_d m \implies n \rightarrow_G^* m$$

Then

$$\left(\rightarrow_{\text{ntscd}}^G \cup \rightarrow_d \cup \rightarrow_{\text{ntsod}}^G \right)^* (M) = \left(\rightarrow_{\text{ntscd}}^{G_{M \not\rightarrow}} \cup \rightarrow_d \right)^* (M)$$

7.3 Weak Order Control Slices

Recall from the discussion of Figure 6.4 on page 87 that slices w.r.t $\rightarrow_{\text{nticd}}$, $\rightarrow_{\text{ntiod}}$ are larger than slices w.r.t \rightarrow_{wod} because the former demand that the slice includes any node n which can make another node in the slice unreachable. This in mind, and given the success of the reduction to $G_{M \nrightarrow}$ in equation (7.2), my following reduction is no longer surprising:

Observation 7.3.1. Let G be any CFG, and M any set of nodes in G , and let

$$G_M = G_{M \nrightarrow}^{*M}$$

be the CFG obtained from G by removing any nodes that cannot reach M , and then any outgoing edges of M . Then

$$(\rightarrow_{\text{wod}}^G)^*(M) = \left(\rightarrow_{\text{nticd}}^{G_M} \right)^*(M)$$

In algorithmic terms, Observation 7.3.1 says that for a given query $(\rightarrow_{\text{wod}}^G)^*(M)$, I can compute $\langle_{\text{SINK}}^{G_M}$ and then (by Algorithm 17 or Lemma 7.1.1) $\left(\rightarrow_{\text{nticd}}^{G_M} \right)^*(M)$.

Remark 7.3.1. When faced with multiple queries M_1, M_2, \dots , for the same CFG G , it is in fact possible to compute $\langle_{\text{SINK}}^{G_M}$ from \langle_{SINK}^G by an incremental algorithm vaguely similarly to the computation of $\langle_{m'_2}$ from \langle_{m_2} in Observation 6.7.8 on page 107, but I do not describe it here. Note that in the extreme case, where G consists of a single sink N , \langle_{SINK}^G is completely useless for the computation of $\langle_{\text{SINK}}^{G_M}$ since then \langle_{SINK}^G is just one big cycle containing every node N , and carries no information about the control structure of G at all.

7.4 Weak Control Closures

Recall (Definition 6.5.2) that a set $S \subseteq N$ was defined to be weakly control closed in G iff all vertices $n \notin S$ reachable from S are S -weakly committing (i.e.: $\text{obs}_S^G(n)$ is at most a singleton) in G . In [Dan+11], the same authors also define the notion of *weakly deciding* nodes:

Definition 7.4.1 (adapting [Dan+11], Definition 50). Given a CFG $G = (N, E)$ and a set $M \subseteq N$, a node $n \notin M$ is M -weakly deciding in G if there exist paths

$$\begin{aligned} \pi_l &= \underbrace{n, \dots, m_l}_{\notin S} \quad \text{such that} \quad m_l \in M \\ \pi_r &= \underbrace{n, \dots, m_r}_{\notin S} \quad \text{such that} \quad m_r \in M \end{aligned}$$

with $\pi_l \cap \pi_r = \{n\}$. Specifically: $m_l \neq m_r$.

Definition 7.4.2. I denote with $\text{WD}_G(M)$ the set of M -weakly deciding nodes in G , and with $\text{WCC}_G(M)$ the smallest set $S \supseteq M$ that is weakly control closed in G .

Lemma 7.4.1 ([LKL18], Property 4). Recall that for a CFG $G = (N, E)$ and $M \subseteq N$, $N^{M \rightarrow^*}$ is the set of nodes reachable from M in G .

$$\text{WCC}_G(M) = (M \cup \text{WD}_G(M)) \cap N^{M \rightarrow^*}$$

Weak control closures and the set of weakly deciding nodes relate to weak order dependence and nontermination-insentitive control dependence as follows:

Observation 7.4.1.

$$(\rightarrow_{\text{wod}}^G)^*(M) = M \cup \text{WD}_G(M)$$

Also

$$\begin{aligned} \text{WCC}_G(M) &= (\rightarrow_{\text{wod}}^G)^*(M) \cap N^{M \rightarrow^*} \\ &= (\rightarrow_{\text{wod}}^{G_M})^*(M) && \text{for } G_M = G^{M \rightarrow^*} \\ &= \left(\rightarrow_{\text{nticd}}^{G_M} \right)^*(M) \cap N^{M \rightarrow^*} && \text{for } G_M = G_{M \nearrow}^{*M} \\ &= \left(\rightarrow_{\text{nticd}}^{G_M} \right)^*(M) && \text{for } G_M = G_{M \nearrow}^{M \rightarrow^* M} \end{aligned}$$

7.5 The Role of $\rightarrow_{\text{nticd}}$ for Nontermination Insensitive Slices

In review of this chapter, I want to comment on the role of nontermination insensitive control dependence $\rightarrow_{\text{nticd}}$, originally defined in [Ran+07].

While it was meant to generalize standard control dependency \rightarrow_{cd} , the fact that $\rightarrow_{\text{nticd}}$ does not capture any dependencies *within* control-sinks begs the question: What *is* the purpose of $\rightarrow_{\text{nticd}}$?

In fact in [Amt08], $\rightarrow_{\text{nticd}}$ was abandoned in favor of \rightarrow_{wod} by its inventors. But in Chapter 6 I showed that “with a little help from” nontermination insensitive order dependence $\rightarrow_{\text{ntiod}}$, nontermination insensitive control dependence $\rightarrow_{\text{nticd}}$ is indeed a sound (and minimal!) basis for nontermination insensitive slicing.

Even more, by moving from G to an appropriate CFG G_M , $\rightarrow_{\text{nticd}}^{G_M}$ *alone* can in fact be used to obtain what I think are the three relevant notions of nontermination insensitive slicing of arbitrary CFG:

$$\begin{aligned} (\rightarrow_{\text{nticd}}^G \sqcup \rightarrow_{\text{ntiod}}^G)^*(M) &= \left(\rightarrow_{\text{nticd}}^{G_M}\right)^*(M) \text{ for } G_M = G_{M \not\rightarrow} \\ (\rightarrow_{\text{wod}}^G)^*(M) &= \left(\rightarrow_{\text{nticd}}^{G_M}\right)^*(M) \text{ for } G_M = G_{M \not\rightarrow}^{*M} \\ \text{WCC}_G(M) &= \left(\rightarrow_{\text{nticd}}^{G_M}\right)^*(M) \text{ for } G_M = G_{M \not\rightarrow}^{M \rightarrow *M} \end{aligned}$$

Due to the fact that all three G_M are of the form $G'_{M \not\rightarrow}$, Lemma 7.1.1 is applicable and I do *not* need to explicitly compute $\rightarrow_{\text{nticd}}^{G_M}$. The only requirement is the corresponding nontermination insensitive postdominance pseudo-forest $\prec_{\text{SINK}}^{G_M}$.

I stress that in this application, the crucial step in the generalization from postdominance trees \prec_{POST} for CFG with unique exit node n_x , to pseudo-forests \prec_{SINK} for arbitrary CFG is not the *pseudo*-part (i.e.:

the presence of cycles in \langle_{SINK} , but the *forest*-part: The transformation $G \mapsto G'_{M \dashv}$ only ever introduce proper *trees* (among them: those with roots M). “Pseudoness” is not important here—in fact, the second and third variant of G_M lead to proper forests, with the each $m \in M$ being a root.

Originally, I specifically designed Algorithm 17 from Appendix C to compute slices in G_M , and it does this adequately. Computation via Lemma 7.1.1, however, is much simpler, and I probably would not have bothered with design of Algorithm 17 but for the fact that at that time, I had not yet discovered Lemma 7.1.1, yet! As is, Algorithm 17 retains useful for the case of multiple queries M_1, M_2, \dots for CFG G without non-trivial sinks.

Summary

- Nontermination insensitive order dependence $\rightarrow_{\text{ntiod}}$ can be of size $\mathcal{O}(|N|^3)$.
- But nontermination insensitive slices can be computed directly from nontermination insensitive postdominance \exists_{SINK} , in a suitable graph G_M obtained from the slicing criterion M .

8 Performance Benchmarks

Atticus told me to delete the adjectives and I'd have the facts.

(Harper Lee — To Kill a Mockingbird)

I evaluated the performance of all new algorithms from Part I of this thesis on three classes of graphs:

1. Control flow graphs of Java methods, as generated by the JOANA system for various third party Java programs
2. Randomly generated graphs $G = (N, E)$ usually with $|E| = 2|N|$, as generated by the standard generator from the JGraphT[NP19] library
3. Variants of “ladder” graphs (see Figure 7.1 on page 119), intended to expose “Bad Case” behavior.

Except for the benchmarks concerned with nontermination insensitive order dependence \rightarrow_{ntiod} , the ladder graphs I use will be unique-exit-node ladder graphs. This sometimes allows me to directly compare with existing algorithms for such graphs.

All benchmarks in this chapter were made on a “desktop workstation” class computer with an Intel i7-6700 CPU at 3.40GHz, and 64 gigabyte RAM. I implemented the algorithms in the Java programming language, and used the OpenJDK Java 9 VM to run them. All benchmarks were run using the Java Microbenchmark Harness JMH[Cor20].

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 which can be finished within 1 second. For example, the data point at $|N| = 21076$, time = 18ms in Figure 8.1a stands for the average of at least ≈ 50 runs of the benchmark that finished within 1 second. On the other hand, the data point in at $|N| = 65000$, time = 88s in Figure 8.1c 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 benchmark in the upper left and upper right of Figure 8.3 suggest that my new algorithm for the computation of nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$ appears to scale “almost linearly” for “average” input CFG, while the original algorithm from [Ran+07] is clearly grows super-linearly for such graphs.

The central outcome of the benchmarks in this chapter is:

1. For “average” CFG, all my algorithms for control dependence variants for arbitrary graphs offer performance “almost linear” in the size of the graph.
2. But for “bad case” CFG, some algorithms perform for decidedly super-linear, and become impractical for very large such graphs.

I also include benchmarks for *timing sensitive* control dependence, *timing* dependence, and *timing sensitive* postdominance, as to be introduced later in part II of this thesis. There the central outcome is:

1. Timing sensitive postdominance can be computed in “almost linear” time for “average” CFG.
2. But even for “average” CFG, timing sensitive control dependence and timing dependence scale super-linearly.



8.1 Nontermination Sensitive Postdominance

In Subsection 5.2.1, I introduced Algorithm 5 for the computation of maximal path postdominance \exists_{MAX} , represented as a pseudo-forest \langle_{MAX} . This algorithm requires the computation of least common ancestors lca_{\langle} in pseudo-forests \langle . I introduced two variants thereof (Algorithm 3 on page 47, and Algorithm 18 on page 396), of which I use the latter.

The Algorithm 5 implements chaotic iteration, by reinserting into a workset those nodes affected by modification to the pseudo-forest. In contrast, the Algorithm 19 on page 397 repeatedly iterates in a fixed node order.

Both these variants do not specify an iteration order (i.e.: Algorithm 5 does not specify which node x to remove from the work set at the start of each iteration, and Algorithm 19 does not specify the initial order of nodes in the workqueue). By default, my implementations 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 Algorithm 5 this means that at the beginning of each iteration, the *rightmost* node (by topological ordering) is removed from the workset.

For Java CFG and randomly generated graphs (neither necessarily with unique exit node), Algorithm 5 (+) and Algorithm 19 (▼) behave similarly (Figure 8.1a and Figure 8.1b). Ladder graphs expose non-linear *bad-case* behavior (Figure 8.1c). This is even more pronounced when additionally, I deliberately choose a bad iteration order (Figure 8.1d).

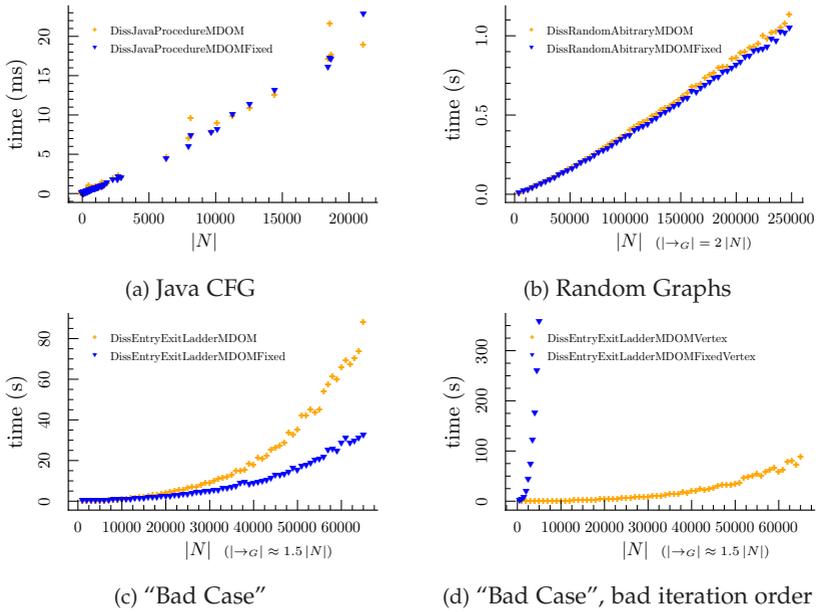


Figure 8.1: Computation of \langle_{MAX} .

8.2 Nontermination Insensitive Postdominance

In Section 5.3, I introduced Algorithm 6 for the computation of sink path postdominance \exists_{SINK} , represented as a pseudo-forest $<_{\text{SINK}}$. Just as before, it uses Algorithm 18 for the computation of least common ancestors $\text{lca}_{<}$.

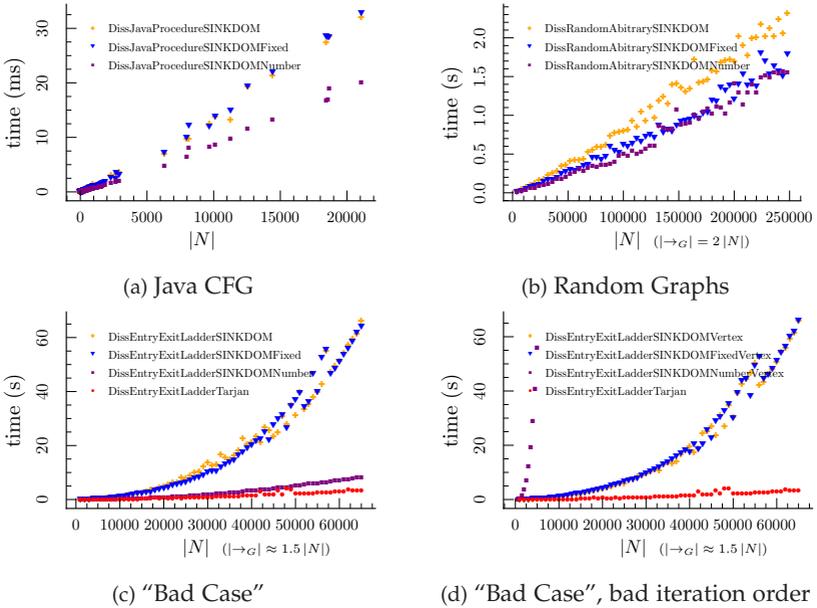
Again, Algorithm 6 implements chaotic iteration, by reinserting into a workset those nodes affected by modification to the pseudo-forest. I also implemented a variant of Algorithm 6 in which the downward fixed point phase repeatedly iterates a workqueue of nodes in a fixed node order. Again, neither variant specifies an iteration order. As before, implementations by defaults orders 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*, which are exactly those components without outgoing edges.

In Section D.2 in the appendix, I introduce Algorithm 24 for the computation of \exists_{SINK} . This algorithm computes least common ancestors $\text{lca}_{<}$ by comparison of postorder numbers, via Algorithm 23. By default, nodes are iterated in reversed-topological order.

For Java CFG (Figure 8.2a) the fixed-iteration order variant of Algorithm 6 (▼) performs on par with the original Algorithm 6 (+). For randomly generated graphs (Figure 8.2b)¹ the variant (▼) appears to perform a bit better than the original (+) for very large graphs, roughly on-par with Algorithm 24 (■).

Using reversed-topological iteration order, ladder graphs (Figure 8.2c) expose non-linear *bad-case* behavior only for Algorithm 6 (+) and its variant (▼). Even with a deliberately bad iteration order, I could not produce much worse performance for these two algorithm (Figure 8.2d). Presumably, the iteration in upward phase SINK_{up} (which is the same both variants) finds a good-enough approximation. On

¹ just as the Java CFG, not necessarily with unique exit node

Figure 8.2: Computation of \prec_{SINK} .

the other hand, Algorithm 24 (■) is affected heavily by a deliberately bad iteration order.

The ladder graphs I use are unique-exit-node ladder graphs, since if I used the ladder graph from Figure 7.1 on page 119 *as shown*, the whole graph would form one big control-sink, and \prec_{SINK} one big cycle, making the performance comparison moot. This also allows me to directly compare with an implementation of the algorithm by Lengauer, Tarjan [LT79] (●).

8.3 Generalized Postdominance Frontiers

In Section 3.2, I introduced Algorithm 1 for the computation generalized postdominance frontiers.

Nontermination Sensitive Control Dependence When instantiated with $<_{\text{MAX}}$, this yields an algorithm for nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$. The benchmarks include the computation time of both Algorithm 1, and $<_{\text{MAX}}$ via Algorithm 19 (▼). I compare with Algorithm 14 from [Ran+07] as shown on page 363 (+).

For Java CFG and randomly generated graphs, Algorithm 14 becomes impractical for moderately sized graphs, while Algorithm 1 performs well even for very large graphs (Figure 8.3, upper left and right).

Ladder graphs expose non-linear *bad-case* behavior even for Algorithm 1 (Figure 8.3c). This cannot be circumvented, since in these ladder graphs, the size of the relation $\rightarrow_{\text{nticd}}$ is quadratic in the number of nodes (similar to the size of relation $\rightarrow_{\text{ntiod}}$ being cubic for the ladder graphs on page 119).

Nontermination Insensitive Control Dependence When instantiated with $<_{\text{SINK}}$, Algorithm 1 yields an algorithm for nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$. The benchmarks include the computation time of both Algorithm 1, and \exists_{SINK} via the variant of Algorithm 19 with fixed iteration order (▼). Since the Algorithm 15 from [Ran+07] is incorrect, I do not compare with it. Instead, I compare with Algorithm 16 as shown on page 380 (+).

For Java CFG and randomly generated graphs, Algorithm 16 becomes impractical for moderately sized graphs, while Algorithm 1 performs well even for very large graphs (Figure 8.4, upper left and right).

Ladder graphs expose non-linear *bad-case* behavior even for Algorithm 1 (Figure 8.4c), which again cannot be circumvented.

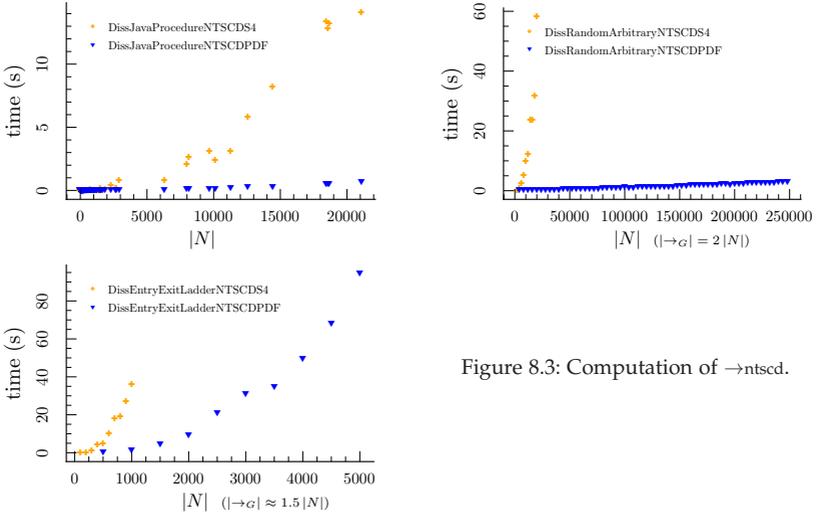


Figure 8.3: Computation of \rightarrow_{ntscd} .

(c) "Bad Case"

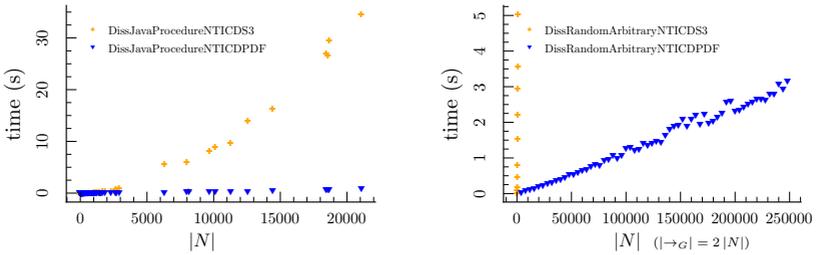
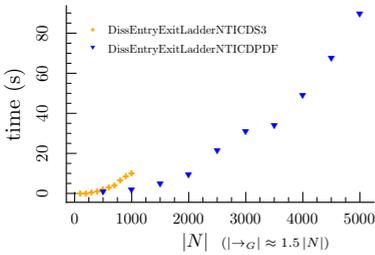


Figure 8.4: Computation of \rightarrow_{nticd} .



(c) "Bad Case"

8.4 Control Slices

Given $\rightarrow_{\text{nticd}}$, corresponding slices $(\rightarrow_{\text{nticd}})^*(M)$ of sets of nodes M can be computed in linear time. The relation $\rightarrow_{\text{nticd}}$, however, can be of size quadratic in the number of nodes. In Appendix C, I give a generalization of an algorithm based on *DJ-Graphs*[SG95] (i.e.: dominance trees enriched by *join* edges) to the pseudo-forests $<_{\text{SINK}}$, enriched with *conditional* edges (Algorithm 17). The algorithm requires only the postdominance pseudo-forest, but not the corresponding control dependence relation.

NTICD Slices (Figure 8.5) For slices via $\rightarrow_{\text{nticd}}$ (+), computation time includes the computation of $\rightarrow_{\text{nticd}}$. The computation time for Algorithm 17 (▼) includes the computation of $<_{\text{SINK}}$. I also implemented a variant of Algorithm 17 that is based on $<_{\text{SINK}}$ as represented by postorder numbers, i.e.: as computed by Algorithm 24 on page 406. Marker ■ include computation time of both.

NTSCD Slices (Figure 8.6) For slices via $\rightarrow_{\text{ntscd}}$ (+), computation time includes the computation of $\rightarrow_{\text{ntscd}}$. Algorithm 17 also works $\rightarrow_{\text{ntscd}}$ slices, by using $<_{\text{MAX}}$ instead of $<_{\text{SINK}}$, and similar in the definition

$$n \rightarrow_C m \Leftrightarrow n \rightarrow_G m \wedge m \notin \text{ipdom}_{\sqsupset_{\text{MAX}}}(n)$$

of conditional edges. The computation time (▼) includes the computation of $<_{\text{MAX}}$.

All slices are w.r.t. randomly selected nodes M , with $|M| = 5$.

Note that in the “bad case” plots, I can show the $|N|$ -axis up to $|N| = 65000$, while before in Section 8.3, I could only up to $|N| = 5000$, since then computation time began to approach and exceed 60s.

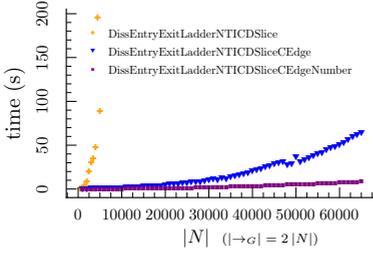
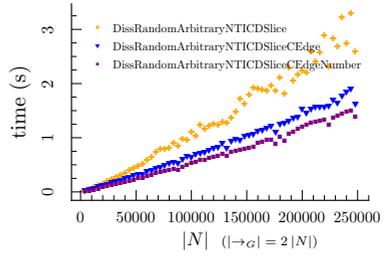
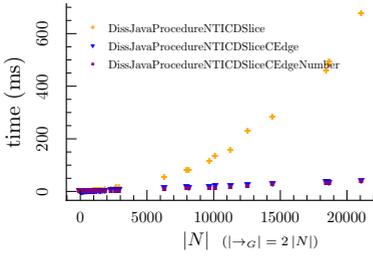


Figure 8.5: Computation of $(\rightarrow_{nticd})^*(M)$

(c) “Bad Case”

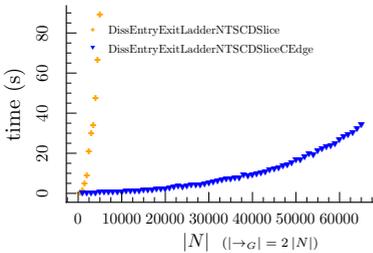
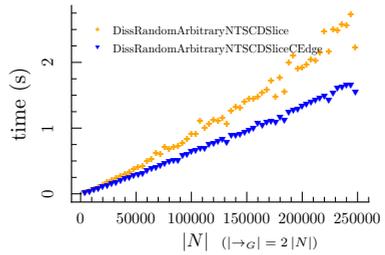
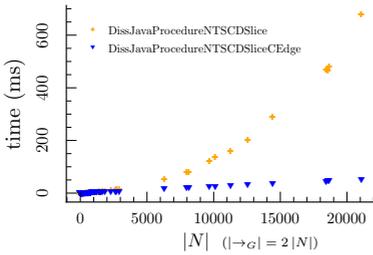


Figure 8.6: Computation of $(\rightarrow_{ntscd})^*(M)$

(c) “Bad Case”

8.5 Nontermination Insensitive Order Dependence

By the slogan

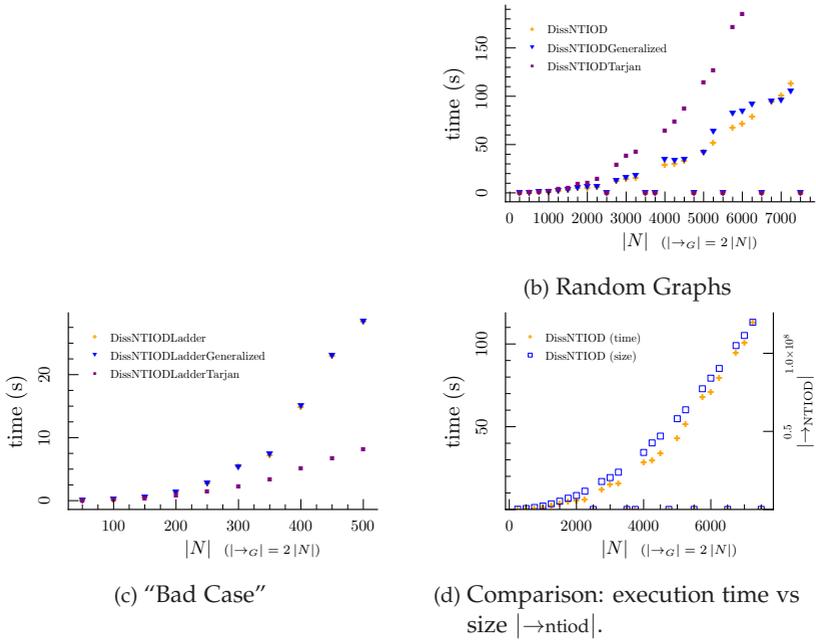
$$\rightarrow_{\text{ntiod}} \approx \left(\sum_{M \in \mathcal{M}} |M| \right) \times \rightarrow_{\text{nticd}}$$

from Subsection 6.7.1, computation of $\rightarrow_{\text{ntiod}}$ for a graph G can be reduced to the computation of $\rightarrow_{\text{nticd}}$ on subgraphs of G . These subgraphs consists of the graphs *control sinks* M (and the those nodes between M and those nodes n of which M are immediate postdominators). Specifically, these subgraphs all have a unique exit node n_x , and hence there it holds that $\rightarrow_{\text{nticd}} = \rightarrow_{\text{cd}}$. This means that classic algorithms [LT79; Cyt+91] can be used. Their computation time form the baseline (■) in Figure 8.7.

Also in Subsection 6.7.1, I introduced two schemes (Observation 6.7.7 shown as +, and Observation 6.7.8 shown as ▼) that avoided full computation of \prec_{SINK} for every node $m \in M$ in some control sink M of G . These are based on phase $\text{SINK}_{\text{down}}$ from Algorithm 19.

The plot in Figure 8.7b contains data points (+, ▼ and ■) near the $|N|$ -axis. These indicate that for the given graph (with size $|N|$), the relation $\rightarrow_{\text{ntiod}}$ is either empty or very small (i.e.: there are no non-trivial control sinks in G , or they are small). In other words: aside from *one* computation of \prec_{SINK} for the whole graph, my algorithms for $\rightarrow_{\text{ntiod}}$ require additional computation only insofar it contains non-trivial control sinks.

This can also be seen in Figure 8.7d, in which the execution time of the computation due to Observation 6.7.7 (+) is compared with the size $|\rightarrow_{\text{ntiod}}|$ (□) of the computed relation. I conjecture that in practice, they are asymptotically equal, i.e.: in practice, Observation 6.7.7 is asymptotically optimal under all algorithms that compute an explicit representation of $\rightarrow_{\text{ntiod}}$.

Figure 8.7: Computation of \rightarrow_{ntiod} .

In randomly generated graphs (Figure 8.7b), the two more complicated schemes (+) and (\blacktriangledown) give an advantage over my baseline scheme (\blacksquare) (Figure 8.7b). In “bad case” ladder graphs (Figure 8.7c), the situation is reversed.

8.6 Nontermination Insensitive Slices

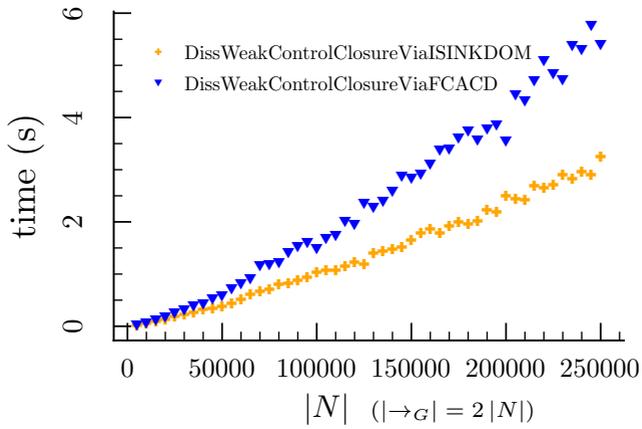
In Section 7.5, I observed that for any given set M of nodes in an arbitrary graph G , I can reduce three different notions of nontermination insensitive slicing to $\rightarrow_{\text{nticd}}$ -slicing in modified graphs G_M , as follows:

$$\begin{aligned} (\rightarrow_{\text{nticd}}^G \cup \rightarrow_{\text{ntiod}}^G)^*(M) &= \left(\rightarrow_{\text{nticd}}^{G_M}\right)^*(M) \text{ for } G_M = G_{M \nrightarrow} \\ (\rightarrow_{\text{wod}}^G)^*(M) &= \left(\rightarrow_{\text{nticd}}^{G_M}\right)^*(M) \text{ for } G_M = G_{M \nrightarrow}^* \\ \text{WCC}_G(M) &= \left(\rightarrow_{\text{nticd}}^{G_M}\right)^*(M) \text{ for } G_M = G_{M \nrightarrow}^{*M} \end{aligned}$$

These equations each yield, together with Lemma 7.1.1, an algorithm for computing M -slices *directly* from $\langle_{\text{SINK}}^{G_M}$, without the need to compute $\rightarrow_{\text{nticd}}^{G_M}$.

With regard to performance, these three reduction behave essentially the same. I show the computation time for weak control slices $\text{WCC}_G(M)$ in Figure 8.8. The best algorithm² other than mine is from [LKL18] (\blacktriangledown), with which I compare my algorithm (\oplus) in randomly generated graphs. The Implementation of the algorithm from [LKL18] is my own, and appears (with respect to constant factors) to be faster than the authors original implementation.

² that I am aware of

Figure 8.8: Computation of $WCC_G(M)$.

8.7 Timing Sensitive Algorithms

Together, Algorithm 11, 10, 9, 10 and 8 implement timing sensitive control dependence $\rightarrow_{\text{tscd}}$. Timing dependence \rightarrow_{td} can be computed via Observation 10.2.1 and Algorithm 8.

Timing Sensitive Control Dependence The benchmarks include the computation time of all sub-algorithms (+). Ladder graphs expose non-linear *bad-case* behavior.

Timing Dependence By Observation 10.2.1, for each graph, I compute $|N|$ transitive reductions $\prec_{\text{TIME}}^{G_m}$ of transitive timing sensitive postdominance $\sqsupseteq_{\text{TIME}}^{G_m}$, for certain subgraphs G_m of the control flow graph G . I use Algorithm 8. The benchmarks include all $|N|$ invocations for each graph (+).

Timing Sensitive Postdominance Pseudo-Forests I also show benchmarks for the computation of computation of \prec_{TIME} only, via Algorithm 8.

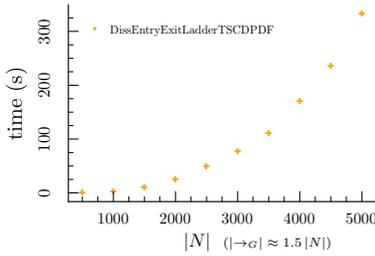
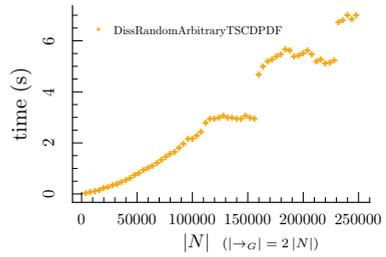
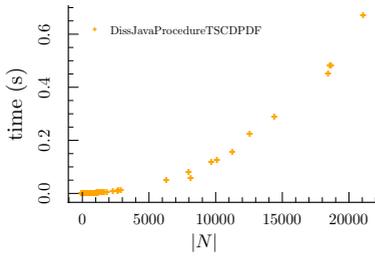


Figure 8.9: Computation of $\rightarrow_{\text{tscd}}$.

(c) "Bad Case"

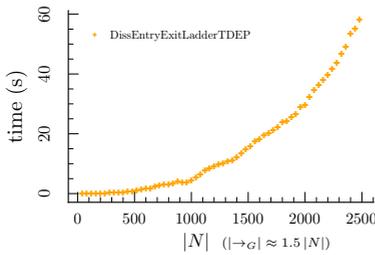
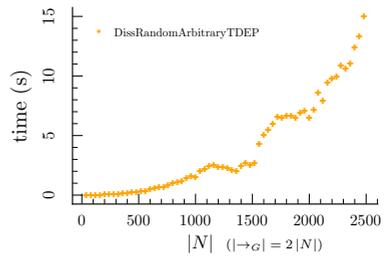
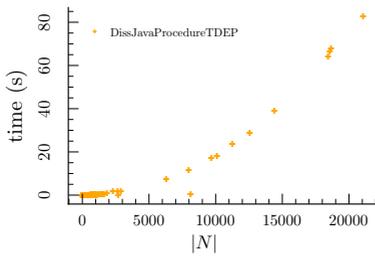


Figure 8.10: Computation of $\rightarrow_{\text{timing}}$.

(c) "Bad Case"

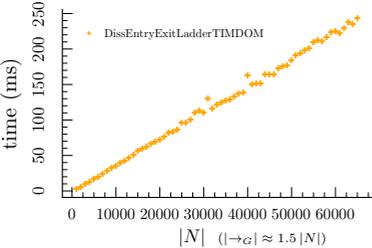
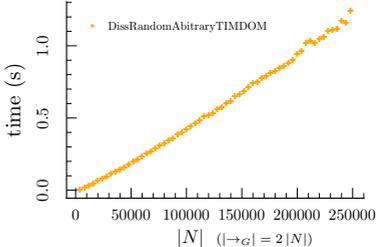
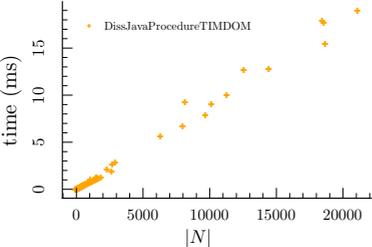


Figure 8.11: Computation of $\langle \text{TIME} \rangle$.

(c) "Bad Case"

Part II

Timing Sensitive Dependency Analysis

Endlich Zeitumstellung zu nehmen (it's time)
Den Wendehals noch schnell umzudrehen
Endlich wieder den Schweinehund ausführen
Zeit umzukehren vor der eigenen Haustür
Höchste Zeit es beim Namen zu nennen
Dass es höchste Zeit ist zum Farbe bekenn', ja

(Dendemann — Zeitumstellung)

9 Timing Sensitive Control Dependence

“Wonderful”, the Flatline said, “I never did like to do anything simple when I could do it ass-backwards.”

(William Gibson — Neuromancer)

In both Section 6.3 and Section 6.8, I established soundness of slicing by virtue of trace based notions of observation. There, I defined a trace t to be a sequence of *partial edges* $(n, n') \in E \cup (N_x \times \{\perp\})$ 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 , or infinite with

$$t = (n_e, n_1), (n_1, n_2), \dots$$

Given a set S of “observable” nodes and a trace t , I defined the S -observation $t|_S$ of t to be the sub-sequence of t containing only edges (n, n') with $n \in S$. In terms of an *attacker model*, this means that I assume an attacker to observe exactly those choices made at nodes $n \in S$. Specifically, I assume that an attacker can observe neither the nodes in a subtrace between observable nodes, nor the *time spent* between two observable nodes (i.e: the *length* of the subtrace between two observable nodes). In other words, I make the assumption that the attacker has no clock! As a result, even programs deemed secure under trace equivalence, i.e.: programs such that

$$t_i \sim_S t_{i'}$$

for any two low-equivalent inputs $i \sim_S i'$ may yet have *external timing leaks*. A most trivial example is shown in Figure 9.1a, with observable



(a) A CFG with external timing leak (b) A CFG without timing leak

Figure 9.1: Dependence of execution time of m_x on n .

nodes $S = \{m, m_x\}$. Regardless of the choice made at n , all inputs i, i' starting in m have the same observable trace

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

consisting of two pseudo-edges and terminating in m_x , hence: $t_i \sim_S t_{i'}$. However, if equipped with a suitably precise clock, and assuming a uniform execution time of one time unit u per edge, an attacker will observe m_x after $3u$ have passed if the input i chooses n'' at n , but only after $5u$ if i chooses n' at n , exposing an external timing leak. This becomes obvious if I annotate each edge in traces $t_i, t_{i'}$ with its execution time, and then compare their S -observation:

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

On the other hand, the program in Figure 9.1b has no timing leak. That is, even if I annotate each edge in the observable trace with its execution time, all inputs i, i' starting in m have the same observable *clocked* trace

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

In this chapter, I fully develop an approach for *timing sensitive* slicing. Unlike nontermination (in)-sensitive slicing, which in the general case required additional ternary notions of dependence, I propose a single (binary) notion of *timing sensitive* control dependence $\rightarrow_{\text{tscd}}$ which will result in slices that are both sound and minimal with regard to observational equivalence of *clocked* traces, justifying the slogan

$$\text{Timing Sensitive IFC} = (\rightarrow_{\text{tscd}} \cup \rightarrow_{\text{data}})^*$$

I will propose an efficient algorithm for the computation of $\rightarrow_{\text{tscd}}$ which will *almost, but not quite* follow the development of algorithms for $\rightarrow_{\text{nticd}}$ and $\rightarrow_{\text{ntscd}}$ I exercised in Chapter 3 and Chapter 5. Specifically, I will

1. Propose a notion $\sqsubseteq_{\text{TIME[FIRST]}}$ of *timing sensitive* postdominance
2. Give a least fixed point characterization of $\sqsubseteq_{\text{TIME[FIRST]}}$
3. Propose a notion $\rightarrow_{\text{tscd}}$ of *timing sensitive* control dependence. It will be based on $\sqsubseteq_{\text{TIME[FIRST]}}$ *almost* the same way that $\rightarrow_{\text{ntscd}}$ is based on \sqsubseteq_{MAX} .

I will (and in general: can) *not* give an algorithm to compute a transitive, reflexive reduction $>_{\text{TIME[FIRST]}}$ of $\sqsubseteq_{\text{TIME[FIRST]}}$. Instead, I propose another notion $\sqsubseteq_{\text{TIME}}$ of *timing sensitive* postdominance that is *almost* the same as $\sqsubseteq_{\text{TIME[FIRST]}}$. Then, I will

5. Give a least fixed point characterization of $\sqsubseteq_{\text{TIME}}$
6. Propose an algorithm to compute a transitive, reflexive reduction $>_{\text{TIME}}$ of $\sqsubseteq_{\text{TIME}}$. The Algorithm will be *almost* the same as Algorithm 5 for $>_{\text{MAX}}$.
7. Propose an auxiliary notion of F_n of “fuel available” at nodes n which will allow me to characterize $\sqsubseteq_{\text{TIME[FIRST]}}$ in terms of $\sqsubseteq_{\text{TIME}}$. Informally:

$$\sqsubseteq_{\text{TIME[FIRST]}} = \sqsubseteq_{\text{TIME}} + F$$

8. Propose a “post processing” algorithm that — given \succ_{TIME} — computes F .
9. Characterize — in terms of \prec_{TIME} and F — for very node n the set $\text{ipdom}_{\exists_{\text{TIME[FIRST]}}}(n)$ of immediate $\exists_{\text{TIME[FIRST]}}$ postdominators of n .
10. Give an algorithm (similar to the generalized Cytron Algorithm from Chapter 3) that — given $\text{ipdom}_{\exists_{\text{TIME[FIRST]}}}$ — computes the timing sensitive postdominance frontiers and hence: $\rightarrow_{\text{tscd}}$.

9.1 Timing Sensitive Control Dependence

In Figure 9.1a, the node m_x does nontermination sensitively postdominate n ($m_x \sqsupseteq_{\text{MAX}} n$) because any maximal path starting in either G -successor of n must contain m_x (i.e.: both $m_x \sqsupseteq_{\text{MAX}} n'$ and $m_x \sqsupseteq_{\text{MAX}} n''$), and so must any maximal path starting in n . Remember that \sqsupseteq_{MAX} was defined via

$$m \sqsupseteq_{\text{MAX}}^G n \Leftrightarrow \forall \pi \in {}_n\Pi_{\text{MAX}}^G. m \in \pi$$

where ${}_n\Pi_{\text{MAX}}^G$ is the set of maximal G -paths starting in n .

In order to account for the different *timing* of the (first) occurrence of m_x in maximal paths starting in n , my following definition is completely natural:

Definition 9.1.1. Let G be any CFG, n, m any nodes in G . Given any path

$$\pi = m_0, m_1, m_2, \dots$$

I 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 $i < k$, I say that m *first* appears in π at position k , and write $m \in_{\text{FIRST}}^k \pi$.

Furthermore, I say that m timing-sensitively postdominates n at position $k \in \mathbb{N}$ in G iff on all maximal G -paths starting in n , m first appears at position k . I omit “in G ” whenever possible, and just say that m timing-sensitively postdominates n iff this is the case for some k . Formally:

$$\begin{aligned} m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k \text{ in } G} n &\Leftrightarrow \forall \pi \in {}_n\Pi_{\text{MAX}}^G. m \in_{\text{FIRST}}^k \pi \\ m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^G n &\Leftrightarrow \forall \pi \in {}_n\Pi_{\text{MAX}}^G. m \in_{\text{FIRST}}^k \pi \text{ for some } k \in \mathbb{N} \end{aligned}$$

Remark 9.1.1. Obviously, given m and n , the k such that $m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k n$ (if it exists!) is unique.

Following Definition 3.1.2 on page 16, but taking into account that $\cdot \sqsupseteq_{\text{TIME}[\text{FIRST}]} \cdot$ is a ternary relation, I can immediately define the following timing sensitive notion of control dependence:

Definition 9.1.2. Let G be any CFG, n, m any nodes in G . Then m is said to be *timing sensitively control-dependent* on n , written $n \rightarrow_{\text{tscd}} m$, if there exists G successors n_l and n_r of n , and some $k \in \mathbb{N}$ such that $m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k$ -post dominates n_l , but *not*: $m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k$ -post dominates n_r . Formally: $n \rightarrow_{\text{tscd}} m \Leftrightarrow$

$$\begin{aligned} & m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k n_l \quad \text{and} \\ \neg & m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k n_r \end{aligned}$$

for some $k \in \mathbb{N}$ and n_l, n_r such that $n \rightarrow_G n_l$ and $n \rightarrow_G n_r$.

Remember from Theorem 5.1.2 that \sqsupseteq_{MAX} is the least fixed point of the following rule system D

$$\frac{}{n \sqsupseteq n} \text{D}^{\text{self}} \quad \frac{\forall p \rightarrow_G x. m \sqsupseteq x \quad p \rightarrow_G^* m}{m \sqsupseteq p} \text{D}^{\text{suc}}$$

in the lattice $(2^{N \times N}, \sqsubseteq)$.

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

Proposition 9.1.1. Let G be a CFG and T_{FIRST} be the rule-system¹

$$\frac{}{n \sqsupseteq^0 n} \text{T}_{\text{FIRST}}^{\text{self}} \quad \frac{\forall p \rightarrow_G x. m \sqsupseteq^k x \quad m \neq p \quad p \rightarrow_G^+ m}{m \sqsupseteq^{k+1} p} \text{T}_{\text{FIRST}}^{\text{suc}}$$

Then $\sqsupseteq_{\text{TIME}[\text{FIRST}]} = \mu \text{T}_{\text{FIRST}}$.

¹ over a ternary relation $\cdot \sqsupseteq \cdot$.

Remark 9.1.2. The condition $p \rightarrow_G^* m$ is redundant for nodes p that have *some* successor x , since I only consider the least (but never: the greatest) fixed point of \top_{FIRST} .

The timing sensitive postdominance for the CFG from the earlier example in Figure 5.1a is shown in Figure 9.3b. Figure 9.3c and Figure 9.3d show the corresponding non-termination sensitive and timing sensitive control dependencies. Note, for example, that $7 \rightarrow_{\text{tscd}} 8$ because a choice $7 \rightarrow_G 11$ can delay node 8, but in contrast: $\neg 7 \rightarrow_{\text{ntsod}}^* 8$, because no choice at node 7 can *prevent* node 8 from being executed. It is *not* the case that, in general, $n \rightarrow_{\text{ntsod}} m$ implies $n \rightarrow_{\text{tscd}} m$. For example: $2 \rightarrow_{\text{ntsod}} 8$, but $\neg 2 \rightarrow_{\text{tscd}} 8$. What *does* hold here is $2 \rightarrow_{\text{tscd}}^* 8$ via $2 \rightarrow_{\text{tscd}} 7 \rightarrow_{\text{tscd}} 8$. In fact, timing sensitive control independence is *transitively* a stricter requirement than non-termination sensitive control independence:

Observation 9.1.1. Let $G = (N, E)$ be any CFG, and $M \subseteq N$ any set of nodes. Then the timing sensitive backward slice of M contains the nontermination sensitive backward slice of M :

$$\begin{aligned} (\rightarrow_{\text{tscd}})^*(M) &\supseteq (\rightarrow_{\text{ntsod}} \cup \rightarrow_{\text{ntsod}})^*(M) \\ &= (\rightarrow_{\text{ntsod}} \cup \rightarrow_{\text{dod}})^*(M) \end{aligned}$$

It is worth noting that the $\rightarrow_{\text{tscd}}$ slice in Observation 9.1.1 does *not* require a timing sensitive analogue of the relation $\rightarrow_{\text{ntsod}}$. Recall that the necessity of either \rightarrow_{dod} or $\rightarrow_{\text{ntsod}}$ was motivated by the canonical irreducible graph from Figure 6.1 on page 66. Essentially the same CFG is shown in Figure 9.2. The problem with $\rightarrow_{\text{ntsod}}$ was that neither m_1 nor m_2 is nontermination sensitively control dependent on n , yet the decision at n determines which node is observed next. In contrast, both m_1 and m_2 are timing-sensitively control dependent on n (e.g.: $n \rightarrow_{\text{tscd}} m_1$ because $m_1 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^1 n'$, but $\neg m_1 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^1 n''$, and also: $m_1 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^2 n''$, but $\neg m_1 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^2 n'$).

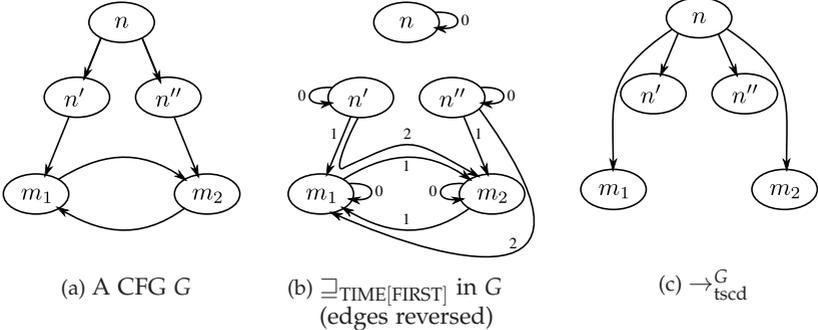
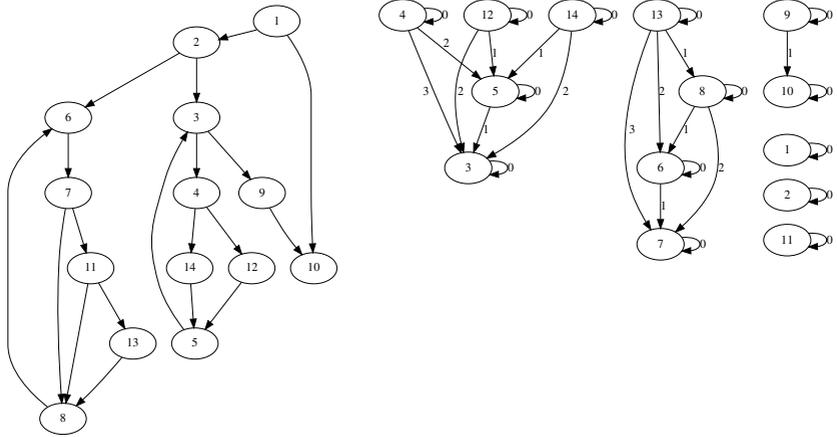
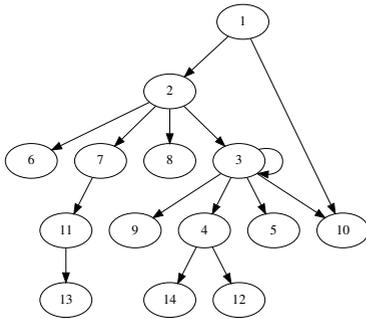


Figure 9.2: The canonical irreducible graph, where neither $n \rightarrow_{\text{ntscd}} m_1$ nor $n \rightarrow_{\text{ntscd}} m_2$.

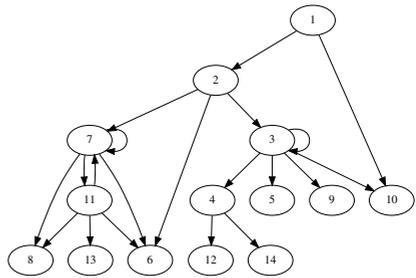


(a) A CFG

(b) Its relation $\sqsupseteq_{\text{TIME}[\text{FIRST}]}^k$ (edges reversed)



(c) Its non-termination sensitive control dependence \to_{ntscd}



(d) Its timing sensitive control dependence \to_{tscd}

Figure 9.3: Timing sensitive postdominance. Edges $n \xrightarrow{k} m$ indicate $m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k n$.

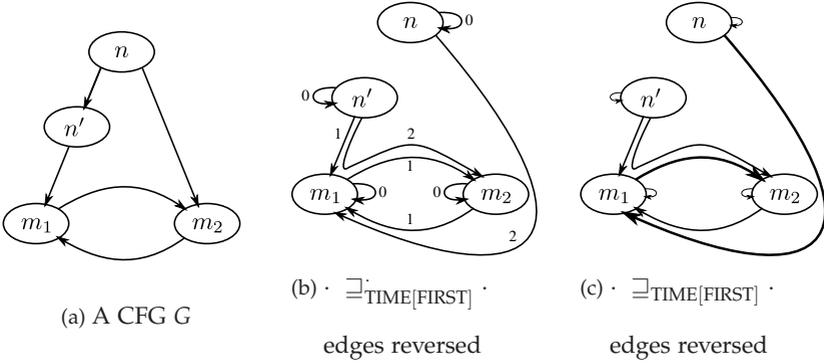


Figure 9.4: An irreducible graph with *intransitive* $\sqsubseteq_{\text{TIME}[\text{FIRST}]}$.

9.2 Timing Sensitive Post Postdominance Frontiers

In order to develop efficient algorithms for the computation of timing sensitive postdominance $\sqsubseteq_{\text{TIME}[\text{FIRST}]}$ and timing sensitive control dependence $\rightarrow_{\text{tscd}}$, I first recall that algorithms I previously developed for *nontermination* sensitive postdominance and control dependence (i.e.: for \sqsubseteq_{MAX} and $\rightarrow_{\text{ntscd}}$) heavily relied on the fact that \sqsubseteq_{MAX} is *transitive*:

1. Transitivity of \sqsubseteq_{MAX} allowed me to efficiently compute and represent \sqsubseteq_{MAX} in form of its transitive reduction $>_{\text{MAX}}$. Here, $<_{\text{MAX}}$ turned out to be a pseudo-forest.
2. Transitivity of \sqsubseteq_{MAX} , and the fact that

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

allowed me to use Algorithm 1 to efficiently compute $\rightarrow_{\text{ntscd}}$ (see Definition 3.2.7) via the nontermination sensitive postdominance frontier $\text{PDF}_{\sqsubseteq_{\text{MAX}}}$.

Disregarding for now the fact that $\rightarrow_{\text{tscd}}$ is defined in terms of the ternary relation $\cdot \sqsubseteq_{\text{TIME[FIRST]}} \cdot$, and not in terms of its (binary) “ $\exists k$. - closure” $\cdot \sqsubseteq_{\text{TIME[FIRST]}}^k \cdot$, let me investigate first if $\cdot \sqsubseteq_{\text{TIME[FIRST]}} \cdot$ is — in general — transitive. To do this, consider the (irreducible) CFG in Figure 9.4a. Here, every maximal path starting in n first reaches m_1 after two steps, hence $m_1 \sqsubseteq_{\text{TIME[FIRST]}} n$. Also, every maximal path starting in m_1 first reaches m_2 after one step, hence $m_2 \sqsubseteq_{\text{TIME[FIRST]}} m_1$. But it is for *no* number k of steps the case that $m_2 \sqsubseteq_{\text{TIME[FIRST]}}^k n$, hence: $\neg m_1 \sqsubseteq_{\text{TIME[FIRST]}} n$. In summary, $\cdot \sqsubseteq_{\text{TIME[FIRST]}} \cdot$ here is *not* transitive.

On the other hand, situations such as that in Figure 9.4 are the *only* in which $\cdot \sqsubseteq_{\text{TIME[FIRST]}} \cdot$ is not transitive:

Observation 9.2.1. Let G be any *reducible* CFG, and write \sqsubseteq for $\cdot \sqsubseteq_{\text{TIME[FIRST]}} \cdot$. Then \sqsubseteq is transitive.

Also, if there exists a unique exit node n_x , then $\cdot \sqsubseteq_{\text{TIME[FIRST]}} \cdot$ is transitive even for irreducible CFG.

Observation 9.2.2. Let G be any CFG with unique exit node n_x , and write \sqsubseteq for $\cdot \sqsubseteq_{\text{TIME[FIRST]}} \cdot$. Then \sqsubseteq is transitive.

Since I know now that $\cdot \sqsubseteq_{\text{TIME[FIRST]}} \cdot$ is “usually” transitive, I am encouraged to work towards a modification of Algorithm 1 which allows me to compute the timing sensitive postdominance frontier $\text{PDF}_{\sqsubseteq_{\text{TIME[FIRST]}}}$. But first, I need to ensure that $\text{PDF}_{\sqsubseteq_{\text{TIME[FIRST]}}}$ will actually allow me to determine $\rightarrow_{\text{tscd}}$. Remember that for $m \neq n$, I simply had

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

which was obvious from the definition of $\rightarrow_{\text{ntscd}}$, which is in terms of the binary relation \sqsubseteq_{MAX} .

In order to obtain the analogous result for $\rightarrow_{\text{tscd}}$, I first need to “conservatively” redefine the notion PDF_{\sqsubseteq} of \sqsubseteq -postdominance in order to

obtain a notion appropriate for non-transitive relations \sqsupseteq . Remember that in Definition 3.2.2 on page 21, I defined for any binary relation \sqsupseteq :

$$\text{PDF}_{\sqsupseteq}(m) = \left\{ n \mid \begin{array}{l} \neg m \text{ 1-}\sqsupseteq n \\ \text{for some } n' \text{ s.t. } n \rightarrow_G n' : m \sqsupseteq n' \end{array} \right\}$$

Syntactically, I will stick with this definition, but I will modify the notion of 1- \sqsupseteq -postdominance. The new definition is

Definition 9.2.1 (1- \sqsupseteq -Postdominance, redefinition). Given a relation $\sqsupseteq \subseteq N \times N$, a node $x \in N$ is said to 1- \sqsupseteq -postdominate z if $x \sqsupseteq z$ and there exists some node y such that

$$x \sqsupseteq y \sqsupseteq z$$

The only change is the new requirement $x \sqsupseteq z$, which of course was redundant up to this chapter, since any relation \sqsupseteq I considered (i.e.: $\sqsupseteq_{\text{POST}}$, \sqsupseteq_{MAX} and $\sqsupseteq_{\text{SINK}}$) was transitive. Implicitly, this change also affects immediate \sqsupseteq -postdominance $\text{ipdom}_{\sqsupseteq}$ — see Definition 3.2.1 on 19.

Theorem 9.2.1. Let G be any CFG, and $n \neq m$ two nodes in G . Then

$$n \in \text{PDF}_{\sqsupseteq_{\text{TIME[FIRST]}}}(m) \Leftrightarrow n \rightarrow_{\text{tscd}} m$$

Proof: In this proof, I write \sqsupseteq for $\sqsupseteq_{\text{TIME[FIRST]}}$, and $x \sqsupseteq^k y$ for $x \sqsupseteq_{\text{TIME[FIRST]}}^k y$, and note that \sqsupseteq is reflexive.

I begin with the forward implication, and assume $n \in \text{PDF}_{\sqsupseteq}(m)$. From $m \neq n$ and since $\neg m \text{ 1-}\sqsupseteq n$, I conclude that $\neg m \sqsupseteq n$. Since $m \neq n$, this just means that for all $k \in \mathbb{N}$ there must exist some G -successor n'' such that

$$\neg m \sqsupseteq^k n''$$

But at the same time, from $n \in \text{PDF}_{\sqsupseteq}(m)$ I obtain some G -successor n' of n with

$$m \sqsupseteq^k n'$$

for some k and hence: $n \rightarrow_{\text{tscd}} m$.

For the reverse implication, I can assume

$$\begin{array}{l} m \sqsupseteq^k n_l \\ \neg m \sqsupseteq^k n_r \end{array}$$

for some G -successors n_l, n_r of n and some $k \in \mathbb{N}$. Since I do have $m \sqsupseteq n_l$, all I need to show is $\neg m \sqsupseteq^1 n$. I assume the opposite in order to derive a contradiction. Given the new definition of $m \sqsupseteq^1 n$, I then have some k' such that $m \sqsupseteq^{k'} n$. Because $m \neq n$, I have $k' > 0$ and both

$$\begin{array}{l} m \sqsupseteq^{k'-1} n_l \quad \text{and} \\ m \sqsupseteq^{k'-1} n_r \end{array}$$

But this means $k = k' - 1$ in contradiction of $\neg m \sqsupseteq^k n_r$. \square

The postdominance frontier algorithm from [Cyt+91], and my generalization of that algorithm in Chapter 3, was based on the separation of the postdominance frontier $\text{PDF}_{\sqsupseteq}(x)$ into its *local* part $\text{PDF}_{\sqsupseteq}^{\text{local}}(x)$, and the parts $\text{PDF}_{\sqsupseteq}^{\text{up}}(z)$ contributed to $\text{PDF}_{\sqsupseteq}(x)$ by nodes z for which x is an immediate postdominator. In order to apply the same idea for timing sensitive postdominance $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$, I can keep the definition for the local part $\text{PDF}_{\sqsupseteq}^{\text{local}}(x)$ (changed only implicitly due to the redefinition of $1\text{-}\sqsupseteq$ -postdominance), which is (see: Definition 3.2.3 on page 21):

$$\text{PDF}_{\sqsupseteq}^{\text{local}}(x) = \left\{ y \mid \begin{array}{l} \neg x \sqsupseteq^1 y \\ y \rightarrow_G x \end{array} \right\}$$

I do, however, have to be more discriminate when considering the up-contributions, since due to the implicit redefinition of immediate postdominance, nodes that are $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$ -“equivalent” (i.e.: nodes

x, x' such that both $x \sqsupseteq_{\text{TIME}[\text{FIRST}]} x'$ and $x \sqsupseteq_{\text{TIME}[\text{FIRST}]} x'$ no longer necessarily are immediate postdominators of the same nodes z . Remember that the definition for immediate postdominators reads:

$$\text{ipdom}_{\sqsupseteq}(z) = \left\{ x \mid \begin{array}{l} x \text{ 1-}\sqsupseteq z \\ \forall x' \in N. x' \text{ 1-}\sqsupseteq z \implies x' \sqsupseteq x \end{array} \right\}$$

Specifically, I do *not* use the Definition 3.2.4 which read:

$$\text{PDF}_{\sqsupseteq}^{\text{up}}(z) = \left\{ y \in \text{PDF}_{\sqsupseteq}(z) \mid \forall x \in \text{ipdom}_{\sqsupseteq}(z). \neg x \text{ 1-}\sqsupseteq y \right\}$$

Instead, I use the following:

Definition 9.2.2 (\sqsupseteq -Postdominance Frontiers: *up* part for a given immediate postdominator x). Given a CFG $G = (N, E)$, a relation $\sqsupseteq \subseteq N \times N$ and nodes $x, z \in N$ such that

$$x \in \text{ipdom}_{\sqsupseteq}(z)$$

the \sqsupseteq -postdominance frontiers *up* part $\text{PDF}_{\sqsupseteq}^{\text{up}}(z, x)$ for z given x is defined by

$$\text{PDF}_{\sqsupseteq}^{\text{up}}(z, x) = \left\{ y \in \text{PDF}_{\sqsupseteq}(z) \mid \neg x \text{ 1-}\sqsupseteq y \right\}$$

Unfortunately, even this change does not give me the same decomposition of $\text{PDF}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}(x)$ as I had before (i.e.: as I had in Lemma 3.2.2 on page 22). Specifically, at some nodes x I can only inherit dominance frontier nodes y (from nodes z s.t. x is a immediate $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$ -postdominator of z) under the additional condition $y \in \sqsubseteq'_x$, for \sqsubseteq'_x as defined as in the following observation:

Observation 9.2.3. Given any CFG $G = (N, E)$ and a node $x \in N$, let

$$\sqsupseteq = \sqsupseteq_{\text{TIME}[\text{FIRST}]}$$

Then:

$$\begin{aligned}
 \text{PDF}_{\sqsupseteq}(x) &= \text{PDF}_{\sqsupseteq}^{\text{local}}(x) \\
 &\cup \bigcup_{\{z \mid x \in \text{ipdom}_{\sqsupseteq}(z), \neg z \in \text{ipdom}_{\sqsupseteq}(x)\}} \text{PDF}_{\sqsupseteq}^{\text{up}}(z, x) \\
 &\cup \bigcup_{\{z \mid x \in \text{ipdom}_{\sqsupseteq}(z), z \in \text{ipdom}_{\sqsupseteq}(x)\}} \text{PDF}_{\sqsupseteq}^{\text{up}}(z, x) \cap \sqsubseteq'_x
 \end{aligned}$$

where $\sqsubseteq'_x = \{y \mid \text{for some } y' \text{ s.t. } y \rightarrow_G y' : x \sqsupseteq y'\}$.

In contrast to the corresponding result (Lemma 3.2.2 on page 22) for transitive postdominance relations, I do not offer a formal proof of Observation 9.2.3. I also do not attempt to give general conditions under which Observation 9.2.3 may hold for (intransitive) postdominance relations other than $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$. I also do not offer any such conditions for the two following simplifications of $\text{PDF}_{\sqsupseteq}^{\text{local}}(x)$ and $\text{PDF}_{\sqsupseteq}^{\text{up}}(z, x)$. I do want to note, however, that while $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$ is in general not transitive, it *does* fulfill the other two earlier conditions of being *closed under \rightarrow_G* and *lacking joins*.

Observation 9.2.4. Given any CFG G , let

$$\sqsupseteq = \sqsupseteq_{\text{TIME}[\text{FIRST}]}$$

Then \sqsupseteq is closed under \rightarrow_G , and

$$\text{PDF}_{\sqsupseteq}^{\text{local}}(x) = \left\{ y \mid \begin{array}{l} \neg x \in \text{ipdom}_{\sqsupseteq}(y) \\ y \rightarrow_G x \end{array} \right\}$$

Observation 9.2.5. Given any CFG G , let

$$\sqsupseteq = \sqsupseteq_{\text{TIME}[\text{FIRST}]}$$

$$\begin{array}{c}
 \frac{\neg x \in \text{ipdom}_{\sqsubseteq}(y) \quad y \rightarrow_G x}{y \in \text{PDF}_{\sqsubseteq}(x)} \text{PDF}^{\text{local}} \\
 \\
 \frac{x \in \text{ipdom}_{\sqsubseteq}(z) \quad \neg x \in \text{ipdom}_{\sqsubseteq}(y) \quad \neg z \in \text{ipdom}_{\sqsubseteq}(x) \quad y \in \text{PDF}_{\sqsubseteq}(z)}{y \in \text{PDF}_{\sqsubseteq}(x)} \text{PDF}_1^{\text{up}} \\
 \\
 \frac{y \rightarrow_G y' \quad \neg x \in \text{ipdom}_{\sqsubseteq}(y) \quad x \sqsubseteq y' \quad x \in \text{ipdom}_{\sqsubseteq}(z) \quad z \in \text{ipdom}_{\sqsubseteq}(x) \quad y \in \text{PDF}_{\sqsubseteq}(z)}{y \in \text{PDF}_{\sqsubseteq}(x)} \text{PDF}_2^{\text{up}}
 \end{array}$$

Figure 9.5: A rule system for $\text{PDF}_{\sqsubseteq_{\text{TIME[FIRST]}}}$, writing \sqsubseteq for $\sqsubseteq_{\text{TIME[FIRST]}}$.

Then \sqsubseteq lacks joins and is closed under \rightarrow_G , and given any $x \in \text{ipdom}_{\sqsubseteq}(z)$:

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

Rephrasing observations 9.2.3, 9.2.4 and 9.2.5, I obtain a characterization of $\text{PDF}_{\sqsubseteq_{\text{TIME[FIRST]}}}$ as the least fixed point of the monotone functional defined by the rules in Figure 9.5.

In order to obtain an efficient algorithm for $\text{PDF}_{\sqsubseteq_{\text{TIME[FIRST]}}}$, I thus need to develop

1. An efficient algorithm for the computation of (an efficient representation of) $\text{ipdom}_{\sqsubseteq_{\text{TIME[FIRST]}}}$

2. A way to replace the explicit *generate-and-test*

$$y \rightarrow_G y', x \sqsubseteq_{\text{TIME[FIRST]}} y'$$

of some y' such in the rule PDF_2^{up} by an efficient (implicit) check.

9.3 Transitive Timing Sensitive Postdominance

When I developed algorithms for $\text{ipdom}_{\sqsupseteq_{\text{SINK}}}$ and $\text{ipdom}_{\sqsupseteq_{\text{MAX}}}$, I made use of the fact that there, I had

$$\text{ipdom}_{\sqsupseteq}^* = \sqsupseteq$$

and $\text{ipdom}_{\sqsupseteq}$ was easily derived from any transitive reduction $>$ of \sqsupseteq (e.g.: Observation 5.2.2 on 44).

With $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$, things are different: for example, in the CFG from Figure 9.4a, repeated in Figure 9.6a, I have

$$\text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}^* \neq \sqsupseteq_{\text{TIME}[\text{FIRST}]}$$

as is evident from $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$ and $\text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}$ shown in the same figure. Note, for example, that

$$\begin{aligned} m_2 &\in \text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}(m_1) \text{ and} \\ m_1 &\in \text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}(m_2), \text{ but} \\ &\neg m_2 \sqsupseteq_{\text{TIME}[\text{FIRST}]} m_1 \end{aligned}$$

As mentioned before, this very same CFG demonstrates that $\cdot \sqsupseteq_{\text{TIME}[\text{FIRST}]} \cdot$ is — in general — not transitive, and hence cannot be efficiently represented by some transitive reduction $>$. Similarly, the ternary relation $\cdot \sqsupseteq_{\text{TIME}[\text{FIRST}]} \cdot$ is — in general — not transitive. Here, by transitive I mean the following:

Definition 9.3.1. A ternary relation $\cdot \sqsupseteq \cdot \subseteq N \times \mathbb{N} \times N$ is *transitive* if

$$\begin{aligned} \text{whenever} \quad & m \sqsupseteq^k m' \quad \text{and} \quad m' \sqsupseteq^{k'} n \\ \text{then i also have} \quad & m \sqsupseteq^{k+k'} n \end{aligned}$$

In contrast to $\cdot \sqsupseteq_{\text{TIME}[\text{FIRST}]} \cdot$, the following ternary relation *is* transitive:

Definition 9.3.2. I say that m timing-sensitively and *transitively* post-dominates n at position k in G iff on all maximal G -paths starting in n , m appears at position k . I omit “in G ” whenever possible. Formally:

$$m \sqsubseteq_{\text{TIME}}^{k \text{ in } G} n \Leftrightarrow \forall \pi \in {}_n\Pi_{\text{MAX}}^G. m \in^k \pi$$

Note that the only difference between this definition of $\sqsubseteq_{\text{TIME}}$ and the Definition 9.1.1 of $\sqsubseteq_{\text{TIME}[\text{FIRST}]}$ is the use of \in^k instead of \in_{FIRST}^k , i.e.: $\sqsubseteq_{\text{TIME}}$ only requires m to appear at position k in π , while $\sqsubseteq_{\text{TIME}[\text{FIRST}]}$ additionally requires k to be the *first* position in π at which m appears.

Observation 9.3.1. Given any CFG G , the ternary relation $\cdot \sqsubseteq_{\text{TIME}} \cdot$ is transitive.

Just as for $\sqsubseteq_{\text{TIME}[\text{FIRST}]}$, I can give a least fixed point characterization of $\sqsubseteq_{\text{TIME}}$ in the underlying lattice $(2^{N \times N \times N}, \subseteq)$.

Proposition 9.3.1. Let G be a CFG and T be the rule-system²

$$\frac{}{n \sqsubseteq^0 n} T^{\text{self}} \quad \frac{\forall p \rightarrow_G x. m \sqsubseteq^k x \quad p \rightarrow_G^+ m}{m \sqsubseteq^{k+1} p} T^{\text{suc}}$$

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

The only difference to T_{FIRST} is the omission of the requirement $m \neq p$ in Proposition 9.3.1.

² over a ternary relation $\cdot \sqsubseteq \cdot \cdot$.

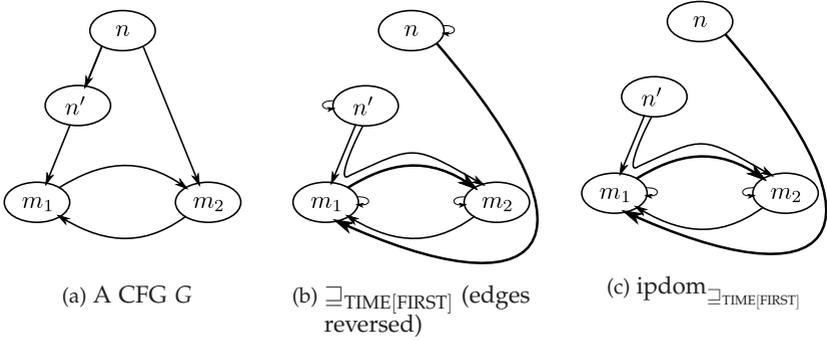


Figure 9.6: An irreducible graph with *intransitive* $\cdot \sqsupseteq_{\text{TIME}[\text{FIRST}]}$.

Unlike $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$, the relation $\sqsupseteq_{\text{TIME}}$ is not finite, so I cannot naively implement $\sqsupseteq_{\text{TIME}}$ to obtain a “reference” implementation. For example, in the CFG from Figure 9.6a, I have

$$\begin{array}{ll}
 m_1 \sqsupseteq_{\text{TIME}}^2 n & \\
 & m_2 \sqsupseteq_{\text{TIME}}^3 n \\
 m_1 \sqsupseteq_{\text{TIME}}^4 n & \\
 & m_2 \sqsupseteq_{\text{TIME}}^5 n \\
 \dots & \dots
 \end{array}$$

I can, however, use \top to obtain (for any $n \in \mathbb{N}$) a reference implementation for $k \leq n$, i.e. for:

$$\sqsupseteq_{\text{TIME}} \cap N \times \mathbb{N}^{\leq n} \times N$$

using an explicit representation.

The fact that $\sqsupseteq_{\text{TIME}}$ is transitive suggests that I can efficiently represent the full relation $\sqsupseteq_{\text{TIME}}$ using a *transitive reduction* $>$ of $\sqsupseteq_{\text{TIME}}$:

Definition 9.3.3. A ternary relation

$$\cdot > \cdot \subseteq N \times \mathbb{N} \times N$$

is called a *transitive reduction* of a ternary relation

$$\cdot \sqsupseteq \cdot \subseteq N \times \mathbb{N} \times N$$

iff \sqsupseteq is equal to the reflexive, transitive closure of $>$, and $>$ is minimal w.r.t. this property, i.e.:

1. For any $n, m \in N$ and any $k \in \mathbb{N}$,

$$m \sqsupseteq^k n \Leftrightarrow \begin{array}{l} m >^{k_1} m' >^{k_2} m'' \dots >^{k_c} n \quad (c \geq 0) \\ \text{for some nodes } m', m'', \dots \text{ such that } k = \sum k_i \end{array}$$

where the case $c = 0$ is understood to mean $k = 0$ and $m = n$, and

2. $>$ has *minimal size* $|\!>| \leq |\!>'|$ among all such relations $>'$.

In the following, given a ternary relation $>$ (or similar), I will write³ $n <^k m$ for $m >^k n$. The ternary relation $<_{\text{TIME}}$ corresponding to a ternary transitive reduction $>_{\text{TIME}}$ of $\sqsupseteq_{\text{TIME}}$ for the example CFG is shown in Figure 9.7.

In order to obtain an efficient algorithm for the computation of some transitive reduction $>$ of $\sqsupseteq_{\text{TIME}}$, the following property is crucial:

Proposition 9.3.2. Let G be any CFG, and $>$ a transitive reduction of $\cdot \sqsupseteq_{\text{TIME}} \cdot$ (and $<$ its inverse, as just defined). Then $<$ is a (\mathbb{N} -) labeled pseudo-forest, i.e. for any node n , whenever both $n <^k m$ and $n <^{k'} m'$, I have $k = k'$ and $m = m'$.

³ in contrast, I will *never* write $<^k$ or $>^k$ for the k th iteration of some *binary* relation $<$ or $>$



Figure 9.7: An irreducible graph with $\text{intransitive} \cdot \sqsupseteq_{\text{TIME}[\text{FIRST}]}$.

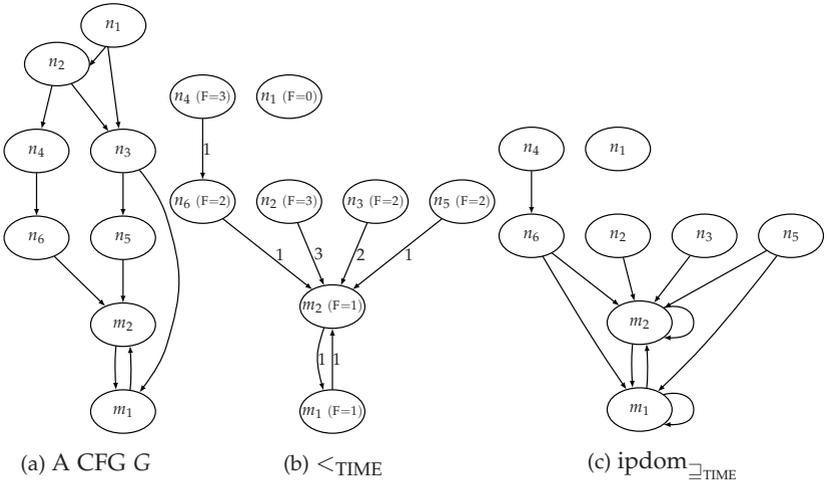


Figure 9.8: An irreducible graph with $\text{intransitive} \cdot \sqsupseteq_{\text{TIME}[\text{FIRST}]}$.

In order for $\sqsupseteq_{\text{TIME}}$ to be of any use in obtaining an algorithm for $\rightarrow_{\text{tscd}}$, I must now (by Observation 9.2.3) explain how to derive $\text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}$ from $\sqsupseteq_{\text{TIME}}$.

In Figure 9.8a, I show a somewhat more interesting CFG. First note that $m_1 \sqsupseteq_{\text{TIME}}^3 n_3$, but *not*: $m_1 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^3 n_3$ (nor: $m_1 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k n_3$ for any other k), since I can reach m_1 from n_3 in just $1 \neq 3$ G-

step $n_3 \rightarrow_G m_1$. In terms of “following edges” in $<_{\text{TIME}}$, I can say: the $<_{\text{TIME}}$ -path $n_3 <_{\text{TIME}}^2 m_2 <_{\text{TIME}}^1 m_1$ is valid with respect to $\sqsubseteq_{\text{TIME}[\text{FIRST}]}$ only up to m_2 .

On the other hand, consider the G -successors n_3 and n_4 of n_2 . Both the $<_{\text{TIME}}$ -paths $n_3 <_{\text{TIME}}^2 m_2$ and $n_4 <_{\text{TIME}}^1 m_6 <_{\text{TIME}}^1 m_2$ are valid up to m_2 , and have the same total “cost” 2, from which I can conclude $m_2 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{3=1+2} n_2$ (adding to the common cost 2 the cost 1 that is due to going from n_2 to any of its G -successor n_3 or n_4).

In general, I make the following observation on the relation of $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$ and $\sqsupseteq_{\text{TIME}}$:

Observation 9.3.2. Let G be any CFG. Then for every node n there exists a number $F_n \in \mathbb{N}$ (the “amount of fuel available at n ”) such that for all nodes m and any “distance” $k \in \mathbb{N}$ such that $m \sqsupseteq_{\text{TIME}}^k n$, I have

$$m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k n \iff k \leq F_n$$

In the following, I just write F_n for the least such number.

Let me try to justify the colloquial phrase “amount of fuel available at n ” by rephrasing Observation 9.3.2 in terms of a transitive reduction $>_{\text{TIME}}$ of $\sqsupseteq_{\text{TIME}}$: When “starting” at node n with a tank filled with F_n units of fuel, and if following edges $<_{\text{TIME}}^f$ consumes f units of fuel, the nodes m such that $m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k n$ for some k are exactly those nodes I can reach without running out of fuel. Given such m , the number k is exactly the amount of fuel consumed before reaching m .

Rephrasing my remarks on Figure 9.8b, I first note that $m_1 \sqsupseteq_{\text{TIME}}^{k=3} n_3$ (and that $k = 3$ is the *least* such k), but *not*: $m_1 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k=3} n_3$ (nor: $m_1 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k=k} n_3$ for any other k), since $k = 3 > 2 = F_{n_3}$. In other words: attempting to reach m_1 from n_3 in $<_{\text{TIME}}$ fails by running out of fuel, since after starting with $F_{n_3} = 2$ and traversing

the edge $n_3 <_{\text{TIME}}^2 m_2$ my tank is empty, and I cannot traverse the edge $m_2 <_{\text{TIME}}^1 m_1$ to reach m_1 .

On the other hand, I have $m_2 \sqsupseteq_{\text{TIME}}^3 n_2$ and $3 \leq 3 = F_{n_2}$, from which I can conclude $m_2 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^3 n_2$.

The following observation states that the function $n \mapsto F_n$ on the nodes of some CFG is easily recoverable from its restriction to nodes “entering” some $<_{\text{TIME}}$ -cycle M , i.e.: nodes $n \in N_M$ where (similar to the N_M in Subsection 6.1.1):

$$N_M = \{ n \in N \mid n \notin M, \exists m \in M, k \in \mathbb{N}. n <_{\text{TIME}}^k m \}$$

In Figure 9.8b, the only cycle is $M = \{m_1, m_2\}$, and $N_M = \{n_2, n_3, n_5, n_6\}$.

Observation 9.3.3. Let G be any CFG, $>_{\text{TIME}}$ a transitive reduction of $\sqsupseteq_{\text{TIME}}$, and $n \neq m$ nodes such that $n <_{\text{TIME}}^k m$. Then either

- a) n is a node “entering” some $<_{\text{TIME}}$ -cycle M (i.e.: $n \in N_M$) and $m \in M$, or
- b) neither n nor m are in any $<_{\text{TIME}}$ -cycle, and $F_n = F_m + k$, or
- c) both n and m are in some $<_{\text{TIME}}$ -cycle M , and

$$\begin{aligned} & m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k n \\ \text{as well as } & n \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{F_m} m \end{aligned}$$

Remark 9.3.1. In item c) of Observation 9.3.3, the sum $k + F_m$ is exactly the *circumference* c_M of the $<_{\text{TIME}}$ -cycle M , i.e. the sum over all the distance in M :

$$c_M = \sum_{m \in M, m <_{\text{TIME}}^k m'} k$$

Note also that for any $m_1, m_2 \in M$ such that $m_1 \neq m_2$, I *always* have

$$m_1 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k m_2$$

as well as

$$m_2 \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k'} m_1$$

for some k, k' with $k + k' = c_M$.

Let me try to rephrase Observation 9.3.3 in colloquial terms. Assume I want to *start a trip* in \prec_{TIME} , starting at node n_0 . The total fuel available for my trip then is F_{n_0} , and traversing some edge $n \prec_{\text{TIME}}^k m$ will cost k units of fuel. By Observation 9.3.2, whenever I reach some node m after spending k units of fuel, I learn $m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k n_0$. Now, item b) and c) in Observation 9.3.3 tell me that

- b) if I start my trip at some node n_0 that is not in any \prec_{TIME} -cycle M , then I will never run out of fuel, *unless* I reach some node $n \in N_M$ from which I will “enter” some \prec_{TIME} -cycle M . Then at that node n , I am guaranteed to have exactly F_n units of fuel remaining (which may allow me to visit some, but not necessarily all nodes $m \in M$).
- c) if I start my trip at in some \prec_{TIME} -cycle M (i.e.: $n_0 \in M$), then I will run out of fuel *just short of coming back to n_0* , i.e. at the node $m \in M$ such that $m \prec_{\text{TIME}} n_0$. In other words, I will visit exactly the nodes in M , and every node exactly once.

Being lazy as I am, this means that after fueling up my empty tank at n_0 , I do not even need to worry about my fuel consumption, unless I visit some border city $n \in N_M$. At that point, some friendly road sign⁴ will assure me: “You have F_n units of fuel remaining”. Additionally, another road sign will tell me: “You have enough fuel to visit: $\{m_1, m_2, \dots\} \subseteq M$ ”. Only *this* road sign will have been put up by $\text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}$:

⁴ valid no matter where exactly I started my journey!

Observation 9.3.4. Given any CFG G , let

$$\sqsupseteq = \sqsupseteq_{\text{TIME}[\text{FIRST}]}$$

Then for any node n

- if $n \in N_M$ for some \prec_{TIME} -cycle M , then

$$\text{ipdom}_{\sqsupseteq}(n) = \{ m_c \mid n \prec_{\text{TIME}}^{k_1} m_1 \prec_{\text{TIME}}^{k_2} \dots \prec_{\text{TIME}}^{k_c} m_c, \sum_i k_i \leq F_n \}$$

where the c are understood to be ≥ 1 .

- if $n \in M$ for some \prec_{TIME} -cycle M , then

$$\text{ipdom}_{\sqsupseteq}(n) = \begin{cases} \emptyset & \text{if } M = \{n\} \\ M & \text{otherwise} \end{cases}$$

- otherwise

$$\text{ipdom}_{\sqsupseteq}(n) = \{ m \mid n \prec_{\text{TIME}}^k m \}$$

which is, by virtue of \prec_{TIME} being a (labeled) pseudo-forest, at most a singleton.

In colloquial terms: outside of \prec_{TIME} -cycles M and border cities N_M , the sign put up by $\text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}$ will just tell me which node m I will visit next (and not to worry about fuel).

9.4 Algorithms for Timing Sensitive Control Dependence

I ended the previous section by giving a (easy-to-implement) scheme to derive $\text{ipdom}_{\sqsupset_{\text{TIME}}[\text{FIRST}]}$ from \langle_{TIME} and the *fuel*-mapping $F = n \mapsto F_n$. Now, I need to explain how to

1. efficiently compute \langle_{TIME}
2. efficiently compute F (from \langle_{TIME})
3. use the previous result from Observation 9.2.3 to obtain an efficient algorithm for $\text{PDF}_{\sqsupset_{\text{TIME}}[\text{FIRST}]}$ and hence (by Theorem 9.2.1):
 $\rightarrow_{\text{tscd}}$.

9.4.1 An Algorithm for \langle_{TIME}

This first task is — in principle — the easiest. I merely need to modify the *least common ancestor* Algorithm 3 from page 47 to support *counting* of the cost of \langle_{TIME} -paths, and then modify the \sqsupset_{MAX} -Algorithm 5 to use the modified least common ancestor algorithm. The result of the latter is shown in Algorithm 8. The most significant modifications are **highlighted**. I need to omit the check $z \neq x$ because I have to differentiate between an exit node x (i.e.: a node x without G -successor), and a “one node sink” x (i.e.: a node x with only G -successor x). For the former I only have $x \sqsupset_{\text{TIME}}^0 x$, while for the latter, I also have $x \sqsupset_{\text{TIME}}^1 x$, $x \sqsupset_{\text{TIME}}^2 x$, \dots . Instead of finding the least common ancestors in some (binary) pseudo-forest, I need to find the least common ancestor in a \mathbb{N} -labeled (ternary) pseudo-forest, by a call to lca_{\langle} . Here, lca_{\langle} on a set (of successors y of x , and their distance 1 from x) is implemented by iterating over Algorithm 9 (similar to Algorithm 4 on page 48).

The *least common ancestor* for \mathbb{N} -labeled pseudo forests like \langle_{TIME} is defined with respect to the amount k^n, k^m of “amount of fuel already spent” before reaching nodes n, m .

Definition 9.4.1. Let $<$ be a \mathbb{N} -labeled pseudo-forest, n, m two nodes in $<$, and $k^n, k^m \in \mathbb{N}$ two natural numbers. If (z, k) is the least (by comparison of k) pair such that both

$$\begin{aligned} n <^{k_1} \dots <^{k_c} z \quad \text{with} \quad k = k^n + \sum_i k_i \\ \text{and} \quad m <^{k'_1} \dots <^{k'_c} z \quad \text{with} \quad k = k^m + \sum_i k'_i \end{aligned}$$

then (z, k) is the *least common ancestor* of n, m in $<$ with respect to initial “fuel consumption” k^n, k^m , and I write

$$\text{lca}_<((n, k^n), (m, k^m)) = (z, k)$$

I do not attempt to rigorously defend the Algorithm 9 for the computation of $\text{lca}_<((n, k^n), (m, k^m))$. I merely note that the check

$$n' \in \pi_n \wedge n' \notin \pi_m \wedge |\text{KS}_m[n]| > 1$$

is needed to guarantee termination whenever π_n has reached (at n') a different $<$ -cycle than π_m is in. The check

$$\max \text{KS}_n[n] > \min \text{KS}_m[n] \quad \wedge \quad \min \text{KS}_n[n] < \max \text{KS}_m[n]$$

terminates the search whenever both π_n and π_m have reached the same $<$ -cycle M , but are “out of phase” with regard to amount of “fuel consumed” before reaching M , and the circumference c_M of M . The use of a map KS that contains multiple numbers k for a given node n once a $<$ -cycle M is reached is certainly not most efficient. I could instead make use of the fact that for such nodes n , I have some k_n such that eventually $k_n + i \cdot c_M \in \text{KS}_n(n)$ for *all* i .

Also note that Algorithm 8 could (by virtue of the two assertions) be optimized by never re-inserting a node n into the workset once $\text{ITIMEDOM}[n] \neq \perp$. Presumably, a workset-free algorithm (similar to Algorithm 19 for $<_{\text{MAX}}$ on page 397) is possible and may be preferable.

Input : A CFG G

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

Output: A transitive reduction $>_{\text{TIME}}$ of $\sqsupseteq_{\text{TIME}}$

begin

for $x \in N$, $\{z \mid x \rightarrow_G z\} = \{z\}$, $z \neq x$ **do**

$\text{IDOM}[x] \leftarrow (z, 1)$

end

TIME_{up}

return IDOM

end

Procedure TIME_{up}

$\text{workset} \leftarrow \text{COND}_G$

while $\text{workset} \neq \emptyset$ **do**

$x \leftarrow \text{remove}(\text{workset})$

$(z, k) \leftarrow \text{lca}_<(\{(y, 1) \mid x \rightarrow_G y\})$

assert $(z, k) \neq \text{IDOM}[x] \Rightarrow (z, k) \neq \perp$

assert $(z, k) \neq \text{IDOM}[x] \Rightarrow \text{IDOM}[x] = \perp$

if $(z, k) \neq \text{IDOM}[x]$ **then**

$\text{workset} \leftarrow \text{workset} \cup \{n \in \text{COND} \mid n \neq x, \exists n \rightarrow_G y. y <^* x\}$

$\text{IDOM}[x] \leftarrow (z, k)$

end

end

end

Algorithm 8: An algorithm for the computation of $<_{\text{TIME}}$. Here, $y <^* x$ means: $x = y \vee y <^{k_1} \dots <^{k_c} x$

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

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

Output: $\text{lca}_{\prec}((n_0, k_0^n), (m_0, k_0^m))$ 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 \prec -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. $k^n = \max_{\text{KS}_n[n]}$ and for any such n
 $\text{KS}_n[n] = \{k_0^n + \sum_i k_i \mid n_0 \prec^{k_1} \dots \prec^{k_e} n \text{ in } \pi_n\}$

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

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

if $n \notin \pi_m$ **then switch** $\text{IDOM}[n]$ **do**

case \perp **do return** \perp

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

if $n' \in \pi_n \wedge n' \notin \pi_m \wedge |\text{KS}_m[m]| > 1$ **then return** \perp

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

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

end

end

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

if $\max_{\text{KS}_n[n]} > \min_{\text{KS}_m[n]} \wedge \min_{\text{KS}_n[n]} < \max_{\text{KS}_m[n]}$ **then return** \perp

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

case \perp **do return** \perp

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

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

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

end

end

end

Algorithm 9: A timing sensitive *least common ancestor* algorithm.

9.4.2 An Algorithm for F

My algorithm for the computation of F from $<_{\text{TIME}}$ is based on the following fixed point characterization of F.

Observation 9.4.1. Let $G = (N, E)$ be any CFG, and F be the rule system

$$\frac{}{F(n) \geq 0} \text{F}^{\text{lin}}$$

$$\frac{m \sqsupseteq_{\text{TIME}}^{k_m} n \quad F(n) \geq k_m \quad m' \sqsupseteq_{\text{TIME}}^k m \quad k = \min\{k \mid m' \sqsupseteq_{\text{TIME}}^k m\} \quad m' \neq n \quad \forall n \rightarrow_G x. m' \sqsupseteq_{\text{TIME}}^{k_m+k-1} x \wedge F(x) \geq k_m + k - 1}{F(n) \geq k_m + k} \text{F}^{\text{trans}}$$

Then $F = \mu F$ for the corresponding monotone functional F over the pointwise lattice on $N \rightarrow \mathbb{N}$.

The corresponding Algorithm 10 will

- remember for each node n with $F(n) = k_m$ the current *witness* m such that $m \sqsupseteq_{\text{TIME}}^{k_m} n$.
- obtain the next mode $m' \neq n$ such that $m' \sqsupseteq_{\text{TIME}}^k m$ and $k = \min\{k \mid m' \sqsupseteq_{\text{TIME}}^k m\}$ simply by following one edge $m <_{\text{TIME}}^k m'$ in $<_{\text{TIME}}$.
- implement the check for successors x of n by looking for a cycle-free $<_{\text{TIME}}$ -path from x to m' with cost $k_m + k - 1$.

Furthermore, by Observation 9.3.3, it is possible to replace the check $F(x) \geq k_m + k - 1$ by only checking for fuel on “border” nodes⁵ n' on the \prec_{TIME} -path from x to m' , i.e.: to check

$$k_m + k - 1 - k_{n'} \leq F(n')$$

for nodes n' such that $n' \in N_M$ for some \prec_{TIME} -cycle M , where $k_{n'}$ is the amount of fuel spent when reaching node n' . But this means that *it is enough to compute the restrictions $F|_{N_M}$ of F to border nodes $n' \in N_M$ for \prec_{TIME} -cycles M* . Even if I am interested in the value F_n for *every node*, then these can be recovered from the restrictions $F|_{N_M}$ by Observation 9.3.3, and I do not need to compute them explicitly in Algorithm 10. In other words: Algorithm 10 only iterates over such border nodes $n \in N_M$, where M is the set of \prec_{TIME} -cycles M , and

$$N_M = \bigcup_{M \in \mathbb{M}} N_M$$

9.4.3 An Algorithm for $\text{PDF}_{\sqsupset_{\text{TIME}}[\text{FIRST}]}$

Now that I have devised algorithms for \prec_{TIME} and F , I am *almost* ready to use the characterization of $\text{ipdom}_{\sqsupset_{\text{TIME}}}$ from Observation 9.3.4 and the rule system in Figure 9.5 on page 178 to give an algorithm for timing sensitive postdominance frontiers $\text{PDF}_{\sqsupset_{\text{TIME}}[\text{FIRST}]}$.

I *could* use the rule system in Figure 9.5 *as is*, since the test

$$x \sqsupset_{\text{TIME}[\text{FIRST}]} y' \text{ in rule } \text{PDF}_2^{\text{up}}$$

can be answered by \prec_{TIME} and F (using the very definition of F , i.e.: Observation 9.3.2). But this would result in an inefficient algorithm

⁵ of which there is at most one on each cycle-free \prec_{TIME} -path

Input : A CFG $G = (N, E)$
Input : The labeled pseudo forest \prec_{TIME} for G
Input : The set N_{M} of “border” nodes n' in \prec_{TIME}
Output: A map $F : N_{\text{M}} \rightarrow \mathbb{N} \times N$ s.t. for all $n \in N_{\text{M}}$: $F[n] = (F_n, m)$
for some “witness” m

```

for  $n \in N_{\text{M}}$  do
  |  $F[n] \leftarrow (0, n)$ 
end
changed  $\leftarrow$  true
while changed do
  | changed  $\leftarrow$  false
  | for  $(n, (k_m, m)) \in F, m \prec_{\text{TIME}}^k m', m' \neq n$  do
    | valid $_{m'}$   $\leftarrow$  true
    | for  $n \rightarrow_G x$  do if valid $_{m'}$  then
      | let  $n_0 \prec_{\text{TIME}}^{k_0} \dots \prec_{\text{TIME}}^{k_c} n_{c+1}$  be the cycle-free  $\prec_{\text{TIME}}$ -path
        | where  $n_0 = x$  and  $n_{c+1} = m'$ 
        | if  $\sum k_i = k_m + k - 1$  then
          | if  $\exists n_j. n_j \in N_{\text{M}}$  and  $k_m + k - 1 - \sum_{i < j} k_i > F[n_j]$  then
            | | valid $_{m'}$   $\leftarrow$  false
            | end
          | end
        | else
          | | valid $_{m'}$   $\leftarrow$  false
          | end
        | end
      | end
    | if valid $_{m'}$  then
      | |  $F[n] \leftarrow (k_m + k, m')$ 
      | | changed  $\leftarrow$  true
    | end
  | end
  | end
end
return  $F$ 

```

Algorithm 10: An algorithm for the computation of F from \prec_{TIME} .
Here, $y \prec^* x$ means: $x = y \vee y \prec^{k_1} \dots \prec^{k_c} x$

whenever the number of nodes in \prec_{TIME} -cycles and hence: the number of nodes x, y such that both

$$x \in \text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}(z) \text{ and } x \in \text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}(z)$$

is large.

Taking a closer look at rule PDF_1^{up} and rule PDF_2^{up} (and still writing \sqsupseteq for $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$), I notice that

1. nodes y are propagated unconditionally from $\text{PDF}_{\sqsupseteq}(z)$ to $\text{PDF}_{\sqsupseteq}(x)$ by PDF_1^{up} unless both x and z are in some common \prec_{TIME} cycle M .
2. if x, z are in some common \prec_{TIME} cycle M , and the $y \in \text{PDF}_{\sqsupseteq}(z)$ originated from some application of rule $\text{PDF}^{\text{local}}$ under substitution $[y \mapsto y, x \mapsto y']$ (i.e.: I have: $\neg y' \in \text{ipdom}_{\sqsupseteq}(y)$ and $y \rightarrow_G y'$) such that also $y' \in M$, then $y' \prec_{\text{TIME}} x$ and hence (by Observation 9.3.3): $x \sqsupseteq_{\text{TIME}[\text{FIRST}]} y'$. But this means that I can propagate y from $\text{PDF}_{\sqsupseteq}(z)$ to $\text{PDF}_{\sqsupseteq}(x)$ by rule PDF_2^{up} .

The resulting algorithm uses a map $\text{DF} : N \rightarrow N \leftrightarrow \text{Bool}$, where \leftrightarrow indicates a partial map, and upon completion, y is in the $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$ -postdominance frontier of z iff $D[x]$ is defined for z :

$$z \in \text{PDF}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}(x) \iff z \in \text{dom}(D[x])$$

During computation (and writing \sqsupseteq for $\sqsupseteq_{\text{TIME}[\text{FIRST}]}$)

$$\begin{aligned} D[z][y] = \mathbf{false} & \text{ indicates } y \in \text{PDF}_{\sqsupseteq}(z) \\ D[z][y] = \mathbf{true} & \text{ indicates } y \in \text{PDF}_{\sqsupseteq}(z) \\ & \mathbf{and } y \in \text{PDF}_{\sqsupseteq}(x) \text{ for } x \in \text{ipdom}_{\sqsupseteq}(z) \end{aligned}$$

In other words: $D[z][y]$ indicates whether y is to be propagated.

Propagation follows $\text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}$ and hence: \prec_{TIME} . But \prec_{TIME} is a pseudo forest, so it can be iterated by starting in leaves, and continuing towards (and: within!) roots. This is either done explicitly, or im-

plicitly by using a priority queue, with each node assigned a priority based on a reverse depth first search starting in the roots. The priority queue based Algorithm 11 even works for arbitrary node-numbering (although then not necessarily efficiently).

Input : A CFG $G = (N, E)$
Input : Any numbering $\# : N \rightarrow \mathbb{N}$
Input : Immediate postdominators $\text{ipdom}_{\sqsupset} = \text{ipdom}_{\sqsupset_{\text{TIME}[\text{FIRST}]}}$
Input : The set $N_{\mathbb{M}}$ of “border” nodes in \langle_{TIME}
Data: A priority queue Q ordered by the numbering $\#$
Output: $\text{PDF}_{\sqsupset_{\text{TIME}[\text{FIRST}]}}$ represented as a map $\text{DF} : N \rightarrow N \leftrightarrow \text{Bool}$

```

 $Q \leftarrow \emptyset$ 
for  $x \in N$ ,  $y \rightarrow_G x$ ,  $\neg x \in \text{ipdom}_{\sqsupset}(y)$  do
  |  $\text{DF}[x][y] \leftarrow \text{true}$ 
  |  $Q \leftarrow Q \cup \{x\}$ 
end
while  $Q \neq \emptyset$  do
  |  $z \leftarrow \text{remove}(Q)$  s.t.  $z\# = \max_{z \in Q} z\#$ 
  | for  $x \in \text{ipdom}_{\sqsupset}(z)$ ,  $(y, \text{true}) \in \text{DF}[z]$ ,  $\neg x \in \text{ipdom}_{\sqsupset}(y)$  do
  | |  $\text{DF}_{x,y} \leftarrow \text{DF}[x][y]$  (may be  $\perp$ )
  | |  $\text{DF}'_{x,y} \leftarrow \text{DF}_{x,y} \vee (z \notin N_{\mathbb{M}})$ 
  | | assert  $\text{DF}'_{x,y} \neq \perp$ 
  | | if  $\text{DF}'_{x,y} \neq \text{DF}_{x,y}$  then
  | | |  $\text{DF}[x][y] \leftarrow \text{DF}'_{x,y}$ 
  | | |  $Q \leftarrow Q \cup \{x\}$ 
  | | end
  | end
end
return  $\text{DF}$ 

```

Algorithm 11: Computation of $\text{PDF}_{\sqsupset_{\text{TIME}[\text{FIRST}]}}$

9.5 Soundness and Minimality of Timing Sensitive Control Dependence

I finish this chapter by making the empirical observation that timing sensitive empirical control dependence $\rightarrow_{\text{tscd}}$ is both sound and minimal with regard to *clocked* traces. Given any sequence

$$t = x_0, x_1, x_2, \dots$$

its clocked sequence is just

$$t^{\circledast} = x_0^{\circledast} [0], x_1^{\circledast} [1], x_2^{\circledast} [2], \dots$$

For clocked traces t^{\circledast} (i.e.: sequences of clocked (pseudo)-edges in G), the S -observable clocked trace $t^{\circledast}|_S$ is defined by removing occurrences of unobservable clocked (partial) edges in the obvious way (similar to Definition 6.3.4 on page 80).

Observation 9.5.1 (Soundness of $\rightarrow_{\text{tscd}}$). Let G be any CFG, and $M \subseteq N$ a set of nodes (the slicing criterion). Let $S = (\rightarrow_{\text{tscd}})^*(M)$ be the timing sensitive backward slice w.r.t M . Then, for any inputs i, i' such that

$$i \sim_S i'$$

I have

$$t_i^{\circledast}|_S = t_{i'}^{\circledast}|_S$$

Observation 9.5.2 (Minimality of $\rightarrow_{\text{tscd}}$). Let G be any CFG, and $M \subseteq N$ a set of nodes (the slicing criterion). Let $S = (\rightarrow_{\text{tscd}})^*(M)$ be the timing sensitive backward slice w.r.t M . Then, for any $n \in S, n \notin M$ and $S' = S \setminus \{n\}$, there exist inputs i, i' such that

$$i \sim_{S'} i'$$

but:

$$t_i^{\circledast}|_{S'} \neq t_{i'}^{\circledast}|_{S'}$$

9.6 Timing Sensitive Control Dependence in Graphs with Unique Exit Node

My algorithms for timing sensitive control dependence $\rightarrow_{\text{tscd}}$ in arbitrary graphs from Section 9.4 are much more complicated than the corresponding algorithms for the computation of nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$. For $\rightarrow_{\text{ntscd}}$, I just needed Algorithm 5 (page 50), and then could directly use Algorithm 1 (page 28). For $\rightarrow_{\text{tscd}}$, on the other hand, I needed not only Algorithm 8 and the more complicated variant Algorithm 11 of Algorithm 1, but also Algorithm 10 for the computation of F.

The reason was that in graphs that are irreducible and lack a unique exit node n_x , the relation $<_{\text{TIME[FIRST]}}$ may not be transitive. There, only the relation $<_{\text{TIME}}$ is guaranteed to be transitive. But what about other graphs? For graphs *with* unique exit node n_x , or reducible graphs, Algorithm 10 is indeed not necessary, and the original Algorithm 1 is adequate for the computation of $\rightarrow_{\text{tscd}}$.

This is true simply due to the following observations:

Observation 9.6.1. Let G be a CFG with unique exit node n_x . Then

$$m \sqsupseteq_{\text{TIME}} n \iff m \sqsupseteq_{\text{TIME[FIRST]}} n$$

Also (writing \sqsupseteq for $\sqsupseteq_{\text{TIME[FIRST]}}$): $\text{ipdom}_{\sqsupseteq}^* = \sqsupseteq$

Observation 9.6.2. Let G be a reducible CFG. Then

$$m \sqsupseteq_{\text{TIME}} n \iff m \sqsupseteq_{\text{TIME[FIRST]}} n$$

Also (writing \sqsupseteq for $\sqsupseteq_{\text{TIME[FIRST]}}$): $\text{ipdom}_{\sqsupseteq}^* = \sqsupseteq$

Then by Observation 9.2.2, Observation 9.2.4 and Observation 9.2.5, $\sqsubseteq_{\text{TIME}[\text{FIRST}]}$ admits an efficient PDF partitioning (Definition 3.2.7 on page 25), and I can use Algorithm 1 on any transitive reduction

$$\prec_{\text{TIME}[\text{FIRST}]} = \prec_{\text{TIME}}$$

of

$$\sqsubseteq_{\text{TIME}[\text{FIRST}]} = \sqsubseteq_{\text{TIME}}$$

But such a reduction is computed by Algorithm 11.

Summary

- Timing sensitive control dependence $\rightarrow_{\text{tscd}}$ can be obtained by a modification of nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$.
- No timing sensitive analogue of nontermination sensitive order dependence is necessary, not even in graphs without unique exit node.
- An algorithm for timing sensitive control dependency can be obtained by a modification of algorithms for nontermination sensitive control dependence.
- This algorithm is complicated for graphs without unique exit nodes, but for graphs with unique exit node a simple algorithm is sufficient.

10 Timing Dependence

As far as I'm concerned, if something is so complicated that you can't explain it in 10 seconds, then it's probably not worth knowing anyway.

(Bill Watterson — The Indispensable Calvin and Hobbes)

Timing sensitive control dependence $n \rightarrow_{\text{tscd}} m$ conflates two kind of dependencies:

- Decisions made at node n that decide whether node m is executed *at all* (i.e.: nontermination sensitive control dependence)
- Decisions made at node n that decide *when* node m is executed (i.e.: timing dependence).

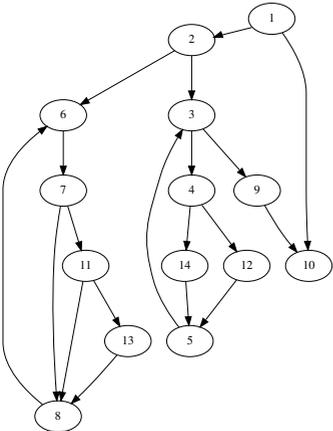
Consider again the graph G repeated in Figure 10.1a. For example, nodes 10 and 9 are both *nontermination sensitively control dependent* on node 3 (since by always successor node 4, node 3 can prevent them from ever executing). They are also *timing dependent* on the same node 3 (since *when* node 9 and 10 are executed is determined by how often node 3 chooses successor node 4 before eventually choosing successor node 9).

In contrast, node 5 is *nontermination sensitively control dependent* on node 3, but *not* timing dependent on node 3, since whenever node 5 executes after node 3, exactly three units of time will have passed.

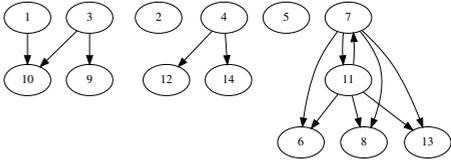
Also, node 8 is timing dependent on node 7, but *not* nontermination sensitively control dependent on node 7.

In this chapter, I propose a new notion \rightarrow_{td} of *timing dependence* that isolates the latter kind of dependence. By a reduction to timing sensitive postdominance $\exists_{\text{TIME[FIRST]}}$, this chapter's slogan will be

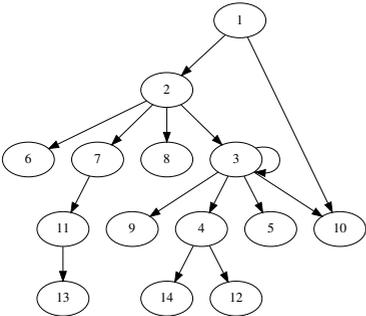
$$\rightarrow_{\text{td}} \approx |N| \times <_{\text{TIME[FIRST]}}$$



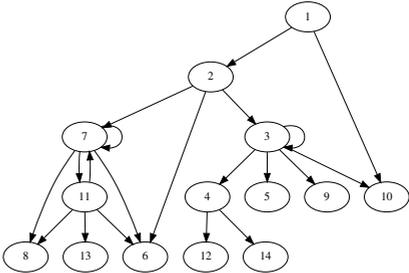
(a) A CFG



(b) Timing Dependence \rightarrow_{td} for the CFG



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



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

Figure 10.1: Timing Dependence

which expresses that timing dependence \rightarrow_{td} in a graph $G = (N, E)$ can be reduced to the computation of $\exists_{\text{TIME}[\text{FIRST}] }^{G_m}$ (represented by its transitive reduction $\prec_{\text{TIME}[\text{FIRST}] }^{G_m}$) for some graph G_m derived from G , for each node $m \in N$.

10.1 Timing Dependence

Definition 10.1.1. Let G be any CFG, and n, n', m nodes in G . I say that n' timing-sensitively postdominates n at position $k \in \mathbb{N}$ on paths towards m , and write

$$n' \sqsupseteq_{\text{TIME}[\text{FIRST tow. } m]}^k \text{in } G n$$

iff on all G -paths starting in n that contain node m , n' first appears at time k . I omit “in G ” whenever possible, and write

$$n' \sqsupseteq_{\text{TIME}[\text{FIRST tow. } m]}^k \text{in } G n$$

if this is the case for *some* k . Formally:

$$n' \sqsupseteq_{\text{TIME}[\text{FIRST tow. } m]}^k \text{in } G n \Leftrightarrow \forall \pi \in {}_n\Pi^G. m \in \pi \rightarrow n' \in_{\text{FIRST}}^k \pi$$

$$n' \sqsupseteq_{\text{TIME}[\text{FIRST tow. } m]}^k \text{in } G n \Leftrightarrow n' \sqsupseteq_{\text{TIME}[\text{FIRST tow. } m]}^k \text{in } G n \text{ for some } k \in \mathbb{N}$$

Proposition 10.1.1. Let G be any CFG, and n, n', m nodes in G . Then if m is unreachable in G from n , I have

$$n' \sqsupseteq_{\text{TIME}[\text{FIRST tow. } m]}^k \text{in } G n \text{ for all } k \in \mathbb{N}$$

If, on the other hand, m is reachable from n in G , let

$$G_m = G_{m \not\rightarrow}^* m$$

be the graph obtained from G by removing all nodes that cannot reach m , and deleting all outgoing edges of m . Then n' first appears at time k in all maximal G -paths starting in n that contain node m iff n' first appears at time k in all maximal G_m -paths starting in n , i.e.:

$$n' \sqsupseteq_{\text{TIME}[\text{FIRST tow. } m]}^k \text{in } G n \Leftrightarrow n' \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k \text{in } G_m n$$

To understand the upcoming definition of timing dependence \rightarrow_{td} , imagine for a conditional node n that for all G -successors x, x' of n I have

$$\begin{aligned} n' &\stackrel{k \text{ in } G}{\equiv}_{\text{TIME}[\text{FIRST tow. } m]} x \\ \text{and } n' &\stackrel{k \text{ in } G}{\equiv}_{\text{TIME}[\text{FIRST tow. } m]} x' \end{aligned}$$

for *some* other node n' .

Then even if the timing of a node m after the execution of n is not constant, i.e.: even if I have

$$\neg m \stackrel{G}{\equiv}_{\text{TIME}[\text{FIRST}]} n$$

still any decision made at n cannot possibly *influence* the timing of m , since the timing of the other node n' after n is constant on paths towards m , and I always must reach n' before m . Hence any influence on the timing of m must happen at or after n' .

Definition 10.1.2. Let $G = (N, E)$ be any graph, and $m \neq n \in N$ be two nodes. Then m is *timing dependent* on n , and I write

$$n \rightarrow_{\text{td}}^G m$$

if there exists no node n' such that for *some* $k \in \mathbb{N}$ and *all* G -successors x of n , on all maximal G -paths starting in x that contain node m , n' first appears at time k , i.e.:

$$n \rightarrow_{\text{td}}^G m \iff \neg \exists n' \in N, k \in \mathbb{N}. \forall n \rightarrow_G x. n' \stackrel{k \text{ in } G}{\equiv}_{\text{TIME}[\text{FIRST tow. } m]} x$$

Observation 10.1.1. Timing dependence \rightarrow_{td} is in fact the timing sensitive part of timing-sensitive control dependence $\rightarrow_{\text{tscd}}$, i.e.: for all graphs G and sets of nodes M ,

$$(\rightarrow_{\text{td}} \cup \rightarrow_{\text{ntscd}})^*(M) = (\rightarrow_{\text{tscd}})^*(M)$$

10.2 Computation of Timing Dependence

Note that for every $m \in N$, the graph

$$G_m = G_{m \rightarrow^* m}$$

is a graph with unique exit node m , and hence:

$$\sqsubseteq_{\text{TIME}}^{G_m} = \sqsubseteq_{\text{TIME}[\text{FIRST}]}^{G_m}$$

An Algorithm for \rightarrow_{td} is available directly from the Algorithm 8 for \langle_{TIME} on page 190, Proposition 10.1.1, and the following observation:

Observation 10.2.1. Let $G = (N, E)$ be any graph, and $m \neq n \in N$ be two nodes. Then

$$\begin{aligned} n \rightarrow_{\text{td}}^G m &\iff \neg \exists n' \neq n. n' \sqsubseteq_{\text{TIME}[\text{FIRST}]}^{G_m} n \\ &\iff \neg \exists n' . n \langle_{\text{TIME}}^{G_m} n' \end{aligned}$$

for any transitive reduction $\langle_{\text{TIME}}^{G_m}$ of $\sqsubseteq_{\text{TIME}}^{G_m}$.

By this observation, I require a computation of $\langle_{\text{TIME}}^{G_m}$ for every node $m \in N$. Informally:

$$\rightarrow_{\text{td}} \approx |N| \times \langle_{\text{TIME}[\text{FIRST}]}$$

Summary

- In timing sensitive control dependence $\rightarrow_{\text{tscd}}$, timing and control dependence are entangled.
- “Pure” timing dependence \rightarrow_{td} can be obtained from timing sensitive control dependence $\rightarrow_{\text{tscd}}$.

11 Timing Stratification

Time, Dr. Freeman? Is it really that time again? It seems as if you only just arrived.

(Valve — G-Man, Half-Life 2)

The two CFG in Figure 9.1 — repeated in Figure 11.1a and 11.1b — were the canonical example of CFG with and with and without timing leak. There, I assumed the set $S = \{m, m_x\}$ of observable nodes, and that the traversal of each CFG edge $n \rightarrow_G m$ took a unit amount $1u$ of time. The second CFG G' can be thought of as a *stratification* of the first CFG G : G' is obtained from G by adding two additional dummy nodes (and corresponding edges) “on the right”¹.

¹ and if i was dealing with edge-labeled CFG, I would label the new edges with an appropriate form of no-op

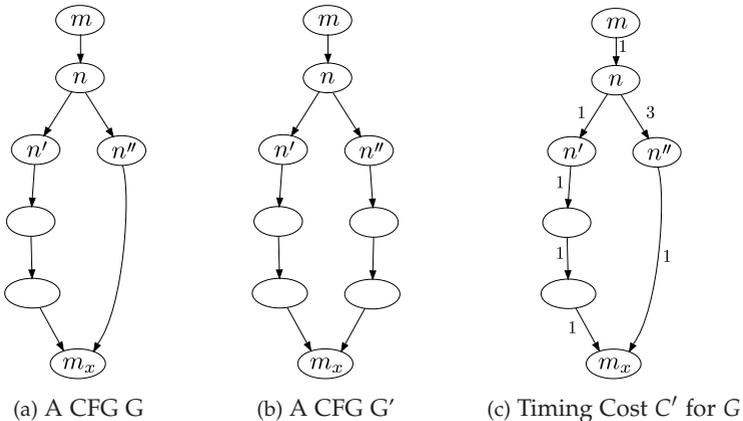


Figure 11.1: Dependence of execution time of m_x on n .

The same effect can be achieved by dropping the assumption that every edge in the CFG has a uniform execution time 1, and instead assume a timing *cost model* C , i.e.: a function

$$C : E \rightarrow \mathbb{N}^+$$

mapping each edge $(u, v) \in E$ of a CFG $G = (N, E)$ to a strictly positive natural number $C(u, v) > 0$: the amount of time spent traversing the edge $u \rightarrow_G v$. For example, I can *stratify* the CFG G in Figure 11.1a by using the timing cost model

$$C'(u, v) = \begin{cases} 3 & \text{for } (u, v) = (n, n'') \\ 1 & \text{otherwise} \end{cases} = \mathbb{1} [(n, n'') \mapsto 3]$$

shown in Figure 11.1c. Here, I assumed the initial timing cost model

$$C = \mathbb{1} := (u, v) \mapsto 1$$

Given an initial timing cost model C for some CFG G , a cost model $C' \geq C$ can be interpreted as a request for a compiler to insert appropriate additional nodes and “no-op”-edges to obtain a CFG G' as in the example, and otherwise keep the initial cost model C . However, making the modified C' explicit has the advantage of leaving the CFG unmodified, making the development in this chapter much simpler.

In this chapter, I will demonstrate how any “common” CFG— i.e.: any CFG that has a unique exit node n_x , or otherwise at least is reducible — can be *stratified*, in a precise sense which will justify the informal slogan

Stratification + Timing Sensitivity = Nontermination Sensitivity

This stratification will be “global” in the sense that it will stratify the timing at every conditional node n , independent of whether the choice at n is relevant w.r.t. a given set S of observable nodes. In other words: it will introduce some unnecessary (given some specific S) increase of

timing cost. Despite these two restrictions, this chapter 11 will serve as the foundation of the application in chapter 12, in which I will demonstrate how to transform out timing leaks in *arbitrary* CFG, and also avoid unnecessary increase of timing cost.

11.1 Timing Sensitive Control Dependence for Arbitrary CFG with Cost Model

It is straightforward to extend the relevant definitions and algorithms for timing sensitive control dependence $n \rightarrow_{\text{tscd}}^G m$ of CFG G with the implicit timing cost model $\mathbb{1}$, to timing sensitive control dependence $n \rightarrow_{\text{tscd}}^{G[C]} m$ of CFG G with explicit timing cost model C . Essentially, I merely need to replace (implicit) occurrences of the cost “1” of some edge $n \rightarrow_G m$ with “ $C(n, m)$ ”.

Definition 11.1.1 (Generalization of Definition 9.1.1 on page 167). Let G be any CFG, C a timing cost model for G , and n, m any nodes in G . Given any path

$$\pi = m_0, m_1, m_2, \dots$$

I say that m appears in π at time k iff $m = m_i$ and

$$k = \sum_{0 \leq j < i} C(m_j, m_{j+1})$$

In this case, I write $m \in^{k[C]} \pi$.

If additionally, $m_i \neq m$ for all $j < i$, I say that m *first* appears in π at time k , and write $m \in_{\text{FIRST}}^{k[C]} \pi$.

Furthermore, I say that $m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k[C]} \text{-postdominates } n$ in G iff on all maximal G -paths starting in n , m first appears at time k . I omit “in G ” whenever possible, and say that $m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^C \text{-post dominates } n$ iff $m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k[C]} \text{-post dominates } n$ for some $k \in \mathbb{N}$. Formally:

$$\begin{aligned} m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k[C] \text{ in } G} n &\Leftrightarrow \forall \pi \in {}_n \Pi_{\text{MAX}}^G. m \in_{\text{FIRST}}^{k[C]} \pi \\ m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^C \text{ in } G n &\Leftrightarrow \forall \pi \in {}_n \Pi_{\text{MAX}}^G. m \in_{\text{FIRST}}^{k[C]} \pi \text{ for some } k \in \mathbb{N} \end{aligned}$$

Remark 11.1.1. Obviously, Definition 11.1.1 really is a generalization of Definition 9.1.1, i.e.:

$$m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^k n \iff m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k[1]} n$$

I define Timing sensitive control dependence for a CFG G with timing cost model C just as expected, but I must not forget to account for the timing cost of traversing an edge $n \rightarrow_G n_l$ or $n \rightarrow_G n_r$ in G .

Definition 11.1.2 (Generalization of Definition 9.1.2 on page 168). Let G be any CFG, C a timing cost model for G , and n, m any nodes in G . Then m is said to be *timing sensitively control-dependent* on n under cost model C , written $n \rightarrow_{\text{tscd}}^{G[C]} m$ or just $n \rightarrow_{\text{tscd}}^C m$, if there exists G successors n_l and n_r of n , and some $k \in \mathbb{N}$ such that for the unique k_l, k_r that satisfy

$$k = k_l + C(n, n_l) \quad \text{and} \quad k = k_r + C(n, n_r)$$

$m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k_l[C]}$ -post dominates n_l , but *not*: $m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k_r[C]}$ -post dominates n_r .

Formally: $n \rightarrow_{\text{tscd}}^C m \Leftrightarrow$

$$\begin{aligned} & m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k_l[C]} n_l \quad \text{and} \\ \neg & m \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{k_r[C]} n_r \end{aligned}$$

for some $k \in \mathbb{N}$ and n_l, n_r such that $n \rightarrow_G n_l$ and $n \rightarrow_G n_r$, where k_l, k_r are defined (given k, n_l, n_r) as above.

Remark 11.1.2. Obviously, Definition 11.1.2 really is a generalization of Definition 9.1.2, i.e.:

$$n \rightarrow_{\text{tscd}} m \iff n \rightarrow_{\text{tscd}}^{\mathbb{1}} m$$

With the example of this generalization it still requires some care, but is otherwise routine to generalize all remaining notions algorithms

from Chapter 9 to CFG and arbitrary timing cost model C . Specifically, not only can I generalize definitions

$$\text{and} \quad \begin{array}{l} m \rightarrow_{\text{tscd}} n \quad \text{to} \quad n \xrightarrow{C}_{\text{tscd}} m \\ m \sqsubseteq_{\text{TIME}[\text{FIRST}]}^k n \quad \text{to} \quad m \sqsubseteq_{\text{TIME}[\text{FIRST}]}^{k[C]} n \end{array}$$

but also:

$$\begin{array}{l} T_{\text{FIRST}} \quad \text{to} \quad T_{\text{FIRST}}^C \\ m \sqsubseteq_{\text{TIME}}^k n \quad \text{to} \quad m \sqsubseteq_{\text{TIME}}^{k[C]} n \\ T \quad \text{to} \quad T^C \\ F_n \quad \text{to} \quad F_n^C \\ F \quad \text{to} \quad F^C \end{array}$$

as well as algorithms 10 to F_n^C and 8 to \langle_{TIME}^C .

I can also generalize Algorithm 11 to $\text{PDF}_{\sqsubseteq_{\text{TIME}}^C}$, but *this* generalization is not completely straight-forward. I describe it in Appendix E.

11.2 Timing-Stratification

Given a CFG G and a timing cost model C for G , to what extent can I hope to stratify the timing behavior of G ? Ideally, I would want all (otherwise) observably equivalent executions of G to also take the same amount of time, but this is clearly impossible in the presence of (unobservable) *loops* in G . The best I can hope to do is to find a modified cost model $C' \geq C$ such that for any observably equivalent executions *in which all loops execute the same number of times*, the executions also take the same total amount of time. But I already have appropriate notions to characterize this difference:

- Timing-sensitive postdominance $\sqsupseteq_{\text{TIME[FIRST]}}$ and timing sensitive control-dependence $\rightarrow_{\text{tscd}}$ are sensitive to *all* timing.
- Nontermination sensitive postdominance \sqsupseteq_{MAX} and nontermination sensitive control-dependence² $\rightarrow_{\text{ntsod}}$ are sensitive to possible non-termination (hence: sensitive to the number of times a loop may execute), but otherwise *insensitive* to timing.

Definition 11.2.1. Let G be any CFG, and C' a timing cost model for G . Then I say that the timing of G under C' is *stratified* if for all sets $M \subseteq N$ of nodes, the timing sensitive backwards slice of M under C' is equal to the nontermination sensitive backward slice of M , i.e.:

$$(\rightarrow_{\text{ntsod}} \cup \rightarrow_{\text{ntsod}})^* = \left(\rightarrow_{\text{tscd}}^{C'} \right)^*$$

The timing of all CFG in Figure 11.2 is stratified under their respective cost model as shown. Note that in Figure 11.2a, I do *not* need to delay the edge $n_e \rightarrow_G n_r$ (for example, by setting $C(n_e, n_r) = 4$), because I already have $n_e \rightarrow_{\text{ntsod}} n_r$ due to the fact that continuing from n_l , the node n_r not necessarily has to be executed. On the other (left) hand, the delays at n and n' stratify this “**if ... then if ...**” region.

² and decisive order dependence $\rightarrow_{\text{dod}} = \rightarrow_{\text{ntsod}}$

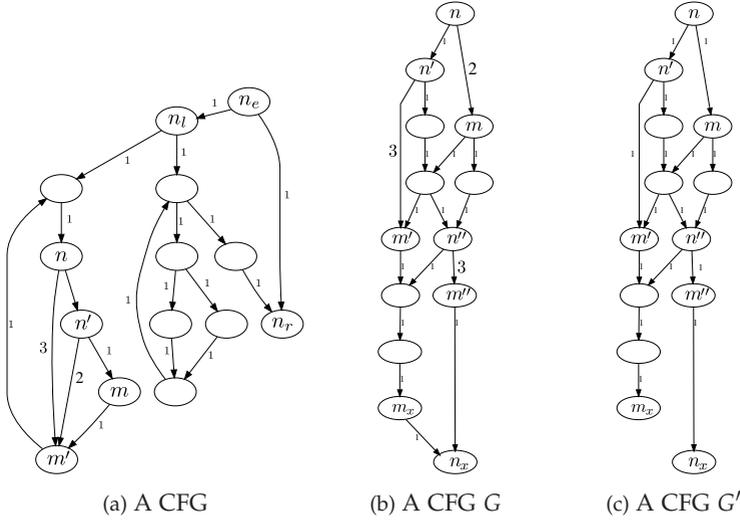
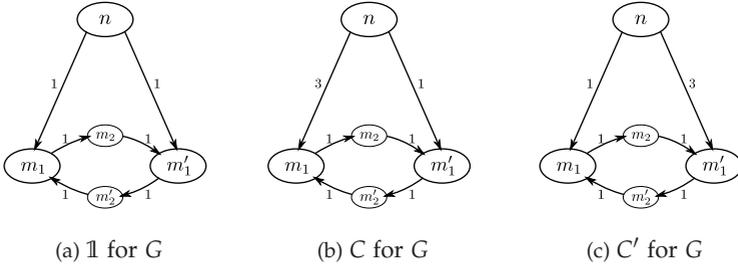


Figure 11.2: Some CFG with stratified timing cost model

In Figure 11.2b, the delay $C(n', m') = 3$ at edge $n' \rightarrow_G m'$ is necessary, but *not* in order to attempt to establish $m' \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{C,3} n'$, which does *not* hold (and *cannot* hold for any C or any k , because $\neg m' \sqsupseteq_{\text{MAX}} n'$). Instead, this delay is necessary to establish $n_x \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{C,k} n'$ for $k = 7$ and the unique exit node n_x , which I must establish (for some k), since $n_x \sqsupseteq_{\text{MAX}} n'$. From an information flow point of view, I must establish this and the other two delays because the execution (time) of n_x might be deemed observable.

On the other hand in Figure 11.2c, which is obtained from Figure 11.2b by removing the edge $m_x \rightarrow_G n_x$, already the default timing cost model $\mathbb{1}$ stratifies the CFG. This example is crucial for information flow control, since removal of this edge certainly does *not* change any observation if n_x is *unobservable*³. In this case, I do *not* need to delay

³ except for the fact that one will observe the pseudo-edge (m_x, \perp) instead of (m_x, n_x) , but this does *not* inform an observer about the *input* to a given execution


 Figure 11.3: Impossibility of Stratification of a CFG G

any edges! I will make use of this technique (removing certain edges, depending on a given set of *observable* nodes) in Chapter 12.

Can I expect to find a stratification C' for every arbitrary CFG (in the sense of Definition 11.2.1)? Consider the CFG in Figure 11.3. It is easy to see that n controls the order *neither* of m_1, m_2 nor of m'_1, m'_2 , since, for example, m_1 *always* occurs before m_2 (when starting in n). So for $M = \{m_1, m_2\}$ and $M' = \{m'_1, m'_2\}$ i have

$$\begin{aligned}
 n &\notin (\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{ntsod}})^*(M) = M \\
 \text{and } n &\notin (\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{ntsod}})^*(M') = M'
 \end{aligned}$$

On the other hand, I have both $n \rightarrow_{\text{tscd}} m_1$ and $n \rightarrow_{\text{tscd}} m'_1$, so under the default timing cost model $\mathbb{1}$:

$$\begin{aligned}
 n &\in \left(\rightarrow_{\text{tscd}}^{\mathbb{1}}\right)^*(M) = M \cup \{n\} \\
 \text{and } n &\in \left(\rightarrow_{\text{tscd}}^{\mathbb{1}}\right)^*(M') = M' \cup \{n\}
 \end{aligned}$$

Can I find *some* timing cost model to stratify G ? I can certainly find a cost model C such that $\neg n \rightarrow_{\text{tscd}}^C m_1$ (as shown in Figure 11.3b). I can also find a different cost model C such that $\neg n \rightarrow_{\text{tscd}}^C m'_1$ (Figure 11.3c).

But I cannot find a cost model C' such that both hold at the same time, for this would require there to be some $k = C'(n, m_1)$ such that $n \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{C',k} m_1$, and some $k' = C'(n, m'_1)$ with $n \sqsupseteq_{\text{TIME}[\text{FIRST}]}^{C',k'} m'_1$. In other words, a cost model C' would need to simultaneously solve the two following equations:

$$\begin{aligned} C'(n, m'_1) + C'(m'_1, m'_2) + C'(m'_2, m_1) - C'(n, m_1) &= 0 \\ -C'(n, m'_1) + C'(m_1, m_2) + C'(m_2, m'_1) + C'(n, m_1) &= 0 \end{aligned}$$

But this is impossible, since all timing cost model are required to be *strictly positive* (i.e.: $C' > 0$).

11.3 An Algorithm for Timing Stratification

I have just shown that it is impossible to find a stratification for *every* CFG (and initial timing cost model C). Nevertheless, in this section I will show an algorithm that *will* return a stratification when applied to a CFG G in which no nodes are (indecisively, nontermination sensitively) order dependent, i.e.: $\rightarrow_{\text{ntsod}}^G = \emptyset$. Note that this class of CFG includes both of the following classes:

1. CFG that are reducible
2. CFG with unique exit node n_x

The key to this algorithm is the observation that I can always stratify such CFG by only delaying CFG-edges $n \rightarrow_G n'$ at *conditional nodes* $n \in \text{COND}$. Given such a node n , I will usually delay *some* but not all of its outgoing edges, going from $C(n, n')$ to $C'(n, n')$ such that $C' \geq C$. I urge the reader to confirm that all examples of stratified CFG shown so far follow this scheme!

The next observation is equally important: whenever timing sensitive post-dominance fails *only due to timing*, i.e.: whenever

$$\begin{array}{l} m \sqsupseteq_{\text{MAX}} n \\ \text{but } \neg m \sqsupseteq_{\text{TIME}}^C n \end{array}$$

then this can inductively reduced to a timing cost model $C' \geq C$ such that for some G -successors n_l, n_r of n I have

$$\begin{array}{l} m \sqsupseteq_{\text{TIME}}^{k_l[C']} n_l \quad \text{and} \quad m \sqsupseteq_{\text{TIME}}^{k_r[C']} n_r \\ \text{but } \underbrace{C'(n, n_l) + k_l}_{=:K_l} < \underbrace{C'(n, n_r) + k_r}_{=:K_r} \end{array}$$

Now the idea for my algorithm is obvious: in this case, just update C' at (n, n_l) to $K_r - k_l$

$$C' [(n, n_l) \mapsto K_r - k_l]$$

The resulting Algorithm 12 is similar to Algorithm 8. The Algorithm 20 (shown on page 398 in the appendix) that computes a least common ancestor of two successor nodes of some node x is only ever called if x has *some* nontermination sensitive immediate postdominator. It also computes a “correction” ΔC of the current timing model C' such that the common ancestor z returned is a timing sensitive least common ancestor *under the timing model* $C' + \Delta C$ (where addition is pointwise).

The computation in Algorithm 12 initializes the \mathbb{N} -labeled pseudo-forest $<$ with a transitive reduction $<_{\text{TIME}}^C$ ⁴. It proceeds by updating $<$ and the timing cost model C' until finally, $<$ is (structurally) a transitive reduction of *maximal-path*⁵ postdominance \sqsupseteq_{MAX} , and (except for self-edges) equal to the transitive reduction $<_{\text{TIME}}^{C'}$ under the *updated* cost model C' . The computation always terminates.

Observation 11.3.1. Let G be a CFG such that no nodes are (indecisively, nontermination sensitively) order dependent, i.e.: $\rightarrow_{\text{ntsod}}^G = \emptyset$. Let C be any timing cost model of G , and C' the timing cost model computed by Algorithm 12. Then $C' \geq C$, and the timing of G is stratified under C' .

Remember that for G be a CFG with unique exit node n_x , I have $\rightarrow_{\text{ntsod}}^G = \emptyset$. The same holds for reducible G .

⁴ which can be obtained from Algorithm 26 on page 415, which itself is just a minor modification of Algorithm 8 on page 190

⁵ i.e.: nontermination sensitive, but not timing sensitive

Input : A CFG G

Input : A initial timing cost model C

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

Output: A timing cost model $C' \geq C$

begin

```

 $C' \leftarrow C$ 
 $< \leftarrow <_{\text{TIME}}^C$ 
 $\text{workset} \leftarrow N_{\text{COND}} \leftarrow \text{COND}_G \cap \{n \mid n <_{\text{MAX}} m \text{ for some } m\}$ 
while  $\text{workset} \neq \emptyset$  do
   $x \leftarrow \text{remove}(\text{workset})$ 
   $(z, k, \Delta C) \leftarrow \text{lca}_<(\{(y, C'(x, y)) \mid x \rightarrow_G y\})$  via alg. 20
  if  $(z, k, \Delta C) \neq \perp$  then
    if  $(z, k) \neq \text{IDOM}[x]$  then
       $\text{IDOM}[x] \leftarrow (z, k)$ 
       $\text{workset} \leftarrow$ 
         $\text{workset} \cup \{n \in N_{\text{COND}} \mid n \neq x, \exists n \rightarrow_G y. y <^* x\}$ 
    end
     $C' \leftarrow C' + \Delta C$ 
  end
end
assert  $\forall n \neq m. \forall k. n <^k m \Leftrightarrow n <_{\text{TIME}}^{k[C']} m$ 
assert  $\forall n \neq m. n <^* m \Leftrightarrow m \exists_{\text{MAX}} n$ 
return  $C'$ 

```

end

Algorithm 12: An efficient algorithm for the computation of a timing stratification $C' \geq C$ of some CFG with initial timing cost model C . Here, $y <^* x$ means: $x = y \vee y <^{k_1} \dots <^{k_c} x$

Summary

- Up to a difference in cost model, timing sensitive control dependence $\rightarrow_{\text{tscd}}$ and nontermination insensitive control dependence $\rightarrow_{\text{ntscd}}$ are essentially the same, in a large class of graphs.
- The timing model that witnesses this correspondence can be computed by a simple algorithm.

Part III

Timing Sensitive Software Security

Es ist nicht genug, zu wissen, man muß auch anwenden;
es ist nicht genug, zu wollen, man muß auch tun.

(Johann Wolfgang von Goethe — Wilhelm Meisters Wanderjahre)

12 Transforming Out Timing Leaks in Arbitrary CFG

If we could perceive time as it really was — what reason would grammar professors have to get out of bed?

(Rosalind and Robert Lutece — Bioshock Infinite)

When transforming out timing leaks, the general goal is — given a program *and* some form of policy that specifies limits on the allowed flow of information — to transform the program such that the transformed program is secure with regard to the given policy. In existing approaches (e.g.: in [Aga00; Mol+06; BRW06]) as well as in the approach I introduce in this chapter, this is only possible if the policy allows information flow due to the number of execution of the programs loops. For example in a simple $\{L, H\}$ -lattice based policy: only if loop predicates whose execution time is L-observable do not depend on H input.

The canonical example of such an approach is [Aga00]. There, the policy consists of a mapping from *left-expressions* (which, aside from program variables, also consists of expressions obtained by dereferencing record-variables, or by indexing array-variables) to security levels L and H. The transformation is syntax-tree-directed, and specified relative to the result of an initial *timing insensitive* information flow analysis, which in turn is specified in form of a security type system. Specifically, the transformation will only successfully transform a loop **while** $e C$ if the type-system inferred the expression e to be of security type L, and *fail* otherwise. For branches **if** $e C_1$ **else** C_2 , timing is harmonized by *cross-copying* parts of one branch into the other.

This somewhat analogous to my algorithm for the timing stratification of CFG from Chapter 11. Imagine a security policy in form of a mapping lvl that maps CFG-nodes to security levels L, H. I deem a program secure if the timing-sensitive backwards-slice of L nodes does not contain H nodes. Then timing stratification is

1. specified relative to the result of *timing insensitive* postdominators \exists_{MAX}^1
2. successful only up to nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$, because there is some node n with $\text{lvl}(n) = \text{H}$ in the timing-sensitive backwards slice of L nodes in the *transformed program*² if and only if such n is in the nontermination sensitive backward slice of L nodes.

The two major *differences* are:

1. Instead of transforming the program (i.e.: the CFG), timing stratification merely change the timing cost model C to $C' \geq C$. But as said before, this then can be understood as a directive to a compiler to insert additional **skip/no-op** instructions, according to the difference $C' - C$ of the stratified timing cost model and the original cost model.
2. Timing stratification is *independent* of the given information flow policy. Instead, it anticipates *all* possible policies, by attempting to harmonize timing-sensitive slicing with nontermination sensitive slicing for *all* slicing criteria M .

I will attack the second difference any minute now, by a simple CFG-preprocessing. But first, I remark with regard to the first difference that the approach [Aga00] and similar techniques (see, e.g., [MS15] for a comparison) have the advantage that they do not require an explicit initial timing cost model C . Indeed, such a model may be hard to get hold of for execution models with under-specified timing behavior (such as the Java Virtual Machine), or otherwise timing models that are at least hard to obtain statically. Just think, for example, of x86-machine code, and the effect of complex memory-cache hierarchies. The transformation in [Aga00] (and others) somewhat alleviate

¹ since in Algorithm 12, i only try so stratify conditional nodes x for which there exists a immediate \exists_{MAX} -postdominator

² i.e.: under the resulting timing cost model C'

this problem by *cross-copying*³ program segments, such that — if so required — both branches essentially execute the same code, and hence can be hoped to have the similar timing behavior (or even the same timing behavior, if the timing model is very simple). In contrast, my approach makes use of the timing cost model C to gauge the execution time of all branches, and equalize (*stratify*) their timing behavior artificially. I suspect — but did not try to establish empirically — that my approach can lead to less runtime overhead in the transformed programs. But again: my approach *does* require an initial timing model C .

Approaches in the tradition of [Aga00] are based on some form of *cross copying*: If a branch condition is H-dependent, code mimicking the behavior of the left branch is appended to the right branch, *and* code mimicking the behavior of the right branch is appended to the left branch. My approach instead is again based not on *cross-copying*, but on the modification of static timing cost model C . To me it appears very difficult to even attempt to apply *cross-copying* or similar techniques in *arbitrary* CFG, since these ideas explicitly rely on the fact that the program model (syntax tree) is *structured* by definition (specifically, it has no arbitrary jumps of control flow). Indeed note that the approach in [Wu+18] which works *not* on the syntax tree but on the CFG still requires the CFG to be of a particular structure. The authors do claim an algorithm to transform CFG into this form *by inserting new variables, and introducing new conditional nodes branching on these new variables*, but it is not clear to me whether their algorithm works for arbitrary CFG. Nor is clear to me the extend to which their method degenerates to the Böhm-Jacopini construction[BJ66] in the general case. Presumably, it is possible to directly apply *cross-copying* or similar techniques to *reducible* CFG which can be obtained by *node splitting*. In general, however this may lead to exponential growth [CFT03]. Though I did not attempt it, what *may* be possible is to apply *cross-copying* (or similar) techniques not to the CFG G , but to some

³ or similar

CFG $G'(M)$ dependent on the slicing criterion M^4 that is “reducible wherever it matters”, and then apply the cross-copying result back to the original CFG G .

⁴ for example: $G' = G_{M/\rightarrow}$

12.1 An Naive Algorithm

My goal of this section is to give a first naive algorithm that — given an *arbitrary* CFG, an arbitrary timing cost model C , and a slicing criterion (i.e.: a set of nodes) M — finds a timing cost model $C' \geq C$ such that for this criterion M

$$(\rightarrow_{\text{ntscd}} \cup \rightarrow_{\text{ntsod}})^*(M) = \left(\rightarrow_{\text{tscd}}^{C'}\right)^*(M)$$

In that case I will say that C' transforms out timing leaks *relevant to M* of G under C *up to timing leaks due to loops*⁵.

My algorithm for transforming out timing leaks is merely an application of the Algorithm 12 for timing stratification. It is based on observations somewhat similar to those that lead to the reduction of nontermination insensitive slicing with $\rightarrow_{\text{ntiod}}$ and $\rightarrow_{\text{nticd}}$ to nontermination insensitive slicing with just $\rightarrow_{\text{nticd}}$ in Section 7.4.

1. The Algorithm 12 for timing stratification only fails to compute a timing cost model C' such that the timing is stratified for *all* slicing criteria, if there is some $<_{\text{TIME-cycle}}$ (remember the example from Figure 11.3 on page 215).
2. But such $<_{\text{TIME-cycle}}$ s (at least: those relevant to a given slicing criterion M) can be broken by *deleting in G* all outgoing edges of nodes $m \in M$.
3. Additionally, as exemplified in Figure 11.2c on page 214, deleting such edges will eliminate the need to delay (some) edges which during *stratification* were only delayed in anticipation of *arbitrary* slicing criteria M .

The transformation then simply works by *stratifying* not the original CFG G under the original timing cost model C , but instead the CFG $G_M := G_{M \nrightarrow}$ obtained from G by removing all edges leaving M . The

⁵ remember that I cannot hope to transform out timing leaks due to loops

stratification is done under the restriction $C_M = C|_{E_M}$ of the original cost model C to the edges E_M still present in G_M . The result is a timing cost model $C'_M \geq C_M$ for G_M , which can simply be extended to a timing cost model $C' \geq C$ for G by using the original timing cost $C(m, m')$ for edges $m \rightarrow_G m'$ missing in G_M .

Observation 12.1.1. Let $G = (N, E)$ be an *arbitrary* CFG, $M \subseteq N$ any set of nodes (the slicing criterion), and C any timing cost model of G . Then for $G_M := (N, E_M) := G_{M \nearrow}$ and $C_M = C|_{E_M}$, let $C'_M \geq C_M$ be the result of running Algorithm 12 on G_M under C_M , and

$$C'(n, m) := \begin{cases} C'_M(n, m) & \text{if } n \rightarrow_{G_M} m \\ C(n, m) & \text{otherwise} \end{cases}$$

Then C' transforms out timing leaks *relevant to M* of G under C up to *timing leaks due to loops*, i.e.:

$$\left(\rightarrow_{\text{ntscd}}^G \sqcup \rightarrow_{\text{ntsod}}^G \right)^* (M) = \left(\rightarrow_{\text{tscd}}^{G[C']} \right)^* (M)$$

Remark 12.1.1. Remember that by Observation 7.2.1, I also have

$$\left(\rightarrow_{\text{ntscd}}^G \sqcup \rightarrow_{\text{ntsod}}^G \right)^* (M) = \left(\rightarrow_{\text{ntscd}}^{G_{M \nearrow}} \right)^* (M)$$



(a) A CFG without timing leak for $M = \{m\}$

(b) Removal of Timing Leaks via Observation 12.1.1

Figure 12.1: Dependence of execution time of m_x on n .

12.2 A More Precise Algorithm

The reduction from G to $G_M = G_{M \nearrow}$ works in the sense of Observation 12.1.1, but introduces unnecessary delays. The simplest example is shown in Figure 12.1. There, the timing leaks transformation shown in Figure 12.1b makes an unnecessary (relative to $M = \{m\}$) delay at n . The transformation here delays an edge at the conditional node n *even though* the node m is not even timing dependent on n .

Such unnecessary delays are avoided by the following Algorithm 13. Apart from to the transition of G to G_M , it is merely a small modification to Algorithm 12. The crucial difference is **highlighted**. The algorithm only ever considers conditional nodes n such that n is in the timing sensitive backward slice of M , but not in the nontermination sensitive backward-slice of M . For the example in Figure 12.1a, which shows a CFG G with initial timing cost model $C = \mathbb{1}$, and slicing criterion $M = \{m\}$, Algorithm 13 returns the unchanged model $C' = C$.

Observation 12.2.1. Let $G = (N, E)$ be an *arbitrary* CFG, $M \subseteq N$ any set of nodes (the slicing criterion), and C any timing cost model of G .

Input : A CFG $G = (N, E)$
Input : A initial timing cost model C for G
Input : A slicing criterion M
Data: A \mathbb{N} labeled pseudo-forest $<$, represented as a map
 $\text{IDOM} : N \leftrightarrow N \times \mathbb{N}$ s.t. $\text{IDOM}[n] = (m, k)$ iff $n <^k m$
Output: A timing cost model $C' \geq C$
begin
 $G_M \leftarrow (N, E_M) \leftarrow G_{M \nrightarrow}$
 $C_M \leftarrow C|_{E_M}$
 $S \leftarrow \left(\rightarrow_{\text{tscd}}^{G_M[C_M]} \right)^* (M) \setminus \left(\rightarrow_{\text{ntscd}}^{G_M} \right)^* (M)$
assert $S = \left(\rightarrow_{\text{tscd}}^{G[C]} \right)^* (M) \setminus \left(\rightarrow_{\text{ntscd}}^G \cup \rightarrow_{\text{ntsod}}^G \right)^* (M)$
 $C'_M \leftarrow C_M$
 $< \leftarrow <_{\text{TIME}}^{G_M[C_M]}$
 $\text{workset} \leftarrow N_{\text{COND}} \leftarrow \text{COND} \cap \{ n \mid n <_{\text{MAX}}^{G_M} m \text{ for some } m \} \cap S$
while $\text{workset} \neq \emptyset$ **do**
 $x \leftarrow \text{remove}(\text{workset})$
 $(z, k, \Delta C) \leftarrow \text{lca}_{<}(\{ (y, C'(x, y)) \mid x \rightarrow_{G_M} y \})$ **via** alg. 20
if $(z, k, \Delta C) \neq \perp$ **then**
if $(z, k) \neq \text{IDOM}[x]$ **then**
 $\text{IDOM}[x] \leftarrow (z, k)$
 $\text{workset} \leftarrow$
 $\text{workset} \cup \{ n \in N_{\text{COND}} \mid n \neq x, \exists n \rightarrow_G y. y <^* x \}$
end
 $C'_M \leftarrow C'_M + \Delta C$
end
end
 $C' \leftarrow C|_{E \setminus E_M} \cup C'_M$
return C'
end

Algorithm 13: An efficient algorithm transforming out timing leaks, given a slicing criterion M for in some CFG with initial timing cost model C . Here, $y <^* x$ means: $x = y \vee y <^{k_1} \dots <^{k_c} x$

Let $C' \geq C$ be the result of running Algorithm 13 on G under C . Then C' transforms out timing leaks *relevant to M* of G under C up to timing leaks due to loops, i.e.:

$$\left(\rightarrow_{\text{ntscd}}^G \cup \rightarrow_{\text{ntscd}}^G \right)^* (M) = \left(\rightarrow_{\text{tscd}}^{G[C']} \right)^* (M)$$

The timing leaks transformation $C' \geq C$ obtained from running Algorithm 13 on G under C is never worse than the timing stratification C'' for $G_{M \nrightarrow}$ obtained from Algorithm 12 along the lines of Observation 12.1.1.

Observation 12.2.2. Let C', C'' as above. Then $C' \leq C''$.

But is the timing leaks transformation C' in some global sense optimal? I cannot answer this question for the reason that I did not attempt formulate such a criterion! The only reason I know that Algorithm 13 is never worse than Algorithm 12 is that both only ever delay timing at conditional nodes n and their G -successors, and hence both limit themselves to the same class of timing cost model transformations. But in general, there is no reason why a somehow globally optimal transformation should not insert delays elsewhere, possibly at *join* nodes. In such cases (say: C''), an edge-wise comparison of the resulting timing cost model can no longer be expected to yield $C'' \geq C'$ or $C' \geq C''$, even though one of these might intuitively be better than the other.

Summary

- Given a set of observable nodes and fixed cost model C , any non-loop timing leaks can be in principle be “transformed out”.

13 Micro-Architectural Dependencies

And by the help of Microscopes, there is nothing so small, as to escape our inquiry; hence there is a new visible World discovered to the understanding.

(Robert Hooke — Micrographia)

Timing based attacks often rely on difference of timing behavior due to *micro-architectural* state, which may in turn depend on private data. Intuitively, given a concrete machine implementing a given (“higher-level”, more abstract *macro-architectural*) execution semantics, micro-architectural state is that part of the concrete machine which is “invisible” to the higher level execution semantics. For example, the micro-architectural state of modern CPU includes the state of data- and code *caches* or the state of the CPU’s *instruction pipeline* (fetch, decode, execute, etc). The micro-architectural state of two different CPU implementing the same instruction set architecture (i.e.: the same *macro-architectural* semantic) may differ considerably. Consider, for example, AMD and Intel CPUs implementing the x86 instruction set architecture. Still ideally, a x86 program’s behavior is the same on both AMD and Intel CPUs, *except for the timing behavior*¹.

In this chapter, I will develop a new method for the dependence analysis of arbitrary micro-architectural state for single-threaded execution models, in form of a new notion $n \rightarrow_{\mu d} m$ of *micro-architectural*

¹ and, of course, difference of behavior due to multi-processing, which I do not cover in this chapter

dependence between nodes n, m . After an introductory example in Section 13.1, I introduce in Section 13.2 the necessary technical framework, and those assumption on the micro-architectural state model necessary for the dependency analysis. In Section 13.6, I define micro-architectural dependence by a reduction to non-termination insensitive control dependence $\rightarrow_{\text{nticd}}$. In Section 13.4, I discuss limitations of my approach, and finally in Section 13.5 I explain how micro-architectural dependence $n \rightarrow_{\mu\text{d}} m$ and timing sensitive control dependence $\rightarrow_{\text{tscd}}$ can be combined to obtain timing sensitive slices that are sensitive to timing channels due to differences in micro-architectural state.

As a running example, I will augment a simple “standard” semantics for CFGs with a simple micro-architectural state, comprised of a simple one-level data cache. I try to make plausible — but give no formal proof — that my method does indeed for work for arbitrary micro-architectural state. For the example cache-semantics, I validated my claims by extensive random testing.

13.1 Introduction

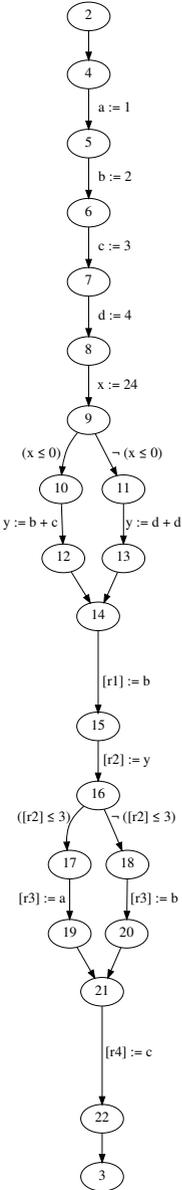
Consider the control flow graph in Figure 13.1a on page 236. Its edges are labeled with assignments and guards that refer to (cachable) variables a, b, \dots , and uncachable registers $r1, r2, \dots$

I want to know whether the choice made at node $n = 9$ influences the execution time of the reads of variables b and y at nodes 14, 15. I 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 13.1b, I 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).

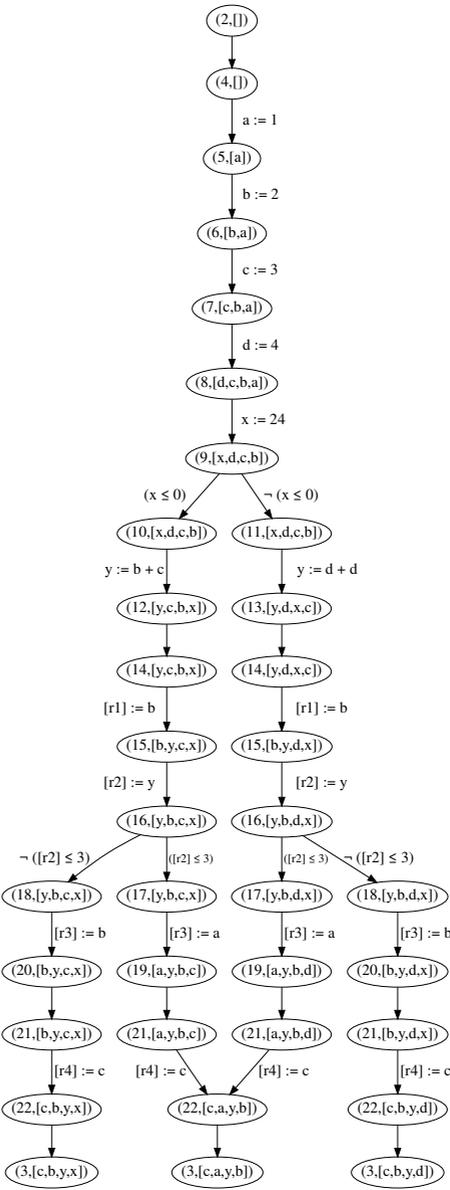
In the example, executions can reach the control node $m = 15$ at cache states represented by either $[b, y, c, x]$, or by $[b, y, d, x]$. Which of these (abstract) cache states is reached is determined by the macro-architectural choice made at $n = 9$. But it is easy to see that the execution time of 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 $15 \xrightarrow{r2:=y} 16$ at m .

For the read of variable b at node $m = 14$, on the other hand, 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



(a) Control Flow Graph



(b) Cache Aware Abstract Executions

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 $(21, [b, y, d, x])$).

On the other hand, 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 13.1b.

Remark 13.1.1. With only 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 control flow graph G , i.e. I will later say:

$$\rightarrow_{\mu d}^G = \{(9, 21), (9, 14)\}$$

13.2 Control Flow Graphs

In this chapter, I will assume programs given in form of *control flow graphs*, i.e. graphs (N, E) with *labeled* edges $n \xrightarrow{l}_G m \in E$. Labels $l \in L$ may be either assignments or guards. Assignments read and write *macro-architectural* state σ_M (e.g.: values of program variables). Guards evaluate macro-architectural state, and then either allow control flow to pass, or not (without changing any macro-architectural state σ_M).

Definition 13.2.1. Formally, I assume some set Σ_M of possible macro-architectural states σ_M , and labels $L = A \cup G$. For each label l I assume a macro-architectural (concrete) interpretation l^M of l such that for $a \in A$, a^M is a function

$$a^M : \Sigma_M \rightarrow \Sigma_M$$

and for $g \in G$, g^M is a function

$$g^M : \Sigma_M \rightarrow \{\mathbf{false}, \mathbf{true}\}$$

A control flow graph $G = (N, E)$ is said to be (two-branch) deterministic (with regard to macro-architectural state) if

1. whenever $n \xrightarrow{a}_G m$, then this is the only edge leaving n in G
2. whenever $n \xrightarrow{g}_G m$ and $n \xrightarrow{g'}_G m'$, then either $m = m'$ and $g = g'$, or $m \neq m'$ and the interpretation $(g')^M$ of the label g' is the negation of the interpretation g^M of the label g : $(g')^M = \neg \circ g^M$ and.

I assume all control flow graphs to be deterministic in this sense.

The macro-architectural small-steps semantics of control flow graph G is then simply given by

$$\frac{n \xrightarrow{a}_G m \quad \sigma'_M = a^M(\sigma_M)}{(n, \sigma_M) \rightarrow (m, \sigma'_M)} \text{Ass}$$

$$\frac{n \xrightarrow{g}_G m \quad g^M(\sigma_M) = \mathbf{true}}{(n, \sigma_M) \rightarrow (m, \sigma_M)} \text{Guard}$$

Later, it will be convenient to assume the macro-architectural to consists of *program variables* $v \in \text{Var}$, and then assume some functions $\text{use}, \text{def} : L \rightarrow 2^{\text{Var}}$, with which I can then define the *use* and *def* sets of a node $n \in N$:

$$\begin{aligned} \text{use}(n) &= \bigcup_{n \xrightarrow{l}_G m} \text{use}(l) \\ \text{def}(n) &= \bigcup_{n \xrightarrow{l}_G m} \text{def}(l) \end{aligned}$$

Definition 13.2.2. I assume a set Σ_μ of micro-architectural states $\sigma_\mu \in \Sigma_\mu$. For each label l , I assume a (concrete) micro-architectural interpretation l^μ of l such that l^μ is a function

$$l^\mu : \Sigma_M \times \Sigma_\mu \rightarrow \Sigma_\mu$$

I also assume for each label l a (concrete) timing cost interpretation l^\ominus of l such that l^\ominus is a function

$$l^\ominus : \Sigma_\mu \rightarrow \mathbb{N}$$

that indicates how much time a transition along an edge labeled with l will take in some micro-architectural state σ_μ .

The full small-steps semantics of a control flow graph G is then given by

$$\frac{n \xrightarrow{a}_G m \quad \sigma'_M = a^M(\sigma_M) \quad \sigma'_\mu = a^\mu(\sigma_M, \sigma_\mu) \quad \Delta t = a^\ominus(\sigma_\mu)}{(n, \sigma_M, \sigma_\mu, t) \rightarrow (m, \sigma'_M, \sigma'_\mu, t + \Delta t)} \text{Ass}$$

$$\frac{n \xrightarrow{g} m \quad g^M(\sigma_M) = \mathbf{true} \quad \sigma'_\mu = g^\mu(\sigma_M, \sigma_\mu) \quad \Delta t = g^\circ(\sigma_\mu)}{(n, \sigma_M, \sigma_\mu, t) \rightarrow (m, \sigma_M, \sigma'_\mu, t + \Delta t)} \text{Guard}$$

Note that this semantics is deterministic if G is (in the macro-architectural sense), since only the macro-architectural state decides which CFG edge is traversed.

Remark 13.2.1. The micro-architectural state will often (to some extent) *track* the macro-architectural state. For example, the micro-architectural state of a data cache will include copies the value of those (macro-architectural) variables that are currently cached. Hence the concrete micro-architectural state space for a given program (i.e.: for a control flow graph with designated initial node) can become arbitrarily large.

But since my approach will require me to compute (a representation of) all micro-architectural state of a given program, I require an *abstract* representation $\sigma_\mu^\#$ of micro-architectural states (with a corresponding abstract interpretation of labels l), which still allows me to assign a definite timing cost to each label l .

Definition 13.2.3 (Abstract Micro-Architectural Semantics). Let α be a (abstraction) function

$$\alpha : \Sigma_\mu \rightarrow \Sigma_\mu^\#$$

from (concrete) micro-architectural states σ_μ to some abstraction $\alpha(\sigma_\mu)$ in some set $\Sigma_\mu^\#$ of abstract micro-architectural states.

I write

$$\gamma(\sigma_\mu^\#) = \{ \sigma_\mu \mid \alpha(\sigma_\mu) = \sigma_\mu^\# \}$$

Recall that in the full concrete small-step semantic (Definition 13.2.2), the choice which control flow graph edge to traverse next depends only on the macro-architectural state. Also, given such a choice (e.g.: $n \xrightarrow{l} m$), the concrete micro-architectural choice is deterministic.

The abstract micro-architectural small-steps semantics of a control flow graph G is then just

$$\frac{n \xrightarrow{l}_G m \quad \alpha(\sigma_\mu) = \sigma_\mu^\# \quad \sigma'_\mu = l^\mu(\sigma_M, \sigma_\mu) \quad \alpha(\sigma'_\mu) = \sigma_\mu^{\prime\#}}{(n, \sigma_\mu^\#) \xrightarrow{l} (m, \sigma_\mu^{\prime\#})} \text{Label}$$

Even for a fixed choice $n \xrightarrow{l}_G m$, the abstract micro-architectural small-steps semantics may be indeterministic. For example, $\sigma_\mu^{\prime\#}$ may differ for different σ_M . For an example, see Section 13.6.

Remark 13.2.2. In the data-cache example, a concrete cache σ_μ (mapping *some* program variables to values, and ordering these variables by recency of use) can be abstractly represented by the list $\alpha_{\text{cache}}(\sigma_\mu)$ of variables currently in the cache. Using a different abstraction, a concrete cache σ_μ can also be abstractly represented by the set $\alpha_{\text{incache}}^{\text{use}(m)}(\sigma_\mu)$ of variables used at some node m that are in cache at state σ_μ .

In Figure 13.1b, I showed the abstract micro-architectural small-steps semantics for the control flow graph in Figure 13.1a, for all abstract configurations reachable from the abstract configuration $(2, \square)$. I used the abstraction α_{cache} .

Definition 13.2.4. A function

$$\alpha^\natural : \Sigma_\mu \rightarrow \Sigma_\mu^\natural$$

from (concrete) micro-architectural states σ_μ to some abstraction $\alpha^\natural(\sigma_\mu)$ in some set Σ_μ^\natural of abstract micro-architectural states is said to respect timing cost for label l if

for all σ_μ, σ'_μ with $\alpha^\natural(\sigma_\mu) = \alpha^\natural(\sigma'_\mu)$ I have:

$$l^\natural(\sigma_\mu) = l^\natural(\sigma'_\mu)$$

Remark 13.2.3. The abstraction α_{cache} respects timing cost for all labels, while the coarser abstraction $\alpha_{\text{incache}}^{\text{use}(m)}$ respects all labels at control flow graph edges leaving m . This is because I assume that the time it takes to access a variable x does only depend on whether x is currently cached or not, but not on the concrete *value* of x (or its position in the cache).

I require one property of abstractions α . Intuitively, the abstraction must *not* conflate two different micro-architectural states σ_μ^1 and σ_μ^2 if these may lead to different execution times from configurations consisting of these two micro-architectural states. This is necessary because if whenever such micro-architectural states are conflated by α i.e.:

$$\alpha \left(\sigma_\mu^1 \right) = \alpha \left(\sigma_\mu^2 \right) = \sigma_\mu^\#$$

then these may introduce *join nodes* $(m, \sigma_\mu^\#)$ in the graph obtained from the abstract micro-architectural semantics. But my notion of micro-architectural dependencies will then judge this *join node*² to be independent from choices made before at conditional nodes (n, \dots) whose successors have all joined in $(m, \sigma_\mu^\#)$, which is unsound if executions continuing from there may have different timing behavior, possibly due the actual choice made at n .

Definition 13.2.5. An abstraction function α is said to respect timing cost for all possible execution if for any two micro-architectural states σ_μ^1 and σ_μ^2 that are conflated by α , any two executions starting at full concrete configurations which differ only in these two micro-architectural states are indiscernible “up to micro-architectural state”.

Formally: For all micro-architectural states σ_μ^1 and σ_μ^2 such that

$$\alpha \left(\sigma_\mu^1 \right) = \alpha \left(\sigma_\mu^2 \right)$$

² and following nodes

and all nodes n , macro-architectural states σ_M and points in time $t \in \mathbb{N}$, whenever

$$\text{and } \begin{pmatrix} n, \sigma_M, \sigma_\mu^1, t \end{pmatrix} \rightarrow \dots \rightarrow \begin{pmatrix} m, \sigma_M^1, \sigma_\mu^{1'}, t_1 \end{pmatrix}$$

$$\begin{pmatrix} n, \sigma_M, \sigma_\mu^2, t \end{pmatrix} \rightarrow \dots \rightarrow \begin{pmatrix} m, \sigma_M^2, \sigma_\mu^{2'}, t_2 \end{pmatrix}$$

with no other occurrences of m in the transition sequences, then $\sigma_M^1 = \sigma_M^2$ and $t_1 = t_2$, as well as $\alpha(\sigma_\mu^{1'}) = \alpha(\sigma_\mu^{2'})$.

Remark 13.2.4. In Definition 13.2.5, $\sigma_M^1 = \sigma_M^2$ must already hold simply due to the fact that both the choice of successor-nodes and the macro-architectural successor state are independent from micro-architectural state (and the fact that the control flow graph is assumed to be deterministic). $t_1 = t_2$ must already hold if α respects timing cost for all labels.

13.3 Micro-Architectural Dependencies

My ultimate goal is to define, given a control flow graph $G = (N, E)$, a binary relation $\rightarrow_{\mu d}^G$ on nodes $n, m \in N$ such that

$$n \rightarrow_{\mu d}^G m$$

if the (macro-architectural) choice made at n influences the timing behavior at node m , if m is executed (after n)³.

Since the timing behavior at node m depends (by Definition 13.2.2) only on the *micro*-architectural state, I need to determine if the choice made at n influences micro-architectural state at node m . In other words: for those labels l at control flow graph edges $m \xrightarrow{l} m'$ leaving m , I am interested in whether the choice made at n influences the timing behavior l^{\otimes} of micro-architectural states σ_{μ} at m .

The general idea that I develop in this chapter is to define $\rightarrow_{\mu d}^G$ in terms of nontermination insensitive control dependence $\rightarrow_{\text{nticd}}$, but not on the control flow graph G , but on a graph derived from the abstract micro-architectural small steps semantics

$$(n, \sigma_{\mu}^{\#}) \xrightarrow{l} (n', \sigma_{\mu}^{\prime\#})$$

as defined in Definition 13.2.3, and previously exemplified in Figure 13.1b on page 236.

The informal slogan

$$\rightarrow_{\mu d}^G \approx |N| \times \rightarrow_{\text{nticd}}^{\#}$$

³ The question whether m is executed at all after n , and whether this depends on the choice made at n , is already answered by control dependence and related notions.

expresses that $\rightarrow_{\mu d}^G$ can be obtained from $|N|$ applications of nontermination-insensitive control dependence $\rightarrow_{\text{nticd}}$ on graphs $G_{\#}$ derived from G (and $m \in N$): One application for each node $m \in N$.

Definition 13.3.1. Fix some deterministic control flow graph $G = (N, E)$ and some (initial) node $n^0 \in N$ in G as well as some (initial) concrete micro-architectural state $\sigma_{\mu}^0 \in \Sigma_{\mu}$.

Let $N^{\#}$ be the set of configurations $(n, \sigma_{\mu}^{\#})$ reachable from $(n^0, \alpha(\sigma_{\mu}^0))$ in the abstract micro-architectural small steps semantics (Definition 13.2.3).

Then the graph $G_{\alpha} = (N_{\alpha}, E_{\alpha})$ consists of the nodes

$$N_{\alpha} = \{ (n, \gamma(\sigma_{\mu}^{\#})) \mid (n, \sigma_{\mu}^{\#}) \in N^{\#} \} \subseteq N \times 2^{\Sigma_{\mu}}$$

and L -labeled edges

$$E_{\alpha} = \{ ((n, \gamma(\sigma_{\mu}^{\#})), l, (m, \gamma(\sigma_{\mu}^{\#}))) \mid (n, \sigma_{\mu}^{\#}) \in N^{\#}, (n, \sigma_{\mu}^{\#}) \xrightarrow{l} (m, \sigma_{\mu}^{\#}) \}$$

From now on, I assume G_{α} to be finite. Also from now on, I fix a node $m \in N$ in the control flow graph G , and let α^{\circledast} be an abstraction coarser than α that respects timing cost for all labels l at edges leaving $m \in G$.

I write $M^{\#} = \{(m, \sigma_{\mu}^{\#}) \in N^{\#}\}$ for the set of abstract configurations at the fixed node m , and

$$M = \{ (m, \gamma(\sigma_{\mu}^{\#})) \mid (m, \sigma_{\mu}^{\#}) \in M^{\#} \}$$

for the set of corresponding nodes in G_{α} .

Then I define

$$G_{\alpha, m} = (G_{\alpha}^{\rightarrow * M})_{M \nrightarrow}$$

as the graph obtained from G_α by removing all nodes that cannot reach M , and all edges leaving M . I also define the graph

$$G_{\alpha,m,\alpha^\circledast}$$

as the graph obtained from $G_{\alpha,m}$ by merging for each $\sigma_\mu^\circledast \in \Sigma_\mu^\circledast$ those nodes (m, Σ) at m such that for all $\sigma_\mu \in \Sigma$, σ_μ is represented in Σ_μ^\circledast by σ_μ^\circledast (i.e.: $\alpha^\circledast(\sigma_\mu) = \sigma_\mu^\circledast$).

In Figure 13.1b on page 236, I previously showed the graph G_α for start node $n^0 = 2$, $\alpha = \alpha_{\text{cache}}$, and the initially empty cache σ_μ^0 . Now in Figure 13.2a on page 247, I show the graph $G_{\alpha,m,\alpha^\circledast}$ for $m = 14$ and $\alpha^\circledast = \alpha_{\text{incache}}^{\text{use}(m)}$, and in Figure 13.2b I do the same for $m = 21$.

For example in G_α , I had two nodes $(14, \gamma([y, c, b, x]))$ and $(14, \gamma([y, d, c, x]))$. On the other hand in $G_{\alpha,14,\alpha^\circledast}$, the abstraction

$$\alpha^\circledast = \alpha_{\text{incache}}^{\text{use}(14)} = \alpha_{\text{incache}}^{\{b\}}$$

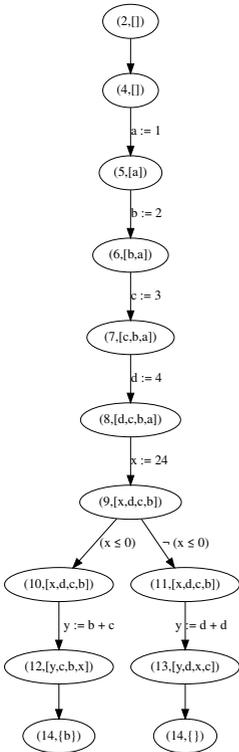
considers only whether the variable b is in the cache. Hence I there have two nodes $(14, \{b\})$ and $(14, \{\})$.

Similarly in G_α , I had four nodes $(21, \gamma([b, y, c, x]))$, $(21, \gamma([a, y, b, c]))$, $(21, \gamma([a, y, b, d]))$ and $(21, \gamma([b, y, d, x]))$. On the other hand in $G_{\alpha,21,\alpha^\circledast}$, the abstraction

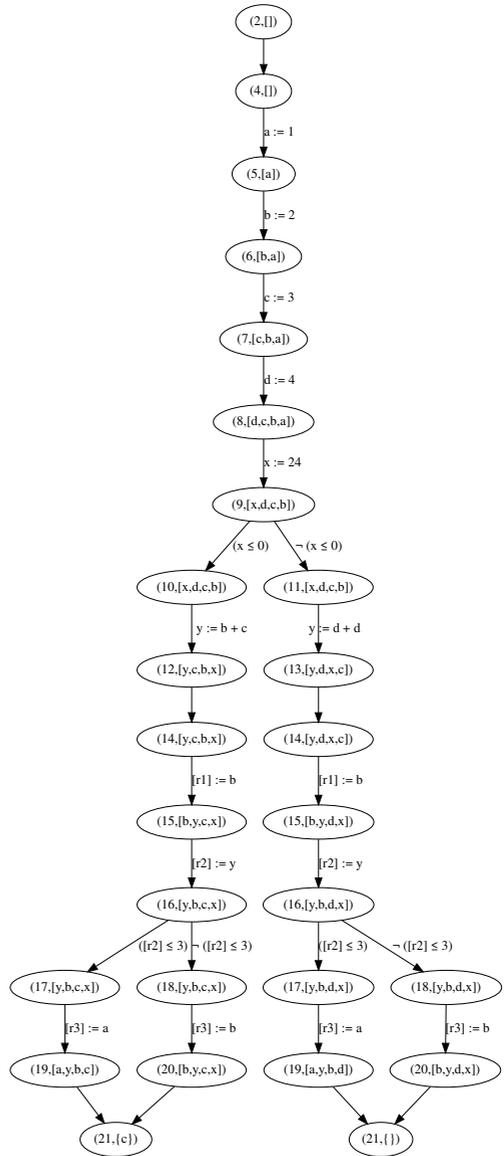
$$\alpha^\circledast = \alpha_{\text{incache}}^{\text{use}(21)} = \alpha_{\text{incache}}^{\{c\}}$$

considers only whether the variable c is in the cache. Hence I there have two nodes $(21, \{c\})$ and $(21, \{\})$.

Evidently, those nodes (m, Σ) for which $m = 14$, are control dependent on the node $(9, [x, d, c, b])$ in $G_{\alpha,14,\alpha^\circledast}$, but no other node. Similarly, those nodes (m, Σ) for which $m = 21$, are control dependent on the



(a) The Graph $G_{\alpha,14,\alpha^{\text{fin}}}$



(b) The Graph $G_{\alpha,21,\alpha^{\text{fin}}}$

same node $(9, [x, d, c, b])$ in $G_{\alpha, 21, \alpha^{\text{tr}}}$, but no other node. This is just as expected, since I had predicted:

$$\rightarrow_{\mu d}^G = \{(9, 21), (9, 14)\}$$

Let me now make some general observations for the nontermination insensitive postdominance relation in $\sqsupseteq_{\text{SINK}}^G$ in the graph $G_{\alpha, m, \alpha^{\text{tr}}}$ for a fixed node m .

- The transitive reduction $<_{\text{SINK}}^G$ is a forest. Its roots include the nodes (m, Σ) , i.e.: those nodes (n', Σ) with $n' = m$. All roots are singular nodes (i.e.: there are only trivial control-sinks).
- The conditional nodes (n, Σ) in $G_{\alpha, m, \alpha^{\text{tr}}}$ correspond to either conditional nodes n in G , or nodes at which the concrete macro-architectural state σ_M may influence the abstract micro-architectural successor state $\sigma_{\mu}^{\#}$ (see Section 13.6 for an example).
- Because all outgoing edges of M were deleted in $G_{\alpha, m}$, for each node (n, Σ_n) , at most one node (m, Σ) postdominates (n, Σ_n) :

$$\left| \{ (m, \Sigma) \mid (m, \Sigma) \sqsupseteq_{\text{SINK}}^G (n, \Sigma_n) \} \right| \leq 1$$

- If some (m, Σ) postdominates some (n, Σ_n) , then all concrete executions first reaching a state represented by (n, Σ_n) , and then next reaching a state at m , have the same timing behavior at m . This is because all $\sigma_{\mu} \in \Sigma$ have the same timing behavior (see Definition 13.2.4).
- Micro-architectural dependency are *not* meant to include “normal” control dependencies. Hence it is intentional that even if some (m, Σ) postdominates some (n, Σ_n) in $G_{\alpha, m, \alpha^{\text{tr}}}$, choices at n may still decide whether a concrete execution reaches m at all (i.e.: m may still be *control dependent* on n in G). This is be-

cause $G_{\alpha,m,\alpha^\circ}$ only consists of nodes that “eventually” (in the nontermination insensitively sense) must reach m (at *some* micro-architectural state σ_μ). This was achieved by taking $G_\alpha \xrightarrow{*} M$.

- If *no* node (m, Σ) postdominates a given node (n, Σ_n) , but some other node $(n', \Sigma_{n'})$ *does* postdominate (n, Σ_n) , i.e., if:

$$(n', \Sigma_{n'}) \stackrel{G_{\alpha,m,\alpha^\circ}}{\exists_{\text{SINK}}} (n, \Sigma_n)$$

(with $n' \neq m$, $n' \neq n$), then any decision made at n cannot possibly influence the micro-architectural state at m , since all concrete executions from states represented by (n, Σ_n) have then joined again at a state represented by $(n', \Sigma_{n'})$ *before* reaching m .

But this last remark just describes the converse of

$$\rightarrow_{\text{ntind}}^{G_{\alpha,m,\alpha^\circ}} = \left(\rightarrow_{\text{nticd}}^{G_{\alpha,m,\alpha^\circ}} \right)^*$$

(see Definition 7.1.1 on page 122, and Observation 7.1.1). Hence I *could* just define⁴

$$n \rightarrow_{\mu\text{d}}^G m \iff \exists (n, \Sigma_n), (m, \Sigma) \in N_{\alpha,m,\alpha^\circ} \cdot (n, \Sigma_n) \rightarrow_{\text{ntind}}^{G_{\alpha,m,\alpha^\circ}} (m, \Sigma)$$

and obtain a sound notion of micro-architectural dependencies. But it turns out that this notion would be too coarse, because it distinguishes too many abstract states (n, Σ_n) .

To see this, consider Figure 13.3b on page 251, which shows the graph $G_{\alpha,26,\alpha^\circ}$ for the control flow graph G in Figure 13.3a. The question is if $m = 26$ is micro-architecturally dependent on $n = 10$, i.e. if

$$10 \rightarrow_{\mu\text{d}}^G 26$$

⁴ with a misuse of notation: the names n, m are not meant to be bound by \exists , but only the names Σ, Σ_n

If I were to answer this directly by consulting $G_{\alpha,26,\alpha}^{\text{nticd}}$, I would have to affirm this, since I have

$$(10, \sqcup) \xrightarrow{G_{\alpha,26,\alpha}^{\text{nticd}}} (18, \sqcup) \xrightarrow{G_{\alpha,26,\alpha}^{\text{nticd}}} (26, \sqcup) \text{ hence } (10, \sqcup) \xrightarrow{G_{\alpha,26,\alpha}^{\text{ntind}}} (26, \sqcup)$$

But in reality, node $n = 10$ can *not* influence whether x is in cache at node $m = 26$. This is because any concrete execution at a concrete configuration represented by $(18, [a, z, d, c])$ will make the same choices as a concrete execution at a concrete configuration represented by $(18, [a, z, c, d])$ ⁵: if $(18, [a, z, d, c])$ proceeds along the edge labeled $l = 0 \leq r_0 \cdot r_1$ (to $(19, [a, z, d, c])$), then $(18, [a, z, c, d])$ will *also* proceed along an edge labeled l (to $(19, [a, z, c, d])$). But it is now easy to see that $(19, [a, z, c, d])$ and $(19, [a, z, d, c])$ behave equivalently with regard to the relevant cache state at $m = 26$, since from both these abstract configuration, execution must reach a configuration at which the relevant variable c is *not* in cache, i.e.:

$$(26, []) \sqsubseteq_{\text{SINK}}^{G_{\alpha,26,\alpha}^{\text{nticd}}} (19, [a, z, c, d])$$

and $(26, []) \sqsubseteq_{\text{SINK}}^{G_{\alpha,26,\alpha}^{\text{ntind}}} (19, [a, z, d, c])$

Similarly, for the other edge $l' = \neg l$ at $(18, [a, z, c, d])$ and $(18, [a, z, d, c])$, I have for their respective successors at 20:

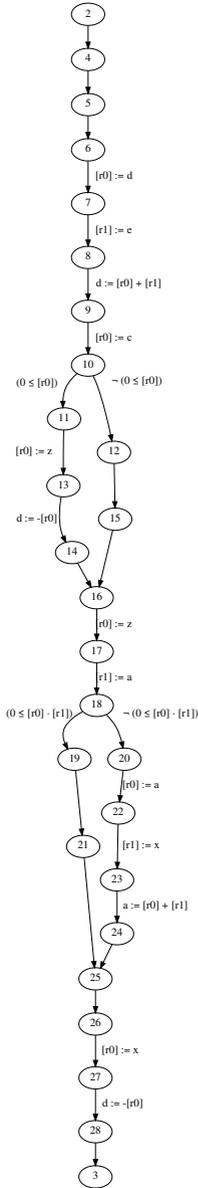
$$(26, [x]) \sqsubseteq_{\text{SINK}}^{G_{\alpha,26,\alpha}^{\text{nticd}}} (20, [a, z, c, d])$$

and $(26, [x]) \sqsubseteq_{\text{SINK}}^{G_{\alpha,26,\alpha}^{\text{ntind}}} (20, [a, z, d, c])$

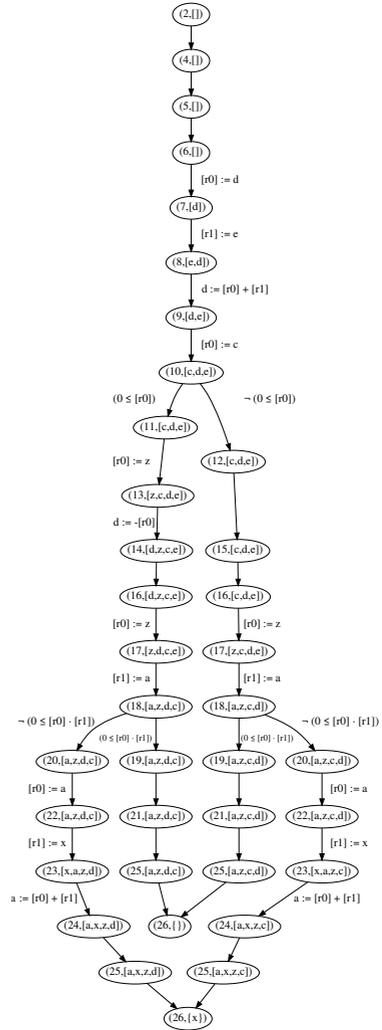
In other words, both successor configurations must reach node $m = 26$ at a configuration in which x is in the cache.

In summary: while the macro-architectural node $n = 10$ does indeed influence in whether execution reaches node 18 at $(18, [a, z, d, c])$ or at $(18, [a, z, c, d])$, this choice at $n = 10$ is irrelevant for the (timing

⁵ note the different order of c, d



(a) A Control Flow Graph G



(b) The Graph $G_{\alpha,26,\nu^\mu}$

relevant) micro-architectural state at node $m = 26$. In this sense (with regard to node $m = 26$), $(18, [a, z, d, c])$ and $(18, [a, z, c, d])$ are *equivalent*: if in a concrete configuration abstracted by $(18, [a, z, d, c])$ and with macro-architectural state σ_M , the concrete execution proceeds towards a configuration at $m = 26$ with x in cache, then any concrete execution from a concrete configuration abstracted by $(18, [a, z, c, d])$ with the same macro-architectural state σ_M will *also* proceed towards a configuration at $m = 26$ with x in cache.

What I need then is a general notion \equiv of equivalence between nodes $(n, \Sigma_n), (n', \Sigma_{n'})$ in $G_{\alpha, m, \alpha^{\circledast}}$.

Intuitively, one might want to make use specifically of the fact that in this example, the abstract cache states $[a, z, d, c]$ and $[a, z, c, d]$ are in some sense equivalent at configurations with control flow graph node $n' = 18$. This *can* be done, but one has to be very careful to find the correct equivalence here, since whether $[a, z, c, d]$ and $[a, z, d, c]$ “behave equivalently” depends very much their position in, and the structure of the control flow graph. Specifically: it depends on the set of variables accessible between their position and nodes at m . But these considerations are completely unnecessary, since I can easily read off the only kind of equivalence that matters to me *directly* from the structure of $G_{\alpha, m, \alpha^{\circledast}}$. All that matters is whether “up to decisions made at conditional nodes” (which are always *independent* of micro-architectural state), two nodes must reach an equivalent micro-architectural state at m . But this information is readily available in the postdominance relation $\sqsupseteq_{\text{SINK}}^{G_{\alpha, m, \alpha^{\circledast}}}$. I do *not* need to further inspect the sets Σ at nodes $(x, \Sigma) \in N_{\alpha, m, \alpha^{\circledast}}$.

Definition 13.3.2. For a fixed control flow graph node $m \in N$, write $G_{\#} = (N_{\#}, E_{\#})$ for $G_{\alpha, m, \alpha^{\circledast}}$. Let $M = \{(m, \Sigma) \mid (m, \Sigma) \in N_{\#}\}$ be the set of nodes in $G_{\#}$ that represent a configuration at control flow graph node m . Also note that whenever

$$(n, \Sigma_n) \xrightarrow{I}_{G_{\#}} (x, \Sigma_x)$$

then also

$$(n, \Sigma'_n) \xrightarrow{l}_{G_{\#}} (x, \Sigma'_x)$$

for some Σ'_x . I.e.: for any $G_{\#}$ -node (n, Σ'_n) with the same G -node n , but (possibly) different set Σ'_n of abstract micro-architectural state, any edge labeled l to some node with the G -node x is matched by some edge also labeled l towards a node with the same G -node x .

Write E both for the following rule system, and the corresponding monotone functional.

$$\frac{}{(n, \Sigma_n) \equiv (n, \Sigma_n)} \text{Base}$$

$$\frac{(m, \Sigma) \sqsupseteq_{\text{SINK}}^{G_{\#}} (n, \Sigma_n) \quad (m, \Sigma) \sqsupseteq_{\text{SINK}}^{G_{\#}} (n, \Sigma'_n)}{(n, \Sigma_n) \equiv (n, \Sigma'_n)} \sqsupseteq_{\text{SINK}}$$

$$\frac{\forall l. \left((n, \Sigma_n) \xrightarrow{l}_{G_{\#}} (x, \Sigma_x) \wedge (n, \Sigma'_n) \xrightarrow{l}_{G_{\#}} (x, \Sigma'_x) \Rightarrow (x, \Sigma_x) \equiv (x, \Sigma'_x) \right) \quad n \neq m}{(n, \Sigma_n) \equiv (n, \Sigma'_n)} \text{Suc}$$

Then define

$$G_{\alpha, m, \alpha}^{\equiv}$$

to be the graph obtained from $G_{\#} = G_{\alpha, m, \alpha}^{\circ}$ by merging all nodes which *cannot be distinguished* by the rule system E, i.e.: by merging all nodes equivalent w.r.t \equiv , where

$$\equiv = \nu E$$

is the greatest fixed point of E. Here, the greatest fixed point is to be taken with respect to the subset order \subseteq , and

$$\top = \{ ((n, \Sigma_n), (n, \Sigma'_n)) \mid (n, \Sigma_n) \in N_{\#}, (n, \Sigma'_n) \in N_{\#} \}$$

Rule Base states that every $G_{\#}$ -node (n, Σ_n) is equivalent to itself. Rule \exists_{SINK} states that two $G_{\#}$ -nodes (n, Σ_n) and (n, Σ'_n) with a common control flow graph node n are equivalent if they are postdominated by the same $G_{\#}$ -node (m, Σ_m) , with m being the fixed G -node for which $G_{\#} = G_{\alpha, m, \alpha^{\boxtimes}}$. The last rule Suc states that two $G_{\#}$ -nodes (n, Σ_n) and (n, Σ'_n) with a common control flow graph node n are equivalent if for each label l , all l -successors with a common control flow graph node x are equivalent.

Now finally, I can define micro-architectural dependencies.

Definition 13.3.3. Fix some deterministic control flow graph $G = (N, E)$ and some node $n^0 \in N$ in G as well as some (concrete) micro-architectural state $\sigma_{\mu}^0 \in \Sigma_{\mu}$.

Let $m \in N$ be some node in G , and let $G_{\alpha, m, \alpha^{\boxtimes}}^{\equiv}$ be as in Definition 13.3.2. Then I define for any node $n \in N$:⁶

$$n \rightarrow_{\mu \text{d}}^G m \iff \exists (n, \Sigma_n), (m, \Sigma) \in N_{\alpha, m, \alpha^{\boxtimes}}^{\equiv} \cdot (n, \Sigma_n) \rightarrow_{\text{ntind}}^{G_{\alpha, m, \alpha^{\boxtimes}}^{\equiv}} (m, \Sigma)$$

For each $m \in N$, computation of $\cdot \rightarrow_{\mu \text{d}}^G m$ is efficient. Specifically, since the definition is of the form

$$\left(\rightarrow_{\text{nticd}}^{(G_{\#}^{\equiv})_{M \not\rightarrow}} \right)^* (M)$$

Lemma 7.1.1 from page 135 applies, and I need to compute neither $\rightarrow_{\text{ntind}}$ nor $\rightarrow_{\text{nticd}}$.

Also in my experiments, computation of $G_{\alpha, m, \alpha^{\boxtimes}}^{\equiv}$ from $G_{\alpha, m, \alpha^{\boxtimes}}$ using a chaotic fixed point iteration in topological order is efficient, even when using a naive explicit representation of the equivalence classes in $G_{\alpha, m, \alpha^{\boxtimes}} / \equiv$.

⁶ where again, only Σ_n, Σ , but not m, n are meant to be bound by \exists

Of course in practice, computation of $\cdot \rightarrow_{\mu d}^G m$ is only feasible if the computation of G_α is, from which then $G_{\alpha, m, \alpha^\circ}$ is easily obtained for each $m \in M$.

13.4 Limitations of Micro-Architectural Dependencies

The requirement of α respecting timing for all possible executions (Definition 13.2.5) is quite strong. Obeying it will lead to huge graphs G_α in some applications. In the micro-architectural cache architecture, for some programs the computation is feasible only for relatively small cache sizes.

By Definition 13.2.2, I assume that the execution time along an edge labeled l depends on the micro-architectural state only. In many CPU architectures, however, there exist instructions whose execution time will depend on the actual *value* of some operand, which is intuitively part of the macro-architectural. For example, the execution of trigonometric operations (\sin, \cos, \dots) and integer division div usually depends on the actual value of the operand. To what extent is this a limitation of my approach? I give three answers:

1. In principle, this is not a limitation at all, since the concrete micro-architectural state σ_μ can always be modeled to include the whole macro-architectural state.
2. However in practice — respecting timing for all possible executions — this *does* appear to be a severe limitation, since I will then be forced to track the actual values of variables in the *abstract* micro-architectural state $\sigma_\mu^\#$.
3. Then again in practice, I can just ignore these particular kind of execution time dependency *in the micro-architectural dependency graph* $\rightarrow_{\mu d}^G$, and defer the obligation to account for these kind of dependencies to the *timing sensitive control dependence* $\rightarrow_{\text{tscd}}$, and data dependence $\rightarrow_{\text{data}}$. I can do this by a trivial modification of the underlying control flow graph G . Consider an edge $n \xrightarrow{r_0 := x \text{ div } y}_G m$, with m being the only G -successor of n . In-

producing two fresh nodes m_l, m_r , I can just replace this edge by four new edges

$$\begin{aligned} n &\xrightarrow{r_0 := x \text{ div } y}_{G'} m_l \rightarrow_{G'} m \\ \text{and } n &\xrightarrow{r_0 := x \text{ div } y}_{G'} m_r \rightarrow_{G'} m \end{aligned}$$

obtaining a new control flow graph G' , and assigning two different costs to these edges in the underlying timing cost model C' for G' . Then m will be timing sensitively control dependent on n in G' under C'

$$n \xrightarrow{\text{tscd}}_{G'[C']} m$$

and if n is data-dependent on some other node n' , then n' will be included in the backward slice of m , as required.

In the next section, I explain how a similar construction can be used to integrate the micro-architectural dependencies $n \xrightarrow{\mu d}_G m$ into a full timing sensitive slice, in general.

13.5 Timing Dependence for Micro-Architectural Dependencies

In a labeled control flow graph $G = (N, E)$, a micro-architectural dependency $n \rightarrow_{\mu d}^G m$ indicates that the time some step at m takes to execute may depend on a (macro-architectural) choice made at n . On the other hand, a *timing* sensitive control dependence $n \rightarrow_{\text{tscd}}^G m$ indicates: when *and* if some node m is executed may depend on a choice made at n . But that a step at node m may take different amount of time *at all* is usually not encoded in the structure of the control flow graph, and hence invisible to timing sensitive control dependence $n \rightarrow_{\text{tscd}}^G m$. Recall the control flow graph G from Figure 13.1a, repeated in Figure 13.4a on page 259.

I previously found that

$$\rightarrow_{\mu d}^G = \{(9, 21), (9, 14)\}$$

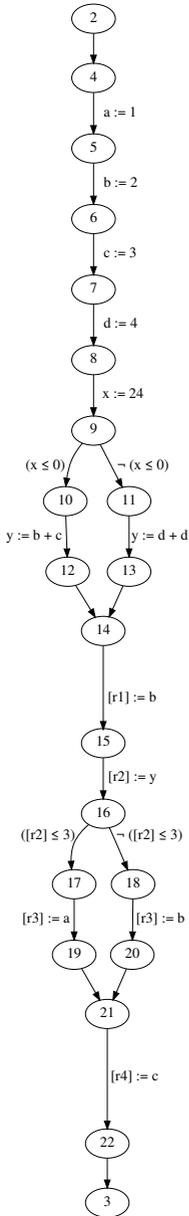
But if I wanted to find out whether the execution of the whole program depended on the choice made at $n = 9$, and did this by naively applying timing sensitive control dependence⁷ to G , I might query the timing-sensitive backward slice of the end node $M = \{3\}$, finding that

$$\left(\rightarrow_{\text{tscd}}^G \cup \rightarrow_{\text{data}}^G \right)^* (M) = \emptyset$$

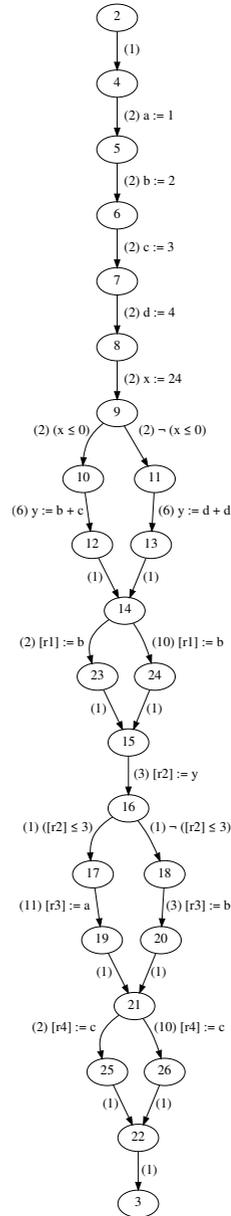
and erroneously conclude that the program has constant execution time (specifically: is independent of choices at $n = 9$).

The solution, of course, is to make different timing behavior due to micro-architectural state visible in the control flow graph, by modifying G at those nodes m at which execution may reach micro-architectural states with different timing behavior. These nodes can

⁷ using, for example, the default timing cost model $\mathbb{1}$, or any timing cost model that assigns some constant timing cost to each variable access



(a) CFG G



(b) Micro-Architecture Aware CFG G'

be identified by inspection of G_α , by gathering for each G -node m all G_α -nodes

$$\left(m, \sigma_\mu^{1\#}\right), \left(m, \sigma_\mu^{2\#}\right), \dots$$

and then checking if for any outgoing labeled edge, there may be different timing in the concrete semantics for any two $\sigma_\mu^i, \sigma_\mu^j$.

At such nodes m , I duplicate edges to fresh artificial nodes immediately behind m and its former successors. I call the resulting graph G' the micro-architecture aware control flow graph for G .

Consider Figure 13.4b at page 259, which shows the result G' of this transformation, together with an explicit timing cost model C' . There, a cache-miss is assumed to take 10 units of time, while a cache-hit takes 2 units⁸. 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.⁹ 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. On the other hand, node 15 still has only one successor, reached with timing cost $3 = 2 + 1$ (cache access plus register access), since I found that there the variable y is always in cache (by inspection of Figure 13.1b).

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

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

Note that even if G is deterministic, G' usually is not. This is no problem, because I can still use the micro-architectural dependencies

⁸ memory writes are assumed to always take 2 units of times, and register accesses take 1 unit of time

⁹ In 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. I need to do this because in my notion of graphs, there can be no multi-edges, and since I required cost models C to be *strictly* positive.

$\rightarrow_{\mu d}^G$ (and data dependence $\rightarrow_{\text{data}}^G$) from the original graph G , and only use G' for timing sensitive control dependence $\rightarrow_{\text{tscd}}^{G'[C']}$.

Observation 13.5.1 (Soundness of Micro-Architectural Dependence for a Data Cache Micro-Architecture). Assume the micro-architecture consisting of the fully associative data cache of size $c \in \mathbb{N}$ with least-recently used eviction strategy, as used in the examples of this chapter.¹⁰ As in the examples, choose the abstractions

$$\alpha = \alpha_{\text{cache}} \quad \text{and} \quad \alpha^{\circledast} = \alpha_{\text{incache}}^{\text{use}(m)}$$

Let $G = (N, E)$ be any deterministic labeled control flow graph, $n^0 \in N$, and σ_{μ}^0 be the initially empty cache. Also let σ_M^0 be the initially empty, partial mapping from variables in G to values, and let $t^0 = 0$ be the starting time.

Assume – starting at n^0 – a linear sequence of “initialization” nodes n_y with outgoing edges labeled $y:=?$ for each variable y in G . Let $G' = (N', E')$ be the micro-architecture aware control flow graph for G , with associated timing cost model C' . Let $m \in N'$ be any non-initialization node, and let

$$S = \left(\rightarrow_{\text{tscd}}^{G'[C']} \cup \rightarrow_{\text{data}}^{G'} \cup \rightarrow_{\mu d}^G \right)^* (\{m\})$$

be the micro-architectural, timing sensitive backwards slice of m .

Assume two *inputs* i_1, i_2 to G , i.e.: two mappings from variables y to some initial value $i_1(y)$ or $i_2(y)$, respectively. Also assume i_1, i_2 to coincide on all variables y for which the initialization node n_y is in the slice S .

Obtain G_1 from G by replacing all labels $y:=?$ at initialization nodes n_y by $y:=i_1(y)$, and obtain G_2 similarly.

¹⁰ for the purpose of gathering experimental evidence, I chose $c \in \{4, 8\}$.

Then the timed, micro-architectural execution sequences for G_1 and G_2 (as defined by the full small-steps semantics in Definition 13.2.2) starting in (n^0, σ_M^0, t^0) coincide, after each removing all configurations (n, σ_M, t) for which $n \notin S$.

13.6 Arrays

In the preceding sections, I (implicitly) made two simplifying assumptions: every program variable resides in a distinct *memory block*, and the memory block of each variable access is statically known. This allowed me to represent abstract caches by lists of variables. The first assumption can easily be done away with: I merely need to respect static, non-injective mapping from variables to memory blocks (such that different variables can share a memory block), and then represent abstract caches by lists of *memory blocks*. Assuming a set $\mathbb{B} = \{b_0, b_1, \dots\}$ of memory blocks, an abstract cache $\sigma_\mu^\#$ is then just a list of memory blocks. For example: $\sigma_\mu^\# = [b_5, b_1, b_7, b_8]$.

The simplest setting in which the second assumption no longer holds is one in which programs contain *array* accesses such as $l = r := a[i]$ with statically unknown indices. Since arrays can span *multiple* memory blocks, the abstract micro-architectural transition from an abstract cache $\sigma_\mu^\#$ along a label l is no longer deterministic. Consider, for example, the abstract cache $\sigma_\mu^\# = [b_5, b_1, b_7, b_8]$ and assume that each memory block is a size of 64 bytes. Assume that the byte-array a is aligned with memory block b_1 , and assume that i is statically known to be in the range $0 \dots 255$. Then the control flow graph edge $n \xrightarrow{G} m$ will induce *four* abstract micro-architectural transitions:

$$\begin{aligned} \left(n, \sigma_\mu^\# \right) &\xrightarrow{l} \left(m, [b_1, b_5, b_7, b_8] \right) && \text{for } i \in 0 \dots 63 && \text{in } \sigma_M \\ \left(n, \sigma_\mu^\# \right) &\xrightarrow{l} \left(m, [b_2, b_5, b_1, b_7] \right) && \text{for } i \in 64 \dots 127 && \text{in } \sigma_M \\ \left(n, \sigma_\mu^\# \right) &\xrightarrow{l} \left(m, [b_3, b_5, b_1, b_7] \right) && \text{for } i \in 128 \dots 191 && \text{in } \sigma_M \\ \left(n, \sigma_\mu^\# \right) &\xrightarrow{l} \left(m, [b_4, b_5, b_1, b_7] \right) && \text{for } i \in 192 \dots 255 && \text{in } \sigma_M \end{aligned}$$

The framework for micro-architectural dependencies from still applies, but *as is* introduces unnecessary imprecision. The reason is that

rule Suc on page 253 for the equivalence relation \equiv is too strict, which results in the graph

$$G_{\alpha, m, \alpha}^{\equiv}$$

having a postdominance relation of too small a size.

Recall that rule Suc requires for two nodes (n, Σ_n) and (n, Σ'_n) in the graph $G_{\alpha, m, \alpha}^{\equiv}$ that for all labels l , every l -transition from (n, Σ_n) is matched by a l -transition from (n, Σ'_n) . In the example, this also requires l -transitions corresponding to $i \in 0 \dots 127$ to be matched by l -transitions corresponding to, e.g., $i \in 192 \dots 255$.

But it is enough to only require l -transitions corresponding to $i \in 0 \dots 127$ to be matched by l -transitions corresponding to $i \in 0 \dots 127$, and l -transitions corresponding to $i \in 64 \dots 127$ to be matched by l -transitions corresponding to $i \in 64 \dots 127$, etc.

In order to do this, I first modify the abstract micro-architectural semantics from Definition 13.2.3 on 240 to include in the transition labels also that part of the macro-architectural state σ_M on which the micro-architectural transition depends. Then rule Suc remains unmodified, except that now the variable l must be read to range over pairs $l_0, \sigma_M^\#$ of control flow graph labels l_0 and abstractions of macro-architectural states.

Definition 13.6.1 (Abstract Micro-Architectural Semantics, Modified). In addition to α as in Definition 13.2.3, overload the name α to also stand for a (abstraction) function

$$\alpha : \Sigma_M \rightarrow \Sigma_M^\#$$

from (concrete) macro-architectural states σ_M to some abstraction $\alpha(\sigma_M)$ in some set $\Sigma_M^\#$ of abstract macro-architectural states. I demand that the abstraction is compatible with the micro-architectural label transition function l^μ , i.e. that: for any σ_M^1, σ_M^2 and $\sigma_\mu^1, \sigma_\mu^2$ with $\alpha(\sigma_M^1) = \alpha(\sigma_M^2)$ and $\alpha(\sigma_\mu^1) = \alpha(\sigma_\mu^2)$ I have

$$\alpha\left(l^\mu\left(\sigma_M^1, \sigma_\mu^1\right)\right) = \alpha\left(l^\mu\left(\sigma_M^2, \sigma_\mu^2\right)\right)$$

Then the modified abstract micro-architectural small-steps semantics of a control flow graph G is just

$$\frac{n \xrightarrow{L}_G m \quad \alpha(\sigma_\mu) = \sigma_\mu^\# \quad \sigma'_\mu = I^\mu(\sigma_M, \sigma_\mu) \quad \alpha(\sigma'_\mu) = \sigma'^{\#\mu}}{\left(n, \sigma_\mu^\#\right) \xrightarrow{L, \alpha(\sigma_M)} \left(m, \sigma'^{\#\mu}\right)} \text{Label}$$

In the example, the relevant abstract macro-architectural states are $\{i \in 0 \dots 63'', \dots\} \subseteq \Sigma_M^\#$, and I have

$$\begin{aligned} \left(n, \sigma_\mu^\#\right) &\xrightarrow{L, "i \in 0 \dots 63''} \left(m, [b_1, b_5, b_7, b_8]\right) \\ \left(n, \sigma_\mu^\#\right) &\xrightarrow{L, "i \in 64 \dots 127''} \left(m, [b_2, b_5, b_1, b_7]\right) \\ \left(n, \sigma_\mu^\#\right) &\xrightarrow{L, "i \in 128 \dots 191''} \left(m, [b_3, b_5, b_1, b_7]\right) \\ \left(n, \sigma_\mu^\#\right) &\xrightarrow{L, "i \in 192 \dots 255''} \left(m, [b_4, b_5, b_1, b_7]\right) \end{aligned}$$

Remark 13.6.1. In order to limit the size of resulting transition relation, in practice it is necessary to establish tight static bounds on index variables such as i in $r := a[i]$ whenever possible.

Summary

- Micro-architectural dependencies $\rightarrow_{\mu d}$ expose timing channels that arise due to the micro-architecture of the executing CPU.
- They can be computed using nontermination insensitive slices, and an additional greatest fixed point computation that is driven by nontermination insensitive postdominance \exists_{SINK} .

14 Cache Timing Attacks on AES256

Dazzle!

(Dazzle — Dota 2)

Cache timing attacks on implementations of cryptographic primitives allow an attacker to recover cryptographic keys and/or plain text messages by observing only the execution time of encryption or decryption operations. Even if (during an computation of the cryptographic operation) the flow of control is independent from keys and plain text¹, the execution time may depend on these due to the effects of data caches. Differences in execution time can then be used to infer keys (e.g., [Ber05; BM06]).

In this chapter, I employ micro-architectural dependencies to analyze cache-based timing channels in implementations of the AES256 block cipher. As expected, my analysis discovers timing channels in naive *substitution table* based implementations. In two more sophisticated implementations, my analysis can proofs the absence of cache based timing channels.

¹ i.e., even if these do not affect decisions at branch- and loop-predicates

14.1 AES256 Encryption

A crucial operation during AES256 encryption is the “S-Box” substitution step, which substitutes values in the current computation state by their value in a constant, publicly known substitution table. The most natural *default* implementation is by simple array lookup, as shown in Figure 14.6a. Here, `state` is an array with 16 entries that holds the current encryption state, i.e.: the encrypted message. The array `sbox` is the constant, publicly known substitution table with 256 entries. In this implementation, the execution time is affected by the current state (and hence: by plain text and key), because the value `r2` affects which part of the substitution table `sbox` is accessed. Since in common CPU micro-architectures the array `sbox` will span multiple memory blocks², the value `r2` in one iteration may influence whether the read of `r3` in a later iteration is served from the data cache or from the main memory, and hence the execution time. In the micro-architecture aware control flow graph, the exit node (i.e.: the total execution time) is timing sensitively control dependent on the node corresponding to line 4, which in turn is data dependent on the key and the plain text.

Additional timing dependencies, also due to micro-architectural dependencies, exists. For example, there exists a dependency chain from the exit node to the plain text input node via a micro-architectural dependency from the node corresponding to line 4 in Figure 14.6a to the node corresponding to line 5 in Figure 14.7a).

Figure 14.1a) shows the results of cache- and timing sensitive slicing w.r.t the exit node of the default AES256 implementation, for different assumed cache sizes. They were computed in my prototype implementation of micro-architectural dependencies in the Haskell programming language. The first column indicates the assumed cache size, in number of cache-lines. For simplicity, each scalar program variable was assumed to occupy a distinct memory block, while arrays (all of size 256, with 8-bit entries) were assumed to span exactly four

² i.e., regions of memory that are associated with the same cache line

memory blocks. Cache lines are assumed to fit exactly one memory block, effectively resulting in a cache line size of 64 bytes. The second column shows the size of the abstract micro-architectural transition graph G_α , measured in the number of nodes N_α . The third column shows the size of the micro-architectural dependency relation. The fourth column shows the size of the backward-slice (as described in Observation 13.5.1 on page 261) of the control flow graphs exit node. The next two columns show execution time of the analysis, and the required amount of memory. The last column indicates whether the analysis could prove the absence of timing channels, i.e.: whether the slice contains *no* key or plain-text input node.

The given computation time includes the whole analysis: computation of G_α and the micro-architectural aware control flow graph G' , as well as data-dependencies, timing-sensitive control dependencies, and micro-architectural dependencies, and slicing. The computation time is dominated by the computation of the graphs

$$G_{\alpha, m, \alpha}^{\equiv}$$

required for micro-architectural dependencies. All times in this chapter were measured on a high end “computation server” class PC with an Intel Xeon Gold 6230 CPU at 2.10 GHz base frequency with 512GB RAM.

For cache sizes of 12 cache lines or more, the analysis did not finish for this AES256 implementation. For 12 cache lines, computation exhausted the maximal memory (500GB). I expect that an analysis implementation more memory-efficient than my prototype Haskell implementation is possible, but since the size of the graph G_α appears to grow exponentially in the number of cache lines, I must expect this to extend the range of testable cache sizes *for this input program* only by a little.³

³ But on the other hand, the phase of “exponential” growth of G_α may end within reach of a more efficient analysis-implementation. Just peek ahead to Figure 14.1b). There, the same Haskell analysis is run for a different input program (concretely: a different

Together, Figure 14.6, 14.7, 14.8, 14.9 , 14.10 and 14.11 form an implementation of AES256 encryption for one block of plain text data. Where applicable, the left subfigure shows the “naive” default implementation with an explicit “S-Box” substitution table `sbox`.

There are two common approaches to avoid cache timing channels in AES256 implementations. The first approach is to employ “pre-caching” of relevant cache lines. My Implementation of this approach is shown in the right hand sides of the aforementioned figures, and explained in Section 14.2. The second approach is to avoid table lookup depending on secret values altogether, by encoding the “S-Box” substitution in a constant time Boolean program (Figure 14.2 and Section 14.3).

AES256 implementation). But for that program, the graphs G_α are much smaller, *and* the growth appears to slow at 16 cache lines, and then stop at 32.

| # c.l. | $ N_\alpha $ | $ \mu d $ | $ S $ | time (s) | memory (kb) | passed |
|--------|--------------|-----------|-------|----------------|-------------|--------|
| 2 | 1132 | 8 | 695 | 1 | 30 968 | ✗ |
| 4 | 8349 | 39 | 695 | 3 | 35 040 | ✗ |
| 6 | 89378 | 132 | 695 | 47 | 268 488 | ✗ |
| 8 | 529225 | 324 | 695 | 396 | 2 097 396 | ✗ |
| 10 | 2627483 | 463 | 695 | 20955 | 62 865 540 | ✗ |
| 12 | 18444663 | | | did not finish | | |

(a) Default AES265 Implementation

| # c.l. | $ N_\alpha $ | $ \mu d $ | $ S $ | time (s) | memory (kb) | passed |
|--------|--------------|-----------|-------|----------|-------------|--------|
| 2 | 1189 | 21 | 703 | 1 | 39 128 | ✗ |
| 4 | 2901 | 66 | 717 | 3 | 36 152 | ✗ |
| 6 | 6320 | 85 | 427 | 4 | 36 952 | ✗ |
| 8 | 14570 | 50 | 62 | 5 | 65 780 | ✓ |
| 10 | 21136 | 44 | 62 | 7 | 84 184 | ✓ |
| 12 | 47593 | 96 | 72 | 30 | 229 676 | ✓ |
| 14 | 110128 | 67 | 74 | 103 | 981 228 | ✓ |
| 16 | 146792 | 101 | 58 | 166 | 1 814 764 | ✓ |
| 18 | 149650 | 45 | 54 | 134 | 2 077 944 | ✓ |
| 20 | 149696 | 18 | 44 | 90 | 1 919 220 | ✓ |
| 22 | 149712 | 18 | 40 | 60 | 1 627 332 | ✓ |
| 24 | 149718 | 18 | 40 | 60 | 1 921 264 | ✓ |
| 26 | 149724 | 18 | 40 | 60 | 1 646 780 | ✓ |
| 28 | 149736 | 18 | 40 | 60 | 2 101 312 | ✓ |
| 30 | 149748 | 18 | 40 | 60 | 1 688 752 | ✓ |
| 32 | 151009 | 18 | 40 | 60 | 1 961 124 | ✓ |
| | | | | ... | | |
| 100 | 151009 | 18 | 40 | 60 | 1 881 268 | ✓ |
| 200 | 151009 | 18 | 40 | 60 | 1 511 584 | ✓ |
| 500 | 151009 | 18 | 40 | 60 | 1 795 240 | ✓ |

(b) With *pre-caching*

Figure 14.1: Cache- and Timing Dependencies

14.2 Pre-Caching

For architectures with big data caches, pre-caching is easy: one simply pre-caches *all* memory blocks associated with program data at the start of the program. For very small data caches, on the other hand, pre-caching must be employed more strategically, with respect not only to the usage of cache-lines, but also to the actual cache size. I am not aware of any automatic process to do this. But micro-architectural and timing can automatically verify the correctness of manual pre-caching for a given cache size.

In Figure 14.6 and 14.7, I surrounded access to the `sbox` and the `skey` arrays⁴ by loads that pre-cache the corresponding array, and then after the access establish a fixed position of all relevant memory blocks in the assumed LRU cache order. In Figure 14.10, the additional loads establish fixed cache-positions for the state memory blocks after the key-expansion. In Figure 14.11, the first additional loads force variables w_i near to the front of the LRU cache — which they are anyway, if the `if` branch is not taken. Similarly, the next two additional loads “emulate” the memory access taken in the first branch of the second `if` statement. The last additional load pre-caches access to `skey`.

As can be seen from Figure 14.1b), my analysis can automatically proof the absence of timing channels due to data caches, for micro-architecture with 8 or more lines of data cache. For this program, my analysis takes up to 166 seconds, and up to \approx 2GB of memory. The reason that the analysis time decreases at 18 cache lines is that my implementation can soundly skip the computation of the set $\{ n \mid n \rightarrow_{\mu d} m \}$ for those nodes m at which memory accesses always take the same amount of time (Observation 14.2.1), as per the observation immediately below. But for larger caches, this can be true for more nodes m .

⁴ the latter of which is holding the *expanded* round key

Observation 14.2.1. Let m be any node in the control flow graph G . If for all two nodes $(m, \Sigma), (m, \Sigma')$ in the graph⁵ G_α at control node m the rule

$$\frac{\sigma_\mu \in \Sigma \quad \sigma'_\mu \in \Sigma' \quad \begin{array}{l} m \xrightarrow{l}_G m' \\ (m, \sigma_M, \sigma_\mu, 0) \xrightarrow{l} (m, \sigma_M, \sigma'_\mu, \Delta t) \\ (m, \sigma_M, \sigma'_\mu, 0) \xrightarrow{l} (m, \sigma_M, \sigma'_\mu, \Delta t') \end{array}}{\Delta t = \Delta t'}$$

is admissible, then

$$\{n \mid n \rightarrow_{\mu d} m\} = \emptyset$$

Micro-Architectural dependencies do not only confirm that the employed pre-caching is sufficient, but the also helped me to conclude where it was necessary: I determined all additional loads in the pre-caching AES256 by manual inspection of the graph G_α as well as micro-architectural dependencies $\rightarrow_{\mu d}$, starting from the default implementation and then iterating.

⁵ Definition 13.3.1 on page 245

```

1  x0 := (state[0] & 128 ^ state[1] & 128 >> 1) ^ state[2] & ...
2  x1 := (state[0] & 64 << 1 ^ state[1] & 64) ^ state[2] & ...
3  ...
4  y14 := x3 ^ x5
5  y13 := x0 ^ x6
6  ...
7  t2 := y12 & y15
8  t3 := y3 & y6
9  ...
10 z0 := t44 & y15
11 z1 := t37 & y6
12 ...
13 t46 := z15 ^ z16
14 t47 := z10 ^ z11
15 ...
16 s0 := t59 ^ t63
17 s6 := t56 ^ ~t62
18 state[0] := ((s0 & 128 ^ s1 & 128 >> 1) ^ s2 & 128 >> 2) ^ ...
19 state[1] := ((s0 & 64 << 1 ^ s1 & 64) ^ s2 & 64 >> 1) ^ ...
20 ...

```

Figure 14.2: Constant Time: SUB^{ct}

14.3 Constant Time S-Box Substitution

In [BP10], the authors present a Boolean function that implements “S-Box” substitution purely by Boolean operations on the (bit-representation) of input. Figure 14.2 gives an idea. As expected, my analysis proves the absence of timing channels if in the default implementation, the array based “S-Box” substitution is replaced by that of Figure 14.2. As can be seen from Figure 14.3 (with the same columns as before in Figure 14.1), this holds for *all* assumed cache sizes.

| # c.l. | $ N_\alpha $ | $ \mu d $ | $ S $ | time (s) | memory (kb) | passed |
|--------|--------------|-----------|-------|----------|-------------|--------|
| 2 | 3717 | 7 | 51 | 3 | 220 344 | ✓ |
| 4 | 7062 | 27 | 60 | 5 | 285 832 | ✓ |
| 6 | 18946 | 102 | 67 | 15 | 221 372 | ✓ |
| 8 | 61399 | 156 | 64 | 57 | 275 576 | ✓ |
| 10 | 195541 | 155 | 64 | 204 | 785 580 | ✓ |
| 12 | 322495 | 146 | 63 | 304 | 1 321 208 | ✓ |
| 14 | 419209 | 141 | 62 | 371 | 1 640 572 | ✓ |
| 16 | 509983 | 137 | 60 | 431 | 1 957 040 | ✓ |
| 18 | 580357 | 132 | 59 | 491 | 2 115 768 | ✓ |
| 20 | 613723 | 125 | 58 | 492 | 2 178 216 | ✓ |
| 22 | 647089 | 125 | 58 | 531 | 2 378 896 | ✓ |
| 24 | 680455 | 125 | 58 | 542 | 2 521 324 | ✓ |
| 26 | 713821 | 125 | 58 | 605 | 2 334 844 | ✓ |
| 28 | 747187 | 125 | 58 | 598 | 2 816 188 | ✓ |
| 30 | 780553 | 125 | 58 | 685 | 2 819 240 | ✓ |

(c) Constant Time S-Box

Figure 14.3: Cache- and Timing Dependencies

14.4 Validation

By Figure 14.1, I know that for the *pre-caching* implementation, execution time is independent from plain text and key for the assumed cache sizes starting at 8 cache lines. Since in AES256 encryption, key and plain-text are the programs *only* inputs, this implementation then is indeed a *constant time* implementation for 8 cache lines. Using an interpreter for the concrete micro-architectural, I validated these results for randomly chosen key and plain-text inputs. For such inputs, encryption took 46578 units of time in the assumed timing model.

For the implementation with constant time “S-Box” substitution (Figure 14.3), this holds even for 2 cache lines. There, encryption takes a constant 211189 units of time.

But to which extend is my analysis precise? By Figure 14.1, it cannot proof the default “naive” implementation constant-time for any cache size, and indeed it is not, as can be seen from Figure 14.4. It shows the execution time histograms for 2 to 12 cachelines, for 1000 random key and plain-text inputs. For the pre-caching implementation, my analysis cannot proof constant-time for 2 to 6 cache lines, but only for 8 and more cache lines. And indeed for cache-sizes 2 to 6, the pre-caching implementation is not constant time, as can be seen from Figure 14.5.

Summary

- Micro-architectural dependencies $\rightarrow_{\mu d}$ can proof the absence of timing channels in AES256 implementations, under the assumption of a simple data cache micro-architectures.
- For one implementation, micro-architectural dependencies can only be computed for small assumed cache sizes.

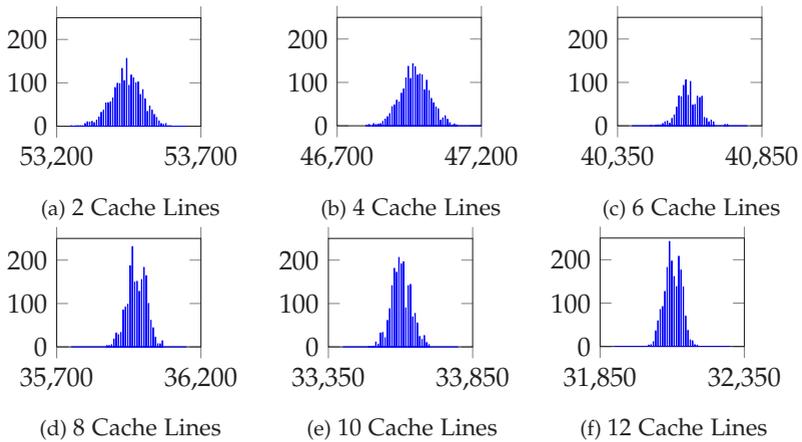
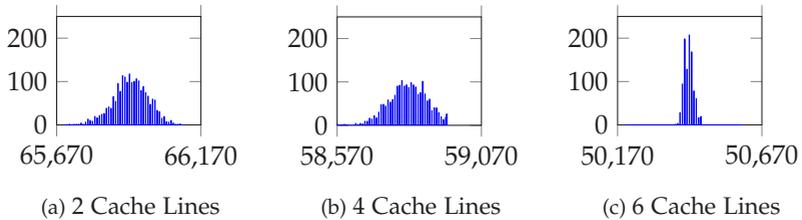


Figure 14.4: Execution Time: Default Implementation

Figure 14.5: Execution Time: With *pre-caching*

| | |
|---|---|
| <pre> 1 2 for r1 : [0, 1 .. 15] 3 r2 := state[r1] 4 r3 := sbox[r2] 5 state[r1] := r3 6 end 7 </pre> | <pre> load sbox[0,64,128,192] for r1 : [0, 1 .. 15] r2 := state[r1] r3 := sbox[r2] state[r1] := r3 end load sbox[0,64,128,192] </pre> |
|---|---|

(a) Default: SUB^{def}(b) With *pre-caching*: SUB^{Pc}

Figure 14.6: AES256: S-Box substitution

| | |
|--|--|
| <pre> 1 2 for r1 : [0, 1 .. 15] 3 r2 := state[r1] 4 r4 := (r1 + 0) 5 r3 := skey[r4] 6 r2 := (r2 ^ r3) 7 state[r1] := r2 8 end 9 </pre> | <pre> load skey[0,64,128,192] for r1 : [0, 1 .. 15] r2 := state[r1] r4 := (r1 + 0) r3 := skey[r4] r2 := (r2 ^ r3) state[r1] := r2 end load skey[0,64,128,192] </pre> |
|--|--|

(a) Default: $\text{ADDRROUND}_0^{\text{def}}$ (b) With *pre-caching*: $\text{ADDRROUND}_0^{\text{pc}}$

Figure 14.7: AES256: Adding the Round Key

```

1 shiftRowsTmp := state[1]
2 state[1] := state[5]
3 state[5] := state[9]
4 state[9] := state[13]
5 state[13] := shiftRowsTmp
6 shiftRowsTmp := state[2]
7 state[2] := state[10]
8 state[10] := shiftRowsTmp
9 shiftRowsTmp := state[6]
10 state[6] := state[14]
11 state[14] := shiftRowsTmp
12 shiftRowsTmp := state[15]
13 state[15] := state[11]
14 state[11] := state[7]
15 state[7] := state[3]
16 state[3] := shiftRowsTmp

```

(a) SHIFT

Figure 14.8: AES256: Shifting Rows

```

1 MIX0
2 MIX4
3 MIX8
4 MIX12

```

(a) MIX

```

1 a0 := state[o + 0]
2 a1 := state[o + 1]
3 a2 := state[o + 2]
4 a3 := state[o + 3]
5 b0 := (a0 << 1) ^ (27 & ((a0 >> 7) * 255))
6 b1 := (a1 << 1) ^ (27 & ((a1 >> 7) * 255))
7 b2 := (a2 << 1) ^ (27 & ((a2 >> 7) * 255))
8 b3 := (a3 << 1) ^ (27 & ((a3 >> 7) * 255))
9 r0 := (((b0 ^ a1) ^ b1) ^ a2) ^ a3
10 r1 := (((a0 ^ b1) ^ a2) ^ b2) ^ a3
11 r2 := (((a0 ^ a1) ^ b2) ^ a3) ^ b3
12 r3 := (((a0 ^ b0) ^ a1) ^ a2) ^ b3
13 state[o + 0] := r0
14 state[o + 1] := r1
15 state[o + 2] := r2
16 state[o + 3] := r3

```

(b) MIX₀

Figure 14.9: AES256: Mixing Columns

```

1 EXPAND
2
3 ADDROUND0
4 for r5 : [1, 2 .. 13]
5   SUB
6   SHIFT
7   MIX
8   ADDROUNDr5 << 4
9 end
10 SUB
11 SHIFT
12 ADDROUND14 << 4

```

(a) Default

```

1 EXPAND
2 load state[0,64,128,192]
3 ADDROUND0
4 for r5 : [1, 2 .. 13]
5   SUB
6   SHIFT
7   MIX
8   ADDROUNDr5 << 4
9 end
10 SUB
11 SHIFT
12 ADDROUND14 << 4

```

(b) With *pre-caching*

Figure 14.10: AES256: Main Loop

```

1  n := 1
2  for i : [0, 1 .. 31]
3    skey[i] := key[i]
4  end
5  for o : [32, 36 .. 236]
6    w0 := skey[o-4]
7    w1 := skey[o-3]
8    w2 := skey[o-2]
9    w3 := skey[o-1]
10   if (o % 32 == 0) then
11     rotateTmp := w0
12     w0 := w1
13     w1 := w2
14     w2 := w3
15     w3 := rotateTmp
16     SUB4
17     w0 := w0 ^ rcon[n]
18
19     n := n + 1
20   end
21   if (o % 32 == 16) then
22     SUB4
23
24   end
25
26   end
27
28   skey[o+0] := skey[o-32]^w0
29   skey[o+1] := skey[o-31]^w1
30   skey[o+2] := skey[o-30]^w2
31   skey[o+3] := skey[o-29]^w3
32 end

```

(a) Default: EXPAND^{def}

```

n := 1
for i : [0, 1 .. 31]
  skey[i] := key[i]
end
for o : [32, 36 .. 236]
  w0 := skey[o-4]
  w1 := skey[o-3]
  w2 := skey[o-2]
  w3 := skey[o-1]
  if (o % 32 == 0) then
    rotateTmp := w0
    w0 := w1
    w1 := w2
    w2 := w3
    w3 := rotateTmp
    SUB4
    w0 := w0 ^ rcon[n]
    load w1, w2, w3
    n := n + 1
  end
  if (o % 32 == 16) then
    SUB4
  else
    load w0, w1, w2, w3
    load sbox[0,64,128,192]
  end
  load skey[0,64,128,192]
  skey[o+0] := skey[o-32]^w0
  skey[o+1] := skey[o-31]^w1
  skey[o+2] := skey[o-30]^w2
  skey[o+3] := skey[o-29]^w3
end

```

(b) With *pre-caching*: EXPAND^{pc}Figure 14.11: AES256: Expanding the *Session Key*

15 Approximate Cache Dependencies

Das eigentliche Problem ist es ja nur, die richtigen Definitionen zu finden. Alles andere ist danach meist trivial.

(Joachim Cuntz (paraphrasiert) — Vorlesung Analysis III)

In the application of micro-architectural dependencies for simple data caches in Chapter 14, the computation of micro-architectural dependencies was practical for the two “constant time” AES256 implementations. But for the *naive* default AES256 implementation, the analysis was impractical for even moderate cache sizes (Figure 14.1a). The reason was that the size of the intermediate graph G_α appeared to grow exponentially in the assumed cache size.

To alleviate this problem, in this section I develop an approximation to micro-architectural dependencies $\rightarrow_{\mu d}$ specifically for a simple cache architecture with least recently used eviction strategy. Since this approximation is best understood as an analogue of traditional data dependencies $\rightarrow_{\text{data}}$ in labeled control flow graphs, I first review their definition in (Section 15.1). Then in Section 15.2, I introduce *local cache-cache* dependencies, which are those dependencies analog to the dependencies implicit in the def and use sets of standard data dependence. In Section 15.3, I describe dependencies from macro-architectural states to cache state (*local state-cache* dependencies). Then in Section 15.4, I describe the cache state analogue of a standard data dependence slice. Finally in Section 15.5, I show how this can be used to obtain an approximation to micro-architectural dependencies $\rightarrow_{\mu d}$ for cache micro-architectures.

15.1 Data Dependence

Intuitively, a node m in a labeled control flow graph G is directly standard data dependent on a node n via variable x if a definition of variable x at node n may reach a use of variable x at node m . In this case, I write $n \xrightarrow{x}_{\text{data}} m$. Conversely, a use of x at node m is not directly data dependent on a definition of x at node n if on all control flow paths from n to m , this definition of x is “killed” by another definition of x at some other intermediate node.

Data dependence is typically computed “value-insensitively”. If I understand the control flow graph to be an abstraction of all possible executions of the represented program, then each control flow graph node n abstractly represents all concrete configurations with control state n , but otherwise *arbitrary* variable states σ (e.g.: arbitrary mappings $\sigma : V \rightarrow \mathbb{N}$, for some set of program variables V , and a domain \mathbb{N} of variable values). For standard data dependence, no attempts are made to consider only those concrete states σ at n which are possible in actual execution of the program, and no attempt is made to interpret the expressions used to define variables.¹ All that is considered are the corresponding use and def sets at control flow edges.²

Under this abstraction, data dependence can then simply be computed from the use and def sets by a standard (forward) data flow analysis. There implicitly, at each control flow edge labeled with l that is leaving a node n , every variable $x \in \text{def}(l)$ is considered to locally depend on every variable $y \in \text{def}(l)$. This is an approximation, since, for example, in the assignment $x := a * \emptyset$, the variable x does semantically *not* depend on a . Another example is the assignment $x := a * b$, in which the local dependence of variable x on variable a as approximated by def and use sets does not hold semantically in programs for which $\sigma(b) = 0$ for all σ at node n .

¹ For example, an assignment $x := a * b$ is treated no differently than $x := a + b$.

² For example: $\text{use}(x := a * b) = \text{use}(x := a + b) = \{a, b\}$, and $\text{def}(x := a * b) = \{x\}$.

15.2 Local Cache-Cache Dependencies

In order to obtain an approximation to micro-architectural dependencies $\rightarrow_{\mu d}$ for simple data cache micro-architectures, I use a new notion of *local* cache-cache dependencies that can be thought of as an analogue of the local dependencies implicit in def and use sets for standard data dependence.

For reasonably precise results, I need for every control flow graph node n an approximation of all possible cache states at n . Similar to Section 13.6 earlier, a concrete cache state σ_{μ} for a data cache with k cache lines is just a list of memory blocks with up to k entries. The next-to-be-evicted memory appears at the end. For example, $[b_5, b_1, b_7, b_8]$ represents a cache with memory blocks $b_i \in \mathbb{B}$, and block b_8 the next to be evicted.

Abstract Caches

In this chapter, I use a cache abstraction similar to that from [Doy+15]. An abstract cache state $\sigma_{\mu}^{\#}$ is a mapping from memory blocks to a set of possible positions in the concrete cache, i.e.: a mapping

$$\sigma_{\mu}^{\#} : \mathbb{B} \rightarrow 2^{\mathbb{K}}$$

where the set $\mathbb{K} = \{0, \dots, k-1\} \cup \{\infty\}$ is comprised of natural numbers $< k$, and the symbol ∞ which indicates that a corresponding memory block may be not in the cache. I write Σ_{μ} for the set of all such mappings $\sigma_{\mu}^{\#}$. A mapping from control flow graph nodes n to abstract caches $\sigma_{\mu}^{\#}(n)$ such that $\sigma_{\mu}^{\#}(n)$ soundly approximates all concrete caches possible at n is then available by abstract interpretation of control flow graph edges $n \xrightarrow{l}_G m$, and a forward data flow analysis.

In an abstract cache state $\sigma_\mu^\#$, the possible positions of two memory blocks are not necessarily disjoint. For example, the abstract cache $\sigma_\mu^\#$ with

$$\begin{array}{c|c|c} b_0 & b_1 & b_2 \\ \hline \{0, 1\} & \{0, 1\} & \{1, 2, \infty\} \end{array}$$

indicates that memory block b_0 is guaranteed to be in the cache, and may be in the 0th or the 1st position in the concrete cache. The same holds for b_1 , while memory block b_2 is indicated to be either in the 1st or 2nd position in the concrete cache, or not in the cache at all.

The set of all concrete caches represented by this abstract cache then is

$$[b_0, b_1, b_2], [b_0, b_1], [b_1, b_0, b_2], [b_1, b_0]$$

Note that by the constraint that any concrete cache position cannot contain two memory blocks, here the 0th and 1st position really must contain either b_0 or b_2 , but never b_2 .

In order to define notions of dependencies between memory blocks in abstract cache states, I do *not* defer to this (“precise”) notion of concretization. Instead I use an approximate notion I call *pseudo-concretization*. The advantage is that with regard to pseudo-concretizations, these dependencies can be computed efficiently (as later described in Observation 15.2.1 and Observation 15.3.1).

By pseudo-concretization I mean the Cartesian product of the concretization of each individual memory block, i.e., for each abstract cache $\sigma_\mu^\#$ the set of functions

$$\sigma_{\text{pseudo}} : \mathbb{B} \rightarrow \mathbb{K}$$

that are compatible with $\sigma_\mu^\#$:

$$\sigma_{\text{pseudo}}(b) \in \sigma_\mu^\#(b)$$

for every $b \in \mathbb{B}$. For the previous example abstract cache $\sigma_\mu^\#$, this set is (with each row specifying one pseudo-concretization σ_{pseudo}):

| b_0 | b_1 | b_2 |
|-------|---------|----------|
| 0 | 0 | 1 |
| 0 | 0 | 2 |
| 0 | 0 | ∞ |
| 0 | 1 | 1 |
| | \dots | |
| 1 | 0 | 1 |
| | \dots | |
| 1 | 1 | 1 |
| | \dots | |

The LRU eviction strategy can be naturally extended to pseudo-concretizations σ_{pseudo} , by assigning to each control flow graph label l a function l^μ that maps macro-architectural state σ_M and a pseudo-cache state σ_{pseudo} to a successor pseudo-cache state σ'_{pseudo} .

Local Cache-Cache Dependencies

When slicing w.r.t standard data dependence, due to the local dependence of (for example) variable x on variable y in an assignment such as $x := y * 5$, variable x “inherits” all dependencies of variable y . The following definition formalizes a similar idea for caches.

Definition 15.2.1 (Local Cache-Cache Dependence). Let $\sigma_\mu^\#$ be an abstract cache state, and l control flow graph edge label. Let $b \in \mathbb{B}$ be any memory block, and $b' \neq b$ be a memory block that may be accessed by label l . Then I say that b is locally cache-cache dependent

on b' in $\sigma_\mu^\#$ via transition l if there exists pseudo-concretizations σ_{pseudo}^1 and σ_{pseudo}^2 of $\sigma_\mu^\#$ and some macro-architectural state σ_M with

$$\sigma_{\text{pseudo}}^1(b) \neq \sigma_{\text{pseudo}}^1(b') \quad \text{and} \quad \sigma_{\text{pseudo}}^2(b) \neq \sigma_{\text{pseudo}}^2(b')$$

such that

- (a) σ_{pseudo}^1 and σ_{pseudo}^2 coincide on all memory blocks except b' , and
- (b) $\sigma_{\text{pseudo}}^{1'}(b) \neq \sigma_{\text{pseudo}}^{2'}(b)$,

where

$$\begin{aligned} \sigma_{\text{pseudo}}^{1'} &= l^\mu \left(\sigma_M, \sigma_{\text{pseudo}}^1 \right) \\ \sigma_{\text{pseudo}}^{2'} &= l^\mu \left(\sigma_M, \sigma_{\text{pseudo}}^2 \right) \end{aligned}$$

are the pseudo-concrete caches arising from transition along edge l .

In that case, I write

$$b' \xrightarrow{l}_{\text{local}[\sigma_\mu^\#]} b$$

For example, in the abstract cache $\sigma_\mu^\#$

| | | |
|---------|------------|------------|
| b | b' | x |
| $\{1\}$ | $\{0, 2\}$ | $\{0, 2\}$ |

and a CFG edge labeled $l = r := b'$ for a variable b' residing in block b' , I do have $b' \xrightarrow{l}_{\text{local}[\sigma_\mu^\#]} b$. This is because in the (pseudo)-concrete cache with b' in position 0 and block x in position 2, the transition along l does *not* change the cache state (of any block), while in the pseudo-concrete cache with block b' in position 2 and block x in position 0, transition along l *does* change the cache state (specifically: of block b , but also of block x), by putting b' in position 0, and moving block b from position 1 to position 2.

On the other hand, in the abstract cache $\sigma_\mu^\#$

| | | |
|---------|------------|---------|
| b | b' | \dots |
| $\{3\}$ | $\{0, 2\}$ | \dots |

I do *not* have $b' \xrightarrow{l}_{\text{local}[\sigma_\mu^\#]} b$ for this label $l = r := b'$.

I can equip the set \mathbb{K} of abstract cache positions with an ordering \leq by defining $a \leq \infty$ for all $a \in \mathbb{K}$, and use the order \leq on natural numbers otherwise. With this, local cache-cache dependency can be simplified as follows:

Observation 15.2.1. Let l, b, b' and $\sigma_\mu^\#$ as in Definition 15.2.1. Let

$$b'_{\min} = \min_{a \in \sigma_\mu^\#(b')} a \quad b'_{\max} = \max_{a \in \sigma_\mu^\#(b')} a$$

Then

$$b' \xrightarrow{l}_{\text{local}[\sigma_\mu^\#]} b \iff b'_{\min} < a < b'_{\max} \text{ for some } a \in \sigma_\mu^\#(b)$$

15.3 Local State-Cache Dependence

Local cache-cache dependencies from the previous section capture how dependencies from one memory block b' transfer to dependencies on other memory blocks b' . Now I formally define how dependencies of memory blocks initially arise from choices due to *macro-architectural* state.

These dependencies arise whenever at a labeled control flow graph edge $n \xrightarrow{l}_G m$ and an abstract cache state $\sigma_\mu^\#$, the macro-architectural state σ_M “chooses” which of multiple possible memory blocks b' is accessed. The canonical example is an array access $r2 := a[r1]$ at which the value of register $r1$ chooses which one of those several memory blocks over which the array a spans is actually accessed.³ If then the resulting cache location of a memory block b depends on the choice made by the macro-architectural state σ_M , I say that there is a local state-cache dependence to b via l . The formal definition is:

Definition 15.3.1 (Local State-Cache Dependence). Let $\sigma_\mu^\#$ be an abstract cache state, and l a control flow graph edge label. Let $b \in \mathbb{B}$ be any memory block. For each macro-architectural state σ_M , let $b'(\sigma_M)$ be the memory block accessed during transition along l from σ_M . Let B' be the set of all such memory blocks b' .

Then I say that in the abstract cache state $\sigma_\mu^\#$, block b is locally state-cache dependent (on the macro-architectural state) via l if there exists a pseudo-concretization σ_{pseudo} of $\sigma_\mu^\#$ and two macro-architectural states σ_M^1 and σ_M^2 relevant to b and σ_{pseudo} such that

$$\sigma_{\text{pseudo}}^{1'}(b) \neq \sigma_{\text{pseudo}}^{2'}(b)$$

³ For simplicity, I assume that “dynamically” (i.e.: in the concrete semantics), for each label l at most one memory block is accessed. For example, I assume that no assignment like $r3 := a[r1] + b[r2]$ occur, but instead only assignments like $r3 := a[r1]; r3 := r3 + b[r2]$.

where

$$\begin{aligned}\sigma_{\text{pseudo}}^1 &= l^\mu \left(\sigma_{M'}^1, \sigma_{\text{pseudo}} \right) \\ \sigma_{\text{pseudo}}^2 &= l^\mu \left(\sigma_{M'}^2, \sigma_{\text{pseudo}} \right)\end{aligned}$$

are the pseudo-concrete caches arising from transition along edge l

By σ_M^1 and σ_M^2 being relevant to b and σ_{pseudo} I mean that block b must be accessed in the transition along l one of the macro-architectural states $\sigma_{M'}^i$, or

1. $b'(\sigma_M^1) = b$ or $b'(\sigma_M^2) = b$, or
2. for both $i \in \{1, 2\}$

$$\sigma_{\text{pseudo}}(b'(\sigma_M^i)) = \infty \text{ or } \sigma_{\text{pseudo}}(b'(\sigma_M^i)) \neq \sigma_{\text{pseudo}}(b)$$

For local state-cache dependencies involving b , l and $\sigma_\mu^\#$, I write

$$\begin{aligned} & \xrightarrow{l}_{\text{choice}[\sigma_\mu^\#]} b \\ \text{and later also } n & \xrightarrow{l}_{\text{choice}[\sigma_\mu^\#]} b \text{ if } n \xrightarrow{l}_G m \text{ and } \sigma_\mu^\# \text{ appears at } n.\end{aligned}$$

Just as before for local cache-cache dependencies, local state-cache dependencies for a memory block b are determined by the possible cache positions of certain memory blocks b' :

Observation 15.3.1. Let l, b, B' and $\sigma_\mu^\#$ as in Definition 15.3.1.

First, observe that if l never accesses any memory block b' , or accesses the same memory block for every macro-architectural state σ_M , then b does not depend on any choice (of memory block).

Let for all $b' \in B'$:

$$b'_{\min} = \min_{a \in \sigma_\mu^\#(b')} a \quad b'_{\max} = \max_{a \in \sigma_\mu^\#(b')} a$$

Then $\xrightarrow{l}_{\text{choice}[\sigma_\mu^\#]} b \iff$

$b_{\min}^1 < a < b_{\max}^2$ for some $a \in \sigma_\mu^\#(b)$ and some $b^1, b^2 \in B'$ with $b^1 \neq b^2$,
or $b \in B'$ and B' is not a singleton set

15.4 Transitive Cache Dependencies

Local state-cache dependence from Section 15.3 “creates” dependencies of memory blocks on macro-architectural state, while local cache-cache dependence from Section 15.2 describes how dependencies from one memory block must be locally propagated to other memory blocks.

From the following rules, I obtain for each control flow graph node m and each memory block b at node m the set of all other nodes n on which b depends, in the sense that the macro-architectural state at node n may (transitively) influence the position of block b in the cache at node m . The resulting relation $\rightarrow_{\text{cache}}$ can be considered an analogue of the transitive closure of standard data dependence $\rightarrow_{\text{data}}$.

Definition 15.4.1. For every control flow graph node n , let $\sigma_{\mu}^{\#}(n)$ be a sound approximation of all concrete cache states at n .⁴

A dependency of some memory block b is said to be killed in some abstract cache state $\sigma_{\mu}^{\#}$ if $\sigma_{\mu}^{\#}(b)$ is a singleton.

I write $\rightarrow_{\text{cache}}$ for the least solution \rightarrow of the rule system

$$\frac{n \xrightarrow{l}_G m \quad n \xrightarrow{l}_{\text{choice}[\sigma_{\mu}^{\#}(n)]} b}{n \rightarrow (m, b)} \text{CHOICE}$$

$$\frac{n \rightarrow (m, b) \quad m \xrightarrow{l}_G m' \quad \sigma_{\mu}^{\#}(m')(b) \text{ is not a singleton}}{n \rightarrow (m', b)} \text{TRANS}$$

$$\frac{m \xrightarrow{l}_G m' \quad b' \xrightarrow{l}_{\text{local}[\sigma_{\mu}^{\#}(m)]} b \quad \sigma_{\mu}^{\#}(m')(b) \text{ is not a singleton}}{n \rightarrow (m', b)} \text{LOCAL}$$

⁴ This is easily obtained by in standard data flow framework, by an abstract interpretation of the micro-architectural semantics for control flow graph edge labels l .

Then for memory block b and control flow graph nodes m and n , I say that at m , memory block b depends on n if

$$n \rightarrow_{\text{cache}} (m, b)$$

15.5 Approximate Cache Dependencies

I want to use transitive cache dependencies $n \rightarrow_{\text{cache}} (m, b)$ to obtain an approximation of micro-architectural dependencies as defined in Chapter 13. Transitive cache dependencies $n \rightarrow_{\text{cache}} (m, b)$ capture dependencies arising from macro-architectural choices at array accesses like $r2 := a[r1]$, but they do not capture influence arising from macro-architectural choices at conditional control flow graph nodes. In order to obtain a sound approximation, I hence need to respect (nontermination insensitive) control dependencies.

Definition 15.5.1. Let $\sigma_\mu^\#$ be an abstract cache state, and l control flow graph edge label. Let $b \in \mathbb{B}$ be any memory block. Then I say that b may be *modified* in $\sigma_\mu^\#$ by transition l if b may be accessed by label l , or if there exists a memory block b' that may be accessed by label l , and some cache positions

$$\begin{aligned} a &\in \sigma_\mu^\#(b) \text{ and} \\ a' &\in \sigma_\mu^\#(b') \end{aligned}$$

with $a < a'$.

In that case, I write

$$\begin{aligned} &\xrightarrow{l}_{\text{mod}[\sigma_\mu^\#]} b \\ \text{and also } n &\xrightarrow{l}_{\text{mod}[\sigma_\mu^\#]} b \text{ if } n \xrightarrow{l}_G m \text{ and } \sigma_\mu^\# \text{ appears at } n. \end{aligned}$$

Definition 15.5.2. Let n, m be two nodes in control flow graph G . Let B be the set of cache lines b that may be accessed along control flow graph edges $m \xrightarrow{l}_G m'$ such that in $\sigma_\mu^\#(n)$, cache line b is not either guaranteed to be in the cache, or guaranteed to be *not* in the cache, i.e.:

$$\neg \left(\infty \notin \sigma_\mu^\#(n)(b) \vee \sigma_\mu^\#(n)(b) = \{\infty\} \right)$$

and write

$$G_m = G_{m \not\rightarrow}^{\rightarrow * m}$$

Then I say that m is approximately cache micro-architecture dependent on n , and write $n \rightarrow_{\mu d}^{\#} m$, if

$$\begin{aligned} n &\in (\rightarrow_{\text{nticd}}^{\text{Gm}})^*(N') && \text{for } N' = \{n' \mid n' \rightarrow_{\text{cache}}(m, b), b \in B\} \\ \text{or } n &\in (\rightarrow_{\text{nticd}}^{\text{Gm}})^*(N') \setminus N' && \text{for } N' = \{n' \mid n' \xrightarrow{l}_{\text{mod}[\sigma_{\mu}^{\#}(n')]} b, b \in B\} \end{aligned}$$

Remark 15.5.1. Note that since (micro-architectural) cache state can depend on macro-architectural choices, but macro-architectural choices cannot depend on cache-state, I need *not* consider an iterated slice

$$(\rightarrow_{\text{nticd}}^{\text{Gm}} \cup \rightarrow_{\text{cache}})^*(m)$$

or similar.

Approximate cache micro-architecture dependency really is a sound approximation of micro-architecture dependence for cache micro-architectures:

Observation 15.5.1.

$$n \rightarrow_{\mu d} m \implies n \rightarrow_{\mu d}^{\#} m$$

15.6 Improving the Precision

By rule TRANS for the computation of transitive cache dependencies $\rightarrow_{\text{cache}}$ in the previous section, a dependency $n \rightarrow (m, b)$ is only ever “killed” along a control flow graph edge $m \xrightarrow{l}_G m'$ if the approximation $\sigma_{\mu}^{\#}(m')(b)$ is a singleton set. But consider the following program, assuming that the array `a` spans 4 memory blocks, and also assume a cache size of $k = 4$.

```
ra1 := a[r1]; rx := x; ry := y; ra2 := a[r2]; ru := u; rv := v;
```

Assume that scalar variables `u, v, x, y` reside in distinct memory blocks. Then the choice⁵ at the first access `ra1 := a[r1]`; to array `a` cannot possibly influence the position of any memory block corresponding to array `a` at the end of the program: Due to the accesses to the four scalar variables, the first choice is followed by four memory accesses to distinct memory blocks. Hence for a cache of size 4, the first choice must have been “evicted” by the end of the program. Only the second choice at `ra2 := a[r2]`; impacts the position of memory blocks corresponding to array `a` at the end of the program.

But the approximations $\sigma_{\mu}^{\#}(m)(a)$ of memory blocks `a` in the span of array `a` are never a singleton in between the first choice and the end of the program. Hence by the rules from the previous section, the dependency from the first choice is never killed.

In my prototype implementation, in order to obtain more precise transitive cache dependencies, I track for each dependency its “age”, and then kill dependencies after traversal of control flow graph edges $m \xrightarrow{l}_G m'$ not only if the approximation of `b` at `m'` is a singleton, but also if since it’s creation, the dependency must have been *evicted* by k accesses to other memory blocks that each must have pushed block `b` towards the back of the cache. Observation 15.5.1 still holds.

⁵ due to the value of `r1`

15.7 Approximation in AES256 Implementations

In order to evaluate the loss of precision in the approximation Observation 15.5.1, I applied it to the three AES256 implementations from Chapter 14. The results are shown in Figure 15.1, 15.2 and 15.3.

As expected for the “naive” default Implementation, the approximation can report possible timing channels not only for cache sizes up to 10 cache lines, also for much bigger assumed cache sizes. The more precise analysis previously (in Figure 14.1 on page 271) could not be run for cache sizes bigger than 10 cache lines.

For the *pre-caching* implementation, the approximate analysis can proof absence of timing channels only for 10, 14 or more cache lines. The more precise analysis could show this also for 8 and 12 lines.

For the implementation with constant time “S-Box” computation, the approximate analysis can proof absence of timing channels for all assumed cache sizes.

The running time and memory requirements for all analysis runs were negligible for small cache sizes, but become more significant for high number of cache lines in the implementation with constant time “S-Box” computation (Figure 15.3).

| # c.l. | $ \mu^{\#} d $ | $ S $ | time (s) | memory (kb) | passed |
|--------|----------------|-------|----------|-------------|--------|
| 2 | 29 | 695 | 1 | 34 812 | X |
| 4 | 126 | 695 | 1 | 33 780 | X |
| 6 | 558 | 695 | 1 | 38 964 | X |
| 8 | 639 | 695 | 1 | 40 936 | X |
| 10 | 617 | 695 | 1 | 35 836 | X |
| 12 | 529 | 695 | 1 | 38 908 | X |
| 14 | 507 | 695 | 1 | 42 928 | X |
| 16 | 557 | 695 | 1 | 68 588 | X |
| 18 | 557 | 695 | 1 | 80 824 | X |
| 20 | 557 | 695 | 1 | 76 776 | X |
| 22 | 311 | 695 | 1 | 75 764 | X |
| 24 | 311 | 695 | 1 | 75 744 | X |
| 26 | 311 | 695 | 1 | 76 780 | X |
| 28 | 311 | 695 | 1 | 75 772 | X |
| 30 | 311 | 695 | 1 | 94 152 | X |
| 32 | 311 | 695 | 1 | 90 108 | X |
| ... | | | | | |
| 100 | 251 | 695 | 1 | 97 224 | X |
| 200 | 251 | 695 | 1 | 96 152 | X |
| 500 | 251 | 695 | 2 | 94 256 | X |
| 1000 | 251 | 695 | 2 | 95 176 | X |
| 2000 | 251 | 695 | 3 | 90 052 | X |

Figure 15.1: Default AES265 Implementation

| # c.l. | $ \mu^{\#} d $ | $ S $ | time (s) | memory (kb) | passed |
|--------|----------------|-------|----------|-------------|--------|
| 2 | 63 | 703 | 1 | 37 864 | ✗ |
| 4 | 219 | 717 | 1 | 40 916 | ✗ |
| 6 | 292 | 717 | 1 | 39 956 | ✗ |
| 8 | 175 | 705 | 1 | 38 856 | ✗ |
| 10 | 260 | 73 | 1 | 40 928 | ✓ |
| 12 | 337 | 432 | 1 | 38 560 | ✗ |
| 14 | 304 | 79 | 1 | 41 956 | ✓ |
| 16 | 296 | 79 | 1 | 46 020 | ✓ |
| 18 | 316 | 83 | 1 | 42 924 | ✓ |
| 20 | 316 | 83 | 1 | 48 008 | ✓ |
| 22 | 32 | 40 | 1 | 44 004 | ✓ |
| 24 | 32 | 40 | 1 | 50 132 | ✓ |
| 26 | 32 | 40 | 1 | 44 156 | ✓ |
| 28 | 32 | 40 | 1 | 42 912 | ✓ |
| 30 | 32 | 40 | 1 | 48 012 | ✓ |
| 32 | 32 | 40 | 1 | 47 020 | ✓ |
| ... | | | | | |
| 100 | 32 | 40 | 1 | 51 180 | ✓ |
| 200 | 32 | 40 | 1 | 51 000 | ✓ |
| 500 | 32 | 40 | 1 | 51 232 | ✓ |
| 1000 | 32 | 40 | 1 | 50 996 | ✓ |
| 2000 | 32 | 40 | 1 | 50 140 | ✓ |

Figure 15.2: With *pre-caching*. Imprecision is highlighted.

| # c.l. | $ \mu^\# d $ | $ S $ | time (s) | memory (kb) | passed |
|--------|--------------|-------|----------|-------------|--------|
| 2 | 17 | 52 | 3 | 206 760 | ✓ |
| 4 | 77 | 60 | 3 | 205 764 | ✓ |
| 6 | 260 | 67 | 3 | 230 316 | ✓ |
| 8 | 222 | 65 | 3 | 290 620 | ✓ |
| 10 | 208 | 64 | 3 | 254 976 | ✓ |
| 12 | 208 | 64 | 3 | 267 264 | ✓ |
| 14 | 206 | 63 | 3 | 226 104 | ✓ |
| 16 | 176 | 60 | 3 | 256 996 | ✓ |
| 18 | 176 | 60 | 3 | 272 384 | ✓ |
| 20 | 162 | 59 | 3 | 241 628 | ✓ |
| 22 | 148 | 58 | 3 | 247 804 | ✓ |
| 24 | 148 | 58 | 4 | 297 940 | ✓ |
| 26 | 148 | 58 | 4 | 289 680 | ✓ |
| 28 | 148 | 58 | 4 | 248 796 | ✓ |
| 30 | 148 | 58 | 4 | 249 816 | ✓ |
| 32 | 148 | 58 | 4 | 269 252 | ✓ |
| | | | ... | | |
| 100 | 148 | 58 | 15 | 1 995 740 | ✓ |
| 200 | 157 | 57 | 47 | 5 969 956 | ✓ |
| 500 | 157 | 57 | 103 | 5 969 836 | ✓ |
| 1000 | 157 | 57 | 194 | 5 962 748 | ✓ |
| 2000 | 157 | 57 | 376 | 5 962 720 | ✓ |

Figure 15.3: Constant Time S-Box

15.8 Related Work

Both the generic micro-architectural dependencies from Chapter 13, as well as (LRU data-cache specific) approximate cache dependencies from this chapter extend standard *dependency graphs* to realize a timing sensitive static information flow analysis sensitive to timing effects due to the (modeled) CPU micro-architecture. To the best of my knowledge, these two are the first such analyses based on dependency graphs.

The perhaps best-known static analysis for cache based timing channels is the *CacheAudit* system[Doy+15]. Based on approach of *counting* the set of possible observations (e.g., [KRB09]), *CacheAudit* also allows for the *quantification* of information leaks. With respect to timing leaks, *CacheAudit* counts the number of different possible execution times in a sound approximation of all possible execution paths. This approximation relies on a sound approximation not only of the possible cache state, but also of macro-architectural (e.g., heap and register) state. In order to obtain useful approximations of the possible execution times in all possible execution paths, the analysis must be able to statically bound the number of executions of each loop. In contrast, neither my generic micro-architectural dependencies, nor my approximate cache dependencies requires such a bound⁶. Two principal advantages of *CacheAudit* over my approaches is a) that it provided a quantification of information leakage for insecure program, and b) it can quantify information leakage not only for timing based attackers (i.e.: those who observe the total execution time), but also *access* based attackers (who monitor the access time of individual statements in *their own programs* running in parallel to the attacked program, in a shared cache environment), and *trace based* attackers who can monitor the sequence of cache hit and miss events. On the other hand, my approach provides not a quantification, but an *explanation* of timing leaks, in form

⁶ Arguably, this is not a crucial advantage of my approach, since in the most relevant class of programs, i.e.: implementation of cryptographic programs, such bounds can usually be inferred or provided

of the involved micro-architectural dependencies (as well as standard control- and data-dependencies). In Chapter 14, I used these dependencies to derive a pre-caching AES256 implementation that did *not* require all involved memory blocks to be in the data cache at all times, but instead also works (i.e.: is without timing leaks) for smaller data cache sizes.

Aside from a LRU caches, the *CacheAudit* system also implements FIFO and *pseudo*-LRU replacement strategies. These also immediately fall into my generic micro-architectural framework, but would require new analogues of Observation 15.3.1 and Observation 15.2.1 for *approximate* cache dependencies.

Other approaches to static timing sensitive information flow analysis sensitive to cache effects are based on product programs resulting from self-composition (e.g., [Alm+16]), or symbolic execution ([Cha+19; Bro+19]). A short bibliography to approaches not based on static program analysis can be found, e.g, in [Doy+15], chapter 8.

Summary

- Micro-architectural dependencies for simple data cache micro-architectures can be approximated by an analogue of data dependencies.
- The approximation is reasonably precise, and allows for the analysis of large modeled cache sizes.

16 Timing Sensitivity in Concurrent Programs

Never tell me the odds!

(Han Solo — Star Wars)

In [Gif12] and [GS15], the authors present dependency graph based algorithms for static information flow control in concurrent programs. The underlying security criterion is **Low Security Observational Determinism**, which requires *any* two executions to be completely indistinguishable by a “low” observer (who can only observe execution of certain “low” program points).

For many applications, this criterion is prohibitively strict, since even most simple programs *without any secret input at all* do not fulfill it. Consider the programs in Figure 16.1. Statements print_l indicate (observable) output on a low channel l . The program on the left can have the observable trace [42, 17] or [17, 42] and hence is LSOD only if one assumes a deterministic scheduler, and not LSOD otherwise. The program on the right is not LSOD even for (some) deterministic schedulers (assume, for example, a round-robin scheduler).

The static checks in [Gif12] and [GS15] — if they succeed — guarantee LSOD for *any* scheduler. The authors also give static checks for a **Relaxed** notion of LSOD. Both notions imply **Probabilistic Noninterference**.

In [Bre+16] and [Bis+18b], we improved on both these static checks. There, we assume the scheduler to be

1. program-state independent, i.e.: a scheduling decision does not depend on, e.g., the state of variables, or the program counter of the threads to be scheduled.
2. stateless, i.e.: there is no other (scheduler-internal) state on which scheduling decision may depend

| | |
|---|---|
| <pre> 1 void main(): 2 fork thread_1(); 3 fork thread_2(); 4 void thread_1(): 5 print_t(42); 6 void thread_2(): 7 print_t(17); </pre> | <pre> 1 void main(): 2 l := read_t; 3 fork thread_1(); 4 fork thread_2(); 5 void thread_1(): 6 if (l < 10) {skip; skip; skip} 7 print_t(42); 8 void thread_2(): 9 if (l >= 10) {skip; skip; skip} 10 print_t(17); </pre> |
|---|---|

Figure 16.1: Some Programs

- probabilistic, in the sense that a scheduling decision may be random, with some underlying distribution.

We called such schedulers *probabilistic*. The canonical probabilistic scheduler is the *uniform* scheduler, which chooses one of the current n threads, each with probability $1/n$. Other probabilistic schedulers are possible. For example, given a *static* assignment of priority to (classes of) threads, a probabilistic scheduler may (after normalization) choose a given thread with probability p_i , where p_i is the probability of the class which the i th thread belongs to. All empirical results in this chapter are based on the uniform scheduler.

In this chapter, after reviewing probabilistic non-interference (Section 16.1) and a short general discussion of timing leaks (Section 16.2), I first review in Section 16.3 the improved relaxed LSOD criterion as described in [Bis+18b]. That section also includes formal definition of all auxiliary notions required in this chapter.

Then in Section 16.4, I describe the statistical test I used to empirically validate all of this chapter's criteria for probabilistic noninterference. In Section 16.5, I explain how the criterion from Section 16.3 can be “less precise” than even Giffhorns original LSOD criterion from [Gif12] and [GS15].

Finally in Section 16.6, I propose a new criterion for probabilistic non-interference, based on *timing dependence* \rightarrow_{td} as defined in Section 10.1.

16.1 Probabilistic Noninterference

Probabilistic Noninterference is based on the probability $P_i(t)$ of a given trace t to occur under input i to the given program. Given an otherwise deterministic program semantic, this probability is determined by the probabilistic scheduler. I sidestep any issues of actual *existence* of such a probability measure for a given program by simply demanding that for any given program, there exists a maximal length $k \in \mathbb{N}$ such that for any input i , any trace t for input i is at most of length k . This will coincide with my empirical observations, which will be based on programs in a (concurrent) For language, i.e.: a language in which all loops will execute a number of times fixed at the beginning of the loop.¹ A less clumsy treatment of this issue can be found in [Bis+18a].

For a fixed probabilistic scheduler, a given program, and any input i to this program, let P_i be the corresponding probability measure. Let $T(i)$ be the set of possible traces for input i . Assume then a relation \sim_L on traces t , such that $t \sim_L t'$ if t, t' are deemed (low)-observational equivalent, and write $[t]_L$ for the equivalence class of t . Assume a similar relation \sim_L on inputs.

Definition 16.1.1 (Definition 6, [Bis+18b]). Let i, i' be inputs; let $\Theta = T(i) \cup T(i')$. Probabilistic Noninterference holds iff

$$i \sim_L i' \implies \forall t \in \Theta. P_i([t]_L) = P_{i'}([t]_L)$$

Remark 16.1.1. Requiring all traces of a program to be of maximal size k also sidesteps any issues of finding an appropriate relation $t \sim_L t'$. For example, in [GS15] and [Bis+18b], its definition is non-trivial in order to account for differences of two traces due to *infinite delay*. Also remember my development of a related notion in Section 6.6.

¹ i.e.: the language allows statements like `For h C`, which executes h iterations of C , where h is the value of the variable h upon reaching the statement. In a concurrent setting, such statements are implemented using a fresh *thread instance local* loop counter variable.

16.2 Observability of Internal Timing Leaks

All criteria in [Gif12; GS15; Bre+16] and [Bis+18b] assume a (externally!) timing *insensitive* notion of observation. For example, in the program on the left of Figure 16.1, the observer is assumed to really observe either $[42, 17]$ or $[17, 42]$. He is assumed to *not*, for example, observe timed traces such as $[42 \odot [3], 17 \odot [4]]$ or $[17 \odot [3], 42 \odot [4]]$.

Similarly in the program on the right of Figure 16.1. Assume, for a moment, that the read $\iota := \text{read}_i$ in line 2 was not from a L (public) channel, but from H (secret) channel. Note that here, if the observer *did* observe timed traces, he might easily conclude from observing, e.g., $[42 \odot [4], 17 \odot [10]]$ that $\iota < 10$ *must not* have been true, since in executions in which $\iota < 10$ *does* hold, 42 can not be observed at time 4, but only as early as time 7 (assuming a simple timing model that takes one unit of time to evaluate each expression and statement). He would be able to directly observe an *external* timing leak.

But also when observing the *untimed* traces $[42, 17]$, an observer will (tentatively²) infer that $\iota < 10$ was *less likely* to have held than not, since in the class of executions in which $\iota < 10$, the observation $[17, 42]$ is significantly *more* likely than the observation $[42, 17]$, while in the class of executions in which *not* $\iota < 10$, the observation $[17, 42]$ is significantly *less* likely than the observation $[42, 17]$.

Here, the delay of $\text{print}_i(42)$ due to $\iota < 10$ caused an *internal* timing leak, *made visible* to the server by the concurrent execution of $\text{print}_i(42)$ and $\text{print}_i(17)$. The statement $\text{print}_i(17)$ was also delayed (due to $\iota \geq 10$), but the program would still have *not* been probabilistically noninterferent even without line this delay in 9.

In this chapter, I too assume a (externally!) timing *insensitive* notion of observation. But as I have just shown, I still have to be sensitive

² The observer would gain even more confidence in his conclusion that $\iota < 10$, if it possible for him to observe multiple executions of this program (with input a priori unknown to him, but known to him to be *constant*).

to *internal* timing, since this concurrency can make this visible even to timing-insensitive observers.

16.3 The RLSOD Criterion

In this section, I review the (improved) relaxed LSOD criterion from [Bis+18b]. There, we also described its derivation from the criteria in [GS15], which I do not repeat here.

In addition to control- and data-dependence, the criterion is based on the following new notions required in *concurrent* programs:

1. The **May Happen in Parallel** relation m_1 MHP m_2 .
2. Inter-Thread dependencies $n \xrightarrow{x}_{\text{inter}} m$, a form of concurrent data dependencies (along variable x), based on either MHP, or a **May Happen in Before** relation MHB.
3. The *common dynamic ancestor* $cda(m_1, m_2)$ of two nodes.

The May Happen in Parallel Relation

Consider a deterministic labeled control flow graph $G = (N, E)$ with labels l as in Section 13.2, together with a binary relation $\xrightarrow{\text{spawn}}$ on N of *spawn edges*. For simplicity, assume G to consist of a numbered set of disconnected parts called *threads*, and each j th thread with an (entry) node n^j , i.e. a node n^j such that all nodes m in thread j are reachable from n^j , but with no G -predecessors. Also assume that G consists of *trivial* control sinks only³, and assume that spawn-edges $n \xrightarrow{\text{spawn}} m$ only enter entry nodes $m = n^j$ for some thread (number) j . I call the 0 th thread with entry node n^0 the *main* thread.

Aside from the global variable state σ and *thread instance local* variable state σ_i , I also assume a global state i for *input-channels*. Labels $a \in A$ then also include operations that operate on input- and output-channels. For example: labels for read statements that consume from

³ such that $\rightarrow_{\text{nticd}}^G$ alone is an appropriate notion of control-dependence, and $\rightarrow_{\text{nticd}}$ is not needed

i the next value available at some channel (yielding i'), and labels for print statements.

The single-threaded small step semantics then is

$$\frac{n \xrightarrow{a}_G m \quad (\sigma', \sigma'_l, i') = a^M(\sigma, \sigma_l, i)}{(n, \sigma, \sigma_l, i) \xrightarrow{a} (m, \sigma', \sigma'_l, i')} \text{State}$$

$$\frac{n \xrightarrow{g}_G m \quad g^M(\sigma, \sigma_l) = \mathbf{true}}{(n, \sigma, \sigma_l, i) \xrightarrow{g} (m, \sigma, \sigma_l, i)} \text{Guard}$$

while the full concurrent semantics then is

$$\frac{(n_l, \sigma, \sigma_l, i) \xrightarrow{l} (m_l, \sigma', \sigma'_l, i')}{\begin{array}{c} \frac{((n_1, \sigma_1), \dots, (n_l, \sigma_l), \dots, (n_{l+1}, \sigma_{l+1}), \dots, (n_k, \sigma_k)), \sigma, i}{\xrightarrow{l, l'} ((n_1, \sigma_1), \dots, (m_l, \sigma'_l), (n'_1, \epsilon), \dots, (n'_K, \epsilon), (n_{l+1}, \sigma_{l+1}), \dots, (n_k, \sigma_k)), \sigma', i'} \end{array}}$$

where n'_1, \dots, n'_K are the nodes n'_j such that $n_l \xrightarrow{\text{spawn}} n'_j$, and ϵ is some fixed initial thread-instance local state.

A configuration $\kappa = ((n_1, \sigma_1), \dots, (n_k, \sigma_k), \sigma)$ is said to be at n if $n = n_l$ for some l . A configuration may contain multiple *instances* of a thread j , i.e.: multiple nodes n_l such that n_l is in the j th disconnected part of G .

Given an initial state σ^0 and input i^0 , a node n may happen in parallel to node m if there is a sequence

$$([n^0], \sigma^0, i^0) \Rightarrow \dots \Rightarrow ((n_1, \sigma_1), \dots, (n_k, \sigma_1)), \sigma, i)$$

such that $n = n_l, m = n_{l'}$ with $l \neq l'$. This notion is symmetric.

I assume some sound, symmetric approximation MHP of this notion, i.e. a symmetric relation MHP such that if n may happen in parallel m , then n MHP m .

Remark 16.3.1. Lacking a notion of *procedures*, this notion of MHP is (vacuously) not calling-context sensitive. It is also not thread-context sensitive, i.e.: it does not differentiate between two occurrences of a single node n in some configuration κ based on the sequences of spawn nodes used to insert these occurrences into κ . This information would be readily available at configurations if my semantics was based, for example, on *execution trees*[Gaw+11]. It can also be approximated by *thread invocation analysis*[Gif12].

Also, if I used execution trees, I could define a more refined *may concurrently happen before* relation. Intuitively, a node n may concurrently happen before m if a configuration at n may happen before a configuration at m , with m occurring in a “different” thread instance than that of the occurrence of n , unless the thread instance of m was (transitively) spawned by the thread instance of n (and after n).

This is difficult to define in my semantics. I don’t want do consider a node n to concurrently happen before m only because in some sequence is first at n and then at m , because this may also be the case simply because m follows n in a single thread instance.

But note that since my semantics lacks any form of *synchronization*, my *may happen in parallel* notion approximates a proper *may concurrently happen before* notion, since if n may concurrently happen before m , then certainly also n may happen in parallel to node m . Just consider the execution sequence in which the thread instance of n remains at n , and only nodes from other thread instances are chosen to proceed, until m is reached from these other thread instances.

Inter-Thread Dependence

Intuitively, the value of a variable x defined at node n may be read at a node m in a thread instance different than that of n , if n may

concurrently happen before m . As just argued, I can substitute *may concurrently happen before* by *may happen in parallel*, and use its static approximation MHP.

Definition 16.3.1. For nodes n, m ,

$$\begin{aligned} n \xrightarrow{x}_{\text{inter}} m &\iff x \in \text{def}(n) \cap \text{use}(m) \text{ and } n \text{ MHP } m \\ \text{and } n \rightarrow_{\text{inter}} m &\iff \exists x. n \xrightarrow{x}_{\text{inter}} m \end{aligned}$$

Inter-thread dependence $n \xrightarrow{x}_{\text{inter}} m$ does *not* account for the flow of data from a definition at node n to those nodes m that are (transitively) spawned from the thread instance in which n occurred. An example is the definition $\text{!} := \text{read}_i$ of variable ! in the right hand side of Figure 16.1 in line 2, to its use **if** ($\text{!} < 10$) in line 6.

In order to account for such flows, I need — in this chapter — standard data dependence $\rightarrow_{\text{data}}$ to include flow along spawn-edges $\xrightarrow{\text{spawn}}$, i.e.: for $G = (N, E)$, I use the data dependence relation for

$$G^{\text{spawn}} := \left(N, E \cup \xrightarrow{\text{spawn}} \right)$$

In the standard data-flow framework used to compute $\rightarrow_{\text{data}}^{G^{\text{spawn}}}$, then, definitions flow along edges $n \xrightarrow{\text{spawn}} m$ unmodified (i.e.: the transformer is the identity function).

Remark 16.3.2. Inter-Thread dependence $n \xrightarrow{x}_{\text{inter}} m$ is called *interference dependence* in [Kri98] and [Gif12].

Common Dynamic Ancestors

In [Bre+16] and [Bis+18b], we introduced the notion of common dynamic ancestors of two nodes m_1, m_2 . A motivating example is shown in Figure 16.2.

Here — somewhat similar to Figure 16.1 (right) — the observable statements $\text{print}_i(17)$ and $\text{print}_i(42)$ are delayed by the loop

```

1  void main():
2    h := readi;
3    for h { skip; }
4    skip;
5    fork thread_1();
6    fork thread_2();
7  void thread_1():
8    printi(42);
9  void thread_2():
10   printi(17);

```

Figure 16.2: The need for *common dynamic ancestors*

`for h { skip; }` at line 3, which executes h skip statements. A (externally) timing sensitive observer would learn the input to h by the execution time of the print statements.

What is *not* similar to Figure 16.1 (right) is that here in Figure 16.2, both $\text{print}_i(17)$ and $\text{print}_i(42)$ are delayed *by the same amount of time*. A (externally) timing insensitive observer will observe either $[42, 17]$ or $[17, 42]$, but neither observation will allow him to infer anything about the secret input value h . Even by repeatedly observing executions of this program, he will observe $[42, 17]$ roughly $2/3$ of the time⁴, and $[17, 42]$ roughly $1/3$ of the time, *independent* of the value of h . In fact, after the delay at line 3, any execution of this program must first pass the skip statement at line 4 before reaching any of the two print statements that otherwise could have made the delay visible externally. The skip statement is a *common dynamic ancestor* of the two print statements.

Definition 16.3.2 (Common dynamic ancestor, [Bis+18b]). Let $n, m, c \in N$ be nodes in $G = (N, E)$, and n^0 the entry node of the main thread. Remember that G^{spawn} is the graph G together with $\xrightarrow{\text{spawn}}$ edges.

⁴ since `thread_1` is forked first

1. c is a common dominator for m_1, m_2 , written $c \sqsupseteq_{\text{CDOM}} (m_1, m_2)$, if c dominates both m_1 and m_2 in G^{spawn} , i.e.: if

$$c \sqsupseteq_{\text{DOM}}^{\text{Gspawn}} m_1 \text{ and } c \sqsupseteq_{\text{DOM}}^{\text{Gspawn}} m_2$$

2. c is a common dynamic ancestor for m_1, m_2 , written $c \sqsupseteq_{\text{CDA}} (m_1, m_2)$, if

$$\begin{aligned} & c \sqsupseteq_{\text{CDOM}} (m_1, m_2) \\ \text{and neither } & c \text{ MHP } m_1 \\ \text{nor } & c \text{ MHP } m_2 \end{aligned}$$

3. If $c \sqsupseteq_{\text{CDA}} (m_1, m_2)$ and $\forall c' \sqsupseteq_{\text{CDA}} (m_1, m_2). c' \sqsupseteq_{\text{DOM}}^{\text{Gspawn}} c$, then c is called an *immediate* common dynamic ancestor. I then write $c = \text{icda}(m_1, m_2)$.

Classification

Based on these three notions specific for concurrent programs, and existing notions of program dependence, in [Bis+18b] we then defined an improved, relaxed LSOD criterion for probabilistic noninterference. In the context of this chapter, I define the concurrent program dependence graph to be

$$\rightarrow_{\text{cpdg}} = \rightarrow_{\text{nticd}}^{G'} \cup \rightarrow_{\text{data}}^{G^{\text{spawn}}} \cup \rightarrow_{\text{inter}} \cup \xrightarrow{\text{spawn}}$$

Here, I treat spawn edges $m \xrightarrow{\text{spawn}} m'$ as dependencies in order to propagate control-dependencies $n \rightarrow_{\text{nticd}}^G m$ from m to m' , since if m is control-dependent on n , then n also decides whether the thread starting in m' executes. In order to make sure that in fact then *all* nodes in the thread starting in m' are dependent on n , I compute control dependence not in the graph G , but in the graph G' obtained

from G by adding an edge from each start node n^j of the j th thread to each *exit node* in thread j .

Then, given some information flow lattice \mathcal{L} , we defined the *classification* of nodes in G with regard to the (user provided) information flow specification ucl (the “user classification”) to be the *least* solution of the rule system RLSOD, as follows:

Definition 16.3.3.

$$\frac{m \in I}{\text{cl}(m) \sqsupseteq \text{ucl}(m)} \text{INPUT} \qquad \frac{n \rightarrow_{\text{cpdg}} m}{\text{cl}(m) \sqsupseteq \text{cl}(n)} \text{CPDG}$$

$$\frac{c = \text{icda}(m_1, m_2) \quad m_1 \text{MHP } m_2 \quad c \rightarrow_{G^{\text{spawn}}}^* n \rightarrow_{G^{\text{spawn}}}^* m_i \quad i \in \{1, 2\}}{\text{cl}(m_i) \sqsupseteq \text{cl}(n)} \text{CDA}$$

The specification ucl is given as a *partial* map from nodes to \mathcal{L} . The domain of ucl is assumed to be partitioned into sets of *input* nodes I and of *output* nodes O . Nodes m in the domain of ucl are specified to be observable at all levels $l \sqsupseteq \text{ucl}(m)$. Specifically for the lattice $\mathcal{L}_2 = L \sqsubset H$, nodes m with $\text{ucl}(m) = L$ are observable at *all* levels, while nodes m with $\text{ucl}(m) = H$ are observable only at level H .

A program is then deemed secure at level $l \in \mathcal{L}$ by the RLSOD criterion (with respect to the specification ucl) if it admits the criterion rule

$$\frac{m \in I \cup O \quad \text{ucl}(m) \sqsubseteq l}{\text{cl}(m) \sqsubseteq l} \text{RLSOD}$$

Remark 16.3.3. If in the rule system RLSOD we instead use the subset-lattice on nodes N and the specification $\text{ucl}^N(m) = \{m\}$, then for each m the solution $\text{cl}(m)$ is a form of backward slice of m .

16.3.1 Observations

In order to connect judgments of the rule RLSOD to probabilistic non-interference, I need make concrete the notion \sim_L of equivalence of *traces*, and the notion $i \sim_L i'$ of equivalence of *inputs*.

Here, an input i is a map from a fixed set of input-channels c^I to *streams* of values. I also assume a fixed set of output-channels c^O .

I extend specifications ucl to channels, i.e.: I demand that ucl maps all input-channels c^I and all output-channels c^O to some security level $l \in \mathcal{L}$.

Two inputs i, i' then are l -equivalent, and I write $i \sim_l i'$ iff i and i' coincide on input-channels classified $\sqsubseteq l$, i.e. iff

$$i(c^I) = i'(c^I)$$

for all input-channels c^I such that $\text{ucl}(c^I) = l'$ for some $l' \sqsubseteq l$.

In need to define observations on execution-traces t , i.e.: observations of sequences

$$\dots \rightsquigarrow ([\dots, (n_i, \sigma_i), \dots], \sigma, i) \xrightarrow{l, \lambda} ([\dots, (n'_i, \sigma'_i), \dots], \sigma', i') \rightsquigarrow \dots$$

in the full concurrent semantics. I demand that the specification ucl for nodes is undefined for nodes n ,⁵ *unless* n has an outgoing control flow graph edge $n \xrightarrow{\lambda} m$ with λ being either an *input*-statement or an *output*-statement, i.e. a label of the form:

$x := \text{read}_{c^I};$ on channel $c = c^I$ with $\text{def}(\lambda) = \{x\}$, $\text{use}(\lambda) = \emptyset$
or $\text{print}_{c^O}(x);$ on channel $c = c^O$ with $\text{def}(\lambda) = \emptyset$, $\text{use}(\lambda) = \{x\}$

In this case, I demand $\text{ucl}(n) = \text{ucl}(c)$.

⁵ i.e.: n is unobservable for *any* observer

The l -observable subtrace t_l of t then is the sequence

$$\dots \xrightarrow{\text{use}} (\sigma_{\text{use}}, n_l, \lambda, \sigma'_{\text{def}}) \xrightarrow{\text{def}} \dots$$

containing only those configurations for which $\text{ucl}(n_l) \sqsubseteq l$. Here, σ_{use} and σ'_{def} are suitable projections of (σ_l, σ) to any variables use (λ) printed via label λ , and of (σ'_l, σ') to the set $\text{def}(l)$ of any variables read from input i via label λ , respectively.

Writing $t \sim_l t'$ whenever $t_l = t'_l$, we showed in [Bis+18b] that program deemed secure with respect to a specification ucl are probabilistically noninterferent.

Theorem 16.3.1 (Corollary 1 in [Bis+18b]). Let G be a program's labeled control flow graph, and ucl a specification G in the security lattice $\mathcal{L} = L \sqsubseteq H$. Then if the least solution cl of rule system RLSOD admits the rule RLSOD, probabilistic noninterference holds for G .

Consider again the program on the left of Figure 16.1. It is *not* low-security observational deterministic, but it is probabilistically noninterferent at level L. This is trivially so, since there is no node classified H. The same holds for the program on the right of Figure 16.1.

For Figure 16.2, the immediate common ancestor of the two print statements is the skip statement in line 4. But no node between this statement and the print statements is (transitively) dependent on the H input in line 2. Hence the program passes the RLSOD criterion, and is probabilistically noninterferent.

16.4 A Statistical Test for Probabilistic Noninterference

In order to facilitate the development of improvements to the RLSOD classification rules, I implemented an ad hoc statistical test for probabilistic noninterference. It consists of

1. A random program generator for a minimal concurrent language For.
2. A compiler from For abstract syntax trees into control flow graphs G and spawn edges $\xrightarrow{\text{spawn}}$.
3. An ad hoc statistical test that for two given inputs i, i' attempts to determine whether i, i' form a counter-example to the supposition that G is probabilistically noninterferent, by repeatedly *executing* G under inputs i and i' until such a determination can be made with confidence.

The language of For statements is made up from

- bounded integer valued program variables, which are either *global* or *thread local*, and have names in Var
- arithmetic expressions $E_{\text{Arithmetic}}$ (over program variables) E_{Boolean} (over arithmetic expressions), with integer constants from Int.
- (static) thread identifiers T, with a designated main thread main.

Statements For $n\ c$ implement loops with constant number of iterations, while the number of iterations of statements For $x\ c$ is determined by the value of variable x immediately before the loop-statements execution. A For-program P then is a map $P : T \rightarrow \text{For}$ from thread identifiers to For-commands.

The compiler from For-programs P to labeled control flow graphs implements a immediate, non-optimizing, syntax-tree directed transla-

tion. Loops `For n c` and `For x c` are both implemented as control flow graph cycles, using fresh thread-local variables as loop counters.

The interpreter for control flow graphs is a direct implementation of the concurrent semantic from Section 16.3.

Given a labeled control flow graph G (obtained from a For-program) and two L-equivalent inputs i, i' , the ad hoc statistical tests must determine (with high confidence) whether inputs i, i' are a counterexample to the supposition that probabilistic noninterference holds for this program, i.e. whether or not it holds that

$$\forall t \in \Theta. P_i([t]_L) = P_{i'}([t]_L) \quad (16.1)$$

for the set $\Theta = T(i) \cup T(i')$ of traces possible under i, i' .

In order to be applicable to automated tests of (many) randomly generated programs, the statistical test must be *two-sided*. By this I mean that, given a set of $n \in \mathbb{N}$ randomly sampled executions of the program under input i , as well as a set of n randomly sampled executions of the program under input i' , the test would ideally

1. Determine (with high confidence) that Equation 16.1 holds, i.e.: *reject* the null hypothesis

$$H_0^{\neq} : P_i \neq P_{i'}$$

2. *or* determine (with high confidence) that Equation 16.1 does *not* hold, i.e.: *reject* the null hypothesis

$$H_0^= : P_i = P_{i'}$$

3. *or*, if it can do neither, “request” more samples (i.e.: more executions of the program under both i and i') until it can.

But *for principal reasons*, I cannot expect a statistical test to ever reject H_0^{\neq} , even if I observe minimal (or even no) differences in the observed

empirical distributions \hat{P}_i and $\hat{P}_{i'}$, because such an observation can always also be explained equally well by two actual underlying distributions P_i and $P_{i'}$ that *do* differ, but only by some infinitesimal amount $\epsilon > 0$. A nonsignificant difference must not be confused with significant homogeneity[Wel10]. Hence all I can hope to do is to reject a modified null-hypothesis

1. $H_0^{\geq \epsilon}$: P_i differs from $P_{i'}$ by at least ϵ , in some metric

Remark 16.4.1. Although presumably possible, I did *not* use a modified dual null hypothesis in the second test, i.e.: I did not need to replace the null-hypothesis $H_0^=$ by some null-hypothesis

2. $H_0^{< \epsilon}$: P_i differs from $P_{i'}$ by less than ϵ , in some metric

In theory, keeping $H_0^=$ has the disadvantage that together with $H_0^{\geq \epsilon}$, it “logically consistent” to *simultaneously* reject $H_0^=$ and $H_0^{\geq \epsilon}$. In practice, this rarely a problem if I choose a very small epsilon, and require very high confidence $(1 - \alpha)$.

I treat P_i (and similarly: $P_{i'}$) as a *multinomial* distribution over all possible observations of the program under inputs i, i' (i.e.: over all possible equivalence classes $[t]_L \in \Theta / \sim_L$), represented by a vector

$$P_i = (p_{i,1}, \dots, p_{i,j}, \dots)$$

with $\sum_j p_{i,j} = 1$, and $p_{i,j}$ being the probability of the j th equivalence class $[t]_L$ of L-equivalent traces⁶ under input i .

The standard approach to reject $H_0^=$ in this setting is (some variant of) Pearson’s Chi-square test. In my setting, however, this test is inappropriate, because it requires a “large” number of observations *in each bin*, i.e.: for each j , it requires a large number of observations of the j th equivalence class of traces in the empirical observation \hat{P}_i . In particular, the often cited requirement of “at least five” observations each is not met for many programs. Also, the “number of bins”, i.e.:

⁶ in a fixed, but arbitrary ordering of these equivalence classes

the number $|\Theta/\sim_L|$ of different possible L-observations are usually high. Hence instead of a Chi-square test, I use a recent test for such *sparse* and *high-dimensional* multinomial distributions due to Plunkett and Park [PP18], which is based on (an unbiased estimator of) the Euclidean distance between P_i and $P_{i'}$.

For the rejection of the null hypothesis $H_0^{\geq \epsilon}$, I use a *goodness of fit* test due to Wellek ([Wel10], Section 9.1). The distance between two multinomial distributions $P_i, P_{i'}$ there is defined as the Euclidean distance

$$0 \leq d(P_i, P_{i'}) = \sqrt{\sum_j |p_{i,j} - p_{i',j}|^2} \leq 1$$

between the corresponding vectors. The null hypothesis hence formally is:

1. $H_0^{\geq \epsilon} : d(P_i, P_{i'}) \geq \epsilon$

Remark 16.4.2. Unlike Plunkett’s and Park’s test, Wellek’s test is a *one* sample test. It is originally designed for the test of goodness-of-fit of distributions $P_i, P_{i'}$, based on one empirical distribution (say: $\hat{P}_{i'}$, obtained from sampling the a priori unknown distribution $P_{i'}$) and one *fully specified reference* distribution (say: P_i), i.e.: one a priori fully specified vector $(p_{i,1}, \dots, p_{i,j}, \dots)$.

Strictly speaking, what I need is a two-sample test, i.e.: a test for goodness-of-fit based on *two* empirical distributions. I shoehorn Wellek’s test into my situation by employing Wellek’s test *twice*, each time treating *one* of the two empirical distributions (say: \hat{P}_i) as the reference distribution (i.e: by assuming $\hat{P}_i = P_i$). I then reject the null hypothesis $H_0^{\geq \epsilon}$ only if both these two instances of Wellek’s test do so.

Presumably, a *proper* two-sample goodness-of-fit test is possible by basing the test-statistics on the “two-sample”-variance estimator $\hat{\sigma}_k^2$ (from [PP18], Equation 9), instead of the “one-sample”-variance estimator v_n^2 (from [Wel10], Equation 9.8), but I did not pursue this.

Definition 16.4.1. The full ad hoc test for probabilistic noninterference for L-equivalent inputs i, i' to a given program works as follows:

1. Choose confidence level α (say: ≤ 0.05), maximal distance ϵ (say: < 0.01), and an initial value k (say: ≥ 10).
2. Sample L-observations of executions of G both for input i and for input i' to obtain a number $n = 2^k$ of total samples, each.
3. If both $H_0^{\geq \epsilon}$ is rejected by the test due to Wellek, and H_0^- is rejected by the test due to Plunkett/Park, then report failure.
4. If H_0^- is rejected by the test due to Plunkett/Park, then finish and report that G is *not probabilistically noninterferent*, with the pair (i, i') being a counterexample.
5. If $H_0^{\geq \epsilon}$ is rejected by the test due to Wellek, then finish and report that the pair (i, i') is *no counterexample* to the claim that G is probabilistically noninterferent.
6. Otherwise, increment k by one and continue at step 2.

In my automated experiments, I chose $\alpha = 0.0000001$, $\epsilon = 0.009$, and initialized $k = 12$ (such that initially, $n = 4096$).

The step 3 is not essential, and could have been left out. If triggered, it indicates evidence for a very small, but non-zero difference

$$0 < d(P_i, P_{i'}) < \epsilon$$

Step 3 could as well just report that G is *not* probabilistically noninterferent (as is done in step item 4). I inserted step 3 merely to determine frequency of situation described in Remark 16.4.1. In my automated experiments, I did indeed observe such a situation, but only *once* for 20000 randomly generated programs (and runs of the statistical test).

The doubling of sample size in each iteration due to step 6 is meant to reduce the probability of type I errors of either test, in lieu of proper *sequential hypothesis testing*.

By running this ad hoc test for two low-equivalent inputs i, i' and 20000 randomly generated For-programs *for which the RLSOD crite-*

tion claims probabilistic noninterference, I was able to validate Theorem 16.3.1.

Observation 16.4.1 (Empirical validation of Corollary 1 in [Bis+18b]). For at least 20000 randomly generated For-programs P (only counting those programs P such that the RLSOD-criterion claims probabilistic noninterference), manual inspection of all those program P for which the ad hoc test from Definition 16.4.1

- failed (via step item 3), or
- reported a counter-example (via step item 4)

revealed those P to be *in fact* probabilistically noninterferent⁷.

For my choice of α, ϵ and initial number of samples $n = 2^k$ as above, the ultimately required number of samples is $n = 2^{12} = 4096$ for the majority of programs, but (rarely) goes up as high as $n = 2^{18} = 262144$. At the same time, false claims of counterexamples are rare. I observed just thirteen in total.

If I choose instead a confidence level of, e.g., $\alpha = 0.01$, then such errors become more common (on the order of roughly one per 100 checked programs).

In total I manually inspected 13 reported programs, all of which turned out to be probabilistically noninterferent, just as was claimed by the RLSOD-criterion in each case.

Efficiency of the Statistical Test

I designed the ad hoc empirical test to gain confidence in the correctness of the RLSOD-criterion from Section 16.3, before the correctness

⁷ i.e.: manual inspection revealed those to be a type I error of the Plunkett/Park test, erroneously rejecting the null hypothesis $H_0^=$

proof from [Bis+18b] was available. I also use this test to check my new timing dependence based criterion coming up in Section 16.6.

But I can gain confidence in correctness by such a test only if I have confidence that the whole process of generating random For-programs and then testing the generated program on *one* pair (i, i') of L-equivalent inputs is indeed capable of exposing *faulty* criteria.

One could argue that what really needed to be done was to run the statistical tests for *all* pairs (i, i') of L-equivalent inputs, which of course is infeasible if not impossible⁸. But in practice, even using two fixed inputs i, i' gives counterexample for faulty criteria relatively quickly. For example, when I tried to validate the *unsound* criterion obtained from RLSOD by omitting the rule CDA, the 232th randomly generated program exposed this unsoundness. This program — which is *not* probabilistically noninterferent, but accepted by this unsound criterion — is shown in Figure 16.3 (left). Here, the *order* of the two observable reads in line 8 and line 19 is influenced by the secret value read in line 2.

⁸ This is different from the situation in Section 6.3, 6.6 and 9.5, in which *due to the simplicity* of the underlying notion of *input*, I *was* able to exhaustively test all equivalent pairs i, i' of inputs.

```
1 void main():
2   z := readi;
3   y := read1;
4   a := 0;
5   fork thread_2();
6   z := z;
7   a := readi;
8   z := read1;
9 void thread_2():
10  if (z*y <= 0)
11    skip;
12  else
13    skip;
14  for z {
15    b := y * a;
16  }
17  y := readi;
18  fork thread_3();
19  b := read2;
20 void thread_3():
21  for 2 {
22    print1(1);
23  }
24  x := -z;
25  x := y * z;
```

Figure 16.3: A randomly generated Example for the necessity of rule CDA

16.5 Imprecision of the RLSOD criterion

The example from Figure 16.2 on page 312 demonstrates how probabilistic noninterference is a less prohibitive notion of information flow security than low security observational determinism, which requires *any* executions for low-equivalent input to be observational indistinguishable (or in other words, it requires deterministic low observations).

Definition 16.5.1. Let i, i' be inputs; let $\Theta = T(i) \cup T(i')$. Low Security Observational Determinism holds iff

$$i \sim_L i' \implies \forall t, t' \in \Theta. t \sim_L t'$$

Since the example from Figure 16.2 is *not* low security observational deterministic, any sound static criterion for LSOD must reject it. On the other hand, this example *is* probabilistically noninterferent, and the RLSOD criterion (from Definition 16.3.3 on page 314) is indeed precise enough to accept it, hence for this instance, the RLSOD for probabilistic noninterference are an improvement over the previous LSOD criterion from [Gif12] and [GS15], which have to (and do) reject it.

In this section, I will first show that the RLSOD do not improve on the LSOD criterion on *every* program. I will then — by employing *timing dependence* \rightarrow_{td} — provide a criterion for probabilistic non-interference that improves on both the LSOD criterion and the RLSOD criterion. The new criterion will accept every program accepted by either of two previous others.

Giffhorns LSOD Criterion

First, I quickly review the LSOD criterion from [Gif12; GS15], in a style of presentation similar to that of the RLSOD criterion. Remember that the RLSOD criterion consisted of a rulesystem RLSOD of which the

least solution cl is then submitted to a check RLSOD, and the program is deemed probabilistically noninterferent if cl admits the rule RLSOD.

In order to facilitate comparison with the RLSOD-criterion, I also present Giffhorns LSOD criterion in this form, but I then require *two* (simple) rule systems:

1. LSOD[cl] whose least solution $cl_{LSOD} : N \rightarrow \mathcal{L}$ simply indicates the security level of nodes in the backward slice of each node.
2. LSOD[R] whose least solution $R \subseteq N$ is the set of nodes (potentially) influenced by a *data race*.

In order to present LSOD[R], I need the notion of *interference-write* dependence.

Definition 16.5.2 ([Gif12], Definition 5.15). For nodes n, m , interference-write-dependence via a variable x between n, m is defined as

$$\begin{aligned} n \overset{x}{\leftrightarrow}_{iw} m &\iff x \in \text{def}(n) \cap \text{def}(m) \text{ and } n \text{ MHP } m \\ \text{and } n \leftrightarrow_{iw} m &\iff \exists x. n \overset{x}{\leftrightarrow}_{iw} m \end{aligned}$$

I also define *interference-read* dependence as follows⁹

$$\begin{aligned} n \overset{x}{\rightarrow}_{ir} m &\iff x \in \text{def}(n) \cap \text{use}(m) \text{ and } n \text{ MHP } m \\ \text{and } n \rightarrow_{ir} m &\iff \exists x. n \overset{x}{\rightarrow}_{ir} m \end{aligned}$$

And ultimately, n, m are said to be in a data race if

$$n \leftrightarrow_{\text{race}} m \iff n \leftrightarrow_{iw} m \vee n \rightarrow_{ir} m$$

⁹ This is just a repetition of Definition 16.3.1 from page 311 of interference dependence $\rightarrow_{\text{inter}}$, but only because that definition is based on the *may happen in parallel* notion, when in different settings it would have better been defined on a *may happen before* notion. Also see Remark 16.5.1

Remark 16.5.1. Unlike interference dependence $\overset{x}{\rightarrow}_{\text{inter}}$, interference write dependence $\overset{x}{\leftarrow}_{\text{iw}}$ is symmetric. If instead of a (symmetric) approximation MHP of *may happen in parallel* behavior, I had a (asymmetric) approximation MHB of *may concurrently happen before* behavior, I would need to demand $n \text{MHB } m \wedge m \text{MHB } n$ in the definition of $\overset{x}{\leftarrow}_{\text{iw}}$.

Similarly, I would then need to demand $n \text{MHB } m \wedge m \text{MHB } n$ in the definition of interference-read dependence $n \overset{x}{\rightarrow}_{\text{ir}} m$, while the definition of interference dependence $n \overset{x}{\rightarrow}_{\text{inter}} m$ would just demand $n \text{MHB } m$.

Definition 16.5.3 ([Bis+18b], Definition 12; following [GS15]). The rule system $\text{LSOD}[\text{cl}]$ is:

$$\frac{m \in I}{\text{cl}_{\text{LSOD}}(m) \sqsupseteq \text{ucl}(m)} \text{INPUT} \quad \frac{n \rightarrow_{\text{cpdg}} m}{\text{cl}_{\text{LSOD}}(m) \sqsupseteq \text{cl}_{\text{LSOD}}(n)} \text{CPDG}$$

Definition 16.5.4 ([Bis+18b], Definition 12; following [GS15]). The rule system $\text{LSOD}[\text{R}]$ is:

$$\frac{m_1 \overset{x}{\leftarrow}_{\text{race}} m_2}{m_2 \in \text{R}} \text{RACE} \quad \frac{n \in \text{R} \quad n \rightarrow_{\text{cpdg}} m}{m \in \text{R}} \text{CPDG}$$

Note that $\text{LSOD}[\text{R}]$ is independent from the user classification ucl .

In the lattice $\mathcal{L}_2 = \{\text{L} \sqsubseteq \text{H}\}$, a program is then judged low (L) observationally deterministic by the LSOD-criterion if (given the least solutions cl_{LSOD} and R of these rule systems), the following three rules are admissible:

$$\frac{m \in O \quad \text{ucl}(m) \sqsubseteq \text{L}}{\text{cl}_{\text{LSOD}}(m) \sqsubseteq \text{L}} \text{LSOD}_1 \quad \frac{m \in O \quad \text{ucl}(m) \sqsubseteq \text{L}}{m \notin \text{R}} \text{LSOD}_2$$

$$\frac{m_1, m_2 \in I \cup O \quad \text{ucl}(m_1), \text{ucl}(m_2) \sqsubseteq \text{L}}{\neg m_1 \text{MHP } m_2} \text{LSOD}_3$$

Note that the propagation of *non-low input* along $\rightarrow_{\text{cpdg}}$ is completely separated (within cl_{LSOD}) from the propagation of *races* along $\rightarrow_{\text{cpdg}}$ (within R). Also completely separated is the treatment of *order conflicts* in rule LSOD_3 , which depends neither on cl_{LSOD} nor on R. Violations of rule LSOD_3 between two L-visible nodes, are always violations of low security observational determinism.

Contrasting these rules with the RLSOD rules, I note that

1. Rules INPUT and CPDG from rule system $\text{LSOD}[\text{cl}]$ are the same as rules INPUT and CPDG from rule system RLSOD.
2. RLSOD has no rule corresponding to rule LSOD_3 . This is because order conflicts prohibited by rule LSOD_3 are *only* violations of probabilistic non-interference if the probability of one of the two nodes (say: m_1) occurring before the other (m_2) is influenced by non-low input.
3. In $\text{LSOD}[\text{R}]$, any race $m_1 \leftrightarrow_{\text{race}} m_2$ is propagated along $\rightarrow_{\text{cpdg}}$, no matter if the probability of m_1 occurring before m_2 in the $m_1 \leftrightarrow_{\text{race}} m_2$ was ever influenced by some non-low input.

In contrast, in rule CDA from RLSOD, two nodes m_1, m_2 with undetermined execution order (as witnessed by the approximation $m_1 \text{MHP } m_2$) lead to propagation of a level $h \sqsupseteq \text{L}$ *only* if previously some node n ¹⁰ has been determined to be influenced by non-L input.

These considerations might lead one to belief that the RLSOD criterion is strictly “more precise” than the LSOD criterion, in the sense that whenever LSOD criterion accepts a given program, then also will the RLSOD criterion. But this is not so, in general. Consider the example on the left of Figure 16.4. It has only one possible low observation (a single output of 42), and hence is not only probabilistically noninterferent, but even low security observationally deterministic. The LSOD criterion is precise enough to detect this, since the $\rightarrow_{\text{cpdg}}$

¹⁰ subject to the *common dynamic ancestor* condition

backwardslice of the only L-output contains no high input (hence condition LSOD₁ is fulfilled), and the program contains no race (hence condition LSOD₂ is fulfilled). There is only one L-observable node n (at the print statement), hence there is no visible order-conflict, and condition LSOD₃ is fulfilled, too.

But the RLSOD criterion *rejects* this program. To see why, acknowledge that for the nodes corresponding to lines 4 and 5, $cl(4), cl(5) = H$ by $\rightarrow_{\text{cpdg}}$ (i.e.: by rule CPDG). Also, I have $cl(7) \sqsupseteq cl(4)$ due to rule CDA, because line 4 can happen after the execution of the common dynamic ancestor $c = 3 = icda(7, 9)$, but before 7.

This particular kind of imprecision in the RLSOD criterion could be argued to be of not too much practical impact, for it requires a parallel thread *without* L observable output; had there been *any* L-observable output in thread 2, then the RLSOD criterion would have *rightfully* rejected the program, since then the delay due to loop in line 4 would have influenced the *relative execution time* of that output in thread 2, and the `print,(42)` statement in line 7.

But also consider the example on the right of Figure 16.4, which is obtained from the example on the left by adding such a L-observable output to thread 2, and also replacing the loop in the main thread by a simple assignment. This example on the right is not low security observationally deterministic, but it *is* probabilistically noninterferent. Yet, the RLSOD criterion rejects it, by the same argument as for the example on the left.

In order to obtain a criterion for probabilistic noninterference more precise than the RLSOD-criterion, I propose to *separate* two kinds of influence of high input on low observations, which were *conflated* in the RLSOD-criterion:

1. The influence of high input on values at other nodes m , and on whether other nodes m are executed or not, as captured by the concurrent program dependence graph.

| | |
|---|--|
| <pre> 1 void main(): 2 h := read_h; 3 fork thread_2(); 4 for h { 5 skip; 6 } 7 print_l(42); 8 void thread_2(): 9 skip; </pre> | <pre> 1 void main(): 2 h := read_h; 3 fork thread_2(); 4 5 h2 := h 6 7 print_l(42); 8 void thread_2(): 9 print_l(17); </pre> |
|---|--|

Figure 16.4: Imprecision of the RLSOD-criterion

2. The influence of high input on the *relative execution order* of pairs (m_1, m_2) of nodes¹¹.

I will do this by use of *two* classifications $cl_{\mathbb{C}}$

1. A map $cl_{\mathbb{C}} : N \rightarrow \mathcal{L}$, from nodes n to security levels $cl_{\mathbb{C}}(n)$, and
2. a partial map $cl_{\mathbb{C}} : N \times N \hookrightarrow \mathcal{L}$, from pairs (m_1, m_2) of nodes to security levels $cl_{\mathbb{C}}(m_1, m_2)$

instead of just one classifications cl .

¹¹ which — lacking methods of *synchronization* — is possible in my program model only due to *delay* depending on high input

16.6 Timing Sensitivity for Probabilistic Noninterference

Consider the program from Figure 16.5. It is *not* probabilistically non-interferent, which can be argued as follows:

1. The value h of the secret input influences the length of the delay in line 8.
2. The length of the delay in line 8 influences the relative execution time of the two assignments to `delay2`, i.e.: it influences the probability of the assignment `delay2 := 0` occurring before the assignment `delay2 := 1000`, as opposed to vice versa.¹²
3. The read of variable `delay2` in line 11 sees the assignment `delay2 := 0` if that was executed after the assignment `delay2 := 1000`. Otherwise, it sees the assignment `delay2 := 1000`. But since the relative execution time of these two assignments was influenced by the value h (via the length of the delay in line 8), then so is the value of `delay2` in line 11.
4. The length of the delay in line 11 influences the relative execution time of the two publicly observable print statements.

Before I can present the rules that rigorously capture this kind of reasoning, I need to extend the notion of timing dependence from Chapter 10 to take into account spawn edges $\xrightarrow{\text{spawn}}$. The intuition here is that if for a spawn edge $m \xrightarrow{\text{spawn}} m'$, the node m is timing dependent on a node n in the control flow graph G , (i.e.: if $n \xrightarrow{\text{td}}^G m$), then also the timing of all nodes reachable from the entry node m' of the spawned thread depends on n .

¹² Line 8 also delays both the `print(42)` statement and the `print(17)` statement, but it delays them both *by the same amount*. Line 8 does not directly influence the *relative* execution time of these two `print` statements.

| | |
|---|---|
| <pre> 1 void main(): 2 delay1 = 0; 3 h := read_i; 4 if (h >= 0) { 5 delay1 := 1000; 6 } 7 fork thread_2(); 8 for delay1 { skip; } 9 delay2 := 0; 10 fork thread_3(); 11 for delay2 { skip; } 12 print_i(42); </pre> | <pre> 13 14 15 16 17 18 19 20 void thread_2(): 21 delay2 := 1000; 22 23 void thread_3(): 24 print_i(17); </pre> |
|---|---|

 Figure 16.5: Interdependence of $cl_{\mathbb{C}}(m)$ and $cl_{\mathbb{C}}(m_1, m_2)$.

Definition 16.6.1. Given an (multi-threaded, labeled) control flow graph $G = (N, E)$ and spawn edges $\xrightarrow{\text{spawn}}$ connecting nodes from G , a node m is timing dependent on a node n in $G^{\text{spawn}} = (N, E \cup \xrightarrow{\text{spawn}})$, and I write $n \xrightarrow[\text{timing}]{G^{\text{spawn}}} m$ or just $n \rightarrow_{\text{timing}} m$ iff:

$$\begin{aligned}
 & n \rightarrow_{\text{td}}^G m \\
 \text{or } & n \rightarrow_{\text{td}}^G m' \xrightarrow{\text{spawn}} m'' \rightarrow_{G^{\text{spawn}}}^* m \text{ for some nodes } m', m''
 \end{aligned}$$

Definition 16.6.2. The classification $cl_{\mathbb{C}}(\cdot) : N \rightarrow \mathcal{L}$ and the classification $cl_{\mathbb{C}}(\cdot, \cdot) : N \times N \leftrightarrow \mathcal{L}$ are the least solution to the *mutually recursive* rule system TIMING consisting of rules

$$\begin{array}{c}
 \frac{m \in I}{cl_{\mathbb{C}}(\sqsupset) \text{ ucl}(m)} \text{ INPUT} \quad \frac{n \rightarrow_{\text{cpdg}} m}{cl_{\mathbb{C}}(m) \sqsupseteq cl_{\mathbb{C}}(n)} \text{ CPDG} \\
 \frac{m_1 \xrightarrow{x}_{\text{inter}} m \quad m_2 \xrightarrow{x}_{\text{inter}} m \quad m_1 \text{ MHP } m_2}{cl_{\mathbb{C}}(m) \sqsupseteq cl_{\mathbb{C}}(m_1, m_2)} \text{ RACE}_1 \\
 \frac{m_1 \xrightarrow{x}_{\text{inter}} m \quad m_2 \xrightarrow{x}_{\text{data}} m \quad m_1 \text{ MHP } m_2}{cl_{\mathbb{C}}(m) \sqsupseteq cl_{\mathbb{C}}(m_1, m_2)} \text{ RACE}_2
 \end{array}$$

and the rule

$$\frac{m_1 \text{MHP } m_2 \quad c = \text{icda}(m_1, m_2)}{c \rightarrow_{G^{\text{spawn}}}^* n \rightarrow_{G^{\text{spawn}}}^* m_i \quad n \rightarrow_{\text{timing}} m_i \quad i \in \{1, 2\}} \text{CDA} \\ \text{cl}_{\mathbb{O}}(m_1, m_2) \sqsubseteq \text{cl}_{\mathbb{O}}(n)$$

Here, $\text{cl}_{\mathbb{O}}(\cdot, \cdot)$ is defined on the set

$$\{ (m_1, m_2) \mid m_1 \text{MHP } m_2 \wedge \text{ucl } m_1 \sqsubseteq l \wedge \text{ucl } m_2 \sqsubseteq l \} \\ \cup \{ (m_1, m_2) \mid \exists x. m_1 \xrightarrow{x}_{\text{iw}} m_2 \}$$

of pairs (m_1, m_2) of nodes that are either both $l \in \mathcal{L}$ observable and may happen in parallel, or are interference-write dependent.

A program is then judged probabilistically noninterferent at level $l \in \mathcal{L}$ by the TIMING-criterion if (given the least solution $\text{cl}_{\mathbb{O}}$ of the rule system TIMING), the following two rules are admissible:

$$\frac{m \in O \quad \text{ucl}(m) \sqsubseteq l}{\text{cl}_{\mathbb{O}}(m) \sqsubseteq l} \text{TIMING}_1 \\ \frac{m_1 \text{MHP } m_2 \quad m_1, m_2 \in O \cup I \quad \text{ucl}(m_1), \text{ucl}(m_2) \sqsubseteq l}{\text{cl}_{\mathbb{O}}(m_1, m_2) \sqsubseteq l \text{ and } \text{cl}_{\mathbb{O}}(m_1, m_2) \sqsubseteq l} \text{TIMING}_2$$

Replaying the argument at the beginning of this chapter, in which I showed that the example from Figure 16.5 is *not* probabilistically noninterferent,

1. was an instance of rule CPDG,
2. was an instance of rule CDA,
3. was an instance of rule RACE₂ and
4. was an instance of rule CDA.

In the solution $\text{cl}_{\mathbb{O}}$ I have both $\text{cl}_{\mathbb{O}}(12) = L$ and $\text{cl}_{\mathbb{O}}(24) = L$ for the two L-observable print statements, indicating that the values printed

by these statements do not depend on high input, and whether these statements are executed (and: how often) also does not depend on high input. Rule TIMING_1 is fulfilled. On the other hand, I *do* have $\text{cl}_{\mathcal{H}}(12, 14) = H$, indicating that *the relative execution order* of these two print statements may depend on high input. The rule TIMING_2 is *not* fulfilled. The program is rejected by the TIMING timing criterion, which could *not* prove it to be probabilistically noninterferent.

I claim that the TIMING criterion is *sound*, i.e.: that any program accepted by the TIMING criterion (given some user classification ucl) is probabilistically noninterferent with regard observations at every security level $l \in \mathcal{L}$, as defined in Subsection 16.3.1. For the simple lattice $\mathcal{L} = \{L \sqsubseteq H\}$, I gathered empirical evidence in support of this claim by use of the statistical test for probabilistic noninterference (Definition 16.4.1 from 320).

Observation 16.6.1 (Soundness of the TIMING criterion, empirical). For at least 20000 randomly generated For-programs P (only counting those programs P such that the TIMING -criterion claims probabilistic noninterference), manual inspection of all those program P for which the ad hoc test from Definition 16.4.1

- failed (via step item 3), or
- reported a counter-example (via step item 4)

revealed those P to be *in fact* probabilistically noninterferent

Since the statistical test may take considerable time (roughly between 5 and 100 seconds for the randomly generated programs), I used in Observation 16.6.1 the same randomly generated programs as before in the empirical validation of the RLSOD criterion (Observation 16.4.1). Those programs that required manual inspection all also passed the RLSOD criterion, so these were the same 13 programs as before.

The two programs from Figure 16.4 are probabilistically noninterferent and accepted by the TIMING criterion, but not by the RLSOD crite-

tion. In general, the TIMING-criterion is more precise than the RLSOD criterion.

Observation 16.6.2 (Precision of the TIMING criterion, relative to the RLSOD criterion). Let cl be the least solution to the rule system RLSOD (Definition 16.3.3) and $\text{cl}_{\mathbb{D}}$ be the least solution to the rule system TIMING (Definition 16.6.2).

Then

$$\begin{aligned} \text{cl}_{\mathbb{D}}(m) &\sqsubseteq \text{cl}(m) \\ \text{and } \text{cl}_{\mathbb{D}}(m_1, m_2) &\sqsubseteq \text{cl}(m_1) \sqcup \text{cl}(m_2) \end{aligned}$$

for all nodes m , and all pairs (m_1, m_2) in the domain of $\text{cl}_{\mathbb{D}}$. Consequently in the lattice $\mathcal{L}_2 = \{\text{L} \sqsubseteq \text{H}\}$ ¹³, if the RLSOD criterion holds at level l , then also does the TIMING criterion.

Remark 16.6.1. Specifically, the two inequations in Observation 16.6.2 hold for the “slicing” subset-lattice $\mathcal{L} = (2^N, \subseteq)$, when initializing each node n with $\text{ucl}(n) = \{n\}$.

¹³ and more generally: all “linear” lattices, in which $l_1 \sqcup l_2 \in \{l_1, l_2\}$

Summary

- For concurrent programs with probabilistic schedulers, probabilistic non-interference the appropriate security property.
- In line with previous research, the execution time is considered *externally* unobservable.
- Still, the relative execution time of memory accesses in concurrent threads makes internal differences in timing externally observable.
- The LSOD criterion, the RLSOD criterion, and the new TIMING criterion guarantee probabilistic non-interference.
- I substantiate this claim by extensive random testing.
- In these tests, probabilistic non-interference is validated by a statistical test.
- The LSOD criterion is very strict. The RLSOD criterion more liberal in some, but not all cases.
- The new TIMING criterion, based on timing dependence $\rightarrow td$, is more liberal than both the LSOD criterion and the RLSOD criterion.

17 Timing Sensitivity with JOANA

“Our product is still totally DeepArcher?”
“Which is ...”
“Like ‘departure’, only you pronounce it DeepArcher?”
“Zen thing,” Maxine guesses.
“Weed thing.”

(Thomas Pynchon — Bleeding Edge)

In Chapter 16, I introduced the new TIMING criterion for concurrent programs. I presented it for control flow graphs in a simple imperative language. But the TIMING criterion can be also be used for the analysis of concurrent programs more complex languages. In fact, I applied the TIMING criterion to a simple case study in the Java programming language, using an implementation in the JOANA¹ system. In this section I show that here, too, the TIMING criterion can improve on the RLSOD criterion.

The JOANA system[HS09; Sne+14] implements information flow analyses for Java programs. For the analysis of sequential programs, JOANA provides an (interprocedural) program dependence graph (“system” dependence graph), consisting of control and data dependencies as well summary dependencies which allow calling-context sensitive slicing. Such slices can be used for the verification of (sequential) non-interference[Was10]. The JOANA program dependence

¹ As explained by Jürgen Graf: “JOANA means Java Object-sensitive ANALysis in case you wondered :)”

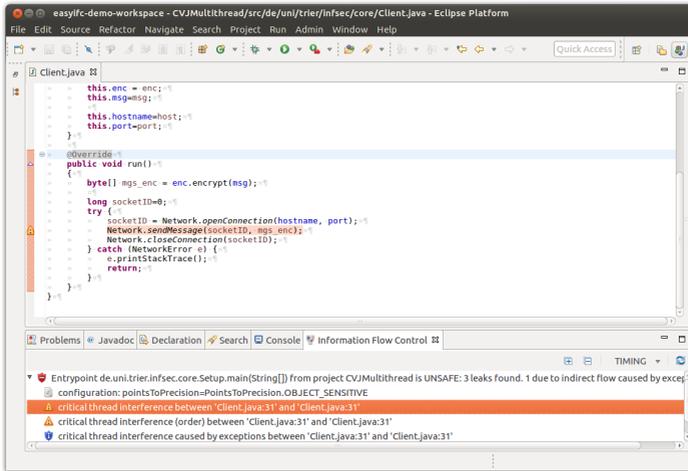


Figure 17.1: The JOANA System

graph for Java programs are described in detail in [Ham09; Gra16]. Slicing Algorithm are described in [Ham09; Gif12]. For concurrent programs, the JOANA program dependence graph is extended with interference dependencies to form the concurrent program dependence graph, as described in the same two theses. Also described there are algorithms for slicing of concurrent programs. The implementation of the LSOD criterion and the computation of precise may-happen-in-parallel information in JOANA is described in [Gif12].

Martin Mohr, Simon Bischof and I implemented the RLSOD criterion and my TIMING criterion based on JOANA dependency graphs. Immediate common ancestors $icda(m_1, m_2)$ and timing dependence $n \rightarrow_{\text{timing}} m$ are implemented as described in Chapter 16 (specifically: they are calling-context insensitive). On the other hand, dependencies due to the concurrent program dependence graph $n \rightarrow_{\text{cpdg}} m$ are computed calling-context sensitively, i.e.: rule CPDG is replaced by an iterated two phase slice (see, e.g., [Gif12]).

17.1 Precision of the TIMING criterion for Java

In a case study in [Bis+18b], we applied the JOANA implementation of the RLSOD criterion to a concurrent Java implementation of a small client/server application. We used the framework for the cryptographic verification of Java Programs due to [KTG12; Küis+14]. In this framework, the concrete implementation of cryptographic primitives (e.g., encryption and decryption) are replaced by *idealized* implementations, in which the content of encrypted messages does *not* depend on the content of plain-text messages. Non-interference of the idealized implementation, together with conventional cryptographic assumptions, then imply computational indistinguishability for the system with a *real* implementation of cryptographic primitives.

In the case study from [Bis+18b], the goal was to prove that a client's choice between two possible plain text messages remains secret to an attacker who observes the network communication between the clients and the server. This is a simplification of a verification goal in e-voting systems. There, the clients choose between two or more alternative candidates, but only the server is allowed to learn the vote (by decryption of the message). The client's choice is modeled as the input to a variable `secret_bit`, resulting in a corresponding user classification $\text{ucl}(\cdot) = H$. In the JOANA system, this is achieved by annotating the corresponding Java code with a `@Source` annotation. We also annotated the *private key* used for decryption of messages with such an annotation. The attackers observation capabilities are modeled by a `@Sink` annotation at the method `Network.sendMessage(...)`.

Using the classification `cl` for the RLSOD criterion inferred by JOANA, and then additional manual inspection, we concluded that any possible information leak must be due to the execution time of the `Encryptor.encrypt()` method. But we could not conclude from RLSOD criterion that the execution time must be independent from the variable `secret_bit`.

On the other hand, the JOANA implementation of the TIMING criterion for the same program reports *no* violation of probabilistic non-interference due to the variable `secret_bit`.² The critical code is

```
for(int i=0; i<msg1.length; ++i) {  
    msg[i] = (secret_bit ? msg1[i] : msg2[i]);  
}
```

in which the plain text message is chosen according to `secret_bit`. The RLSOD criterion must compute $\text{ucl}(n) = H$ for the corresponding control flow graph nodes n in the loop body, and then by rule CDA also conclude $\text{ucl}(m) = H$ also for the node m corresponding to network communication at `Network.sendMessage(...)`. But the TIMING criterion does *not* propagate the classification $\text{cl}_\psi(n) = H$ to the relative timing of any two nodes (m_1, m_2) . To see this, consider the CDA rule for TIMING. There, nodes m_i are *not* timing dependent on any node n in the loop body. Specifically, they are not timing dependent on the choice node implicit in the expression

```
secret_bit ? msg1[i] : msg2[i]
```

² Although just like the RLSOD criterion, it *does* report a violation due to the private encryption key.

17.2 Scalability of the TIMING criterion for Java

In the case study from the previous section, the program consisted of 550 lines of code in 16 classes. Since the JOANA analyses are *whole program* analyses, the code was analyzed together with used components of the Java standard libraries. The total analysis used 2713 megabytes RAM and finished within 2.9 seconds. This includes the computation of the JOANA concurrent program dependence graph, may happen in parallel information, and the computation of the TIMING criterion. The latter took a total of 297 milliseconds, of which the computation of timing dependencies $n \rightarrow_{\text{timing}} m$ required 108 milliseconds, and the computation of the control flow graph “chops” implicit in the TIMING classification rule CDA took 97 milliseconds.

In order to give an idea of the scalability of the TIMING criterion, I also applied it to the *Apache FtpServer*[Fou19]. In addition to the Java standard libraries, the analyzed core of this program makes use of 125818 lines of library source code, and consists of 20645 lines of code³. The total analysis used 87 gigabytes RAM and finished within 10996 seconds ($\approx 3\text{h}03\text{m}$). The computation of the TIMING criterion took a total of 9844 seconds ($\approx 2\text{h}44\text{m}$), of which the computation of timing dependencies $n \rightarrow_{\text{timing}} m$ required 4306 seconds ($\approx 1\text{h}11\text{m}$), while the computation of the control flow graph “chops” implicit in the TIMING classification rule CDA took 5343 seconds ($\approx 1\text{h}29\text{m}$). Within the remaining time (1152 seconds), the computation of the concurrent program dependence graph (including summary edges) graph took 640 seconds, and the may-happen-in-parallel information required 197 seconds.

The *Privacy Crash Cam* is a system developed at the KASTEL Competence Center for Applied Security Technology. I ran the TIMING criterion on its web-service component, which makes available crash-incident related videos from car *dash-cams* to various stakeholders (drivers, other parties involved in an car-accident, law-enforcement).

³ all as counted by the `c1oc` utility[Dan18], ignoring comments

| | total | TIMING | → _{timing} | “chops” | other | → _{cpdg} | MHP |
|--------|----------|----------|---------------------|---------|---------|-------------------|--------|
| apache | 10 996s | 9 844s | 4 306s | 5 343s | 1 152s | 640s | 197s |
| pcc | 121 549s | 103 304s | 25 278s | 77 066s | 18 245s | 9 112s | 3 870s |

Figure 17.2: Run Time of TIMING Based Analysis for Concurrent Java in JOANA

In addition to the Java standard libraries, this web service component makes use of 495607 lines of library source code, and consists of 2894 lines of code. The analysis used 305 gigabytes RAM. See Figure 17.2 for times required for the analysis.

All times in this chapter were measured on a high end “computation server” class PC with an Intel Xeon Gold 6230 CPU at 2.10 GHz base frequency with 512GB RAM.

Summary

- An implementation of the new TIMING criterion for concurrent Java is available in the JOANA system.
- In a simple case study, it improves in precision on the RLSOD criterion of the JOANA system.
- The computation of the TIMING criterion is feasible for a concurrent server application of $\approx 100\,000$ lines of code.

18 Summary and Future Work

Arrakis teaches the attitude of the knife — chopping off what's incomplete and saying: "Now, it's complete because it's ended here."

(Frank Herbert — Dune)

In this thesis, I introduced timing sensitive control dependence $\rightarrow_{\text{tscd}}$ as a natural modification to nontermination sensitive control dependence $\rightarrow_{\text{ntscd}}$. For control flow graphs with unique exit node (also: for all reducible graphs), I can compute $\rightarrow_{\text{tscd}}$ using the generalized control dependence Algorithm 1. For such graphs, timing sensitive postdominance $\sqsupseteq_{\text{TIME[FIRST]}}$ is transitive, and I can compute it by a modification of an algorithm for nontermination sensitive postdominance \sqsupseteq_{MAX} . I used timing sensitive control dependence to support a static information flow analysis for concurrent programs, and an information flow analysis sensitive to timing channels due to micro-architectural effects. For the first analysis, I provided a practical implementation for Java programs in the JOANA system. For the second, I provided a prototype implementation for programs with variables and arrays.

No work is ever truly finished, and mine is no exception. In the following, I want to mention some possible avenues for future work on further improvements of my results.

Calling-Context Sensitivity In the JOANA analysis for concurrent Java programs, neither timing dependence $\rightarrow_{\text{timing}}$, nor the classifications cl_{c} are computed calling-context sensitively. I suspect that

calling-context timing dependence could be achieved, possibly by enriching control flow graphs with timing sensitive “control flow summary edges”: Whenever in a called procedure, timing sensitive postdominance is established between exit and entry node (for some k), equip corresponding call-sites with a “summary” edge with timing cost k . Otherwise, equip call sites with a summary edge with timing cost \top , to be handled appropriately.

Weakening of Timing Sensitivity Timing sensitive control dependence $n \rightarrow_{\text{tscd}} m$ holds unless all successor of a branch node n reach node m after exactly k units of time. Is it possible (and useful?) to consider not only exact matches of time, but also approximate matches? Possibly, by computing intervals $[k_{\min}, k_{\max}]$ of timing costs instead of just one value k ?

“Declassification” of Timing Leaks Pragmatically, a user of timing sensitive analysis will want to ignore timing in selected parts of the program. Perhaps this is possible by allowing him to axiomatically introduce additional edges in the transitive reduction $\langle_{\text{TIME[FIRST]}}$ of timing sensitive postdominance? Note that if this is done at nodes n with no current successor in $\langle_{\text{TIME[FIRST]}}$, it remains a pseudo-tree. Such *timing declassifications* could also be provided by an auxiliary more costly but more precise sound analysis of the relevant program part.

Micro-Architectural / Cache Dependencies for *real* Architectures

In order to be practically useful, (approximate) cache dependencies must be computed for the actual binaries run on actual hardware, using a “usefully realistic” cache model of that hardware. I expect my analysis to work practically unchanged for common instruction set architectures. But it will only be useful together with a precise analysis of memory accesses, i.e.: for each machine operation l , a “as small as possible” static approximation of those memory blocks that may be

accessed by l , which in turn will require a precise analysis of values in, e.g., registers used as offsets in relative address accesses.

New Algorithm for Micro-Architectural dependencies Micro-Architectural dependencies first compute the full (sometimes huge) graph G_α , starting from some initial cache state. Is it possible instead to compute, for each node m , only a relevant part of G_α by going *backwards* from m ?

Approximate Cache Dependencies I introduced approximate cache dependencies only for LRU caches. But I expect similar notions to be possible also for, e.g., *pseudo* LRU caches.

A More Efficient TIMING Criterion As stated, the rule CDA in the TIMING Criterion require the computation of numerous “control flow chops”. I strongly suspect that the explicit computation of these chops can be avoided by use of an appropriate data structure.

Appendices

A Proofs

Very deep. You should send that in to the Reader's Digest. They've got a page for people like you.

(Douglas Adams — The Hitchhiker's Guide to the Galaxy)

A.1 Nontermination (In-)Sensitive Control Dependence in Arbitrary Graphs

Proof of Lemma 3.2.1 on page 20: If $x \neq y$, then $x \sqsupset y \sqsupseteq y$. Also, let $x' \sqsupset \sqsubseteq \sqsupseteq y$. Then $x' \sqsupseteq x$ by transitivity.

For EQ₂[□], I only need to show $\text{ipdom}_{\sqsupseteq}(x) \subseteq \text{ipdom}_{\sqsupseteq}(y)$. Now, for any $x' \in \text{ipdom}_{\sqsupseteq}(x)$, I have $x' \sqsupset \sqsubseteq \sqsupseteq x \sqsupseteq y$. Also, for any y' such that $y' \sqsupset \sqsubseteq \sqsupseteq y$, I have $y' \sqsupset \sqsubseteq \sqsupseteq y \sqsupseteq x$. Because $x' \in \text{ipdom}_{\sqsupseteq}(x)$, I conclude $y' \sqsupseteq x'$, and hence $x' \in \text{ipdom}_{\sqsupseteq}(y)$.

For EQ₃[□], assume $x \neq y$. Then $y \sqsupset x \sqsupseteq z$. Also, if $y' \sqsupset \sqsubseteq \sqsupseteq z$, then $y' \sqsupseteq x \sqsupseteq y$ because $x \in \text{ipdom}_{\sqsupseteq}(z)$. \square

Proof of Lemma 3.2.2 on page 22:

1. $\text{PDF}_{\sqsupseteq}^{\text{local}}(x) \subseteq \text{PDF}_{\sqsupseteq}(x)$ because $x \sqsupseteq x$.
2. $\text{PDF}_{\sqsupseteq}^{\text{up}}(z) \subseteq \text{PDF}_{\sqsupseteq}(x)$ for $x \in \text{ipdom}_{\sqsupseteq}(z)$:
Let $y \in \text{PDF}_{\sqsupseteq}^{\text{up}}(z)$. I have $\neg x \sqsupset \sqsubseteq \sqsupseteq y$ by definition. Also, from $y \in \text{PDF}_{\sqsupseteq}(z)$ I obtain some s such that $y \rightarrow_G s$ and $z \sqsupseteq s$. From $x \sqsupset \sqsubseteq \sqsupseteq z$ and due to reflexivity I am done, because $x \sqsupseteq z \sqsupseteq s$.
3. I show for any x, s such that $(x, s) \in \text{ipdom}_{\sqsupseteq}^*$ that the rule

$$\frac{\neg x \sqsupset \sqsubseteq \sqsupseteq y \quad y \rightarrow_G s \quad x \sqsupseteq s}{y \in \text{PDF}_{\sqsupseteq}^{\text{local}}(x) \vee y \in \bigcup_{\{z \mid x \in \text{ipdom}_{\sqsupseteq}(z)\}} \text{PDF}_{\sqsupseteq}^{\text{up}}(z)}$$

is admissible, by induction on $(x, s) \in \text{ipdom}_{\sqsupseteq}^*$.

I $x = s$.

Let y be as in the rule's premise. Then $y \in \text{PDF}_{\sqsupseteq}^{\text{local}}(x)$ by definition.

II $x \in \text{ipdom}_{\sqsubseteq}(z)$ and $(z, s) \in \text{ipdom}_{\sqsubseteq}^*$ for some $z \in N$.

Let y be as in the rule's premise, i.e.:

$$\neg x \perp\!\!\!\sqsubseteq y \quad y \rightarrow_G s \quad x \sqsubseteq s$$

In order to exploit the induction hypothesis, I want to show:

$$\neg z \perp\!\!\!\sqsubseteq y \quad z \sqsubseteq s$$

In order to obtain a contradiction, assume $z \perp\!\!\!\sqsubseteq y$. From $x \in \text{ipdom}_{\sqsubseteq}(z)$, I also have $x \perp\!\!\!\sqsubseteq z$, i.e.:

$$x \sqsubseteq_{\sqsubseteq} z \sqsubseteq_{\sqsubseteq} y$$

which contradicts $\neg x \perp\!\!\!\sqsubseteq y$, due to the transitivity of \sqsubseteq .

Also, $z \sqsubseteq s$ because $\text{ipdom}_{\sqsubseteq}^* = \sqsubseteq$.

Now, from the induction hypothesis, either:

i. I have $y \in \text{PDF}_{\sqsubseteq}^{\text{local}}(z)$,

in which case I show $y \in \text{PDF}_{\sqsubseteq}^{\text{up}}(z)$ as follows:

First, I conclude $y \in \text{PDF}_{\sqsubseteq}(z)$ as before.

Also, let $x' \in \text{ipdom}_{\sqsubseteq}(z)$. Then $x \sqsubseteq x'$ and $x' \sqsubseteq z$. Assume $x \neq x'$, because otherwise immediately $\neg x' \perp\!\!\!\sqsubseteq y$. If I had $x' \perp\!\!\!\sqsubseteq y$, then:

$$x \sqsubseteq x' \sqsubseteq_{\sqsubseteq} y, \text{ i.e., via reflexivity, } x \sqsubseteq x' \sqsubseteq y,$$

in contradiction with the choice of y .

or

ii. I obtain a node z' such that $z \in \text{ipdom}_{\sqsubseteq}(z')$ and $y \in \text{PDF}_{\sqsubseteq}^{\text{up}}(z')$

In this case, I need to show: $y \in \text{PDF}_{\sqsubseteq}^{\text{up}}(z)$.

- By definition of $\text{PDF}_{\sqsubseteq}^{\text{up}}(z')$, and because $z \in \text{ipdom}_{\sqsubseteq}(z')$, I have $\neg z \perp\!\!\!\sqsubset y$.
- From $y \in \text{PDF}_{\sqsubseteq}^{\text{up}}(z')$ I have $y \in \text{PDF}_{\sqsubseteq}(z')$ by definition. Hence: $z' \sqsubseteq s'$ for some s' such that $y \rightarrow_G s'$. Because $z \in \text{ipdom}_{\sqsubseteq}(z')$, this means that

$$z \sqsubseteq_{\sqsubseteq} z' \sqsubseteq s'$$

This shows $z \sqsubseteq s'$ due to transitivity, and hence: $y \in \text{PDF}_{\sqsubseteq}(z)$.

- It remains to show for arbitrary $x' \in \text{ipdom}_{\sqsubseteq}(z)$: $\neg x' \perp\!\!\!\sqsubset y$. But this follows just as it did in case i.

□

Proof of Lemma 3.2.3 on page 23: Given $y \rightarrow_G x$, I show

$$x \perp\!\!\!\sqsubset y \iff x \in \text{ipdom}_{\sqsubseteq}(y)$$

Assume $x \perp\!\!\!\sqsubset y$, and let x' be any node such that $x' \perp\!\!\!\sqsubset y$. If $x' \neq y$, then $x' \sqsubseteq x$ by $x' \sqsubseteq y$ (transitivity) and rule $\text{CL}^{\rightarrow_G}$. If $x' = y$, let v be some node such that $y \sqsubseteq v \sqsubseteq y$. Then, by rule $\text{CL}^{\rightarrow_G}$ I have $v \sqsubseteq x$, and hence $x' = y \sqsubseteq x$, as well, which proofs $x \in \text{ipdom}_{\sqsubseteq}(y)$.

The reverse implication follows directly from the definitions. □

Proof of Lemma 3.2.4 on page 25: Given $y \in \text{PDF}_{\sqsubseteq}(z)$, I show

$$x \in \text{ipdom}_{\sqsubseteq}(y) \iff \exists x' \in \text{ipdom}_{\sqsubseteq}(z) . x' \perp\!\!\!\sqsubset y$$

The implication \implies is trivial, given the assumption on x .

For the reverse implication \Leftarrow let x' as provided. Because both x, x' are in $\in \text{ipdom}_{\sqsubseteq}(z)$, I obtain $x \sqsubseteq x'$ and then — regardless whether $x = x'$ or not —

$$x \text{ 1-}\sqsubseteq y$$

From this, if also $y \sqsubseteq x$, I immediately conclude $x \in \text{ipdom}_{\sqsubseteq}(y)$ from $\text{EQ}_1^{\sqsubseteq}$. Otherwise, because $\text{ipdom}_{\sqsubseteq}^* = \sqsubseteq$, I obtain via Observation 3.2.1 nodes w, v such that

$w \in \text{ipdom}_{\sqsubseteq}(v)$ and $v \sqsubseteq y$ and $x \sqsubseteq w \sqsubseteq x$ and $\neg v \sqsubseteq x$
but also, because of $\text{EQ}_3^{\sqsubseteq}$

$$x \in \text{ipdom}_{\sqsubseteq}(v)$$

If $z = v$, I conclude $z = y$, because otherwise: $z \sqsupset y \sqsubseteq y$, i.e.: $z \text{ 1-}\sqsubseteq y$, in contradiction to $y \in \text{PDF}_{\sqsubseteq}(z)$. This shows $x \in \text{ipdom}_{\sqsubseteq}(y)$. This is also true if $v = y$.

This means I still have to show $x \in \text{ipdom}_{\sqsubseteq}(y)$ if $z \neq v$ and $y \neq v$. From $y \in \text{PDF}_{\sqsubseteq}(z)$, I obtain a node s such that $y \rightarrow_G s$ and $z \sqsubseteq s$. From $\text{CL}^{\rightarrow_G}$ and $v \sqsubseteq y$ I conclude $v \sqsubseteq s$. Now, I can use NoJoin to infer that either $z \in \text{ipdom}_{\sqsubseteq}(v)$ or $v \in \text{ipdom}_{\sqsubseteq}(z)$. But if $z \in \text{ipdom}_{\sqsubseteq}(v)$, then

$$z \sqsupset v \sqsubseteq y, \text{ i.e.: } z \text{ 1-}\sqsubseteq y$$

in contradiction to $y \in \text{PDF}_{\sqsubseteq}(z)$. If, on the other hand, $v \in \text{ipdom}_{\sqsubseteq}(z)$, specifically: $v \text{ 1-}\sqsubseteq z$, but then $v \sqsubseteq x$ because $x \in \text{ipdom}_{\sqsubseteq}(z)$, in contradiction to the choice of v .

□

A.2 Postdominator Pseudoforests

Proof of Theorem 5.1.1 on page 41: First, I show $\exists_{\text{POST}} \subseteq \nu P$. By the co-induction proof principle, I have to show

$$\exists_{\text{POST}} \subseteq P(\exists_{\text{POST}})$$

So let $m \exists_{\text{POST}} p$. The case $p = m$ is trivial, so assume $p \neq m$. I know that $\forall p \rightarrow_G x. x \exists_{\text{POST}} m$ by definition of $m \exists_{\text{POST}} p$, so: $(m, p) \in P(\exists_{\text{POST}})$.

Now, let me show $\exists_{\text{POST}} \supseteq \nu P$, by assuming $\neg m \exists_{\text{POST}} p$, and showing $(m, p) \notin \nu P$.

There must be some path $p \rightarrow_G^\pi n_x$ to the unique exit node n_x such that $m \notin \pi$. Specifically, $m \neq p$, so I cannot use rule P^{self} to validate $(m, p) \in \nu P$. Let x be the successor of p in π , so that $\pi = p, x, \pi'$. Then I can only validate $(m, p) \in \nu P$ if I can validate $(m, x) \in \nu P$. But I now have a *shorter* path $x \rightarrow_G^{x, \pi'} n_x$ such that $m \notin x, \pi'$, so by iterating I eventually find that I cannot validate $(m, p) \in \nu P$ at all. \square

Proof (Sketch) of Theorem 5.1.2 on page 42: The proof for \exists_{SINK} is similar to the proof of Theorem 5.1.1. For $\mu D \subseteq \exists_{\text{MAX}}$, by the induction proof principle I have to show:

$$D(\exists_{\text{MAX}}) \subseteq \exists_{\text{MAX}}$$

i.e. I have to show:

whenever $\forall p \rightarrow_G x. m \exists_{\text{MAX}} x$ and $p \rightarrow_G^* m$ then also $m \exists_{\text{MAX}} p$

for $m \neq x$, but this follows directly from the definition. For $\exists_{\text{MAX}} \subseteq \mu D$, I assume $m \exists_{\text{MAX}} n$. Let $m \neq n$, because otherwise I have $(m, n) \in \mu D$ by rule D^{self} . Because of $m \exists_{\text{MAX}} n$, all paths from n to m in which m occurs only once (at the end) are cycle-free. Let π be such a path with maximal length among all such paths. Also denote by π_x for each successor x of n a path with maximal length among all

such paths from x to m . Then all π_x are certainly shorter than π , and inductively I can assume $m \sqsupseteq_{\text{MAX}} x$ for all such x . But then $(m, n) \in \mu\text{D}$ by rule D^{succ} . \square

Proof of Lemma 5.2.2 on page 45: If both $v \sqsupseteq z$ and $z \sqsupseteq v$ i'm done by rule $\text{EQ}_1^{\sqsupseteq}$. So by symmetry, it is enough to show

$$\neg z \sqsupseteq v \implies v \in \text{ipdom}_{\sqsupseteq_{\text{MAX}}}(z)$$

So let $>_{\text{MAX}}$ be any transitive reduction of \sqsupseteq_{MAX} , and assume $\neg z \sqsupseteq v$ (i.e., even: $\neg z \sqsupset v$). From $v \sqsupseteq_{\text{MAX}} s$ and because $<_{\text{MAX}}$ is a pseudo forest, i know that there is exactly one sequence

$$v >_{\text{MAX}} \dots >_{\text{MAX}} s$$

such that v appears exactly once. But since $\neg z \sqsupset v$ and $z \sqsupseteq_{\text{MAX}} s$, i know that z must appear in between, i.e. (since $v \neq z$) i have:

$$v >_{\text{MAX}}^* v' >_{\text{MAX}} z >_{\text{MAX}}^* s$$

I also have — because $x \in \text{ipdom}_{\sqsupseteq_{\text{MAX}}}(z)$ — some x' s.t.

$$x' >_{\text{MAX}}^* x >_{\text{MAX}}^* x' >_{\text{MAX}} z$$

But since $<_{\text{MAX}}$ is a pseudo forest, i have $x' = v'$, and also: v must lay on the cycle $x' >_{\text{MAX}}^* x >_{\text{MAX}}^* x'$, i.e.:

$$v' >_{\text{MAX}}^* v >_{\text{MAX}}^* v' >_{\text{MAX}} z$$

\square

A.2.1 Proof of Correctness of the Workset implementation of Algorithm 5

On order to provide — apart from the empirical evidence – proof of correctness, I need some terminology. I want to express that any node

n for which I have not already determined its successor (if any) in \prec_{MAX} either is still in the workset, or will be put there eventually.

Definition A.2.1. Let $\text{workset} \subseteq \text{COND}_G$ and \prec be some pseudo-forest. Then the set of nodes *accessible* from workset is defined as the least fixed point of the rule system

$$\frac{n \in \text{workset}}{n \in \text{workset}^*}$$

$$\frac{w \in \text{workset}^* \quad n \in \text{COND}_G \quad n \rightarrow_G y \quad y \prec^* w}{n \in \text{workset}^*}$$

Also, i define

$$\prec_{\text{workset}} = \prec \cup \left\{ (w, m) \mid w \in \text{workset}^*, m \in \text{ipdom}_{\supseteq_{\text{MAX}}}(w) \right\}$$

and write \prec_{workset}^* for $(\prec_{\text{workset}})^*$.

Proof (Sketch): The algorithm establishes and then upholds the following invariants for \prec (represented by IMDOM) and workset :

$$m \supseteq_{\text{MAX}} n \iff n \prec_{\text{workset}}^* m$$

$$n \prec m \implies m \in \text{ipdom}_{\supseteq_{\text{MAX}}}(n)$$

$$\text{IMDOM}[n] = \perp \wedge \exists m \neq n. m \supseteq_{\text{MAX}} n \implies n \in \text{workset}^*$$

Obviously, upon termination, I have $\prec_{\text{workset}} = \prec$.

Also, the algorithm always terminates: Observe that the choice $a \in \text{lca}_{\prec}(S)$ made by Algorithm 3 is such that \prec -cycles, once established, remain stable, i.e.:

- whenever $\text{IMDOM}[x]$ is about to be updated from some $z \neq \perp$ to z' , i have: $\neg z' <^* x$

Also, once $\text{IMDOM}[x] = z \neq \perp$, future changes of $\text{IMDOM}[x]$ to z' are only possible if x is just outside some $<$ -cycle Z , and z, z' are in that cycle, i.e.: $z <^* z' <^* z$.

But there cannot be an infinite number of such changes, since this would require there to be some nodes $x_1 \neq x_2^1$, $x_i \notin Z$, $\text{ipdom}_{\exists\text{MAX}}(x_i) = Z$ such that

- $\exists x_2 \rightarrow_G y_2. y_2 <^* x_1$
- $\exists x_1 \rightarrow_G y_1. y_1 <^* x_2$

But from $y_2 <^* x_1$ i obtain a path

$$\underbrace{y_2, \dots, x_1}_{\pi_2} \text{ in } G \text{ with } x_1 \notin \pi_2, \text{ and also: } \pi_2 \cap Z = \emptyset,$$

since if I had such a z , I would also have a cycle $\pi_z = z, \dots, z', \dots, z$ in G (for any $z' \in Z$) with $x_1 \notin \pi_z$, because $\neg z <^* x_1$ (and: $\neg z' <^* x_1$) and hence $\neg x_1 \exists_{\text{MAX}} z$ (and $\neg x_1 \exists_{\text{MAX}} z'$). But this contradicts $y_2 <^* x_1$. Similarly, I obtain a path

$$\underbrace{y_1, \dots, x_2}_{\pi_1} \text{ in } G \text{ with } x_2 \notin \pi_1, \text{ and also: } \pi_1 \cap Z = \emptyset.$$

But then, the concatenation of these paths form a G cycle without nodes from Z , in contradiction with $\text{ipdom}_{\exists\text{MAX}}(x_i) = Z$.

□

¹ or: a sequence x_1, \dots, x_n in a similar situaion

A.3 Order Dependence

Proof of Lemma 6.1.2 on page 67: Assume $n \rightarrow_{\text{dod}} (m_1, m_2)$.

- (i) I show $m_2 \exists_{\text{MAX}} m_1$. Let $\pi_1 = m_1, \dots$ be a maximal path. Let n_l be some successor of n in accord with clause (b) of Definition 6.1.2. By definition (clause (a)) I obtain a path $\pi_n = n, n_l, \dots, m_1$, and n_l was chosen s.t. $\neg m_2 \in \pi_n$. Again by clause (a), any extension π of π_n to a maximal path $\pi_n \pi$ must contain m_2 (in it's extension π), and I am done (i.e.: $m_2 \in \pi$) for $\pi = \pi_1$.
- (ii) First, i show $m \exists_{\text{MAX}} m_1$. Assume the opposite. Then I have a maximal path $\pi_1 = m_1, \dots$ with $m \notin \pi_1$. I note that, since $m_2 \exists_{\text{MAX}} m_1$ (see (i)), I have

$$\pi_1 = m_1, \dots, m_2, \pi'_1 \quad \text{with } m \notin \pi'_1 \quad (\text{A.1})$$

From $m \notin \pi_1$ I infer that

$$\forall \pi_n = n, \dots, m_1. m \in \pi_n \quad (\text{A.2})$$

since otherwise I had a maximal path $\pi_n = n, \dots, \pi_1$ with $m \notin \pi_n$, in contradiction to $m \exists_{\text{MAX}} n$. Specifically, since $n \neq m$, given n_l as in clause (b), for any path

$$\pi_l = n_l, \dots, m_1$$

with only one appearance of m_1 , I have $m \in \pi_l$ and $m_2 \notin \pi_l$. Then, given n_r as in clause (c), for any path

$$\pi_r = n_r, \dots, m_2$$

with only one appearance of m_2 I have $m_1 \notin \pi_r$ and $m \notin \pi_r$, since if I had $m \in \pi_r$, I'd have a path $n_r, \dots, m, \dots, m_1$ that does *not* contain m_2 , in contradiction of clause (c).

Now, since $m_1 \exists_{\text{MAX}} n$ and $m_1 \neq n$, *any* maximal extension $\pi = \pi_r \pi'$ to π_r must contain m_1 . But because of (A.2), any such

extension must also contain m , and it must appear in π' , which contradicts (A.1), because $\pi_r \pi'_1$ is a maximal extension to π_r .

$m_1 \sqsupset_{\text{MAX}} m$ then is obvious from (i), due to the fact that $m_1 \neq m_2$, and hence the set $\{m \mid m \sqsupset_{\text{MAX}} m_1\}$ form a cycle (remember that for any transitive reduction \succ_{MAX} of \sqsupset_{MAX} , the graph \prec_{MAX} is a psuedo-forest).

- (iii) I show $\neg m_1 \rightarrow_G^* n$. Assume the opposite. Given n_l as in clause (b), for any path

$$\pi_l = n_l, \dots, m_1$$

with only one appearance of m_1 I have $m_2 \notin \pi_l$. From $m_1 \rightarrow_G^* n$, I infer that any maximal path starting in m_1 must contain m_2 before n since otherwise, I would have a cycle

$$n_l, \dots, m_1, \dots, n, n_l$$

not containing m_2 , contradicting $m_2 \sqsupset_{\text{MAX}} n$ (clause (a)).

But with $m_1 \rightarrow_G^* n$ and $m_2 \rightarrow_G^* m_1$ (see (i)) I also have $m_2 \rightarrow_G^* n$, and likewise infer that any maximal path starting in m_2 must contain m_1 before n . But then, any maximal path starting in m_1 must be of the form

$$\underbrace{m_1, \dots, m_2, \dots, m_1, \dots, m_2, \dots}_{n \notin}$$

contradicting $m_1 \rightarrow_G^* n$.

□

B Nontermination (In-)Sensitive Control Dependence

All your life people will tell you things. And most of the time, probably ninety-five percent of the time, what they'll tell you will be wrong.

(Michael Crichton — The Lost World)

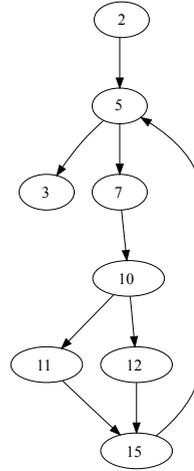
B.1 Analysis of previous Algorithms

Ranganath et al. propose Algorithm 14 for the computation of $\rightarrow_{\text{ntscd}}$ ¹. Their algorithm works by computing — for each $x \in N, n \in \text{COND}_G$ — the set

$$S[x, n] = \{ n \rightarrow_G m \mid x \sqsupseteq_{\text{SINK}} m \}$$

of edges $n \rightarrow_G m$ starting in n such that every maximal path starting in m contains x . The authors also write t_{nm} for the edge $n \rightarrow_G m$.

I added to their algorithm the highlighted parts, which are missing in [Ran+07]. In order to see that these are indeed necessary, consider the graph depicted on the right. Note that it has the unique exit node $n_x = 3$, and that 15 is standard (and hence also both nontermination sensitively and insensitively) control-dependent on 5, so we also expect $5 \rightarrow_{\text{ntscd}} 15$. Specifically, we need $S[15, 5] = \{t_{5,7}\}$.



If the highlighted parts are missing, the sequence $n \in [3, 7, 10, 11, 12, 15, 5]$ is a possible iteration order. I cannot learn that $t_{5,7} \in S[15, 5]$ before I learn that both $t_{10,11}, t_{10,12} \in S[15, 10]$, which I do in the iteration $n = 12$. I do not learn $t_{5,7} \in S[15, 5]$ at that iteration because I did not (and must not!) learn that $t_{5,7} \in S[12, 5]$. But I do not learn $t_{5,7} \in S[15, 5]$ in iteration $n = 15$ (because 15 is not a conditional node, and not $15 \rightarrow_G 15$) or $n = 5$ (trivially), either.

If, on the other hand, in iteration $n = 12$ I put $p = 10$ in the workset, I learn $t_{5,7} \in S[15, 5]$ from $|S[15, 12]| = |\{11, 12\}| = 2$ in a later iteration $n = 10$.

¹ For purposes of comparison only, I temporarily adapt adapt their choice of variable naming.

Input : A CFG $G = (N, E)$

Output: A map NTSCD such that $\text{NTSCD}[n] = \{m \mid n \rightarrow_{\text{ntscd}} m\}$

begin

```

for  $n \in \text{COND}_G, n \rightarrow_G m$  do
  |  $S[m, n] \leftarrow \{t_{nm}\}$ 
  |  $\text{workset} \leftarrow \text{workset} \cup \{m\}$ 
end
while  $\text{workset} \neq \emptyset$  do
  |  $n \leftarrow \text{remove}(\text{workset})$ 
  | if  $\{m \mid n \rightarrow_G m\} = \{m\}$  then
  | | for  $p \in \text{COND}_G, S[n, p] \setminus S[m, p] \neq \emptyset$  do
  | | |  $S[m, p] \leftarrow S[m, p] \cup S[n, p]$ 
  | | |  $\text{workset} \leftarrow \text{workset} \cup \{m\} \cup \{p\}$ 
  | | end
  | | end
  | | if  $\{m \mid n \rightarrow_G m\} = \{m_1, m_2, \dots\}$  then
  | | | for  $m \in N$  do
  | | | | if  $|S[m, n]| = |\{x \mid n \rightarrow_G x\}|$  then
  | | | | | for  $p \in \text{COND}_G \setminus \{n\}, S[n, p] \setminus S[m, p] \neq \emptyset$  do
  | | | | | |  $S[m, p] \leftarrow S[m, p] \cup S[n, p]$ 
  | | | | | |  $\text{workset} \leftarrow \text{workset} \cup \{m\} \cup \{p\}$ 
  | | | | | end
  | | | | end
  | | | end
  | | end
  | end
end
for  $n \in N, m \in \text{COND}_G$  do
  | if  $0 < |S[m, n]| < |\{x \mid n \rightarrow_G x\}|$  then
  | |  $\text{NTSCD}[m] \leftarrow \text{NTSCD}[m] \cup \{n\}$ 
  | end
end

```

end

Algorithm 14: An Algorithm for $\rightarrow_{\text{ntscd}}$. The highlighted parts were missing in [Ran+07], but are crucial for correctness.

Is Algorithm 14 as given here now in fact correct? For now, I can at least make use of a trusted algorithm for $\sqsupseteq_{\text{POST}}$ -control dependence, and validate the second implication of Theorem 4.1.1 empirically:

Observation B.1.1. Let NTSCD be computed via Algorithm 14, and G be a CFG with unique exit node. Then

$$n \rightarrow_{\text{cd}} m \implies m \in \text{NTSCD}[n]$$

Indeed with the previously missing workset updates in place, Algorithm 14 is a standard workset implementation of a least fixed point computation:

Observation B.1.2. Algorithm 14 computes the least fixed point of the monotone functional implicit in the rule system S_4 below:

$$\frac{p \rightarrow_G x}{(p, x) \in S[x, p]} S_4^{\text{suc}} \quad \frac{\{m \mid n \rightarrow_G m\} = \{m\} \quad (p, x) \in S[n, p]}{(p, x) \in S[m, p]} S_4^{\text{lin}}$$

$$\frac{n \neq p \quad |S[m, n]| = |\{y \mid n \rightarrow_G y\}| \quad (p, x) \in S[n, p]}{(p, x) \in S[m, p]} S_4^{\text{cond}}$$

Here, S is understood to be a map $(N \times \text{COND}_G) \rightarrow 2^E$.

Lemma B.1.1. S_4 is *sound*, i.e.

$$S_4 \vdash (p, x) \in S[m, p] \implies m \sqsupseteq_{\text{MAX}} x$$

But S_4 also is *complete* (i.e.: the reversed implication holds for edges $p \rightarrow_G x$) because any fact $m \sqsupseteq_{\text{MAX}} x$ necessarily has a *finite* proof (in S_4).

Proof (Sketch): Soundness follows directly by induction. For completeness, I assume $m \sqsupseteq_{\text{MAX}} p$. For conditional nodes $p \neq m$ this means: $m \sqsupseteq_{\text{MAX}} x$ for all x s.t. $p \rightarrow_G x$. Then either m appears on all the linear segments starting in x — in which case rule S_4^{suc} and S_4^{lin} suffice,

and I stop with a finite proof —, or there exists successors x of p such that from x I can reach the next conditional node p' without visiting m . For such x , I continue to search for proofs that $m \sqsupseteq_{\text{MAX}} p'$. This search must eventually come to an end, because otherwise, I would have found a sequence p, p', \dots, p' of conditional nodes such that m does not necessarily appear on every paths π through p, p', \dots, p' , and I could construct a maximal path π, π, \dots without m , in contradiction of the assumption $m \sqsupseteq_{\text{MAX}} p$.

Ranganath et al. also claim an algorithm for the computation of $\rightarrow_{\text{nticd}}$. I repeat it — together with the fixes w.r.t workset management I proposed for Algorithm 14 — as Algorithm 15. In contrast to Algorithm 14, that algorithm is *not* correct, and there appears to be not obvious fix. First observe that — if we allow the additional highlighted lines, and then again read it as a work-set algorithm that computes the least fixed point of some monotone functional — Algorithm 15 differs from Algorithm 14 only by adding the rule

$$\frac{p \rightarrow_G m \quad (p, x) \in S [p, p]}{(p, x) \in S [m, p]} S_5^{\text{self}}$$

to the rule system S_4 . Let the resulting system be denoted by S_5 .

Input : A CFG $G = (N, E)$

Output: A map NTICD such that $\text{NTICD}[n] = \{m \mid n \rightarrow_{\text{ntscd}} m\}$

begin

for $n \in \text{COND}_G, n \rightarrow_G m$ **do**

$S[m, n] \leftarrow \{t_{nm}\}$

$\text{workset} \leftarrow \text{workset} \cup \{m\}$

end

while $\text{workset} \neq \emptyset$ **do**

$n \leftarrow \text{remove}(\text{workset})$

if $\{m \mid n \rightarrow_G m\} = \{m\}$ **then**

 ... same as Algorithm 14 ...

end

if $\{m \mid n \rightarrow_G m\} = \{m_1, m_2, \dots\}$ **then**

 ... same as Algorithm 14 ...

end

if $|S[n, n]| > 0$ **then**

for $n \rightarrow_G m, m \neq n, S[n, n] \setminus S[m, n] \neq \emptyset$ **do**

$S[m, n] \leftarrow S[m, n] \cup S[n, n]$

$\text{workset} \leftarrow \text{workset} \cup \{m\} \cup \{p\}$

end

end

end

for $n \in N, m \in \text{COND}_G$ **do**

if $0 < |S[m, n]| < |\{x \mid n \rightarrow_G x\}|$ **then**

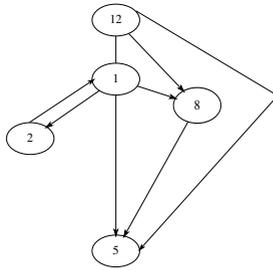
$\text{NTICD}[m] \leftarrow \text{NTICD}[m] \cup \{n\}$

end

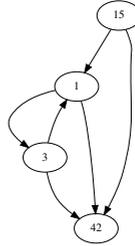
end

end

Algorithm 15: An incorrect Algorithm for $\rightarrow_{\text{nticd}}$. Only the framed part is new (w.r.t. Algorithm 14). The highlighted parts were missing in [Ran+07]. The algorithm is incorrect no matter if these are present or not.



(a) Unsoundness.



(b) \rightarrow_{nticd} requires circular reasoning.

Figure B.1: Problems with Algorithm 15.

Unfortunately, this rule is not even *sound*, i.e. it does *not* hold that

$$S_5 \vdash (p, x) \in S[m, p] \implies m \supseteq_{\text{SINK}} x$$

In order to see this, consider the CFG in Figure B.1a. Note that I expect from a solution S :

$$S[1, 1] = \{(1, 2)\} \quad \text{but not: } (1, 2) \in S[8, 1]$$

because any (sink-bound) path from 2 does contain 1, but there is the sink-bound path 2, 1, 5 from 2 which does *not* contain 8. But rule S_5^{self} allows me to infer just that!

This problem with rule system S_5 does not yet necessarily lead to wrong results *in the resulting map* NTICD. However, a much more serious problem is exemplified in Figure B.1b. Here, Algorithm 15 computes

$$S[42, 1] = \{(1, 42)\} \quad \text{and concludes: } 1 \rightarrow_{nticd} 42$$

while in reality:

$$S[41, 1] = \{(1, 42), (1, 3)\} \quad \text{and} \quad \neg 1 \rightarrow_{nticd} 42$$

To see that in fact $(1,3) \in S [42,1]$ (i.e.: $42 \exists_{\text{SINK}} 3$), just note that $\{42\}$ is the *only* control sink.

Can Algorithm 15 be fixed by adding another rule to S_5 ? To answer this, I explain how I fail to infer $(1,3) \in S [42,1]$ using Algorithm 15 (i.e.: how I fail to infer $S_5 \vdash (1,3) \in S [42,1]$ in the rule system S_5). If I am to infer $S_5 \vdash (1,3) \in S [42,1]$, I need either to infer $S_5 \vdash (1,3) \in S [1,1]$ via rule S_5^{self} , or

$$|S [42,n]| = |\{y \mid n \rightarrow_G y\}| \text{ and } (1,3) \in S [n,1]$$

for some $n \in \text{COND}_G, n \neq 1$. Since (in order to apply rule S_4^{cond}) attempting to infer $S_5 \vdash (1,3) \in S [15,1]$ is pointless, I am left with $n = 3$, and I am required to show $S_5 \vdash (3,42) \in S [42,3]$ (trivial) and $S_5 \vdash (3,1) \in S [42,3]$. Then, I need either to infer $S_5 \vdash (3,1) \in S [3,3]$ via rule S_5^{self} , or

$$|S [42,n']| = |\{y \mid n' \rightarrow_G y\}| \text{ and } (3,1) \in S [n',3]$$

for some $n' \in \text{COND}_G, n' \neq 3$. Since attempting to infer $S_5 \vdash (3,1) \in S [15,3]$ is pointless, I am left with $n' = 1$, and I am required to show $S_5 \vdash (1,42) \in S [42,1]$ (trivial) and $S_5 \vdash (1,3) \in S [42,1]$.

In summary, I need to complete one of the following proof trees:

$$(a) \left\{ \begin{array}{c} \frac{?}{(1,3) \in S [1,1]} \\ \frac{?}{(1,3) \in S [42,1]} \end{array} \quad \cdots \quad \frac{\frac{?}{(3,1) \in S [3,3]}}{(3,1) \in S [42,3]} \quad \cdots \right\} (b)$$

$$\frac{\quad}{(1,3) \in S [42,1]}$$

$$\left. \begin{array}{c}
 \dots \quad \frac{\dots \quad (1,3) \in S [42,1] \quad \dots}{(3,1) \in S [42,3]} \quad \dots \\
 \hline
 (1,3) \in S [42,1]
 \end{array} \right\} (c)$$

The last proof tree (c) is circular, so at most the first two are feasible. But attempting $S_5 \vdash (1,3) \in S [1,1]$, the only sensible rule is S_{4cond} with $n = 3$, which requires me to infer $S_5 \vdash (3,1) \in S [1,3]$ (trivial) and $S_5 \vdash (3,42) \in S [1,3]$, which I cannot because it is false. Similarly, attempting $S_5 \vdash (3,1) \in S [3,3]$, I have to use S_{4cond} with $n = 1$, but this requires me to infer $S_5 \vdash (1,3) \in S [3,1]$ (trivial) and $S_5 \vdash (1,42) \in S [3,1]$, which I cannot.

I have demonstrated how the rule system S_5 is incapable of finitely proving $(1,3) \in S [42,1]$, and thus how Algorithm 15 is incorrect. How can I find a correct algorithm? Refocus on the last “circular” proof tree (c), and recall how I argued for the completeness of S_4 w.r.t \exists_{MAX} : I demonstrated that whenever $m \exists_{MAX} x$, I could find a *finite proof* of this in S_4 . Now, \exists_{SINK} differs from \exists_{MAX} in that it disregards infinite loops such as the loop such as $1 \rightarrow_G 3 \rightarrow_G 1$. In particular, considering *sink-bound* paths starting in x instead of maximal paths starting in x , whenever $\neg m \exists_{SINK} x$, I must be able to find a *finite disproof*, disregarding loops on paths towards x . This suggests that we should believe $(1,3) \in S [42,1]$ since, as considering the circular proof tree (c), it cannot be *finitely disproven*, i.e.: we should define S *co-inductively*, by taking the *greatest* fixed point of some monotone functional (implicit in some rule-system).

Is the rule system S_4 (or: S_5) a suitable rule system for the computation of \rightarrow_{nticd} when read as a *co-inductive* definition of S ? Not quite, since it suffers from three problems.

1. It does *not* enforce that from $(p,x) \in S [m,p]$, it follows that:
 $x \rightarrow_G^* m$.

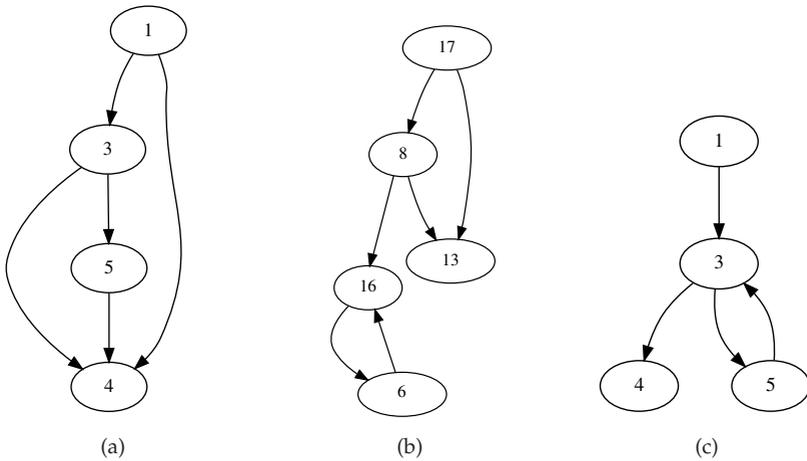


Figure B.2:

2. Following the graph structure (from n to m along $n \rightarrow_G m$) in rules S_4cond and S_4^{in} one step a time is too liberal, allowing too many “self-justifications”.
3. The requirement $n \neq p$ in S_4cond is too strict (while, as i have shown, S_5^{self} , is too liberal).

The problem 1 is demonstrated in the CFG from Figure B.2a. I cannot finitely disprove in S_4 the following assertions, of which only the highlighted is true w.r.t \rightarrow_{nticd} , and all the other assertions $(p, x) \in S[m, p]$ are false because $\neg x \rightarrow_G^* m$.

- $(3, 4), (3, 5) \in S[1, 3]$
- $(1, 3), (1, 4) \in S[1, 1]$
- $(1, 3), (1, 4) \in S[3, 1]$**
- $(3, 4), (3, 5) \in S[3, 3]$

as demonstrated by the following partial derivation trees consistent with S

$$\frac{(1,3), (1,4) \in S [1,1] \quad (3,4) \in S [1,3]}{(3,4) \in S [1,3]}$$

$$\frac{(3,4), (3,5) \in S [1,1] \quad (3,5) \in S [1,3]}{(3,5) \in S [1,3]}$$

$$\frac{(3,4), (3,5) \in S [1,3] \quad \overline{(1,3) \in S [3,1]}^{S_4^{\text{suc}}}}{(1,3) \in S [1,1]}$$

$$\frac{(3,4), (3,5) \in S [1,3] \quad (1,4) \in S [3,1]}{(1,4) \in S [1,1]}$$

$$\frac{(3,4), (3,5) \in S [3,3] \quad (1,4) \in S [3,1]}{(1,4) \in S [3,1]}$$

$$S_4^{\text{suc}} \frac{\overline{(1,3) \in S [3,1]} \quad (1,4) \in S [3,1] \quad (3,4) \in S [1,3]}{(3,4) \in S [3,3]}$$

From such a S , I would have to conclude: $\neg 1 \rightarrow_{\text{nticd}} 3$, when in reality $(1,3) \in S [3,1]$ but *not* $(1,4) \in S [3,1]$, and hence $1 \rightarrow_{\text{nticd}} 3$.

The problem 2 is demonstrated in Figure B.2b. I cannot finitely disprove in S_4 the following false assertions

$$\begin{aligned} (17,8) &\in S [16,17] \\ (17,8) &\in S [6,17] \end{aligned}$$

as demonstrated by the following partial derivation trees consistent with S

$$S_4^{\text{lin}} \frac{(17,8) \in S [6,17]}{(17,8) \in S [16,17]}$$

$$S_4^{\text{lin}} \frac{(17, 8) \in S [16, 17]}{(17, 8) \in S [6, 17]}$$

which is valid even if I — in the light of the *previous* example —, also require

$$(p, x) \in S [m, p] \implies x \rightarrow_G^* m \quad (\text{B.1})$$

But then I *can* disprove, e.g., $(17, 13) \in S [6, 17]$ in “ $S_4 + (\text{B.1})$ ”, and falsely conclude $17 \rightarrow_{\text{nticd}} 6$ (what holds here is: $17 \rightarrow_{\text{nticd}} 8 \rightarrow_{\text{nticd}} 6$).

Note that the problem with rule S_4^{lin} here is that it allows me to “self-validate” the cycle $(17, 8) \in S [6, 17], S [8, 17]$ for paths starting in the edge $17 \rightarrow_G 8$ *without* requiring me to validate that from node 8 (i.e: the “next” conditional after 17 for paths starting in $17 \rightarrow_G 8$), any sink-path reaches that cycle. Similarly, S_4^{lin} allows me to validate an assertion $(p, x) \in S [m, p]$ using *any* other conditional node n (reachable from x), without requiring me to validate

$$|S [m, n']| = |\{y \mid n' \rightarrow_G y\}|$$

for those conditional nodes n' “in between” x and m .

The problem 3 is demonstrate in Figure B.2c. Here, I can disprove the following *true* fact in S_4 :

$$(3, 5) \in S [4, 3]$$

and erroneously conclude $3 \rightarrow_{\text{nticd}} 4$.

B.2 Duality of Nontermination (In-)Sensitivity

Discussing the example in Figure B.1b, I concluded that for the computation of $\rightarrow_{\text{nticd}}$, I need to characterize S co-inductively, rather than inductively. I then proceeded by explaining why the rule set S_4 is not suitable for such a characterization. Now in this section, I will propose a rule system S_3 (and also write S_3 for the corresponding monotone functional) such that

- for $S = \nu S_3$

$$S[m, p] = \{ p \rightarrow_G x \mid m \sqsupseteq_{\text{SINK}} x \}$$

and *at the same time*:

- for $S = \mu S_3$

$$S[m, p] = \{ p \rightarrow_G x \mid m \sqsupseteq_{\text{MAX}} x \}$$

Informally:

nontermination insensitivity = greatest fixed point
 nontermination sensitivity = least fixed point

Before I introduce this system S_3 of rules, I go on a short detour in which I explain this correspondence from the point of view of *safety* and *liveness* properties. It is well-known that liveness properties correspond to least fixed-points, while safety-properties correspond to greatest fixed points. In fact, adopting the standard definitions of *liveness* and *safety* to executions I obtain:

Definition B.2.1 ([AS85]). A set $\mathcal{S} \subseteq N^\omega$ (of sequences called *infinite traces*) is called a *safety* property iff a violation of it is finitely observable and irremediable, i.e. iff whenever $\pi \notin \mathcal{S}$ for some infinite trace π , there already exists some *finite* prefix π_0 of π such that:

- for all infinite traces $\pi' = \pi_0 \dots \in N^\omega$ extending π_0 : $\pi' \notin \mathcal{S}$

Definition B.2.2 ([AS85]). A set $\mathcal{L} \subseteq N^\omega$ (of infinite traces) is called a *liveness* property iff it is always possible (and possibly infinite), i.e. iff for *any* finite trace π_0 , there already exists some *infinite* extension $\pi' = \pi_0 \dots$ of π_0 such that

- $\pi' \in \mathcal{L}$

Note that in this terminology, *traces* are in fact only arbitrary sequences of nodes, not necessarily corresponding to proper paths in some graph G .

Definition B.2.3. I write ${}_x N \subseteq N^\omega$ for the set of infinite traces starting in x . Moreover, let $G = (N, E)$ be some CFG, and $x \rightarrow_G^* m$. Then I define

$$\Pi_{\text{SINK}}^G [x, m] = \{ \pi \in {}_x N \mid \neg (\pi \text{ has a prefix } \pi_0 = x \dots n \text{ with } m \notin \pi_0, \neg n \rightarrow_G^* m) \}$$

$$\Pi_{\text{MAX}}^G [x, m] = \{ \pi \in {}_x N \mid \neg (\hspace{15em} m \notin \pi \hspace{15em}) \}$$

Intuitively, any trace $\pi \notin \Pi_{\text{SINK}}^G [x, m]$ is a counterexample to the claim that node x nontermination insensitively postdominates node m , while any maximal trace $\pi \notin \Pi_{\text{MAX}}^G [x, m]$ is a counterexample to the claim that node x nontermination sensitively postdominates node m . Also intuitively, $\Pi_{\text{SINK}}^G [x, m]$ is a safety property, since I can just expose the prefix π_0 obtained by definition. $\Pi_{\text{MAX}}^G [x, m]$ is trivially a liveness property, since any finite sequence of nodes can be extended to include m (what's not clear is whether this can be done with a proper *path* in some graph G , i.e.: whether this liveness property holds in G).

In order to establish this formally, I make a technical modification to the underlying graph G .

Observation B.2.1. Let $G_0 = (N, E)$ be some CFG, $x, m \in N$ some nodes, $x \rightarrow_G^* m$, and define

$$G = (N, E \cup \{ (n_x, n_x) \mid n_x \text{ is an exit node} \})$$

Then the maximal paths in G are exactly the infinite paths, and

1. $\Pi_{\text{SINK}}^G[x, m]$ is a safety property.
2. $\Pi_{\text{MAX}}^G[x, m]$ is a liveness property.
3. $\forall \pi \in {}_x\Pi_{\text{SINK}}. m \in \pi \iff {}_x\Pi^G \cap N^\omega \subseteq \Pi_{\text{SINK}}^G[x, m]$
4. $\forall \pi \in {}_x\Pi_{\text{MAX}}. m \in \pi \iff {}_x\Pi^G \cap N^\omega \subseteq \Pi_{\text{MAX}}^G[x, m]$

Proof (Sketch): 1. Let $\pi \notin \Pi_{\text{SINK}}^G[x, m]$. The case $\pi \notin {}_xN$ is trivial, so i can assume some prefix $\pi_0 = x \dots n$ of π s.t. $m \notin \pi_0$ and $\neg n \rightarrow_G^* m$. But this then also holds for every π' extending π_0 , so $\pi' \notin \Pi_{\text{SINK}}^G[x, m]$.

2. This is trivially true.
3. First, I show the implication \Rightarrow .

Let $\pi = x \dots$ be some infinite path starting in x . If it is a sink path, then $m \in \pi$, hence $\pi \in \Pi_{\text{SINK}}^G[x, m]$. Otherwise, let let $\pi_0 = x \dots n$ of π be any prefix of π . I can always extend this in G to *obtain* a sink-bound path π' starting with π_0 . But then $m \in \pi'$, so either $m \in \pi_0$, or $n \rightarrow_G^* m$.

Turning to the implication \Leftarrow , assume ${}_x\Pi^G \cap N^\omega \subseteq \Pi_{\text{SINK}}^G[x, m]$, and $\pi \in {}_x\Pi_{\text{SINK}}$. Let $\pi_0 = x \dots n$ be some prefix of π s.t. $n \in S$, with S being the sink that π is bound for. Then either $m \in \pi_0$, and I am done, or $n \rightarrow_G^* m$. But then $m \in S$, and $m \in \pi$, because a sink-bound path visits every $m \in S$ infinitely often.

4. Trivial, given that the maximal paths in G are exactly the infinite paths.

But, again informally,

$$\begin{aligned} \text{liveness} &= \text{least fixed point} \\ \text{safety} &= \text{greatest fixed point} \end{aligned}$$

which is made somewhat more explicit in the setting of the modal μ calculus (see, e.g., [BS07]). So the fact that $\rightarrow_{\text{nticd}}$ is obtainable via greatest fixed points, and $\rightarrow_{\text{ntscd}}$ via least fixed points is — *a posteriori* — perhaps not so surprising, while perhaps the fact that it is obtained from a single functional is.

Remark B.2.1. My characterization of \exists_{MAX} as liveness-, and \exists_{SINK} as safety property is unusual, since the definitions involve the predicate $n \rightarrow_G^* m$, which is a property of the graphs G structure, while liveness-, safety properties usually only involve predicates on the graphs nodes (i.e.: states).

Note also that i did *not* characterize $\rightarrow_{\text{nticd}}$ or $\rightarrow_{\text{nticd}}$ as trace properties. They presumably are not trace properties, but *2-hyper-properties*[CS10].

Returning form this diversion, I will now finally define the rules S_3 . First, I need a concept already implicit in the argument for Lemma B.1.1.

Definition B.2.4. Given any node x , it's *next conditional node*

$$\text{next}_{\text{COND}} [x] = \epsilon n. n \in \text{COND}_G, x \rightarrow_G^\pi n, \pi \cap \text{COND}_G = \{n\}$$

is *the* (if any!) unique conditional node n that is reachable from x without first reaching another conditional node.

The set of nodes on the path to the next conditional node from x (or just the nodes linearly following x if there is no such node) is:

$$\overline{\text{next}_{\text{COND}}} [x] = \begin{cases} \{ m \mid x \rightarrow_G^\pi n, \pi \text{ cycle-free}, m \in \pi \} & \text{next}_{\text{COND}} [x] = n \\ \{ m \mid x \rightarrow_G^* m \} & \text{next}_{\text{COND}} [x] = \perp \end{cases}$$

Similarly,

$$\text{prev}_{\text{COND}}[x] = \{ p \mid p \in \text{COND}_G, p \xrightarrow{\pi}_G x, \pi \cap \text{COND}_G = \{p\}, \pi \neq x \}$$

is the set of set of conditional nodes p that can reach x (in at least one step) without first reaching another conditional node.

Observation B.2.2. Note that for conditional nodes $x \in \text{COND}_G$

- $\text{next}_{\text{COND}}[x] = x$, but
- *not necessarily* $x \in \text{prev}_{\text{COND}}[x]$.

Remark B.2.2. Both $\text{next}_{\text{COND}}$, $\overline{\text{next}_{\text{COND}}}$ and $\text{prev}_{\text{COND}}$ can easily and (simultaneously) be computed in $\mathcal{O}(|\text{COND}_G| \times |N|)$ steps, but in practice usually not much more than $\mathcal{O}(|\text{COND}_G|)$.

Now, I can state S_3 :

Definition B.2.5. With S_3 I denote both the rule system below, and the corresponding monotone functional.

$$\frac{p \xrightarrow{G} x \quad m \in \overline{\text{next}_{\text{COND}}}[x]}{(p, x) \in S[m, p]} S_3^{\text{lin}}$$

$$\frac{p \xrightarrow{G} x \quad n = \text{next}_{\text{COND}}[x] \quad x \xrightarrow{G^*} m \quad |S[m, n]| = |\{y \mid n \rightarrow y\}|}{(p, x) \in S[m, p]} S_3^{\text{cond}}$$

Here, again, S is understood to be a map $(N \times \text{COND}_G) \rightarrow 2^E$.

Observation B.2.3. With regard to the *least* fixed point, S_3 is equivalent to S_4 :

$$\mu S_3 = \mu S_4$$

Remark B.2.3. This is also obvious: S_3^{lin} subsumes S_4^{suc} and S_4^{lin} , but limits its applicability to nodes n in *regions* directly behind x . But this is no limitation w.r.t μS_4 , since there, any finite derivation along linear segments of G starting with an edge $p \rightarrow_G x$ at S_4^{suc} and continuing with S_4^{lin} have to end at the next conditional node $\text{next}_{\text{COND}}[x]$, anyway.

With regard to S_3^{lin} , first note that the requirement $m \rightarrow_G^* x$ is vacuously true in the least fixed point of S_4 . Also, w.r.t least fixed points of S_4 , dropping the requirement $n \neq p$ does not change anything, since for $n = p$, this rules then *requires* $(p, x) \in S[m, p]$ in order to derive $(p, x) \in S[m, p]$.

Moreover, for $n = \text{next}_{\text{COND}}[x]$, I immediately have $(p, x) \in S[n, p]$ via S_3^{lin} , so I can drop that requirement in S_3^{cond} .

Finally, $S = \mu S_4$ has the property that whenever I can proof $(p, x) \in S[m, p]$ with $S_4\text{cond}$, then I can already proof

$$|S[m, n]| = |\{y \mid n \rightarrow y\}|$$

for all conditional nodes n reachable from x (without using $(p, x) \in S[m, p]$), so limiting myself to $n = \text{next}_{\text{COND}}[x]$ does not inhibit my power. Also, for $S = \mu S_4$ I never invoke $S_4\text{cond}$ with n unreachable from x .

B.3 New Algorithms

In this section, I will provide the — to the best of my knowledge — first correct algorithm for the computation of $\rightarrow_{\text{ntcd}}$. Before I proof that it really is enough to compute the *greatest* fixed point of S_3 (Theorem B.3.1), I will first provide the Algorithm 16.

Observation B.3.1. Let S be computed by Algorithm 16. Then

$$S = \nu S_3$$

Remark B.3.1. The initialization is correct because for *any* S consistent with S_3 , $(p, x) \in S [m, p']$ implies $p = p'$, $x \rightarrow_G^* m$ and $p \rightarrow_G x$ — hence, I need not start the iteration as high as $\top = (p, x) \mapsto E$. The work-set is managed correctly, since any refutation of $(n, y) \in S [m, n]$ may invalidate at most those assertions $(p, x) \in S [m, p]$ for which $n = \text{next}_{\text{COND}} [x]$ (see S_3^{cond}), but then it is enough to re-inspect those (m, p) s.t. $p \in \text{prev}_{\text{COND}} [n]^2$. In fact, I only need to re-inspect (m, p) the *first* time i refute *any* $(n, y) \in S [m, n]$. So in Algorithm 16, I can strengthen the **highlighted** check $S'_{mp} \subset S [m, p]$ by replacing it with

$$|S [m, p]| = |\{x \mid p \rightarrow_G x\}| \quad \wedge \quad S'_{mp} \subset S [m, p]$$

in order to reduce the number of iterations.

Incidentally, since the relevant lattice $(\sqsubseteq, N \times \text{COND}_G \hookrightarrow 2^E)$ is constructed from subset-lattice $(\subseteq, 2^E)$, the greatest fixed point computation $S = \nu S_3$ can be replaced by a least fixed computation on the *dual* (monotone) functional

$$\overline{S_3}(S) [m, p] = \overline{S_3} \left((m', p') \mapsto \overline{S [m', p']} \right) [m, p] \quad (\text{B.2})$$

² note the reversal of varibale names n, p w.r.t the algorithms source code

Input : A CFG $G = (N, E)$
Input : Maps $\text{next}_{\text{COND}}$ and $\overline{\text{next}_{\text{COND}}}$
Output : A map NTICD such that $\text{NTICD}[n] = \{m \mid n \rightarrow_{\text{nticd}} m\}$
Data: A map S from pairs of nodes to sets of nodes
Data: A workset $\text{workset} \subseteq N \times \text{COND}_G$
Notation: $|n|$ for $|\{x \mid n \rightarrow_G x\}|$

```

begin
  for  $p \in \text{COND}_G, p \rightarrow_G x, x \rightarrow_G^* m$  do
    |  $S[m, p] \leftarrow S[m, p] \cup \{(p, x)\}$ 
  end
  for  $p \in \text{COND}_G, p \rightarrow_G^* m$  do
    |  $\text{workset} \leftarrow \text{workset} \cup \{(m, p)\}$ 
  end
  while  $\text{workset} \neq \emptyset$  do
     $(m, p) \leftarrow \text{remove}(\text{workset})$ 
     $S'_{mp} \leftarrow \{(p, x) \mid p \rightarrow_G x, m \in \overline{\text{next}_{\text{COND}}}[x]\}$ 
     $\cup \{(p, x) \mid p \rightarrow_G x, n = \text{next}_{\text{COND}}[x], |S[m, n]| = |n|\}$ 
    assert  $S'_{mp} \subseteq S[m, p]$ 
    if  $S'_{mp} \subset S[m, p]$  then
      |  $\text{workset} \leftarrow \text{workset} \cup \{(m, n) \mid n \in \text{prev}_{\text{COND}}[p]\}$ 
    end
  end
end
for  $n \in N, m \in \text{COND}_G$  do
  if  $0 < |S[m, n]| < |n|$  then
    |  $\rightarrow_{\text{nticd}}[m] \leftarrow \rightarrow_{\text{nticd}}[m] \cup \{n\}$ 
  end
end
end
    
```

Algorithm 16: A correct algorithm for $\rightarrow_{\text{nticd}}$. The highlighted check can be optimized.

where for sets $X \subseteq \{(p, x) \mid p \rightarrow_G x\}$ I denote with \overline{X} its complement w.r.t. $\{(p, x) \mid p \rightarrow_G x\}$. If I extend this notion of complement to maps $S : N \times \text{COND}_G \hookrightarrow 2^E$ via

$$\overline{S}[m, p] = \overline{S[m, p]}$$

then definition (B.2) just reads:

$$\overline{S}_3(S) = \overline{S_3(\overline{S})}$$

The following lemma is easily shown along the lines of, e.g. [Fri02], Lemma 20.9.

Lemma B.3.1. The greatest fixed point of S_3 is dual to the least fixed point of \overline{S}_3 , i.e.:

$$\nu S_3 = \overline{\mu \overline{S}_3}$$

The rule system corresponding to \overline{S}_3 is:

$$\frac{p \rightarrow_G x \quad \neg x \rightarrow_G^* m}{(p, x) \in S[m, p]} \overline{S}_3^{\text{unreach}}$$

$$\frac{p \rightarrow_G x \quad \neg m \in \overline{\text{next}_{\text{COND}}}[x] \quad n = \text{next}_{\text{COND}}[x] \quad |S[m, n]| \neq 0}{(p, x) \in S[m, p]} \overline{S}_3^{\text{cond}}$$

To understand the following Lemma, recall the safety-property

$$\{\pi \in {}_x N \mid \neg(\pi \text{ has a prefix } \pi_0 = x \dots n \text{ such that } m \notin \pi_0, \neg n \rightarrow_G^* m)\}$$

from Definition B.2.3 that characterized $m \sqsupseteq_{\text{SINK}} x$.

Lemma B.3.2. For $S = \mu\overline{S}_3$,

$$(p, x) \in S[m, p] \iff \exists n, \pi_0. m \notin \pi_0 \wedge x \rightarrow_G^{\pi_0} n \wedge \neg n \rightarrow_G^* m$$

Proof (Sketch): For the implication \implies , proceed by induction on the derivation of $(p, x) \in S[m, p]$. Case $\overline{S}_3^{\text{unreach}}$ is trivial. For case $\overline{S}_3^{\text{cond}}$, from $|S[m, n]| \neq 0$ obtain some $n \rightarrow_P y$ s.t. — from the induction hypothesis — there exists n', π'_0 with $y \rightarrow_G^{\pi'_0} n'$, $m \notin \pi'_0$, but *not*: $n' \rightarrow_G^* m$. But then immediately for

$$\pi_0 = x, \underbrace{\dots}_{\text{linear}}, n, y, \pi'_0$$

I have: $m \notin \pi_0$, and *not*: $n' \rightarrow_G^* m$.

For the implication \impliedby , assume $m \notin \pi_0 \wedge x \rightarrow_G^{\pi_0} n \wedge \neg n \rightarrow_G^* m$. In all but the trivial cases, let $p' \in \text{prev}_{\text{COND}}[n]$ be the conditional node immediately preceding n in π_0 , and x' the node succeeding that occurrence of p' in π_0 . Then — since there are no conditional nodes in G between p' and n — also $\neg p' \rightarrow_G^* m$, and hence $(p', x') \in S[m, p']$ by rule $\overline{S}_3^{\text{unreach}}$. But then $(p, x) \in S[m, p]$ by following the path π_0 backwards, applying $\overline{S}_3^{\text{cond}}$ at each segment $p'' \rightarrow_G x'', n'' = \text{next}_{\text{COND}}[x'']$.

I immediately obtain:

Theorem B.3.1. The rule system S_3 is correct w.r.t it's greatest fixed point, and $\rightarrow_{\text{nticd}}$. Specifically, let $S = \nu S_3$. Then

$$(p, x) \in S[m, p] \iff x \exists_{\text{SINK}} m$$

Proof: The proof that

$$\exists n, \pi_0. m \notin \pi_0 \wedge x \rightarrow_G^{\pi_0} n \wedge \neg n \rightarrow_G^* m \iff \neg x \exists_{\text{SINK}} m$$

is just as the proof that $\Pi_{\text{SINK}} [x, m]$ characterizes $x \sqsupseteq_{\text{SINK}} m$ in Observation B.2.1. But then the theorem follows immediately from Lemma B.3.2 and Lemma B.3.1. \square

Remark B.3.2. A machine checked proof of Theorem B.3.1 using the Isabelle/HOL proof assistant[NWP02] was prepared by Simon Bischof. Bischofs proof does not invoke the dual functional \bar{S}_3 , and directly invokes co-inductive proof principles (while my proof of Lemma B.3.2 is purely inductive).

Observation B.3.2. Let S be computed by Algorithm 22. Then for all relevant m, p :

$$(\nu S_3) [m, p] = \overline{S [m, p]}$$

C A Slicing Algorithm using C-Edges

I don't even know what I was running for —
I guess I just felt like it.

(J.D. Salinger — The Catcher in the Rye)

Whenever a CFG G has no non-trivial control-sink, the corresponding pseudo-forests \prec_{SINK} are *proper forests*¹. Then by Observation 6.7.4,

$$\rightarrow_{\text{nticd}}^G = \emptyset$$

and I can obtain control slices for arbitrary M from $\rightarrow_{\text{nticd}}^G$. In the special case that there is only one exit node n_x , I have $\rightarrow_{\text{nticd}}^G = \rightarrow_{\text{cd}}^G$, and by an algorithm[SG95] based on *DJ-Graphs*, any backward-slice

$$\left(\rightarrow_{\text{nticd}}^G \right)^* (M) \tag{C.1}$$

can be computed in linear time from \prec_{SINK} without the construction of $\rightarrow_{\text{nticd}}^G$.

In this section, I will demonstrate that the technique from [SG95] to compute Equation C.1 can be generalized to *arbitrary* CFG with (possibly) *multiple* exit nodes n_x, n'_x, \dots and control-sinks.²

The original algorithm is presented with the application of placing ϕ -nodes in CFG with unique entry node n_e . It presupposes the *dominance-tree*, and computes the iterated dominance frontiers of sets M of variable definitions by considering *join-edges* $n \rightarrow_J m$, where

$$\begin{aligned} n \rightarrow_J m &\Leftrightarrow n \rightarrow_G m \wedge \neg n \sqsupset_{\text{DOM}} m \\ &\Leftrightarrow n \rightarrow_G m \wedge n \neq \text{idom}(m) \end{aligned}$$

¹ i.e.: \prec_{SINK} is free of cycles

² but I make no formal claim about the worst-case running time of this generalization

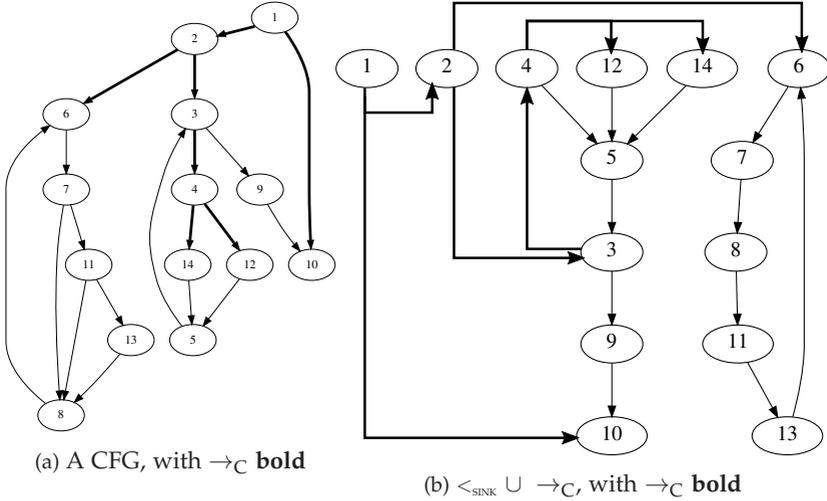


Figure C.1: Conditional Edges \rightarrow_C

In contrast, my application is backward-control-slicing in arbitrary CFG. I presuppose the pseudo forest \prec_{SINK} , and compute the iterated *postdominance* frontiers of sets M of slicing criteria by considering *conditional-edges* $n \rightarrow_C m$, where³

$$n \rightarrow_C m \iff n \rightarrow_G m \wedge m \notin \text{ipdom}_{\sqsupseteq_{\text{SINK}}}(n)$$

That is, my setting is obtained from the original by

1. as usual, considering CFG with edges \rightarrow_G flipped, and then
2. generalizing to arbitrary CFG, by going from $\sqsupseteq_{\text{POST}}$ to $\sqsupseteq_{\text{SINK}}$.

³ it is *not*, in general, the case that $n \rightarrow_C m \iff n \rightarrow_G m \wedge \neg m \sqsupseteq_{\text{SINK}} n$

I recall the example CFG from Figure 5.1a and the corresponding pseudo-forest \prec_{SINK} in Figure C.1. Conditional edges $n \rightarrow_C m$ are **bold** in Figure C.1a, and have also been added to Figure C.1b.

In order to derive a generalized algorithm for the computation of Equation C.1, I will now present the appropriate generalizations of the relevant properties from [SG95]. I will require the following notation:

Definition C.0.1. The set of \exists_{SINK} -ancestors of m is

$$\exists_{\text{SINK}} \uparrow^+ (m) = \{ m'' \mid m'' \exists_{\text{SINK}} m', m' \in \text{ipdom}_{\exists_{\text{SINK}}} (m) \}$$

i.e., for a transitive, reflexive reduction $>_{\text{SINK}}$ of \exists_{SINK}

$$= \{ m'' \mid m <_{\text{SINK}}^+ m'' \}$$

similarly:

$$\exists_{\text{SINK}} \uparrow^* (m) = \{ m'' \mid m'' \exists_{\text{SINK}} m \}$$

$$= \{ m'' \mid m <_{\text{SINK}}^* m'' \}$$

The following observation is to be understood as a generalization of the (key) Lemma 3.1 in [SG95]:

Observation C.0.1. Let G be any CFG, and $n \neq m$. Then

$$n \rightarrow_{\text{nticd}} m \iff n \rightarrow_C m_0 \text{ and } m \exists_{\text{SINK}} m_0 \text{ and } \neg m \exists_{\text{SINK}} n$$

for some node m_0 . With regard to any corresponding pseudo-forest \prec_{SINK} :

$$n \rightarrow_{\text{nticd}} m \iff n \rightarrow_C m_0 <_{\text{SINK}}^* m \text{ and } \neg n <_{\text{SINK}}^* m$$

for some node m_0 .

Up to the flipping of edges in G , the authors in [SG95] assume a unique exit node n_x , and hence can assume \prec_{SINK} to be a proper tree, which allows them to presuppose for each node n it's level $\text{lvl}(n)$, defined as its distance from the root.

Lemma C.0.1 ([SG95], Lemma 3.1, after flipping edges in G). Let G be a CFG with unique exit node n_x , and $n \neq m$. Then $n \rightarrow_{cd} m$ iff there exists a node m_0 in the sub-tree rooted in m with conditional edge $n \rightarrow_C m_0$ and $lvl(n) \leq lvl(m)$.

So in my generalization,

$$\begin{array}{ll} m_0 \text{ in the sub-tree rooted in } m & \text{became } m_0 \prec_{SINK}^* m \\ lvl(n) \leq lvl(m) & \text{became } \neg n \prec_{SINK}^* m \end{array}$$

By Observation C.0.1, I can compute

$$(\rightarrow_{nticd})^*(M)$$

by following $(\rightarrow_C \cup \prec_{SINK})$ -paths backwards from nodes $m \in M$, disregarding reached nodes n in accord with the condition $\neg n \prec_{SINK}^* m$. In the special case of a unique exit node n_x , there is an efficient algorithm to do this because by keeping a (never increasing) *current tree level* l ,

1. no candidate node n ever needs to be visited twice (even if there paths from n to multiple nodes $m, m', \dots \in M$),
2. the algorithm only ever needs to consider nodes n such that $lvl(n) \leq l$, which — unlike $\neg n \prec_{SINK}^* m$ — is a $\mathcal{O}(1)$ check!

So in order to devise an algorithm for the general case, I need a generalized notion of the “pseudo-forest level” of nodes n that allows me to replace checks $\neg n \prec_{SINK}^* m$ by a comparison of pseudo-forest levels.

Recall that in pseudo-forests \prec , roots are either single nodes $\{r\}$, or \prec -cycles R . So the following definition generalizes the notion of tree-level:

Definition C.0.2. For any pseudo forest \prec and any node n , let $R_n \subseteq N$ be the root of the pseudo-tree of n . Then I define

$$lvl^{\prec}(n) = \min_{\pi=n, \dots, r} \min_{r \in R_n} |\pi| \geq 1$$

I also define the *level of the <-successor of n*:

$$\text{lvl}_{\text{next}}^{\lt}(n) = \begin{cases} \text{lvl}(n') & \text{if } n \lt n' \\ 0 & \text{if there is no such } n' \end{cases}$$

Note that for “most” nodes, $\text{lvl}_{\text{next}}^{\lt}(n) = \text{lvl}^{\lt}(n) - 1$, i.e.:

$$\text{lvl}_{\text{next}}^{\lt}(n) < k \iff \text{lvl}^{\lt}(n) \leq k$$

In a backward-traversal of $(\rightarrow_C \cup \lt_{\text{SINK}})$, I obviously cannot “jump trees by” following \lt_{SINK} . Also, as a consequence of the following Observation C.0.2, whenever I reach a candidate node n by backwards-following a conditional edge $n \rightarrow_C m_0 \lt_{\text{SINK}}^* m$:

1. n cannot be in any proper \lt_{SINK} -cycle (and specifically: $\neg n \in R_m$)
2. if n is in a different pseudo tree than m_0 and m , then $\{n\}$ is a root, i.e.: $\neg n \lt_{\text{SINK}}^+ m'$ for *any* node m' , and $\text{lvl}^{\lt_{\text{SINK}}}(n) = 1$.
3. if $n \neq m$ is in the same pseudo-tree as m_0 and m , then $\neg n \lt_{\text{SINK}}^* m$ iff $\text{lvl}_{\text{next}}^{\lt_{\text{SINK}}}(n) < \text{lvl}^{\lt_{\text{SINK}}}(m)$.

But this just means that regardless whether the edge $n \rightarrow_C m_0$ jumps trees or not, I can simply test $\text{lvl}_{\text{next}}^{\lt_{\text{SINK}}}(n) < \text{lvl}^{\lt_{\text{SINK}}}(m)$.

Observation C.0.2. Let G be any CFG with corresponding pseudoforest \lt_{SINK} . Then whenever $n \rightarrow_C m$,

$$\begin{aligned} \exists_{\text{SINK}} \uparrow^*(m) \supset \exists_{\text{SINK}} \uparrow^+(n) \\ \text{and} \quad m \notin \exists_{\text{SINK}} \uparrow^+(n) \end{aligned}$$

Note that the inequation is strict. For n that do have *some* immediate \exists_{SINK} -dominator, this can be understood to mean that conditional edges $n \rightarrow_C m$ “never advance toward the root of n ”, but remain in the same tree. If n does not have any immediate \exists_{SINK} -dominator, Ob-

servation C.0.2 makes no restriction, and $n \rightarrow_C m$ may “jump between pseudo-trees”.

In the example Figure C.2 on page 393, I show a two-tree pseudo-forest $<_{\text{SINK}}$ and highlight for some n the set M_n of nodes m that “are allowed” by Observation C.0.2.

In Algorithm 17, I write just lvl for $lvl^{<_{\text{SINK}}}$ and lvl_{next} for $lvl_{\text{next}}^{<_{\text{SINK}}}$. Unlike the original algorithm for the special case with unique exit node n_x , my general algorithm cannot proceed with a never-increasing global *current level*, because I will sometimes discover new trees n via an edge $n \rightarrow_C m_0$. Instead, I hold fast to a current node m — selected from the priority-queue Q — whose sub-tree I explore, noting edges $n \rightarrow_C m_0$ *en passant*. The check $m'_0 \notin \text{Visited}$ ensures that I never explore a node more than once.

Observation C.0.3. Let G be any CFG, M any set of nodes, and N be computed by Algorithm 17. Then

$$N = (\rightarrow_{\text{nticd}})^* (M)$$

I derived Algorithm 17 by generalizing the algorithm from [SG95] for classical (post-)dominance, which is a nontermination-insensitive notion. But Algorithm 17 works just as well for nontermination sensitive postdominance:

Observation C.0.4. Let G be any CFG, M any set of nodes, and N be computed by Algorithm 17, with input $>_{\text{MAX}}$ instead of $>_{\text{SINK}}$. Then

$$N = (\rightarrow_{\text{ntscd}})^* (M)$$

Closing this section, let me note that since control-sinks do not contribute to $(\rightarrow_{\text{nticd}})^*$, it was to be expected that the original algorithm could be generalized to pseudo-trees. What may not have been obvious is my generalization to (pseudo)-forests. I also fully expect that the algorithm does not only work for $\sqsubseteq_{\text{SINK}}$ and \sqsubseteq_{MAX} , but for a general class of “postdominance like” relations \sqsubseteq . It would be interesting

Input : A set M of nodes (the slice criteria)
Input : A transitive reduction \succ_{SINK} of \exists_{SINK}
Input : The conditional edges \rightarrow_C
Output: The set $(\rightarrow_{\text{nticd}})^*(M)$

```

N ← M
Q ← M
Visited ← ∅
while Q ≠ ∅ do
  m ← remove(Q) s.t. lvl(m) = maxm∈Q lvl(m)
  if m ∉ Visited then
    M0 ← {m}
    while M0 ≠ ∅ do
      m0 ← remove(M0)
      for n →C m0, n ∉ N, lvlnext(n) < lvl(m) do
        N ← N ∪ {n}
        Q ← Q ∪ {n}
      end
      for m'0 <SINK m0, m'0 ∉ Visited do
        M0 ← M0 ∪ {m'0}
      end
      Visited ← Visited ∪ {m0}
    end
  end
end
return N

```

Algorithm 17: Computation of $(\rightarrow_{\text{nticd}})^*(M)$ via \rightarrow_C

to check if, for example, the conditions for *admitting an efficient PDF partitioning* (Definition 3.2.7 on page 25) are already sufficient.

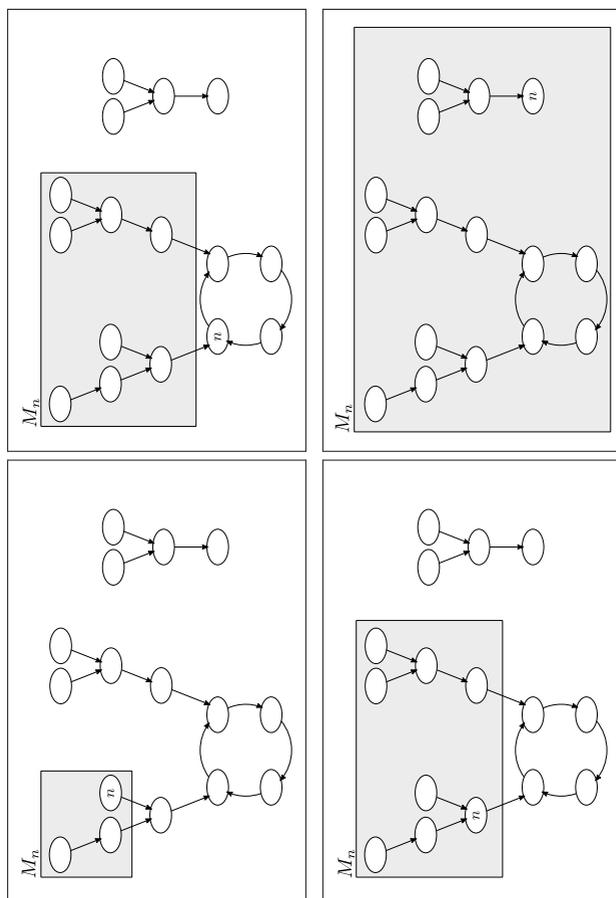


Figure C.2: The region M_n such that $n \rightarrow_C m$ is not prohibited by Observation C.0.2

D Algorithm Variants

An algorithm must be seen to be believed, and the best way to learn what an algorithm is all about is to try it.

(Donald Knuth — The Art of Computer Programming Vol. 1)

A slightly more efficient variant of Algorithm 3 is shown in Algorithm 18. Whenever it is known that the pseudo-forest $<$ is cycle-free, the checks $n' \in \beta_n$ and $m' \in \beta_m$ can be omitted.

Algorithm 19 is an algorithm for the computation of a pseudoforest $<_{\text{MAX}}$. It maintains a work-queue of nodes with fixed iteration order, instead of a work set.

Algorithm 20 and 21 compute a *cost demand*, i.e.: costs such that after adding them to the timing cost of edges towards nodes n_0, m_0 , the returned node is a timing sensitive least common ancestor of n_0, m_0 .

Input : A pseudo-forest \prec , represented as a map
IMDOM : $N \leftrightarrow N$ s.t. $\text{IMDOM}[n] = m$ iff $n < m$.
Input : Nodes n_0, m_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** $\text{lin}(\pi_m m')$

| | **end**

| **end**

end

Algorithm 18: A least common ancestor algorithm variant of Algorithm 3

Input : A CFG G

Data: A pseudo-forest \prec represented as a map $\text{IMDOM} : N \leftrightarrow N$ s.t.

$$\text{IMDOM}[n] = m \text{ iff } n < m$$

Output: A transitive reduction \prec_{MAX} of \exists_{MAX}

begin

for $x \in N$, $\{z \mid x \rightarrow_G z\} = \{z\}$, $z \neq x$ **do**
 | $\text{IMDOM}[x] \leftarrow z$

end

$\text{MAXIMAL}_{\text{up}}$

return IMDOM

end

Procedure $\text{MAXIMAL}_{\text{up}}$

 workqueue $\leftarrow \text{COND}_G$

 oldest $\leftarrow \perp$

while workqueue $\neq \emptyset$ **do**

$x \leftarrow \text{removeFront}(\text{workqueue})$

assert $\text{IMDOM}[x] = \perp$

if oldest = x **then**

 | **return**

end

if oldest = \perp **then**

 | oldest $\leftarrow x$

end

$a \leftarrow \text{lca}(\{y \mid x \rightarrow_G y\})$

$z \leftarrow \begin{cases} \perp & \text{if } a = \perp \vee a = x \\ a & \text{otherwise} \end{cases}$

if $z \neq \perp$ **then**

 | $\text{IMDOM}[x] \leftarrow z$

 | oldest $\leftarrow \perp$

end

else

 | $\text{pushBack}(\text{workqueue}, x)$

end

end

end

Algorithm 19: An efficient algorithm for the computation of some

\prec_{MAX} .

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 $n <^k m$

Input : Numbers $k_0^n, k_0^m \in \mathbb{N}$ and nodes n_0, m_0 such that *some*
 $<_{\text{MAX}}$ -least common ancestor exists

Output : A triple $(a, k, \Delta C)$ such that ΔC is a map with
 $(a, k) = \text{lca}_{<}((n_0, k_0^n + \Delta C[n_0]), (m_0, k_0^m + \Delta C[m_0]))$, or
 \perp if no such triple exists

begin

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

end

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
 | $\text{KS}_n[n] = k_0^n + \sum_i k_i$ for $n_0 <^{k_1} \dots <^{k_c} n$ in π_n

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

 | **if** $m \in \pi_n$ **then let** $\Delta C = |k^m - \text{KS}_n[m]|$ **in**

 | **if** $k^m = \text{KS}_n[m]$ **then return** $(m, [n_0 \mapsto 0, m_0 \mapsto 0])$

 | **if** $k^m < \text{KS}_n[m]$ **then return** $(m, [n_0 \mapsto 0, m_0 \mapsto \Delta C])$

 | **if** $k^m > \text{KS}_n[m]$ **then return** $(m, [n_0 \mapsto \Delta C, m_0 \mapsto 0])$

 | **end**

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

 | **case** \perp **do return** $\text{lin}[\pi_n](\pi_m)$

 | **case** (n', k) **do**

 | **if** $n' \in \pi_n$ **then return** $\text{lin}[\pi_n](\pi_m)$

 | **return** $\text{lca}(\pi_m, \pi_{n'})$ **where** $\pi_{n'} = (n', k^n + k, \text{KS}_n[n' \mapsto k^n + k])$

 | **end**

 | **end**

end

Algorithm 20: A least common ancestor algorithm that computes a timing cost demand, continued in Algorithm 21

```

Function  $\text{lin}[\pi_n](\pi_m)$ 
  Input      : A  $\leftarrow$ -path  $\pi_m = m_0, \dots, m$  ending in  $m$ , represented
                as in Algorithm 20
  Implicit   : A  $\leftarrow$ -path  $\pi_n = n_0, \dots, n$  ending in  $n$ , likewise
  switch  $\text{IDOM}[m]$  do
    case  $\perp$  do return  $\perp$ 
    case  $(m', k)$  do let  $k^{m'} = k^m + k$  in
      if  $m' \in \pi_n$  then let  $\Delta C = |k^{m'} - \text{KS}_n[m']|$  in
        if  $k^{m'} = \text{KS}_n[m']$  then return  $(m', [n_0 \mapsto 0, m_0 \mapsto 0])$ 
        if  $k^{m'} < \text{KS}_n[m']$  then return  $(m', [n_0 \mapsto 0, m_0 \mapsto \Delta C])$ 
        if  $k^{m'} > \text{KS}_n[m']$  then return  $(m', [n_0 \mapsto \Delta C, m_0 \mapsto 0])$ 
      end
      if  $m' \in \pi_m$  then return  $\perp$ 
      return  $\text{lin}(\pi_{m'})$  where  $\pi_{m'} = (m', k^{m'}, \text{KS}_m[m' \mapsto k^{m'}])$ 
    end
  end

```

end

Algorithm 21: A least common ancestor algorithm that computes a timing cost demand (continued from Algorithm 20)

D.1 Another Algorithm for $\rightarrow_{\text{nticd}}$

Owing to Lemma B.3.1, I can give another work-set algorithm (Algorithm 22) for the computation of $\rightarrow_{\text{nticd}}$. I use a variant $\text{prev}_{\text{COND}}^{\rightarrow}$ of $\text{prev}_{\text{COND}}$ that — along with predecessor-conditional nodes n of some node p — also gives me the node(s) x s.t. $n \rightarrow_G x \rightarrow_G^* p$.

Input : A CFG $G = (N, E)$
Input : Maps $\overrightarrow{\text{next}}_{\text{COND}}$ and $\overrightarrow{\text{prev}}_{\text{COND}}$
Output : A map NTICD such that $\text{NTICD}[n] = \{m \mid n \rightarrow_{\text{nticd}} m\}$
Data: A map S from pairs of nodes to sets of nodes
Data: A workset $\text{workset} \subseteq N \times \text{COND}_G \times N$
begin
 for $p \in \text{COND}_G, p \rightarrow_G x, \neg x \rightarrow_G m$ **do**
 $S[m, p] \leftarrow S[m, p] \cup \{(p, x)\}$
 end
 for $p \in \text{COND}_G, m \in N, |S[m, p]| \neq 0, (n, x) \in \overrightarrow{\text{prev}}_{\text{COND}}[p]$ **do**
 $\text{workset} \leftarrow \text{workset} \cup \{(m, n, x)\}$
 end
 while $\text{workset} \neq \emptyset$ **do**
 $(m, p, x) \leftarrow \text{remove}(\text{workset})$
 if $\neg m \in \overrightarrow{\text{next}}_{\text{COND}}[x]$ **and** $\neg (p, x) \in S[m, p]$ **then**
 if $|S[m, p]| = 0$ **then**
 $\text{workset} \leftarrow \{(m, n, x') \mid (n, x') \in \overrightarrow{\text{prev}}_{\text{COND}}[p]\}$
 end
 $S[m, p] \leftarrow S[m, p] \cup \{(p, x)\}$
 end
 end

 for $n \in N, m \in \text{COND}_G$ **do**
 if $0 < |S[m, n]| < |\{x \mid n \rightarrow_G x\}|$ **then**
 $\rightarrow_{\text{nticd}}[m] \leftarrow \rightarrow_{\text{nticd}}[m] \cup \{n\}$
 end
 end
end

Algorithm 22: A least fixed point algorithm for $\rightarrow_{\text{nticd}}$.

D.2 Efficient $\text{lca}_<$ via Postorder Numbers

In Algorithm 5 and Algorithm 6, the intersection of the postdominance sets of all successors of x of a conditional node p , i.e. the computation of the set of nodes m s.t. $m \sqsupseteq p$ via the rule (see: Theorem 5.1.2)

$$\frac{\forall p \rightarrow_G x. m \sqsupseteq x \quad p \rightarrow_G^* m}{m \sqsupseteq p} \text{D}^{\text{suc}}$$

was replaced by it's abstraction in $<$, i.e. by the computation of

$$\text{lca}_<(\{x \mid p \rightarrow_G x\})$$

In [CHK01], the algorithm for the computation of dominance in graphs with unique entry node n_e similarly computes the set of nodes m s.t. $m \sqsupseteq_{\text{DOM}} p$ (by considering all predecessors x of join-nodes p), but employs one additional abstraction:

There, the least common ancestor $\text{lca}_<(\{x \mid x \rightarrow_G p\})$ is also computed by following paths in $<$ and terminating once the path $\pi_n = n_0, \dots, n$ from one predecessor n_0 joins the path π_m from another predecessor m_0 (as is done by Algorithm 3), but the check $n \in \pi_m$ is replaced by an (arithmetic) comparison of the nodes' postorder numbers.

In this subsection, i will demonstrate how this very technique can be adapted to the computation of $<_{\text{SINK}}$.

First, i observe that the phase SINK_{up} in Algorithm 6 can be replaced by a depth-first traversal in the reversed graph G^{-1} . Remember that i need to compute an approximation $>_0$ of $>_{\text{SINK}}$, i.e. a pseudo-forest $>_0$ such that

$$>_0^* \supseteq \sqsupseteq_{\text{SINK}}$$

Observation D.2.1. Let $G = (N, E)$ be any CFG, S_1, \dots, S_n its sinks, and $S = \bigcup S_i$. Choose a distinct node s_i for each S_i . Also choose for each S_i a fixed ordering n_1, \dots, n_k of each S_i .

Let $F \subseteq N \times N$ be the depth-first forest obtained from a search in G^{-1} , iteratively starting from the fixed nodes s_i , and let

$$\begin{aligned} <_0 = & \bigcup_i \{ (n_j, n_{(j+1 \bmod k)}) \mid n_j \in S_i \} \\ & \cup \{ (m, n) \mid (n, m) \in F, m \notin S \} \end{aligned}$$

Then

$$>_0^* \supseteq \exists_{\text{SINK}}$$

With regard to the precision of this $>_0$ compared with that computed by phase SINK_{up} in Algorithm 6, experiments with randomly generated graphs suggest that the result of SINK_{up} is for *most* graphs more precise than that of the depth first search (if measured by the cardinality of $>_0^*$) and leads to quicker termination of phase $\text{SINK}_{\text{down}}$; but this is not true for every graph G .

Next I observe that, given a post-order numbering $\# : N \rightarrow \mathbb{N}^+$ of the depth-first search as obtained in Observation D.2.1, the $\text{lca}_{<}(n, m)$ of any approximation $<$ of \exists_{SINK} with which that numbering is *compatible* can be computed using *arithmetic* comparisons:

Definition D.2.1. Let $<$ be a pseudo-forest with cycles S_i , and $S = \bigcup_i S_i$.

A numbering¹ $\# : N \rightarrow \mathbb{N}$ is called *compatible* with $<$ if

1. for $n \notin S$ and $n < m$: $n^\# < m^\#$
2. for $n \in S_i$ and $n < m$: $m = s_{j+1 \bmod k}$

¹ i.e.: an injective map

given that $n = s_j$ and $S_i = s_1, \dots, s_k$ ordered by the numbering $\#$.

Observation D.2.2. In the situation of Definition D.2.1, and given $n_0, m_0 \in N$, Algorithm 23 computes a number $a^\#$ such that $a \in \text{lca}_<(n, m)$, or 0 if $\text{lca}_<(n, m) = \emptyset$.

Proof (Sketch): Upon terminating with $a^\# \neq 0$, the algorithm has followed two strictly ascending² sequences $\pi_n^\# = n_0^\#, \dots, n^\#$ and $\pi_m^\# = m_0^\#, \dots, m^\#$ such that $n^\# = m^\#$ and $n_0 < \dots < n$ and $m_0 < \dots < m$, i.e.: $n_0 <^* a \wedge m_0 <^* a$. Also, since π_n, π_m are disjunct up to $m = n$, $a < a'$ for any other common ancestor of n_0, m_0 . If, on the other hand, the algorithm terminates with 0, then — up to symmetry — either

- $n^\# = 0, m^\# \neq 0$, i.e.: $\pi_n^\# = n_0^\#, \dots, n'^\#, 0$ and $\pi_m^\# = m_0^\#, \dots, m^\#$ with $\pi_n \cap \pi_m = \emptyset$ and $n' < x$ for no $x \in N$, or
- $\pi_n^\# = n_0^\#, \dots, n^\#$ and $\pi_m^\# = m_0^\#, \dots, m^\#$ with $m < m', m^\# \leq m^\#$ and $n^\# > m^\#$. But because $\#$ is compatible, $m^\# = \max_{m_0 <^* x} x^\# = s_k^\#$, with $s_k \in S_i$. Since $n^\# > m^\#$, there is no node $s \in S_i$ s.t. $n_0 <^* s$, and with $\pi_n \cap \pi_m = \emptyset$, it follows that $\text{lca}_<(n_0, m_0) = \emptyset$.

Ordering each sinks nodes by $\#$, any approximation $<_0$ obtained via Observation D.2.1 is obviously compatible with the corresponding post-order numbering $\#$. If i am to use Algorithm 23 in an algorithm phase similiar to $\text{SINK}_{\text{down}}$, compatibility has to be upheld — but this is simply to the fact that initially, if $n <_0^* m$ and $n \notin S$, not only $n^\# < m^\#$ (compatibility), but also $n \rightarrow_G m$. Since for $a = \text{lca}_<(\{x \mid n \rightarrow_G x\})$ as computed by Algorithm 23 i have $\forall x. x^\# \leq a^\#$, i retain $n^\# < a^\#$ whenever i update $<$ at n with such a least common ancestor of the successors of n .

Empirical Observation D.2.1. Let G be any CFG. Then Algorithm 24 terminates with a result $<_{\text{SINK}}$ s.t. $>_{\text{SINK}}$ is a transitive reduction of \exists_{SINK} .

² due to compatibility

Implicit : A pseudo-forest $<$ with cycles S_i , represented as a map $\text{ISDOM}^\# : \mathbb{N} \rightarrow \mathbb{N}$ s.t. $\text{ISDOM}^\# [n^\#] = m^\#$ iff $n < m$, and $\text{ISDOM}^\# [n^\#] = 0$ if there is no such m , for a postorder numbering $^\# : N \rightarrow \mathbb{N}^+$ compatible with $<$.

Input : Postorder Numbers Nodes $m_0^\#, n_0^\#$ of Nodes m_0, n_0

Output : The Postorder Number of a least common ancestor of n_0, m_0 , or 0 if there is none.

Function $\text{lca}^\# (n^\#, m^\#)$

Input : A Postorder Number $n^\#$ such that $\pi_n = n_0, \dots, n$ is a $<$ -path ending in n

Input : A Postorder Number $m^\#$ such that $\pi_m = m_0, \dots, m$ is a $<$ -path ending in m

if $m^\# = 0 \vee n^\# = 0$ **then return** 0

if $n^\# > m^\#$ **then**

$m'^\# \leftarrow \text{ISDOM}[m^\#]$

if $m'^\# \leq m^\#$ **then**

return 0

($m' \in \pi_m$)

end

else return $\text{lca}^\# (n^\#, m'^\#)$

end

if $n^\# = m^\#$ **then return** $n^\#$

if $n^\# < m^\#$ **then**

$n'^\# \leftarrow \text{ISDOM}[n^\#]$

if $n'^\# \leq n^\#$ **then**

return 0

($n' \in \pi_n$)

end

else return $\text{lca}^\# (n'^\#, m^\#)$

end

end

Algorithm 23: A least common ancestor algorithm for nodes represented by their postorder number

Input : A CFG G

Data : A pseudo-forest \prec , represented as a map $\text{ISDOM}^\# : \mathbb{N} \rightarrow \mathbb{N}$
 s.t. $\text{ISDOM}^\# [n^\#] = m^\#$ iff $n < m$, and $\text{ISDOM}^\# [n^\#] = 0$ if
 there is no such m

Output: A transitive reduction \prec_{SINK} of \exists_{SINK}

begin

$\{S_1, \dots, S_n\} \leftarrow \{S_k \mid S_k \in \text{scc}(G), \neg \exists s \rightarrow_G n. s \in S_k \wedge n \notin S_k\}$

$S \leftarrow \bigcup S_i$

for $1 \leq i \leq n$ **do**

$s_i \leftarrow$ any node in S_i

end

let \prec_0 be obtained by a depth first search in G^{-1} (see:

 Observation D.2.1), compatible with post-order numbering $\#$

for $n <_0 m$ **do**

$\text{ISDOM}^\# [n^\#] \leftarrow m^\#$

end

$\text{SINK}_{\text{down}}^\#$

return ISDOM

end

Procedure $\text{SINK}_{\text{down}}^\#$

repeat

$\text{changed} \leftarrow$ **false**

for $x \in \text{COND}$ **do**

$a^\# \leftarrow \text{lca}^\# (\{y^\# \mid x \rightarrow_G y\})$

$z^\# \leftarrow \begin{cases} 0 & \text{if } a^\# = 0 \\ s_i^\# & \text{if } a \in S_i \\ a^\# & \text{otherwise} \end{cases}$

if $z^\# \neq \text{ISDOM}^\# [x^\#]$ **then**

$\text{ISDOM}^\# [x^\#] \leftarrow z^\#$

$\text{changed} \leftarrow$ **true**

end

end

until $\neg \text{changed}$

end

Algorithm 24: An Algorithm for the computation of some \prec_{SINK} .

The fixed point iteration in Algorithm 24 is the most naive implementation possible. Several worklist based optimizations are possible. Again, one can also make use of the fact that once $\text{ISDOM}^\# [x^\#] = 0$, x no longer needs to be considered. As is, Algorithm 24 can be read as a strict generalization³ of the algorithm in [CHK01].

³ after the transition from G to G^{-1}

E Generalizations for CFG with Timing Cost Model

This is a real frickin' embarrassment.

(Scout — Teamfortress 2)

In Subsection 9.4.3, I described an algorithm (Algorithm 11) for the computation of the timing sensitive postdominance frontier $\text{PDF}_{\sqsupset_{\text{TIME}[\text{FIRST}]}}$ for CFG under the (implicit) uniform timing cost model $\mathbb{1}$ which assigns a duration of 1 unit of time to each edge $n \rightarrow_G m$.

The algorithm was derived from Observation 9.2.3 on page 176 for

$$\sqsupset = \sqsupset_{\text{TIME}[\text{FIRST}]} = \sqsupset_{\text{TIME}[\text{FIRST}]}^{\mathbb{1}}$$

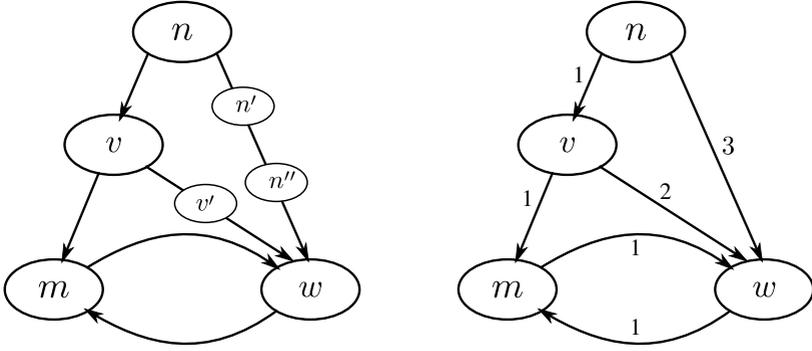
via a corresponding least fixed point characterization in Figure 9.5 on page 178.

I did not justify why it is enough to take the *least* fixed of rule system in Figure 9.5. In fact, Observation 9.2.3 merely characterizes, for each node x , $\text{PDF}_{\sqsupset_{\text{TIME}[\text{FIRST}]}}(x)$ in terms of some sets $\text{PDF}_{\sqsupset_{\text{TIME}[\text{FIRST}]}}^{\text{up}}(x, z)$, and hence ultimately: in terms of $\text{PDF}_{\sqsupset_{\text{TIME}[\text{FIRST}]}}(z)$ of some nodes z . In other words: it is — in general — merely a *mutually recursive* system of equations, and Algorithm 11 merely states that $\text{PDF}_{\sqsupset_{\text{TIME}[\text{FIRST}]}}$ is *some* solution of this system, but — a priori — not necessarily the least.

Consider the CFG in Figure E.1a under the default timing cost $\mathbb{1}$.

For $\sqsupset = \sqsupset_{\text{TIME}[\text{FIRST}]}$, I have

$$\text{ipdom}_{\sqsupset_{\text{TIME}[\text{FIRST}]}}(m) = \text{ipdom}_{\sqsupset_{\text{TIME}[\text{FIRST}]}}(w) = \{m, w\}$$


 (a) A CFG G

 (b) A CFG G_C with timing cost C

Figure E.1: The need for a modification of Algorithm 11.

and so from Observation 9.2.3 i merely obtain

$$\begin{aligned}
 \text{PDF}_{\sqsupseteq}(m) &= \underbrace{\text{PDF}_{\sqsupseteq}^{\text{local}}(m)}_{= \{v\}} \cup \dots \cup \underbrace{\text{PDF}_{\sqsupseteq}^{\text{up}}(m, m) \cap \sqsubseteq'_w}_{\subseteq \text{PDF}_{\sqsupseteq}(m)} \\
 \text{and } \text{PDF}_{\sqsupseteq}(w) &= \underbrace{\text{PDF}_{\sqsupseteq}^{\text{local}}(w)}_{= \emptyset} \cup \dots \cup \underbrace{\text{PDF}_{\sqsupseteq}^{\text{up}}(m, m) \cap \sqsubseteq'_w}_{\subseteq \text{PDF}_{\sqsupseteq}(w)}
 \end{aligned}$$

Note that the inequation $\text{PDF}_{\sqsupseteq}^{\text{up}}(m, m) \cap \sqsubseteq'_w \subseteq \text{PDF}_{\sqsupseteq}(m)$ is already due to the definition

$$\text{PDF}_{\sqsupseteq}^{\text{up}}(m, m) = \{y \in \text{PDF}_{\sqsupseteq}(m) \mid \neg m \perp_{\sqsupseteq} y\}$$

of $\text{PDF}_{\sqsupseteq}^{\text{up}}(m, m)$.

At the same time I observe — by manual inspection of G — that m is timing sensitively dependent on n (i.e.: $n \in \text{PDF}_{\sqsupseteq}(m)$). Since in the equation for $\text{PDF}_{\sqsupseteq}(m)$, n does not appear left from the ellipsis \dots , and the term on the right of the ellipsis is not “productive” (it is, after all, a subset of $\text{PDF}_{\sqsupseteq}(m)$) in a least fixed point computation, the node

n must “reach” $\text{PDF}_{\sqsupseteq}(m)$ via some term *in* the ellipsis. And indeed it does: realizing that $n \in \text{PDF}_{\sqsupseteq}(n'')$, and that $m \in \text{ipdom}_{\sqsupseteq}(n'')$, but $\neg n'' \in \text{ipdom}_{\sqsupseteq}(m)$, the node n reaches $\text{PDF}_{\sqsupseteq}(m)$ via

$$\text{PDF}_{\sqsupseteq}(m) = \dots \cup \text{PDF}_{\sqsupseteq}^{\text{up}}(n'', m) \cup \dots$$

In summary, everything is fine for $\text{PDF}_{\sqsupseteq_{\text{TIME[FIRST]}}}^{\text{ll}} = \text{PDF}_{\sqsupseteq_{\text{TIME[FIRST]}}}^{\text{ll}}$. Even though the equation for $\text{PDF}_{\sqsupseteq}(m)$ is (syntactically) not fully productive whenever $m \in \text{ipdom}_{\sqsupseteq}(m)$ (i.e.: when m is in a non-trivial \langle_{TIME} cycle M), nevertheless all $y \in \text{PDF}_{\sqsupseteq}(m)$ are provided either by $\text{PDF}_{\sqsupseteq}^{\text{local}}(m')$ for some $m' \in M$, or by $\text{PDF}_{\sqsupseteq}(n')$ for some node $n' \notin M$ such that $m \in \text{ipdom}_{\sqsupseteq}(n')$.

In contrast, things are *not* generally fine for $\text{PDF}_{\sqsupseteq_{\text{TIME[FIRST]}}}^{\text{C}}$ for *arbitrary* timing cost model C . Both Observation 9.2.3, and the simplifications Observation 9.2.4 and Observation 9.2.5 of $\text{PDF}_{\sqsupseteq}^{\text{local}}$ and $\text{PDF}_{\sqsupseteq}^{\text{up}}$ *do* hold even for $\sqsupseteq = \sqsupseteq_{\text{TIME[FIRST]}}^{\text{C}}$ and arbitrary G, C . But for $\text{PDF}_{\sqsupseteq_{\text{TIME[FIRST]}}}^{\text{C}}$, it is no longer enough to simply take the least fixed point of this system of equations.

Consider the CFG and the timing cost model C shown in Figure E.1b. Up to *intermediate* nodes v', n', n'' , this is essentially the same CFG as before.

Again, I have

$$\text{ipdom}_{\sqsupseteq_{\text{TIME[FIRST]}}^{\text{C}}}(m) = \text{ipdom}_{\sqsupseteq_{\text{TIME[FIRST]}}^{\text{C}}}(w) = \{m, w\}$$

and so from Observation 9.2.3 I obtain gain, for $\sqsupseteq = \sqsupseteq_{\text{TIME[FIRST]}}^{\text{C}}$,

$$\begin{aligned} \text{PDF}_{\sqsupseteq}(m) &= \underbrace{\text{PDF}_{\sqsupseteq}^{\text{local}}(m)}_{= \{v\}} \cup \dots \cup \underbrace{\text{PDF}_{\sqsupseteq}^{\text{up}}(m, m) \cap \sqsubseteq'_w}_{\subseteq \text{PDF}_{\sqsupseteq}(m)} \\ \text{and } \text{PDF}_{\sqsupseteq}(w) &= \underbrace{\text{PDF}_{\sqsupseteq}^{\text{local}}(w)}_{= \emptyset} \cup \dots \cup \underbrace{\text{PDF}_{\sqsupseteq}^{\text{up}}(m, m) \cap \sqsubseteq'_w}_{\subseteq \text{PDF}_{\sqsupseteq}(w)} \end{aligned}$$

I observe — again by manual inspection, this time of CFG G_C under timing cost model C — that m is timing sensitively dependent on n under timing cost model C (i.e.: $n \in \text{PDF}_{\sqsupseteq}(m)$).

But can I obtain $n \in \text{PDF}_{\sqsupseteq}(m)$ by a least fixed point iteration? In other words, can I obtain $n \in \text{PDF}_{\sqsupseteq}(m)$ purely from the center ellipsis ... for m ? First, I must observe that

$$m \notin \text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}^C(v) = \text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}^C(n) = \{w\}$$

In other words: now there exists *no* node n'' outside of $M = \{m, w\}$ such that $m \in \text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}^C(n'')$. Also, $w \in M$ is *not* timing sensitively control dependent on n under cost model C : $\neg n \in \text{PDF}_{\sqsupseteq}(w)$.

The full equation for m thus reads: $\text{PDF}_{\sqsupseteq}(m) =$

$$\underbrace{\text{PDF}_{\sqsupseteq}^{\text{local}}(m)}_{= \{v\}} \cup \underbrace{\text{PDF}_{\sqsupseteq}^{\text{up}}(w, m) \cap \sqsubseteq'_m}_{\subseteq \text{PDF}_{\sqsupseteq}(w)} \cup \underbrace{\text{PDF}_{\sqsupseteq}^{\text{up}}(m, m) \cap \sqsubseteq'_m}_{\subseteq \text{PDF}_{\sqsupseteq}(m)}$$

which is *not* productive for $n \in \text{PDF}_{\sqsupseteq}(m)$, so I can *not* obtain $\text{PDF}_{\sqsupseteq}(m)$ as a least fixed point of the implied functional (nor as the least fixed point from the simplified rule system Figure 9.5).

I did *not* investigate whether PDF_{\sqsupseteq} can be obtained as the greatest fixed point of Figure 9.5 or some related system. Instead I observe that nodes n missing in the least fixed point computation for $\text{PDF}_{\sqsupseteq}(m)$ must always be nodes at the *border* N_M of the $<_{\text{TIME}}^C$ -cycle M that m is part of. This means that I *can* compute PDF_{\sqsupseteq} as a least fixed point, if I augment the rule system Figure 9.5 with the additional rule

$$\frac{\neg x \in \text{ipdom}_{\sqsupseteq}(y) \quad y \rightarrow_G y' \quad x \in M \quad M \in \mathbb{M} \quad y \in N_M \quad x \sqsupseteq y'}{y \in \text{PDF}_{\sqsupseteq}(x)} \text{PDF}$$

where \mathbb{M} denotes the set of non-trivial $<_{\text{TIME}}^C$ -cycles. The rule is directly derived from the *definition* of PDF_{\sqsupseteq} , but restricted to a *small*

subset of pairs x, y of nodes. Note that since $x \in M$, the condition $x \sqsupseteq y'$ can be replaced by

$$y' \in M \quad \vee \quad y' < \dots < n \wedge x \in \text{ipdom}_{\sqsupseteq}(n) \text{ for some } n$$

where

$$< := \{ (n, m) \mid n <_{\text{TIME}}^{C,k} m, m \notin M \}$$

is obtained from $<_{\text{TIME}}^C$ by deleting all edges into M .

The resulting Algorithm 25 differs from Algorithm 11 only by addition of the highlighted lines.

Input : A CFG $G = (N, E)$
Input : Any numbering $\# : N \rightarrow \mathbb{N}$
Input : Immediate post dominators $\text{ipdom}_{\sqsupseteq} = \text{ipdom}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}}^{\text{C}}$
Input : The transitive reduction $\prec_{\text{TIME}}^{\text{C}}$ of transitive timing sensitive postdominance under timing cost model C
Input : The set \mathbb{M} of $\prec_{\text{TIME}}^{\text{C}}$ -cycles M
Input : The set N_M of corresponding “border” nodes in $\prec_{\text{TIME}}^{\text{C}}$, and $N_M = \bigcup_{M \in \mathbb{M}} N_M$
Data: A priority queue Q ordered by the numbering $\#$
Output: $\text{PDF}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}^{\text{C}}}$ represented as a map $\text{DF} : N \rightarrow N \leftrightarrow \text{Bool}$

$Q \leftarrow \emptyset$

```

for  $M \in \mathbb{M}, |M| > 1$  do
  | for  $y \in N_M, x \in M \setminus \text{ipdom}_{\sqsupseteq}(y), \exists y'. y \rightarrow_G y' \wedge x \sqsupseteq y'$  do
  | |  $\text{DF}[x][y] \leftarrow \text{false}$ 
  | end
end

```

```

for  $x \in N, y \rightarrow_G x, \neg x \in \text{ipdom}_{\sqsupseteq}(y)$  do
  |  $\text{DF}[x][y] \leftarrow \text{true}$ 
  |  $Q \leftarrow Q \cup \{x\}$ 
end
while  $Q \neq \emptyset$  do
  |  $z \leftarrow \text{remove}(Q)$  s.t.  $z^\# = \max_{z \in Q} z^\#$ 
  | for  $x \in \text{ipdom}_{\sqsupseteq}(z), (y, \text{true}) \in \text{DF}[z], \neg x \in \text{ipdom}_{\sqsupseteq}(y)$  do
  | |  $\text{DF}_{x,y} \leftarrow \text{DF}[x][y]$ 
  | |  $\text{DF}'_{x,y} \leftarrow \text{DF}_{x,y} \vee (z \notin N_M)$ 
  | | if  $\text{DF}'_{x,y} \neq \text{DF}_{x,y}$  then
  | | |  $\text{DF}[x][y] \leftarrow \text{DF}'_{x,y}$ 
  | | |  $Q \leftarrow Q \cup \{x\}$ 
  | | end
  | end
end
return  $\text{DF}$ 

```

Algorithm 25: Computation of $\text{PDF}_{\sqsupseteq_{\text{TIME}[\text{FIRST}]}^{\text{C}}}$

Input : A CFG G

Input : A timing cost model C for G

Data: 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 $n <^k m$

Output: A transitive reduction $>_{\text{TIME}}^C$ of $\sqsupseteq_{\text{TIME}}^C$

begin

for $x \in N$, $\{z \mid x \rightarrow_G z\} = \{z\}$ **do**

$\text{IDOM}[x] \leftarrow (z, C(x, z))$

end

TIME_{up}

return IDOM

end

Procedure TIME_{up}

$\text{workset} \leftarrow \text{COND}_G$

while $\text{workset} \neq \emptyset$ **do**

$x \leftarrow \text{remove}(\text{workset})$

$(z, k) \leftarrow \text{lca}_{<}(\{(y, C(x, y)) \mid x \rightarrow_G y\})$

assert $(z, k) \neq \text{IDOM}[x] \Rightarrow (z, k) \neq \perp$

assert $(z, k) \neq \text{IDOM}[x] \Rightarrow \text{IDOM}[x] = \perp$

if $(z, k) \neq \text{IDOM}[x]$ **then**

$\text{workset} \leftarrow \text{workset} \cup \{n \in \text{COND} \mid n \neq x, \exists n \rightarrow_G y. y <^* x\}$

$\text{IDOM}[x] \leftarrow (z, k)$

end

end

end

Algorithm 26: An efficient algorithm for the computation of $<_{\text{TIME}}^C$.

Here, $y <^* x$ is taken to mean: $x = y \vee y <^{k_1} \dots <^{k_c} x$

Bibliography

Ahenny (adj.) — The way people stand when examining other people's bookshelves.

(Douglas Adams — The Deeper Meaning of Liff)

- [Aga00] Johan Agat. "Transforming out Timing Leaks". In: *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL '00. Boston, MA, USA: ACM, 2000, pp. 40–53. ISBN: 1-58113-125-9. DOI: 10.1145/325694.325702.
- [Alm+16] Jose Bacelar Almeida et al. "Verifying Constant-Time Implementations". In: *25th USENIX Security Symposium (USENIX Security 16)*. Austin, TX: USENIX Association, Aug. 2016, pp. 53–70. ISBN: 978-1-931971-32-4.
- [Amt08] Torben Amtoft. "Slicing for modern program structures: a theory for eliminating irrelevant loops". In: *Information Processing Letters* 106.2 (2008), pp. 45–51. ISSN: 0020-0190. DOI: <https://doi.org/10.1016/j.ipl.2007.10.002>.
- [AS85] Bowen Alpern and Fred B. Schneider. "Defining liveness". In: *Information Processing Letters* 21.4 (1985), pp. 181–185. ISSN: 0020-0190. DOI: [https://doi.org/10.1016/0020-0190\(85\)90056-0](https://doi.org/10.1016/0020-0190(85)90056-0).
- [Ask+08] Aslan Askarov et al. "Termination-Insensitive Noninterference Leaks More Than Just a Bit". In: *Computer Security - ESORICS 2008*. Ed. by Sushil Jajodia and Javier Lopez. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 333–348. ISBN: 978-3-540-88313-5.
- [Bar+10] Richard W. Barraclough et al. "A trajectory-based strict semantics for program slicing". In: *Theoretical Computer Science* 411.11 (2010), pp. 1372–1386. ISSN: 0304-3975. DOI: <https://doi.org/10.1016/j.tcs.2009.10.025>.

- [BB05] David Brumley and Dan Boneh. "Remote timing attacks are practical". In: *Computer Networks* 48.5 (2005). Web Security, pp. 701–716. ISSN: 1389-1286. DOI: <https://doi.org/10.1016/j.comnet.2005.01.010>.
- [Ber05] Daniel J. Bernstein. *Cache-timing attacks on AES*. Tech. rep. 2005.
- [Bis+18a] Simon Bischof et al. "Illi Isabellistes Se Custodes Egregios Praestabant". In: *Principled Software Development: Essays Dedicated to Arnd Poetzsch-Heffter on the Occasion of his 60th Birthday*. Ed. by Peter Müller and Ina Schaefer. Springer International Publishing, 2018, pp. 267–282. DOI: 10.1007/978-3-319-98047-8_17.
- [Bis+18b] Simon Bischof et al. "Low-Deterministic Security For Low-Nondeterministic Programs". In: *Journal of Computer Security* 26 (2018), pp. 335–366. DOI: 10.3233/JCS-17984.
- [Bis19] Simon Bischof. *Isabelle Theories about NTxCD*. 2019. URL: <https://pp.ipd.kit.edu/~bischof/ntxcd.html> (visited on 12/13/2010).
- [BJ66] Corrado Böhm and Giuseppe Jacopini. "Flow Diagrams, Turing Machines and Languages with Only Two Formation Rules". In: *Commun. ACM* 9.5 (May 1966), pp. 366–371. ISSN: 0001-0782. DOI: 10.1145/355592.365646.
- [BM06] Joseph Bonneau and Ilya Mironov. "Cache-Collision Timing Attacks Against AES". In: *Cryptographic Hardware and Embedded Systems - CHES 2006*. Ed. by Louis Goubin and Mitsuru Matsui. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 201–215. ISBN: 978-3-540-46561-4.
- [BP10] Joan Boyar and René Peralta. "A New Combinational Logic Minimization Technique with Applications to Cryptology". In: *Experimental Algorithms*. Ed. by Paola Festa. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 178–189. ISBN: 978-3-642-13193-6.

- [BP96] Gianfranco Bilardi and Keshav Pingali. “A Framework for Generalized Control Dependence”. In: *Proceedings of the ACM SIGPLAN 1996 Conference on Programming Language Design and Implementation*. PLDI '96. Philadelphia, Pennsylvania, USA: ACM, 1996, pp. 291–300. ISBN: 0-89791-795-2. DOI: 10.1145/231379.231435.
- [Bre+16] Joachim Breitner et al. “On Improvements Of Low-Deterministic Security”. In: *Principles of Security and Trust - 5th International Conference, POST 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings*. Ed. by Frank Piessens and Luca Viganò. Vol. 9635. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2016, pp. 68–88. DOI: 10.1007/978-3-662-49635-0_4.
- [Bro+19] R. Brotzman et al. “CaSym: Cache Aware Symbolic Execution for Side Channel Detection and Mitigation”. In: *2019 IEEE Symposium on Security and Privacy (SP)*. May 2019, pp. 505–521. DOI: 10.1109/SP.2019.00022.
- [BRW06] Gilles Barthe, Tamara Rezk, and Martijn Warnier. “Preventing Timing Leaks Through Transactional Branching Instructions”. In: *Electronic Notes in Theoretical Computer Science* 153.2 (2006). Proceedings of the Third Workshop on Quantitative Aspects of Programming Languages (QAPL 2005), pp. 33–55. ISSN: 1571-0661. DOI: <https://doi.org/10.1016/j.entcs.2005.10.031>.
- [BS07] Julian Bradfield and Colin Stirling. “12 Modal *mu*-calculi”. In: *Handbook of Modal Logic*. Ed. by Patrick Blackburn, Johan Van Benthem, and Frank Wolter. Vol. 3. Studies in Logic and Practical Reasoning. Elsevier, 2007, pp. 721–756. DOI: [https://doi.org/10.1016/S1570-2464\(07\)80015-2](https://doi.org/10.1016/S1570-2464(07)80015-2).
- [CFT03] Larry Carter, Jeanne Ferrante, and Clark Thomborson. “Folklore Confirmed: Reducible Flow Graphs Are Exponentially Larger”. In: *Proceedings of the 30th ACM*

- SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL '03. New Orleans, Louisiana, USA: ACM, 2003, pp. 106–114. ISBN: 1-58113-628-5. DOI: 10.1145/604131.604141.
- [CH02] Koen Claessen and John Hughes. “Testing Monadic Code with QuickCheck”. In: *SIGPLAN Not.* 37.12 (Dec. 2002), pp. 47–59. ISSN: 0362-1340. DOI: 10.1145/636517.636527.
- [Cha+19] Sudipta Chattopadhyay et al. “Quantifying the Information Leakage in Cache Attacks via Symbolic Execution”. In: *ACM Trans. Embed. Comput. Syst.* 18.1 (Jan. 2019). ISSN: 1539-9087. DOI: 10.1145/3288758.
- [CHK01] Keith D. Cooper, Timothy J. Harvey, and Ken Kennedy. “A simple, fast dominance algorithm”. In: *Software Practice & Experience* (2001).
- [Cor20] Oracle Corporation. *Code Tools: jmh*. 2020. URL: <https://github.com/AlDanial/cloc> (visited on 02/11/2020).
- [CR06] Feng Chen and Grigore Roşu. “Parametric and Termination-Sensitive Control Dependence”. In: *Static Analysis*. Ed. by Kwangkeun Yi. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 387–404. ISBN: 978-3-540-37758-0.
- [CS10] Michael R Clarkson and Fred B Schneider. “Hyperproperties”. In: *Journal of Computer Security* 18.6 (2010), pp. 1157–1210.
- [Cyt+91] Ron Cytron et al. “Efficiently Computing Static Single Assignment Form and the Control Dependence Graph”. In: *ACM Trans. Program. Lang. Syst.* 13.4 (Oct. 1991), pp. 451–490. ISSN: 0164-0925. DOI: 10.1145/115372.115320.
- [Dan+05] Sebastian Danicic et al. “Static program slicing algorithms are minimal for free liberal program schemas”. In: *The Computer Journal* 48.6 (2005), pp. 737–748.

- [Dan+11] Sebastian Danicic et al. “A unifying theory of control dependence and its application to arbitrary program structures”. In: *Theoretical Computer Science* 412.49 (2011), pp. 6809–6842. ISSN: 0304-3975. DOI: <https://doi.org/10.1016/j.tcs.2011.08.033>.
- [Dan18] Al Danial. *Cloc*. 2006 - 2018. URL: <https://github.com/AlDanial/cloc> (visited on 02/07/2020).
- [Doy+15] Goran Doychev et al. “CacheAudit: A Tool for the Static Analysis of Cache Side Channels”. In: *ACM Trans. Inf. Syst. Secur.* 18.1 (June 2015). ISSN: 1094-9224. DOI: 10.1145/2756550.
- [EM17] Martin Erwig and Ivan Lazar Miljenovic. *fgl: Martin Erwig’s Functional Graph Library*. <https://hackage.haskell.org/package/fgl>. 2001–2017.
- [Erw01] Martin Erwig. “Inductive graphs and functional graph algorithms”. In: *Journal of Functional Programming* 11.5 (2001), pp. 467–492. DOI: 10.1017/S0956796801004075.
- [Fou19] The Apache Software Foundation. *Apache FtpServer*. 2019. URL: https://mina.apache.org/ftpserver-project/download_1.1.1.html (visited on 02/07/2020).
- [FOW87] Jeanne Ferrante, Karl J. Ottenstein, and Joe D. Warren. “The Program Dependence Graph and Its Use in Optimization”. In: *ACM Trans. Program. Lang. Syst.* 9.3 (July 1987), pp. 319–349. ISSN: 0164-0925. DOI: 10.1145/24039.24041.
- [Fri02] Carsten Fritz. “Some Fixed Point Basics”. In: *Automata Logics, and Infinite Games: A Guide to Current Research*. Ed. by Erich Grädel, Wolfgang Thomas, and Thomas Wilke. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 359–364. ISBN: 978-3-540-36387-3. DOI: 10.1007/3-540-36387-4_20.

- [Gaw+11] Thomas Martin Gawlitza et al. “Join-Lock-Sensitive Forward Reachability Analysis for Concurrent Programs with Dynamic Process Creation”. In: *Verification, Model Checking, and Abstract Interpretation*. Ed. by Ranjit Jhala and David Schmidt. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 199–213. ISBN: 978-3-642-18275-4.
- [Gif12] Dennis Giffhorn. “Slicing of Concurrent Programs and its Application to Information Flow Control”. PhD thesis. Karlsruher Institut für Technologie, Fakultät für Informatik, 2012.
- [GM03] Roberto Giacobazzi and Isabella Mastroeni. “Non-Standard Semantics for Program Slicing”. In: *Higher-Order and Symbolic Computation* 16.4 (Dec. 2003), pp. 297–339. ISSN: 1573-0557. DOI: 10.1023/A:1025872819613.
- [Gra16] Jürgen Graf. “Information Flow Control with System Dependence Graphs – Improving Modularity, Scalability and Precision for Object Oriented Languages”. PhD thesis. Karlsruher Institut für Technologie, Fakultät für Informatik, Nov. 2016. DOI: 10.5445/IR/1000068211.
- [GS15] Dennis Giffhorn and Gregor Snelting. “A new algorithm for low-deterministic security”. In: *International Journal of Information Security* 14.3 (2015), pp. 263–287. DOI: 10.1007/s10207-014-0257-6.
- [Ham09] Christian Hammer. “Information Flow Control for Java - A Comprehensive Approach based on Path Conditions in Dependence Graphs”. ISBN 978-3-86644-398-3. PhD thesis. Universität Karlsruhe (TH), Fak. f. Informatik, July 2009.
- [Hec20] Martin Hecker. *Properties for: Timing Sensitive Dependency Analysis and its Application to Software Security*. 2020. URL: <https://pp.ipd.kit.edu/~hecker/dissertation/> (visited on 03/03/2010).

- [HS09] Christian Hammer and Gregor Snelting. “Flow-sensitive, context-sensitive, and object-sensitive information flow control based on program dependence graphs”. In: *International Journal of Information Security* 8.6 (Dec. 2009), pp. 399–422. ISSN: 1615-5270. DOI: 10.1007/s10207-009-0086-1.
- [HU73] Matthew S. Hecht and Jeffrey D. Ullman. “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. Boston, Massachusetts: ACM, 1973, pp. 207–217. DOI: 10.1145/512927.512946.
- [Jus+11] Seth Just et al. “Information Flow Analysis for Javascript”. In: *Proceedings of the 1st ACM SIGPLAN International Workshop on Programming Language and Systems Technologies for Internet Clients*. PLASTIC ’11. Portland, Oregon, USA: ACM, 2011, pp. 9–18. ISBN: 978-1-4503-1171-7. DOI: 10.1145/2093328.2093331.
- [Koc+19] Paul Kocher et al. “Spectre Attacks: Exploiting Speculative Execution”. In: *40th IEEE Symposium on Security and Privacy (S&P’19)*. 2019.
- [KRB09] B. Köpf, A. Rybalchenko, and M. Backes. “Automatic Discovery and Quantification of Information Leaks”. In: *2009 30th IEEE Symposium on Security and Privacy (SP)*. Los Alamitos, CA, USA: IEEE Computer Society, May 2009, pp. 141–153. DOI: 10.1109/SP.2009.18.
- [Kri98] Jens Krinke. “Static Slicing of Threaded Programs”. In: *Proceedings of the 1998 ACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering*. PASTE ’98. Montreal, Quebec, Canada: ACM, 1998, pp. 35–42. ISBN: 1-58113-055-4. DOI: 10.1145/277631.277638.

- [KTG12] Ralf Küsters, Tomasz Truderung, and Jürgen Graf. “A Framework for the Cryptographic Verification of Java-like Programs”. In: *Computer Security Foundations Symposium (CSF), 2012 IEEE 25th*. IEEE Computer Society, June 2012.
- [Küs+14] Ralf Küsters et al. “Extending and Applying a Framework for the Cryptographic Verification of Java Programs”. In: *Proc. POST 2014*. LNCS 8424. Springer, 2014, pp. 220–239.
- [Lip+18] Moritz Lipp et al. “Meltdown: Reading Kernel Memory from User Space”. In: *27th USENIX Security Symposium (USENIX Security 18)*. 2018.
- [LKL18] Jean-Christophe Léchenet, Nikolai Kosmatov, and Pascale Le Gall. “Fast Computation of Arbitrary Control Dependencies”. In: *Fundamental Approaches to Software Engineering*. Ed. by Alessandra Russo and Andy Schürr. Cham: Springer International Publishing, 2018, pp. 207–224. ISBN: 978-3-319-89363-1.
- [Lor+14] Steffen Lortz et al. “Cassandra: Towards a Certifying App Store for Android”. In: *Proceedings of the 4th ACM Workshop on Security and Privacy in Smartphones & Mobile Devices*. SPSM '14. Scottsdale, Arizona, USA: ACM, 2014, pp. 93–104. ISBN: 978-1-4503-3155-5. DOI: 10.1145/2666620.2666631.
- [LT79] Thomas Lengauer and Robert Endre Tarjan. “A Fast Algorithm for Finding Dominators in a Flowgraph”. In: *ACM Trans. Program. Lang. Syst.* 1.1 (Jan. 1979), pp. 121–141. ISSN: 0164-0925. DOI: 10.1145/357062.357071.
- [Mil17] Ivan Lazar Miljenovic. *fgl-arbitrary: QuickCheck support for fgl*. <https://hackage.haskell.org/package/fgl-arbitrary>. 2015–2017.
- [Mol+06] David Molnar et al. “The Program Counter Security Model: Automatic Detection and Removal of Control-Flow Side Channel Attacks”. In: *Information Security and Cryptology - ICISC 2005*. Ed. by Dong Ho Won and Se

- ungjoo Kim. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 156–168. ISBN: 978-3-540-33355-5.
- [Mou+15] Leonardo de Moura et al. “The Lean Theorem Prover (System Description)”. In: *Automated Deduction - CADE-25*. Ed. by Amy P. Felty and Aart Middeldorp. Cham: Springer International Publishing, 2015, pp. 378–388. ISBN: 978-3-319-21401-6.
- [MS15] Heiko Mantel and Artem Starostin. “Transforming Out Timing Leaks, More or Less”. In: *Computer Security – ESORICS 2015*. Ed. by Günther Pernul, Peter Y A Ryan, and Edgar Weippl. Cham: Springer International Publishing, 2015, pp. 447–467. ISBN: 978-3-319-24174-6.
- [NP19] Barak Naveh and Stephane Popinet. *JGraphT: A Java Library of Graph Theory Data Structures and Algorithms*. <https://jgrapht.org/>. 2003–2019.
- [NWP02] Tobias Nipkow, Markus Wenzel, and Lawrence C. Paulson. *Isabelle/HOL: A Proof Assistant for Higher-order Logic*. Berlin, Heidelberg: Springer-Verlag, 2002. ISBN: 3-540-43376-7.
- [PC90] A. Podgurski and L. A. Clarke. “A formal model of program dependences and its implications for software testing, debugging, and maintenance”. In: *IEEE Transactions on Software Engineering* 16.9 (Sept. 1990), pp. 965–979. ISSN: 0098-5589. DOI: 10.1109/32.58784.
- [PP18] Amanda Plunkett and Junyong Park. “Two-sample test for sparse high-dimensional multinomial distributions”. In: *TEST* (July 2018). ISSN: 1863-8260. DOI: 10.1007/s11749-018-0600-8.
- [Ran+07] Venkatesh Prasad Ranganath et al. “A New Foundation for Control Dependence and Slicing for Modern Program Structures”. In: *ACM Trans. Program. Lang. Syst.* 29.5 (Aug. 2007). ISSN: 0164-0925. DOI: 10.1145/1275497.1275502.

- [SG95] Vugranam C. Sreedhar and Guang R. Gao. “A Linear Time Algorithm for Placing Φ -nodes”. In: *Proceedings of the 22Nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL '95. San Francisco, California, USA: ACM, 1995, pp. 62–73. ISBN: 0-89791-692-1. DOI: 10.1145/199448.199464.
- [Sne+14] Gregor Snelting et al. “Checking Probabilistic Noninterference Using JOANA”. In: *it - Information Technology* 56 (Nov. 2014), pp. 280–287. DOI: 10.1515/itit-2014-1051.
- [Tar72] Robert Tarjan. “Depth-First Search and Linear Graph Algorithms”. In: *SIAM Journal on Computing* 1.2 (1972), pp. 146–160. DOI: 10.1137/0201010.
- [Tea17] The Coq Development Team. *The Coq Proof Assistant, version 8.7.0*. Version 8.7.0. Oct. 2017. DOI: 10.5281/zenodo.1028037.
- [Was10] Daniel Wasserrab. “From Formal Semantics to Verified Slicing - A Modular Framework with Applications in Language Based Security”. PhD thesis. Karlsruher Institut für Technologie, Fakultät für Informatik, Oct. 2010.
- [Wei81] Mark Weiser. “Program Slicing”. In: *Proceedings of the 5th International Conference on Software Engineering*. ICSE '81. San Diego, California, USA: IEEE Press, 1981, pp. 439–449. ISBN: 0-89791-146-6.
- [Wel10] Stefan Wellek. *Testing Statistical Hypotheses of Equivalence and Noninferiority*. 2nd ed. Boca Raton, FL, USA: CRC Press, 2010. ISBN: 978-1-4398-0819-1.
- [Wol95] Michael Joseph Wolfe. *High Performance Compilers for Parallel Computing*. Ed. by Carter Shanklin and Leda Ortega. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1995. ISBN: 0805327304.

- [Wu+18] Meng Wu et al. “Eliminating Timing Side-channel Leaks Using Program Repair”. In: *Proceedings of the 27th ACM SIGSOFT International Symposium on Software Testing and Analysis*. ISSTA 2018. Amsterdam, Netherlands: ACM, 2018, pp. 15–26. ISBN: 978-1-4503-5699-2. DOI: 10 . 1145 / 3213846 . 3213851.
- [XZ07] Bin Xin and Xiangyu Zhang. “Efficient Online Detection of Dynamic Control Dependence”. In: *Proceedings of the 2007 International Symposium on Software Testing and Analysis*. ISSTA '07. London, United Kingdom: ACM, 2007, pp. 185–195. ISBN: 978-1-59593-734-6. DOI: 10 . 1145 / 1273463 . 1273489.

Beppu (n.) — The triumphant slamming shut
of a book after reading the final page.

(Douglas Adams — The Deeper Meaning of Liff)