
Practical Dynamic Information Flow Control

Abhishek Bichhawat



A dissertation submitted towards the degree
Doctor of Engineering
of the Faculty of Mathematics and Computer Science
of Saarland University

by
Abhishek Bichhawat

Saarbrücken
September, 2017

Day of the Colloquium
Dean of the Faculty

Chair of the Committee
Reviewers

Academic Assistant

22/06/2018

Univ.-Prof. Dr. Sebastian Hack

Prof. Dr. Andreas Zeller

Prof. Dr. Christian Hammer

Dr. Deepak Garg

Prof. Dr. Bernd Finkbeiner

Dr. Marco Patrignani

Abstract

Over the years, computer systems and applications have grown significantly complex while handling a plethora of private and sensitive user information. The complexity of these applications is often assisted by a set of (un)intentional bugs with both malicious and non-malicious intent leading to information leaks. Information flow control has been studied extensively as an approach to mitigate such information leaks. The technique works by enforcing the security property of non-interference using a specified set of security policies. A vast majority of existing work in this area is based on static analyses. However, some of the applications, especially on the Web, are developed using dynamic languages like JavaScript that make the static analyses techniques stale and ineffective. As a result, there has been a growing interest in recent years to develop dynamic information flow analysis techniques. In spite of the advances in the field, dynamic information flow analysis has not been at the helm of information flow security in dynamic settings like the Web; the prime reason being that the analysis techniques and the security property related to them (non-interference) either over-approximate or are too restrictive in most cases. Concretely, the analysis techniques generate a lot of false positives, do not allow legitimate release of sensitive information, support only static and rigid security policies or are not general enough to be applied to real-world applications.

This thesis focuses on improving the usability of dynamic information flow techniques by presenting mechanisms that can enhance the precision and permissiveness of the analyses. It begins by presenting a sound improvement and enhancement of the permissive-upgrade strategy, a strategy widely used to enforce dynamic information flow control, which improves the strategy's permissiveness and makes it generic in applicability. The thesis, then, presents a sound and precise control scope analysis for handling complex features like unstructured control flow and exceptions in higher-order languages. Although non-interference is a desired property for enforcing information flow control, there are program instances that require legitimate release of some parts of the secret data to provide the required functionality. Towards this end, this thesis develops a sound approach to bound information leaks dynamically while allowing information release in accordance to a pre-specified budget. The thesis concludes by applying these techniques to an information flow control-enabled Web browser and explores a policy specification mechanism that allows flexible and useful information flow policies to be specified for Web applications.

Kurzzusammenfassung

Seit Jahren werden Computersysteme und -Anwendungen immer komplexer und verarbeiten eine Unmenge private und sensible Daten. Die Komplexität der Anwendungen trägt neben der Existenz von (un)gewollt eingefügten Software Fehlern zur Weitergabe dieser sensiblen Informationen bei. Information Flow Control (IFC, zu Deutsch Informations-Fluss-Analyse) Mechanismen sind Gegenstand intensiver Forschung um diesem Problem entgegen zu wirken. Grundsätzlich basieren diese Ansätze auf der Anwendung von vordefinierten Sicherheitsregeln, die die Unbeeinflussbarkeit (engl. non-interference) garantieren. Der überwiegende Teil dieser Techniken nutzt statische Analyse zur Erzeugung der Regeln. Dem gegenüber steht die Tatsache, dass Anwendungen, insbesondere im Bereich Web-Anwendungen, in dynamischen Sprachen wie JavaScript entwickelt werden, wodurch rein statische Analysen unzureichend sind. Dynamische Methoden auf der anderen Seite approximieren das Verhalten einer Anwendung und können daher die grundlegende non-interference nicht garantieren. Sie tendieren dazu besonders restriktive Regeln zu erzeugen, wodurch auch der rechtmäßige Zugriff auf Information verweigert wird. Beide Ansätze sind daher nicht zur Anwendung auf Systeme in der realen Welt geeignet.

Das Ziel dieser Arbeit besteht darin die Benutzbarkeit von dynamischen IFC Mechanismen zu verbessern indem Techniken entwickelt werden, die die Genauigkeit und Toleranz steigern. Die Arbeit präsentiert eine korrekte (engl. 'sound') Erweiterung der permissive-upgrade Strategie (eine Standardstrategie für dynamische IFC), die die Toleranz der Strategie verbessert und sie weithin anwendbar macht. Darüber hinaus präsentiere ich eine neue dynamische IFC Analyse, die auch komplexe Funktionen, wie unstrukturierte Kontrollflüsse und Exceptions in Hochsprachen, abbildet. Obwohl Unbeeinflussbarkeit eine wünschenswerte Eigenschaft ist, gibt es Anwendungen, die rechtmäßigen Zugang zu sensiblen Daten benötigen um ihre Funktion zu erfüllen. Um dies zu ermöglichen präsentiert diese Arbeit einen Ansatz, der die ungewollte Weitergabe von Information quantifiziert und anhand eines vordefinierten Grenzwertes freigibt. Diese Techniken wurden in einen Web-Browser integriert, welcher es erlaubt die Definition von flexiblen und nützlichen Informationsflussregeln für Web Anwendungen umzusetzen.

Acknowledgment

My first thanks goes to my primary advisor, Christian Hammer, for believing in me and offering me a research position. Apart from teaching me important lessons in research and shaping my development as a PhD student, he helped me get settled at a new home away from home. He has been highly supportive in whatever initiatives I took. Additionally, I would like to thank my secondary advisor, Deepak Garg, who taught me about how foundations are necessary for building a proper system. Besides, the other not-so-research based talks and advices from both have helped me take further steps in my career. The acknowledgment would be incomplete without mentioning my master's supervisor, Prof. Ramesh Joshi, who introduced me to the wonderful field of research at IIT Roorkee. I am grateful to all my committee members, Bernd Finkbeiner, Andreas Zeller and Marco Patrignani, for agreeing to be a part of my thesis committee and putting up with my requirements.

Naturally, many thanks also go to my closest collaborator and co-author, Vineet Rajani, without whom this thesis would not have been possible and with whom I had really great and fruitful discussions on security and every other topic in the world. Together, we had a lot of fun while doing research and outside of it. I would also like to thank my friends and colleagues — Abhishek Tiwari, Aravinda Krishnan, Himangshu Saikia, Kaustubh Beedkar, and Milivoj Simeonovski who added loads of fun and food to my PhD and travel as part of it. In particular, I want to mention my colleagues David Pfaff, Jenny Hotzkow, and Marie-Therese Walter, with whom I had the pleasure of sharing the office and who helped me in handling all the German documents and procedures, and also writing the German abstract of my dissertation. Special thanks to CISPA secretaries Stephanie Feyahn, Isa Maurer, and Sabine Nermerich, who literally did all the bureaucratic tasks that I had to and helped me in every way possible. I would also like to thank RS3, a special priority program under the DFG, for primarily sponsoring my research and Heiko Mantel and his group for organizing and taking care of all the administrative stuff while we enjoyed the project meetings.

I would also like to acknowledge Karthikeyan Bhargavan, Aseem Rastogi and Kapil Vaswani to have given me the opportunity to work as a research intern with them. It was a wonderful and highly interesting work environment for me and greatly helped me explore other related fields of research. I met and made friends with Rahul Kidambi, Sameer Wagh, and Shruti Tople during these internships who also helped me in making career choices besides all the fun-filled game and movie sessions. Thanks are also due to Jean Yang for offering me a research position before my defense was scheduled and giving me feedbacks on all my endeavors since I have joined her group, in particular, my defense presentation.

I am deeply grateful to my family and friends for their unconditional love, their faith in me, and their advice in all situations. Without them, I would not have been able to go my own way during the years of my PhD and would certainly not stand where I am today.

Contents

I	Introduction and Background	1
1	Introduction	3
1.1	Contributions of the Thesis	5
2	Background and Overview	7
2.1	Information Flow Control	7
2.2	Information Release	9
2.3	Dynamic Information Flow Control	10
II	Generalized Permissive-Upgrade Strategy	13
3	Improved Permissive-Upgrade	15
3.1	Overview	15
3.2	Austin and Flanagan’s Permissive-Upgrade Strategy	18
3.3	Improved Permissive-Upgrade Strategy	20
4	Generalized Permissive-Upgrade	23
4.1	Generalization of the Improved Permissive-Upgrade Strategy	23
4.2	Generalized Permissive-Upgrade on Arbitrary Lattices	24
4.3	Comparison of the Generalization of Section 4.2 with the Generalization of Section 4.1	28

III	Precise Control Scope Analysis	31
5	Dynamic IFC with Unstructured Control Flow and Exceptions	33
5.1	Control Flow Graphs and Post-dominator Analysis	34
5.2	Exceptions and Synthetic Exit Nodes	35
5.3	Formal Model	36
IV	Budget-based Limited Information Release	41
6	Bounding Information Leaks Dynamically	43
6.1	Quantifying Information Leaks	45
6.2	Limited Information Release	46
6.3	LIR Enforcement	48
6.4	Formalization of LIR	56
6.5	Soundness and Decoding Semantics	58
V	Application to an IFC-enabled Web Browser	63
7	Information Flow Policies for Web Browsers	65
7.1	Policy Component for Web Browsers	65
7.2	WebPol policy model	66
7.3	Expressiveness of WebPol	69
8	Policy Implementation and Evaluation in an IFC-enabled Browser	75
8.1	Implementation and Evaluation of WebPol	76
8.2	Implementation and Evaluation of LIR	79
9	Related Work	81
VI	Conclusion and Outlook	85
10	Conclusion	87
11	Future Directions	89
11.1	Evaluating Information Flow Policies on Real-World Websites	89
11.2	Exploring Alternative Granularities for Enforcing Information Flow Control	89

Contents	xi
11.3 Handling Timing Leaks on the Web	90
Appendix	91
A Proofs for Improved and Generalized Permissive Upgrade	93
B Proofs of Precision for Dynamic IFC with Unstructured Control Flow and Exceptions	104
C Proofs for IFC with Unstructured Control Flow and Exceptions	106
D Proofs for Limited Information Release	111
 List of Figures	 121
 List of Tables	 123
 Bibliography	 132

Part I

Introduction and Background

Chapter 1

Introduction

Isn't it amazing how the computer has evolved from a machine that filled a large room to what occupy our offices and pockets today! With the growth in the use of computers and the Internet for almost every application, the amount of data, in particular sensitive data, that programs compute has also increased dramatically. For instance, e-commerce websites have access to the personal details of a user including her address, credit-card details etc. Similarly, a hospital database stores sensitive and private medical data of its patients. While the complexity of such systems has increased manifold in the last few decades, the privacy and confidentiality guarantees still remain questionable [5]. Often leaks occur either due to buggy programs or due to malicious code. Cryptographic techniques protect data by encrypting it but many programs need to operate on confidential data, which requires the data to be available in plaintext form. To assist this, *access control* techniques have been widely used for preventing leakage of confidential data. However, once the authorized users have access to data, there is no control on how the data can be used. The data might be input to a program that writes to publicly-observable locations or outputs data that is accessible to all users.

Consider the Web, for instance. Web applications rely extensively on third-party JavaScript to provide useful libraries, page analytics, advertisements and many other features [78]. In such a *mashup* model, wherein the hosting page and the included scripts share the page's state (called the DOM), by design, all included third-party scripts run with the same access privileges as the hosting page. While some third-party scripts are developed by large, well-known, trustworthy vendors, many other scripts are developed by small, domain-specific vendors whose commercial motives do not always align with those of the web-page providers and users. This leaves sensitive information such as passwords, credit card numbers, email addresses, click histories, cookies and location information vulnerable to inadvertent bugs and deliberate exfiltration by third-party scripts. In many cases, developers are fully aware that a third-party script accesses sensitive data to provide useful functionality, but they are unaware that the script also leaks that data on the side. In fact, this is a widespread problem [59]. The traditional browser security model is based on restricting scripts' access to data, not on tracking how scripts use data. Existing web security standards like the same-

origin policy [19] and web browsers address this problem unsatisfactorily, favoring functionality over privacy. The same-origin policy implemented in all major browsers restricts a web-page and third-party scripts included in it to communicating with web servers from the including web-page's domain only. However, broad exceptions are allowed. For instance, there is no restriction on request parameters in URLs that fetch images and, unsurprisingly, third-party scripts leak information by encoding it in image URLs. The candidate web standard Content Security Policy [3], also implemented in most browsers, allows a page to white list scripts that may be included, but places no restriction on scripts that have been included, thus not helping with the problem above. Quite a few *fine-grained* access control techniques have also been proposed [1, 4, 38, 45, 68, 74, 95, 105]. However, all these techniques enforce only access policies and cannot control what a script does with data it has been provided in good faith, i.e., if a third-party script was allowed access to some data only for local computations, the policy places no restriction on the script for sending it on the network while performing the computation. More broadly, no mechanism based only on access control can solve the problem of information leakage.

The academic community has proposed solutions based on *information flow control* (IFC). It ensures the security of confidential information even in the presence of untrusted and buggy code. The idea is to track the flow of information through the program and prevent any undesired flows based on a security policy. Research has considered static methods such as type checking and program analysis, which verify the security policy at compile time [40, 41, 53, 58, 76, 79, 87, 100], dynamic methods that track information flow through program execution at runtime [10, 13–15, 21, 43, 46, 54, 89, 94, 102], and hybrid approaches that combine both static and dynamic analyses to add precision to the analysis [22, 51, 52, 56, 77, 85, 98] for handling information leaks described above.

Dynamic analysis has the drawback of introducing significant performance overheads at runtime, though it can be more permissive than static analysis methods in certain cases [85]. Although helpful in most scenarios, static analyses are mostly ineffective when working with dynamic languages like JavaScript, which is an indispensable part of the modern Web. The dynamic nature of JavaScript [82, 83] with features like dynamic typing, `eval`, scope-chains, prototype chains etc. makes sound static analysis difficult. Thus, recent research has focused on dynamic analysis for enforcing information flow control especially in languages like JavaScript. Even though research in dynamic information flow control has made significant inroads in the last decade, the applicability of these techniques still remains bleak, *permissiveness* being the major challenge to the practicality of dynamic information flow control.

In recent years, researchers have also explored methods to quantify the amount of sensitive information leaked by a program as an alternative to the qualitative notion of information flow control. The developer determines an “acceptable” upper-bound on the number of bits of sensitive information a program may leak. Programs that leak less than this “acceptable” amount of information are considered secure. The information released is generally quantified as the knowledge gained by an adversary about the sensitive data as a result of the information flow in the program. Specifically, the adversary has knowledge about the possible set of initial values that can be associated with the sensitive data. A flow in the program could reduce the number of possibilities, thereby increasing the knowledge of the adversary about the specific value associated with the sensitive data in that execution of the program. This change in knowledge can be measured as the number of bits of in-

formation of sensitive data released to the adversary [37]. Statically quantifying the information leaked by a program has been well studied in the literature and various static quantitative information flow measures [7, 34, 71, 92, 93] and techniques [16, 33, 34, 42, 64, 93] have been proposed. However, quantifying information leaks in a purely dynamic setting is still largely an open problem (prior work is limited to [73]) as dynamic analysis is generally limited to the current execution of the program in contrast to the static methods, which analyze the program as a whole.

1.1 Contributions of the Thesis

This thesis focuses on improving the usability of dynamic information flow techniques by presenting mechanisms that enhance the precision and permissiveness of the analyses. The main contributions are as follows:

- **Generalized Permissive-Upgrade Strategy.** To improve the permissiveness of dynamic information flow analysis, the thesis presents a sound improvement and enhancement of the permissive-upgrade strategy (a strategy widely used to enforce dynamic information flow control). The development improves the original strategy's permissiveness and applicability by generalizing the approach to an arbitrary security lattice, in place of the two-point lattice considered in prior work.
- **Precise Control Scope Analysis.** Most of the existing work in dynamic information flow control does not handle complex features like unstructured control flow and exceptions. The proposals that handle these features are too conservative and require additional annotations in the program. This thesis presents a sound and precise dynamic control scope analysis for handling these features without requiring any additional annotations from the developer.
- **Bounding Information Leaks Dynamically.** Although non-interference is a desired property for enforcing information flow control, there are program instances that require legitimate release of some parts of the secret data to provide the required functionality. However, in many cases it might not be known upfront as to what information needs to be released. Towards this end, this thesis develops a sound approach to bound information leaks dynamically while allowing information release in accordance with a pre-specified *budget*.
- **Application to a Web Browser.** The thesis concludes by applying these techniques to an information flow control-enabled Web browser. An information flow control system enforces security policies that are a collection of rules for labeling private information sources, generally specified by a policy component in the system. To complement the work on enforcement components in Web browsers, this thesis explores a policy specification mechanism to specify flexible and useful information flow policies for Web applications.

Chapter 2

Background and Overview

2.1 Information Flow Control

Information flow control (IFC) refers to controlling the flow of (confidential) information through a program based on a given security policy. Typically, pieces of information are classified into security *labels* and the policy is a lattice over labels. Information is only allowed to flow up the lattice. For illustration purposes often the smallest non-trivial lattice $L \sqsubset H$ is used, which specifies that public (low, L) data must not be influenced by confidential (high, H) data. Information flow control can be used to provide confidentiality (or integrity) of secret (trusted) information; the work in this thesis is, however, limited to confidentiality guarantees. Roughly the idea behind information flow control is that an adversary can view all the public outputs of a program. By preventing private or sensitive data to flow to public outputs, the adversary does not get any information about the private or sensitive data.

The seminal work by Denning [40–42] defined most of the theoretical ideas pertaining to information flow control. In general, information can flow along many channels. However, this thesis considers two of the most important flows — *explicit* and *implicit* — in deterministic programs. Covert channels like timing or resource usage are beyond the scope of this thesis.

An *explicit flow* occurs as a result of direct assignment, e.g., the statement `public = secret + 1` causes an explicit flow from `secret` to `public`. An *implicit flow* occurs due to the control structure of the program. For instance, consider the program in Listing 2.1. The final value of `y` is equal to the value of `z` even though there is no direct assignment from `z` to `y`. Leaking a bit like this can be magnified into leaking a bigger secret bit-by-bit [12].

The correctness of the approaches enforcing information flow control is often stated in terms of a well-defined security property known as *non-interference* [49], which basically stipulates that high or secret inputs of a program must not influence its low or publicly observable outputs. While non-interference is too strong a property in practice, different variants of the definition are proven.

```

1  x = false, y = false
2  if (not(z))
3    x = true
4  if (not(x))
5    y = true

```

Listing 2.1: Implicit flow from z to y

One such variant of non-interference usually established for information flow control techniques is *termination-insensitive non-interference* [100]. Roughly, a program is termination-insensitive non-interferent if any two terminating runs of the program starting from low-equivalent heaps (i.e., heaps that look equivalent to the adversary assuming that an adversary can observe some part of the heap) end in low-equivalent heaps. Termination-insensitive means that one-bit of leak is tolerable when an adversary checks whether or not the program terminated. In particular, this discounted one-bit leak accounts for termination due to the failure of a runtime security check. Askarov *et al.* [12] show that for programs with intermediate observable outputs, termination-insensitivity may leak more than one bit but also show that this attack is limited to a brute-force attack.

Information flow control approaches either statically analyze the program at compile-time and reject or accept a program based on the security policy, or dynamically analyze the flow of data in a program by either using a modified runtime or with the help of a reference monitor, or employ a combination of both the approaches to shrug off some of the individual limitations of the two approaches. Static approaches normally support restricted policies, and are less permissive in nature. Besides, it becomes extremely complicated to use such methods with dynamically-typed languages and languages that are loaded with dynamic features. Dynamic approaches, on the other hand, add runtime overheads and are less precise as compared to the static approaches. Dynamic information flow control, which is the central theme of this thesis, is described in detail later in the chapter.

2.1.1 Flow-sensitivity

Static analysis techniques are usually flow-insensitive in nature, i.e., they do not account for the ordering of the instructions in the program or the general flow of execution of the program. In other words, all operations should be individually secure to secure the program as a whole. For instance, consider the program snippet:

$$l = h$$

Assume that h is a secret variable labeled H and l is a public variable labeled L . This program gets rejected by a flow-insensitive analysis as the assignment of a secret value to a public variable is considered insecure.

Flow-sensitive analysis improves the permissiveness of static analysis techniques and is mostly used in dynamic analyses. Under a flow-sensitive static analysis, the above program upgrades the labels of l to H indicating the influence of H -labeled h on l . If the program was to output the value of l on a L -observable channel later, the program gets rejected by the analysis. Flow-sensitive analysis also accepts programs with dead-code like the one shown in Listing 2.2, thus improving the

```
1 var x
2 if (not(z))
3   x = true
4   x = false
```

Listing 2.2: Permissiveness of flow-sensitive analysis

permissiveness of the analysis. Some flow-sensitive static analyses have been proposed to enforce information flow control [53, 58] while almost all dynamic approaches are flow-sensitive in nature.

2.2 Information Release

Non-interference prevents all unauthorized flows without regarding the severity of the leaks. However, many practical applications require some fragment of the sensitive data to be released to programs for providing proper functionality or to improve the user-experience. For instance, according to non-interference even a standard log-in application would be considered insecure because it leaks information about the user’s password, which is considered confidential, by either allowing or denying access. As a remedy, researchers have proposed various ways to intentionally release or *declassify* sensitive data through relaxations of non-interference [91], which is generally enforced using annotations like `declassify` in the program. For instance, `declassify(pwd == input)` would treat the result of the password check as public, thus allowing access-information to be released to the user. Most of these approaches require policy specifications by the developer that define what information can be released by the program as in many settings it is difficult to trust third-parties to include correct and appropriate (`declassify`) annotations for information release in the code. Declassification approaches generally ensure that the adversary is unable to declassify or force declassification of sensitive information as per his/her needs. Sabelfeld and Sands [91] present a survey of many such techniques that allow release of information in a secure manner.

Some of these declassification techniques are described below. Cohen [35, 36] in his work on selective dependency used assertions to eliminate certain information paths thereby allowing programs that satisfy the constraint to be accepted. The idea was that information flows to the adversary should not allow her to deduce more information about secrets than is allowed by the assertions. Volpano and Smith [99] present relative secrecy, a property that instead of providing absolute secrecy allows information to be released through specific *match* queries relative to the size of the secret, i.e., if the adversary cannot guess the secret in polynomial time. Giacobazzi and Mastroeni [48] generalize non-interference by modelling what property the adversary can observe about the secrets. Li and Zdancewic’s relaxed non-interference characterizes the information released through downgrading policies [66]. Sabelfeld and Sands [90] proposed the PER model using partial equivalence relations for specifying information flow policies. Sabelfeld and Myers [88] introduced the notion of delimited release, which enables the specification of declassification policy using *escape hatches*. The policy specifies upfront what information can be released through these escape hatches. Localized delimited release [9] requires that an expression is only released after a declassification operation for this expression has appeared in the code. Lux and Mantel [69] propose explicit reference points

to allow flexible specification of what secrets may be declassified at the specific reference points. Robust declassification [103] ensures that only trusted code is allowed to declassify information. Chong and Myers [30] propose a framework for specifying application-specific information release policies. Chong [29] further proposes the concept of required information release, which considers applications that are obligated to release information under certain conditions. Askarov and Sabelfeld introduce the concept of gradual release [8] for allowing release of information at certain release or declassification points. The policy does not allow the adversary knowledge to refine its knowledge of secrets at points other than the release points. Banerjee *et al.* [17] present a way to specify declassification policies that satisfy conditioned gradual release, which is an extension of the gradual release property. Similarly, another security property, *policy controlled release*, that extends gradual release was presented by Rocha *et al.* [84] for declassification in untrusted and legacy programs.

2.3 Dynamic Information Flow Control

Dynamic IFC usually works by tracking taints or labels on individual program values in the language runtime. A label represents a mandatory access policy on the value. A value v labeled ℓ is written v^ℓ .

Flow-sensitive dynamic IFC analysis *propagates* labels as data flows during program execution. *Explicit* flows are generally handled by carrying over the label of the computed value to the variable being assigned. For example, in the statement $x = y + z$, the result of computing $y + z$ will have the label that is a join of the individual labels on y and z , which is the final label of x , i.e., if either of y or z is labeled confidential or H , then the final label of x is also labeled H ¹.

Implicit flows in a flow-sensitive IFC analysis are tracked by maintaining an additional taint, usually called the program counter taint or program context taint or pc , which is an upper bound on the label of all the control dependencies that lead to the current instruction being executed. For example, in the program of Listing 2.1, the value in variable x at the end of line 3 depends on the value in z . If z is labeled H , then at line 3, $pc = H$ because of the branch in line 2 that depends on z . Thus, by tracking pc , dynamic IFC can enforce that x has label H at the end of line 3, thus taking into account the control dependency.

However, simply tracking control flow dependencies via pc is not enough to guarantee absence of information flows when labels are flow-sensitive, i.e., when the same variable may hold values with different labels depending on what program paths are executed. The program in Listing 2.1 is a classic counterexample, taken from [13]. Assume that z is labeled H and x and y are labeled L initially. The final value in y is computed as a function of the value in z . If z contains true^H , then y ends with true^L : The branch on line 2 is not taken, so x remains false^L at line 4. Hence, the branch on line 4 is taken, but $pc = L$ at line 5 and y ends with true^L . If z contains false^H , then similar reasoning shows that y ends with false^L . Consequently, in both cases y ends with label L and its value is exactly equal to the value in z . Hence, an adversary can deduce the value of z by observing y at the end (which is allowed because y ends with label L). So, this program leaks information

¹“ z is labeled H ” actually means “the value in z is labeled H ”. This convention is used consistently.

about z despite correct use of pc .

2.3.1 No-sensitive-upgrade Check

Preventing leaks due to implicit flow in dynamic IFC requires coarse approximation because a dynamic monitor only sees program branches that are executed and does not know what assignments may happen in alternate branches in other executions. One such coarse approximation is the *no-sensitive-upgrade* (NSU) check proposed by Zdancewic [104]. In the program in Listing 2.1, x 's label is upgraded from L to H at line 3 in one of the two executions above, but not the other. Subsequently, information leaks in the other execution (where x 's label remains L) via the branch on line 4. The NSU check stops the leak by preventing the assignment on line 3. More generally, it stops a program whenever a public variable's label is upgraded due to a high pc . This check suffices to provide termination-insensitive non-interference as shown by Austin and Flanagan [13].

2.3.2 Permissive-upgrade and Faceted Execution

To tackle the issue of permissiveness with the no-sensitive-upgrade strategy, Austin and Flanagan proposed the permissive-upgrade strategy [14] and faceted execution [15]; faceted execution being the most permissive of the three. Faceted execution simulates multiple executions simultaneously within a single runtime. They introduce the concept of faceted values that are pairs of values for both low and high observers. With multiple levels, each value in the pair is represented as a pair. When branching on a faceted value, multiple executions are simulated for the different values in the facet. However, the runtime overheads of faceted execution are quite prohibitive for multiple security levels. This thesis, thus, considers only the permissive-upgrade strategy, which is more permissive than the no-sensitive-upgrade strategy and much less performance-intensive compared to faceted execution. Permissive-upgrade is described in detail in Section 3.2.

2.3.3 Other Approaches for Permissiveness

Secure multi-execution [43] is another approach for enforcing non-interference at runtime. Instead of tracking information flow through the program, the approach executes multiple copies of the program with different values of sensitive data. Conceptually, one executes the same code once for each security level (like low and high) with a few constraints. The private data in the low execution are replaced by default values, i.e., the public copy of the program does not see the actual value of the private data but a pre-determined default value, and outputs on an ℓ -labeled channel are permitted only in the ℓ -level execution of the program, i.e., the high execution of the program outputs on the high channel, if any, and the low execution of the program outputs on the low channel, typically the network. The modification of the semantics forces that private data and the outputs resulting from it are visible to only high-level observers. The public-level observers or the adversary observe inaccurate results for the outputs that depend on the private value as the value is replaced by the default value in that execution. This modification forces even unsafe programs to adhere to non-interference. Additionally, secure multi-execution guarantees precision, i.e., the semantics of a secure program is not altered. Thus, for a secure program the outputs are all accurate. Secure

multi-execution normally guarantees termination-insensitive non-interference as the high execution may not terminate in some cases. Flowfox [39] demonstrates secure multi-execution in the context of web browsers. However, executing a program multiple times can be prohibitive for a security lattice with multiple levels [15]. The runtime overhead incurred can be reduced if the executions are run in parallel (which requires more hardware resources), though the program has to be run for all levels irrespective of whether it uses private data or not. In addition to this, secure multi-execution makes declassification complicated as it requires synchronization between different executions [80].

Birgisson *et al.* [27] describe a testing-based approach that adds variable upgrade annotations to avoid halting on the NSU check in an implementation of dynamic IFC for JavaScript [54]. Hritcu *et al.* improve permissiveness by making IFC errors recoverable in the language Breeze [57]. This is accomplished by a combination of two methods: making all labels public (by upgrading them when necessary in a public *pc*) and by delaying exceptions. A different way of handling the problem of implicit flows through flow-sensitive labels is to assign a (fixed) label to each label; this approach has been examined in recent work by Buiras *et al.* in the context of a language with a dedicated monad for tracking information flows [28]. The precise connection between that approach and permissive-upgrade remains unclear, although Buiras *et al.* sketch a technique related to permissive-upgrade in their system, while also noting that generalizing permissive-upgrade to arbitrary lattices is non-obvious. This thesis confirms the latter and shows how it can be done.

Part II

Generalized Permissive-Upgrade Strategy

Chapter 3

Improved Permissive-Upgrade

The no-sensitive-upgrade (NSU) check described earlier provides the basic foundations for sound dynamic IFC. However, terminating a program preemptively because of the NSU check is quite restrictive in practice. For example, consider the program of Listing 3.1, where z is labeled H and y is labeled L . This program potentially upgrades variable x at line 3 under a high pc , and then executes function f when y is `true` and executes function g otherwise. Suppose that f does not read x . Then, for $y \mapsto \text{true}^L$, this program leaks no information, but the NSU check would terminate this program prematurely at line 3. (Note: g may read x , so x is not a dead variable at line 3.)

To improve permissiveness, Austin and Flanagan [14] proposed the permissive-upgrade strategy as a replacement for NSU. However, that development lacks permissiveness in certain cases. This chapter presents the soundness results of the permissive-upgrade strategy with the improvement for further permissiveness in place.

3.1 Overview

This section presents a *formal* description of the no-sensitive-upgrade check. The technical development in this thesis is mostly based on the simple imperative language shown in Figure 3.1. However, the key ideas are orthogonal to the choice of language and generalize to other languages easily. The

```
1 x = false
2 if (not(z))
3   x = true
4 if (y) f() else g()
5 x = false
```

Listing 3.1: Impermissiveness of the NSU strategy

$$\begin{aligned}
e &= \mathbf{n} \mid \mathbf{x} \mid e_1 \odot e_2 \\
c &= \mathbf{skip} \mid \mathbf{x} := e \mid c_1; c_2 \mid \mathbf{if} \ e \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2 \mid \mathbf{while} \ e \ \mathbf{do} \ c \\
\ell &= L \mid M \mid H \mid \dots \\
k, l, m, pc &= \ell
\end{aligned}$$

Figure 3.1: Syntax of the Language

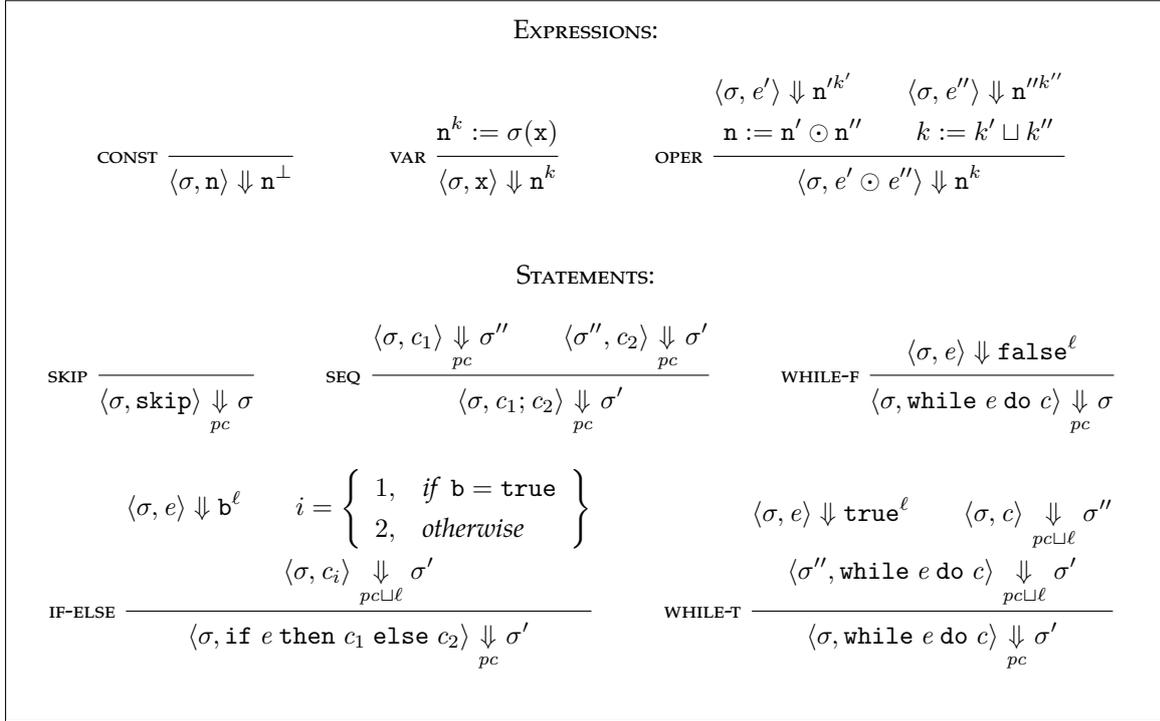


Figure 3.2: Semantics

use of a simpler language is to simplify non-essential technical details. The parts in this thesis that require a more complex language define the additional features in place. The language's expressions include constants or values (\mathbf{n} , \mathbf{b}), variables (\mathbf{x}) and unspecified binary operators (\odot) to combine them. The set of variables is fixed upfront. Labels (ℓ) are drawn from a fixed security lattice. The lattice contains different labels $\{L, M, H, \dots\}$ with a partial ordering between the elements. Join (\sqcup) and meet (\sqcap) operations are defined as usual on the lattice. The program counter label pc is an element of the lattice.

3.1.1 Basic IFC Semantics

The rules in Figure 3.2 define the big-step semantics of the language, including standard taint propagation for IFC: the evaluation relation $\langle \sigma, e \rangle \Downarrow \mathbf{n}^k$ for expressions, and the evaluation relation $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$ for commands. Here, σ denotes a store, a map from variables to labeled values of the

$$\text{ASSN-NSU} \frac{l = \Gamma(\sigma(x)) \quad pc \sqsubseteq l \quad \langle \sigma, e \rangle \Downarrow \mathbf{n}^m}{\langle \sigma, x := e \rangle \Downarrow_{pc} \sigma[x \mapsto \mathbf{n}^{(pc \sqcup m)}]}$$

Figure 3.3: Assignment rule for NSU

form \mathbf{n}^k . \mathbf{b} represents a Boolean constant. For now, labels $k ::= \ell$; this is generalized later when the “partially-leaked” taints are introduced in Section 3.2.

The evaluation relation for expressions evaluates an expression e and returns its value n and label k . The label k is the join of labels of all variables occurring in e (according to σ). The relation for commands executes a command c in the context of a store σ , and the current program counter label pc , and yields a new store σ' . The function $\Gamma(\sigma(x))$ returns the label associated with the value in x in store σ : If $\sigma(x) = \mathbf{n}^k$, then $\Gamma(\sigma(x)) = k$. \perp denotes the least element of the lattice.

The sequencing rule (SEQ) evaluates the command c_1 under store σ and the current pc label; this yields a new store σ'' . It then evaluates the command c_2 under store σ'' and the same pc label, which yields the final store σ' . The IF-ELSE rule evaluates the branch condition e to a Boolean value \mathbf{b} with label ℓ . Based on the value of \mathbf{b} , one of the branches c_1 and c_2 is executed under a pc obtained by joining the current pc and the label ℓ of \mathbf{b} . Similarly, the rules for `while` (WHILE-T and WHILE-F) evaluate the loop condition e and execute the loop command c_1 while e evaluates to `true`. The pc for the loop is obtained by joining the current pc and the label ℓ of the result of evaluating e .

The rule for assignment statements are conspicuously missing from Figure 3.2 because they depend on the strategy used to control implicit flows. The rule for assignment (ASSN-NSU) corresponding to the NSU check is shown in Figure 3.3. The rule checks that the label l of the assigned variable x in the initial store σ is at least as high as pc (premise $pc \sqsubseteq l$). If this condition is not true, the program gets stuck. This is exactly the NSU check described in Section 2.3.1.

3.1.2 Formalization of the No-sensitive-upgrade Check

For establishing and proving the security property of termination-insensitive non-interference (TINI), the observational power of the adversary needs to be defined. An adversary at level ℓ in the lattice is allowed to view all values that have a label less than or equal to ℓ . To prove the security property of non-interference, it is enough to show that when executing a program beginning with two different memory stores that are *observationally equivalent* to an adversary, the final memory stores are also *observationally equivalent* to the adversary. For this, the observational equivalence of two memory stores with respect to an adversary needs to be defined. Store equivalence is formalized as a relation \sim_ℓ , indexed by lattice elements ℓ , representing the adversary.

Definition 1 (Value equivalence). *Two labeled values \mathbf{n}_1^k and \mathbf{n}_2^m are ℓ -equivalent, written $\mathbf{n}_1^k \sim_\ell \mathbf{n}_2^m$, iff either:*

1. $(k = m) \sqsubseteq \ell$ and $\mathbf{n}_1 = \mathbf{n}_2$ or
2. $k \not\sqsubseteq \ell$ and $m \not\sqsubseteq \ell$

This definition states that for an adversary at security level ℓ , two labeled values n_1^k and n_2^m are equivalent iff either ℓ can access both values and n_1 and n_2 are equal, or it cannot access either value ($k \not\sqsubseteq \ell$ and $m \not\sqsubseteq \ell$). In the lattice $L \sqsubset H$, two values labeled L and H are distinguishable for the L -adversary.

Definition 2 (Store equivalence). *Two stores σ_1 and σ_2 are ℓ -equivalent, written $\sigma_1 \sim_\ell \sigma_2$, iff for every variable x , $\sigma_1(x) \sim_\ell \sigma_2(x)$.*

The following theorem states TINI for the NSU check. The theorem has been proved for various languages in the past.

Theorem 1 (TINI for NSU). *With the assignment rule ASSN-NSU from Figure 3.3, if $\sigma_1 \sim_\ell \sigma_2$ and $\langle \sigma_1, c \rangle \Downarrow_{pc}$ σ'_1 and $\langle \sigma_2, c \rangle \Downarrow_{pc}$ σ'_2 , then $\sigma'_1 \sim_\ell \sigma'_2$.*

Proof. Standard, see e.g., [13] □

3.2 Austin and Flanagan’s Permissive-Upgrade Strategy

To allow a dynamic IFC analysis to accept safe executions of programs with variable upgrades due to high pc , Austin and Flanagan proposed a less restrictive strategy called the *permissive-upgrade strategy* [14]. They study this strategy for a two-point lattice $L \sqsubset H$ and their strategy does not immediately generalize to arbitrary security lattices. Whereas NSU stops a program when a variable’s label is upgraded due to assignment in a high pc , permissive-upgrade allows the assignment, but labels the variable as *partially-leaked* or P . The exact intuition behind the partially-leaked label P is the following:

A variable with a value labeled P may have been implicitly influenced by H -labeled values in this execution, but in other executions (obtainable by changing H -labeled values in the initial store), this implicit influence may not exist and, hence, the variable may be labeled L .

The program must be stopped later if it tries to use or case-analyze the variable (in particular, branching on a partially-leaked Boolean variable is stopped). Permissive-upgrade also ensures termination-insensitive non-interference, but is strictly more permissive than NSU. For example, permissive-upgrade stops the leaky program of Listing 2.1 at line 4 when z contains false^H , but it allows the program of Listing 3.1 to execute to completion when y contains true^L .

In the revised syntax of labels, summarized in Figure 3.4, the labels k, l, m on values can be either elements of the lattice (L, H) or P . The pc can only be one of L, H because branching on partially-leaked values is prohibited. The join operation \sqcup is lifted to labels including P . Joining any label with P results in P . For brevity in definitions, they extend the order \sqsubset to $L \sqsubset H \sqsubset P$. However, P is not a new “top” member of the lattice because it receives special treatment in the semantic rules.

$$\begin{array}{ll}
\ell = L \mid H & k \sqcup k = k \\
pc = \ell & L \sqcup H = H \\
k, l, m = \ell \mid P & L \sqcup P = P \\
& H \sqcup P = P
\end{array}$$

Figure 3.4: Syntax of labels including the partially-leaked label P

The rule for assignment with permissive-upgrade is

$$\text{ASSN-PUS} \frac{l := \Gamma(\sigma(\mathbf{x})) \quad \langle \sigma, e \rangle \Downarrow \mathbf{n}^m}{\langle \sigma, \mathbf{x} := e \rangle \Downarrow_{pc} \sigma[\mathbf{x} \mapsto \mathbf{n}^k]}$$

where k is defined as follows:

$$k = \begin{cases} m & \text{if } pc = L \\ m \sqcup H & \text{if } pc = H \text{ and } l = H \\ P & \text{otherwise} \end{cases}$$

The first two conditions in the definition of k correspond to the NSU rule (Figure 3.3). The third condition applies, in particular, when a variable whose initial label is L is assigned with $pc = H$. The NSU check would stop this assignment. With permissive-upgrade, however, the updated variable is labeled P , consistent with the intuitive meaning of P . This allows more permissiveness by allowing the assignment to proceed in all cases. To compensate, any program (in particular, an adversarial program) is disallowed from case analyzing any value labeled P . Consequently, in the rules for `if-then` and `while` (Figure 3.2), the label of the branch condition is of the form ℓ , which does not include P . Thus, assignments under high pc succeed under the permissive-upgrade check but branching or case-analyzing a partially-leaked value is not permitted as that can also leak information.

The noninterference result obtained for NSU earlier can be extended to permissive-upgrade by changing the definition of store equivalence. Because no program can case-analyze a P -labeled value, such a value is equivalent to any other labeled value.

Definition 3. Two labeled values \mathbf{n}_1^k and \mathbf{n}_2^m are equivalent to an adversary at level L , written $\mathbf{n}_1^k \sim_L \mathbf{n}_2^m$, iff either:

1. $(k = m) = L$ and $\mathbf{n}_1 = \mathbf{n}_2$ or
2. $k = H$ and $m = H$ or
3. $k = P$ or $m = P$

Definition 4. Two stores σ_1 and σ_2 are L -equivalent, written $\sigma_1 \sim_L \sigma_2$, iff $\forall \mathbf{x}. \sigma_1(\mathbf{x}) \sim_L \sigma_2(\mathbf{x})$.

Theorem 2 (TINI for permissive-upgrade with a two-point lattice). *With the assignment rule `assn-PUS`, if $\sigma_1 \sim_L \sigma_2$ and $\langle c, \sigma_1 \rangle \Downarrow_{pc} \sigma'_1$ and $\langle c, \sigma_2 \rangle \Downarrow_{pc} \sigma'_2$, then $\sigma'_1 \sim_L \sigma'_2$.*

Proof. See [14]. □

```

1  y = false
2  if (not(x))
3    y = true
4  z = y || x
5  if (not(z))
6    w = true

```

Listing 3.2: Example showing the impermissiveness of the original permissive-upgrade strategy

3.3 Improved Permissive-Upgrade Strategy

The original permissive-upgrade strategy, however, lacks permissiveness; it rejects secure programs like the one shown in Listing 3.2. Consider that x is labeled H and w, y are labeled L . With the original permissive-upgrade strategy, the label of z on line 4 would remain P and the execution would be terminated when branching on z on line 5. With the improvement

$$H \sqcup P = H$$

the analysis can accept such programs while remaining sound. With the improvement, z would be labeled H on line 4, which would allow the execution to branch on line 5, thus, taking the execution to completion. The idea behind the improvement is that an H -labeled value is never observable at L -level. Similarly, the result of any operation involving an H -labeled value is also never observable at L -level. Thus, any operation involving a partially-leaked value and a H -labeled value does not reveal any information to an adversary at level L about the partially leaked value.

The final label k in the assignment rule ASSN-PUS under the improved permissive-upgrade strategy becomes:

$$k = \begin{cases} m & \text{if } pc = L \\ H & \text{if } pc = H \text{ and } l = H \\ P & \text{otherwise} \end{cases}$$

The soundness results of the original permissive-upgrade strategy can be extended to show the soundness of the improved permissive-upgrade strategy. However, a significant difficulty in proving the theorem using the modified notation for the imperative language is that the definition of \sim is not transitive. In [14], the authors resolve this issue by defining a special relation called evolution. The need for evolution is averted here using the auxiliary lemmas listed below. Lemma 2 proves the required result substituting evolution.

Lemma 1 (Expression Evaluation). *If $\langle \sigma_1, e \rangle \Downarrow n_1^{k_1}$ and $\langle \sigma_2, e \rangle \Downarrow n_2^{k_2}$ and $\sigma_1 \sim_L \sigma_2$, then $n_1^{k_1} \sim_L n_2^{k_2}$.*

Proof. By induction on e . □

Lemma 2 (Evolution). *If $pc = H$ and $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$, then $\forall x. \Gamma(\sigma(x)) = P \implies \Gamma(\sigma'(x)) = P$.*

Proof. By induction on the derivation rules and case analysis on the last rule. \square

Lemma 3 (Confinement for improved permissive-upgrade with a two-point lattice). *If $pc = H$ and $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$, then $\sigma \sim_L \sigma'$.*

Proof. By induction on the derivation rules. \square

Theorem 3 (TINI for improved permissive-upgrade with a two-point lattice). *With the assignment rule ASSN-PUS and the modified syntax of Figure 3.4, if $\sigma_1 \sim_L \sigma_2$ and $\langle \sigma_1, c \rangle \Downarrow_{pc} \sigma'_1$ and $\langle \sigma_2, c \rangle \Downarrow_{pc} \sigma'_2$, then $\sigma'_1 \sim_L \sigma'_2$.*

Proof. By induction on c and case analysis on the last step. \square

The detailed proofs are provided in Appendix A.1. Note that the definitions and proofs presented in this chapter are specific to the two-point lattice and with respect to an adversary at level L .

Chapter 4

Generalized Permissive-Upgrade

Although the permissive-upgrade strategy as described in the previous chapters is useful, its development in literature is incomplete so far: Austin and Flanagan’s original paper [14], and the work building on it (Chapter 3), develops permissive-upgrade for *only* a two-point security lattice, containing levels L and H with $L \sqsubset H$, and the new label P . A generalization to a pointwise product of such two-point lattices (and, hence, a powerset lattice) was suggested by Austin and Flanagan in the original paper, but not fully developed. As explained later in Section 4.1, this generalization works for the improved permissive-upgrade strategy and can be proven sound.

However, that still leaves open the question of generalizing permissive-upgrade to arbitrary lattices. It is not even clear hitherto that this generalization exists. This chapter shows by construction that a generalization of permissive-upgrade to arbitrary lattices does indeed exist and that it is, in fact, non-obvious. Specifically, the rule for adding partially-leaked labels and the definition of store (memory) equivalence needed to prove non-interference are reasonably involved.

4.1 Generalization of the Improved Permissive-Upgrade Strategy

Austin and Flanagan point out that permissive-upgrade on a two-point lattice can be generalized to a pointwise product of such lattices. This generalization can also be extended to the improved permissive-upgrade strategy presented in the previous chapter. Specifically, let X be an index set — these indices are called principals in [14]. Let a label l be a map of type $X \rightarrow \{L, H, P\}$ and let the subclass of pure labels contain maps ℓ, pc of type $X \rightarrow \{L, H\}$. The order \sqsubset and the join operation \sqcup can be generalized pointwise to these labels. Finally, the rule `ASSN-PUS` can be generalized pointwise

The content of this chapter is based on the work published as part of the paper, “Generalizing Permissive-Upgrade in Dynamic Information Flow Analysis” [21]

$$\begin{array}{ll}
\ell = L \mid M \mid \dots \mid H & \ell_1 \sqcup \ell_2^* = (\ell_1 \sqcup \ell_2)^* \\
pc = \ell & \ell_1^* \sqcup \ell_2^* = (\ell_1 \sqcup \ell_2)^* \\
k, l, m = \ell \mid \ell^* &
\end{array}$$

Figure 4.1: Labels and label operations

by replacing it with the following rule:

$$\text{ASSN-GPUS} \frac{l := \Gamma(\sigma(\mathbf{x})) \quad \langle \sigma, e \rangle \Downarrow \mathbf{n}^m}{\langle \sigma, \mathbf{x} := e \rangle \Downarrow_{pc} \sigma[\mathbf{x} \mapsto \mathbf{n}^k]}$$

where k is defined as follows:

$$k(a) = \begin{cases} m(a) & \text{if } pc(a) = L \\ H & \text{if } pc(a) = H \text{ and } l(a) = H \\ P & \text{otherwise} \end{cases}$$

In the definition above, a represents a principal from the set X . It can be shown that for any semantic derivation in this generalized system, projecting all labels to a given principal yields a valid semantic derivation in the system with a two-point lattice. This immediately implies non-interference for the generalized system, where observations are **limited** to individual principals.

Definition 5. Two labeled values \mathbf{n}_1^k and \mathbf{n}_2^m are a -equivalent, written $\mathbf{n}_1^k \approx^a \mathbf{n}_2^m$, iff either:

1. $k(a) = m(a) = L$ and $\mathbf{n}_1 = \mathbf{n}_2$ or
2. $k(a) = m(a) = H$ or
3. $k(a) = P$ or $m(a) = P$

Definition 6 (Store equivalence). Two stores σ_1 and σ_2 are ℓ -equivalent, written $\sigma_1 \approx^a \sigma_2$, iff for every variable \mathbf{x} , $\sigma_1(\mathbf{x}) \approx^a \sigma_2(\mathbf{x})$.

Theorem 4 (TINI for permissive-upgrade with a product lattice). With the assignment rule ASSN-GPUS , if $\sigma_1 \approx^a \sigma_2$ and $\langle \sigma_1, c \rangle \Downarrow_{pc} \sigma'_1$ and $\langle \sigma_2, c \rangle \Downarrow_{pc} \sigma'_2$, then $\sigma'_1 \approx^a \sigma'_2$.

Proof. Outlined above. □

This generalization of the two-point lattice to a *product* of such lattices is interesting because a *powerset* lattice can be simulated using such a product. However, this still leaves open the question of constructing a generalization of permissive-upgrade to an *arbitrary* lattice (for instance, lattices like the one shown in Figure 4.3). Such a generalization is developed in the next section.

4.2 Generalized Permissive-Upgrade on Arbitrary Lattices

This section shows by construction the generalization of the permissive-upgrade strategy to arbitrary security lattices. For every element ℓ of the lattice, a new label ℓ^* is introduced which means “partially-leaked ℓ ”, with the following intuition:

$$\begin{array}{c}
\text{ASSN-N} \frac{\langle \sigma, e \rangle \Downarrow \mathbf{n}^m \quad l = \Gamma(\sigma(\mathbf{x})) \quad l = \ell_x \vee l = \ell_x^* \quad pc \sqsubseteq \ell_x \quad k = pc \sqcup m}{\langle \sigma, \mathbf{x} := e \rangle \Downarrow_{pc} \sigma[\mathbf{x} \mapsto \mathbf{n}^k]} \\
\text{ASSN-S} \frac{\langle \sigma, e \rangle \Downarrow \mathbf{n}^m \quad l = \Gamma(\sigma(\mathbf{x})) \quad l = \ell_x \vee l = \ell_x^* \quad pc \not\sqsubseteq \ell_x \quad k = ((pc \sqcup m) \sqcap \ell_x)^*}{\langle \sigma, \mathbf{x} := e \rangle \Downarrow_{pc} \sigma[\mathbf{x} \mapsto \mathbf{n}^k]}
\end{array}$$

Figure 4.2: Assignment rules for the generalized permissive-upgrade

A variable labeled ℓ^* may contain partially-leaked data, where ℓ is a *lower-bound* on the \star -free labels the variable may have in alternate executions.

The syntax of labels is listed in Figure 4.1. Labels k, l, m may be lattice elements ℓ or \star -ed lattice elements ℓ^* . In examples, suggestive lattice element names L, M, H (low, medium, high) are used. Labels of the form ℓ are called \star -free or *pure*. Figure 4.1 also defines the join operation \sqcup on labels. This definition is based on the intuition above. When the two operands of \odot are labeled ℓ_1 and ℓ_2^* , $\ell_1 \sqcup \ell_2$ is a lower bound on the pure label of the resulting value in any execution (because ℓ_2 is a lower bound on the pure label of ℓ_2^* in any run). Hence, $\ell_1 \sqcup \ell_2^* = (\ell_1 \sqcup \ell_2)^*$. The reason for the definition $\ell_1^* \sqcup \ell_2^* = (\ell_1 \sqcup \ell_2)^*$ is similar.

The rules for assignment are shown in Figure 4.2. They strictly generalize the rule `ASSN-PUS` for the two-point lattice, treating $P = L^*$. Rule `ASSN-N` applies when the existing label of the variable being assigned to is ℓ_x or ℓ_x^* and $pc \sqsubseteq \ell_x$. The key intuition behind the rule is the following: If $pc \sqsubseteq \ell_x$, then it is safe to overwrite the variable, because ℓ_x is necessarily a lower bound on the (pure) label of x in this and any alternate execution (see the framebox above). Hence, overwriting the variable cannot cause an implicit flow. As expected, the label of the overwritten variable is $pc \sqcup m$, where m is the label of the value assigned to x .

Rule `ASSN-S` applies in the remaining case — when $pc \not\sqsubseteq \ell_x$. In this case, there may be an implicit flow, so the final label on x must have the form ℓ^* for some ℓ . The question is which ℓ . Intuitively, it may seem that one could choose $\ell = \ell_x$, the pure part of the original label of x . The final label on x would be ℓ_x^* and this would satisfy the intuitive meaning of \star written in the framebox above. Indeed, this intuition suffices for the two-point lattice of Section 3.2 and 3.3. However, for a more general lattice, this intuition is unsound, as illustrated with an example below. The correct label is $((pc \sqcup m) \sqcap \ell_x)^*$.

Example The need for the label $k := ((pc \sqcup m) \sqcap \ell_x)^*$ instead of $k := \ell_x^*$ in rule `ASSN-S` is illustrated below. Consider the lattice of Figure 4.3 and the program of Listing 4.1. Assume that, initially, the variables z, w, x_1, x', x_2, y_1 and y_2 have labels $H, L_1, L_1, L', L_2, M_1$ and M_2 , respectively. Fix the attacker at level L_1 . Fix the value of x_1 at true^{L_1} , so that the branch on line 5 is always taken and line 6 is always executed. Set $y_1 \mapsto \text{false}^{M_1}, y_2 \mapsto \text{true}^{M_2}, w \mapsto \text{false}^{L_1}$ initially. The initial value of z is irrelevant. Consider two executions of the program starting from two stores σ_1 with

```

1  if (x')
2    z = y1
3  else
4    z = y2
5  if (x1)
6    z = x1
7  if (not(x2))
8    z = x2
9  if (z)
10  w = z

```

Listing 4.1: Example explaining rule `ASSN-S`

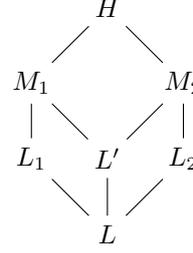


Figure 4.3: Lattice explaining rule `ASSN-S`

$x' \mapsto \text{true}^{L'}$, $x_2 \mapsto \text{true}^{L_2}$ and σ_2 with $x' \mapsto \text{false}^{L'}$, $x_2 \mapsto \text{false}^{L_2}$. Note that as L' and L_2 are incomparable to L_1 in the lattice, σ_1 and σ_2 are equivalent for L_1 .

Requiring $k := \ell_x^*$ in rule `ASSN-S` causes an implicit flow that is observable for L_1 . The intermediate values and labels of the variables for executions starting from σ_1 and σ_2 are shown in the second and third columns of Table 4.1. Starting with σ_1 , line 2 is executed, but line 4 is not, so z ends with false^{M_1} at line 5 (rule `ASSN-N` applies at line 2). At line 6, z contains true^{L_1} (again by rule `ASSN-N`) and line 8 is not executed. Thus, the branch on line 9 is taken and w ends with true^{L_1} at line 10. Starting with σ_2 , line 2 is not executed, but line 4 is, so z becomes true^{M_2} at line 5 (rule `ASSN-N` applies at line 4). At line 6, rule `ASSN-S` applies, but because $k := \ell_x^*$ is assumed in that rule, z now contains the value $\text{true}^{M_2^*}$. As the branch on line 7 is taken, at line 8, z becomes false^{L_2} by rule `ASSN-N` because $L_2 \sqsubseteq M_2$. Thus, the branch on line 9 is not taken and w ends with false^{L_1} in this execution. Hence, w ends with true^{L_1} and false^{L_1} in the two executions, respectively. The attacker at level L_1 can distinguish these two results; hence, the program leaks the value of x' and x_2 to L_1 .

With the correct `ASSN-S` rule in place, this leak is avoided (last column of Table 4.1). In that case, after the assignment on line 6 in the second execution, z has label $((L_1 \sqcup L_1) \sqcap M_2)^* = L^*$. Subsequently, after line 8, z gets the label L^* . As case analysis on a \star -ed value is not allowed, the execution is halted on line 9. This guarantees termination-insensitive non-interference with respect to the attacker at level L_1 .

4.2.1 Termination-insensitive Non-interference (TINI)

To prove non-interference for the generalized permissive-upgrade, equivalence of labeled values relative to an adversary at arbitrary lattice level ℓ needs to be defined. The definition is shown below (Definition 7). Note that clauses (3)–(5) here refine clause (3) of Definition 5 for the two-point lattice. The obvious generalization of clause (3) of Definition 5 — $n_1^k \sim_\ell n_2^m$ whenever either k or m is \star -ed — is too coarse to prove non-interference inductively. For the degenerate case of the two-point lattice, this definition also degenerates to Definition 5 (there, ℓ is fixed at L , $P = L^*$ and only L may be \star -ed).

Definition 7. Two values n_1^k and n_2^m are ℓ -equivalent, written $n_1^k \sim_\ell n_2^m$, iff either

	$w = \text{false}^{L_1}, x_1 = \text{true}^{L_1}, y_1 = \text{false}^{M_1}, y_2 = \text{true}^{M_2}$		
	$x' = \text{true}^{L'}$ $x_2 = \text{true}^{L_2}$	$x' = \text{false}^{L'}$ $x_2 = \text{false}^{L_2}$	
		$k := \ell_x^*$	$k := ((pc \sqcup m) \sqcap \ell_x)^*$
if (x') $z = y_1$ else $z = y_2$	$pc = L'$ $z = \text{false}^{M_1}$	$pc = L'$ $z = \text{true}^{M_2}$	$pc = L'$ $z = \text{true}^{M_2}$
if (x_1) $z = x_1$	$pc = L_1$ $z = \text{true}^{L_1}$	$pc = L_1$ $z = \text{true}^{M_2^*}$	$pc = L_1$ $z = \text{true}^{L^*}$
if (not(x_2)) $z = x_2$	branch not taken	$pc = L_2$ $z = \text{false}^{L_2}$	$pc = L_2$ $z = \text{false}^{L^*}$
if (z) $w = z$	$pc = L_1$ $w = \text{true}^{L_1}$	branch not taken	execution halted
Result	$w = \text{true}^{L_1}$	$w = \text{false}^{L_1}$ (leak)	no leak

Table 4.1: Execution steps in two runs of the program from Listing 4.1, with two variants of the rule ASSN-S

1. $k = m = \ell' \sqsubseteq \ell$ and $\mathbf{n}_1 = \mathbf{n}_2$, or
2. $k = \ell' \not\sqsubseteq \ell$ and $m = \ell'' \not\sqsubseteq \ell$, or
3. $k = \ell_1^*$ and $m = \ell_2^*$, or
4. $k = \ell_1^*$ and $m = \ell_2$ and ($\ell_2 \not\sqsubseteq \ell$ or $\ell_1 \sqsubseteq \ell_2$), or
5. $k = \ell_1$ and $m = \ell_2^*$ and ($\ell_1 \not\sqsubseteq \ell$ or $\ell_2 \sqsubseteq \ell_1$)

Definition 8. Two stores σ_1 and σ_2 are ℓ -equivalent, written $\sigma_1 \sim_\ell \sigma_2$, iff for every variable x , $\sigma_1(x) \sim_\ell \sigma_2(x)$.

This definition is obtained by constructing (through examples) an extensive transition graph of pairs of labels that may be assigned to a single variable at corresponding program points in two executions of the same program. The starting point is label-pairs of the form (ℓ, ℓ) . This characterization of equivalence is both sufficient and necessary. It is sufficient in the sense that it allows us to prove TINI inductively. It is necessary in the sense that example programs can be constructed that end in states exercising every possible clause of this definition. Appendix A.2 lists these examples.

Using the above definition of equivalence of labeled values, TINI can be proven for the generalized permissive-upgrade strategy presented above. A significant difficulty in proving the theorem is that the definition of \sim_ℓ is not transitive unlike the previous definition of \sim . Detailed proofs of all the lemmas and the theorems are presented in Appendix A.3.

Lemma 4 (Expression evaluation). *If $\langle \sigma_1, e \rangle \Downarrow \mathbf{n}_1^{k_1}$ and $\langle \sigma_2, e \rangle \Downarrow \mathbf{n}_2^{k_2}$ and $\sigma_1 \sim_\ell \sigma_2$, then $\mathbf{n}_1^{k_1} \sim_\ell \mathbf{n}_2^{k_2}$.*

Proof. By induction on e . □

Lemma 5 (\star -preservation). *If $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$ and $\Gamma(\sigma(x)) = \ell^\star$ and $pc \not\sqsubseteq \ell$, then $\Gamma(\sigma'(x)) = \ell'^\star$ and $\ell' \sqsubseteq \ell$.*

Proof. By induction on the derivation rule. □

Corollary 1. *If $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$ and $\Gamma(\sigma(x)) = \ell^\star$ and $\Gamma(\sigma'(x)) = \ell'$, then $pc \sqsubseteq \ell$.*

Proof. Immediate from Lemma 5. □

Lemma 6 (pc -lemma). *If $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$ and $\Gamma(\sigma'(x)) = \ell$, then $\sigma(x) = \sigma'(x)$ or $pc \sqsubseteq \ell$.*

Proof. By induction on the derivation rule. □

Corollary 2. *If $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$ and $\Gamma(\sigma(x)) = \ell^\star$ and $\Gamma(\sigma'(x)) = \ell'$, then $pc \sqsubseteq \ell'$.*

Proof. Immediate from Lemma 6. □

Using these lemmas, the standard confinement lemma and non-interference can be proven.

Lemma 7 (Confinement Lemma). *If $pc \not\sqsubseteq \ell$ and $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$, then $\sigma \sim_\ell \sigma'$.*

Proof. By induction on the derivation rule. □

Theorem 5 (TINI for generalized permissive-upgrade for arbitrary lattices). *If $\sigma_1 \sim_\ell \sigma_2$ and $\langle \sigma_1, c \rangle \Downarrow_{pc} \sigma'_1$ and $\langle \sigma_2, c \rangle \Downarrow_{pc} \sigma'_2$, then $\sigma'_1 \sim_\ell \sigma'_2$.*

Proof. By induction on c . □

4.3 Comparison of the Generalization of Section 4.2 with the Generalization of Section 4.1

Two distinct and sound generalizations of the permissive-upgrade strategy for the two-point lattice have now been described: The generalization of the improved permissive-upgrade to pointwise products of two-point lattices or, equivalently, to powerset lattices as described in Section 4.1, and the generalization to arbitrary lattices described in Section 4.2. Since both the generalizations apply to powerset lattices, an obvious question is whether one is more permissive than the other on such lattices. The generalization of permissive-upgrade to pointwise lattices described in Section 4.1 can be more permissive than the generalization described in Section 4.2 for powerset lattices in certain cases as shown by the example below. The reason for this permissiveness is that the generalization of permissive-upgrade to pointwise lattices tracks finer taints, i.e., it tracks partial leaks for each principal separately.

```

1  if (y)
2    x = z
3  if (z)
4    x = z
5  if (x)
6    z = x

```

Listing 4.2: Example where generalization of permissive-upgrade to pointwise lattices is more permissive than the generalization to arbitrary lattices

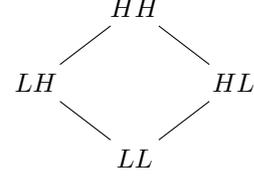


Figure 4.4: A powerset/product lattice

Example The powerset lattice of Figure 4.4 is used for illustration purpose. This lattice is the pointwise lifting of the order $L \sqsubset H$ to the set $S = \{L, H\} \times \{L, H\}$. For brevity, this lattice's elements are written as LL, LH , etc. When generalization from Section 4.1 is applied to this lattice, labels are drawn from the set $\{L, H, P\} \times \{L, H, P\}$. These labels are concisely written as LP, HL , etc. For the generalization from Section 4.2 to arbitrary lattices, labels are drawn from the set $S \cup S^*$. These labels are written LH, LH^* , etc. Note that LH^* parses as $(LH)^*$, not $L(H^*)$ (the latter is not a valid label in the generalization applied to this lattice). Consider the program in Listing 4.2. Assume that x, y and z have initial labels LL, HL and LH , respectively and that the initial store contains $y \mapsto \text{true}^{HL}, z \mapsto \text{true}^{LH}$, so the branches on lines 1 and 3 are both taken. The initial value in x is irrelevant but its label is important. Under the generalization from Section 4.1, x obtains label $((HL) \sqcup (LH)) \sqcap (LL)^* = LL^*$ at line 2 by rule ASSN-S . At line 4, the same rule applies but the label of x remains LL^* . When the program tries to branch on x at line 5, it is stopped. In contrast, under generalization of permissive-upgrade to pointwise lattices, this program executes to completion. At line 2, the label of x changes to PH by rule ASSN-GPUS . At line 4, the label changes to LH because pc and the label of z are both LH . Since this new label has no P , line 5 executes without halting. Hence, for this example, generalization of permissive-upgrade to pointwise lattices is more permissive than the generalization to arbitrary lattices presented in Section 4.2.

An example for which the generalization of permissive-upgrade to arbitrary lattices is more permissive than generalization to pointwise lattices in the case of powerset lattices was not found. But the generalization presented in Section 4.2 is more general than the product construction (Section 4.1) when applied to arbitrary lattices (and hence, applicable to a broader set of lattices) as it is unclear whether or how the improved permissive upgrade strategy generalizes to arbitrary lattices. By developing this generalization, this work makes permissive-upgrade applicable to arbitrary security lattices like other IFC techniques.

Part III

Precise Control Scope Analysis

Chapter 5

Dynamic IFC with Unstructured Control Flow and Exceptions

This chapter presents a mechanism to prevent leaks due to implicit flows in the presence of unstructured control constructs like `break`, `return-in-the-middle` and exceptions.

Implicit flow corresponds to *control dependence* in program analysis, where a predicate governs which program path is executed and leaks information through the control flow of the program. For sound analysis and to avoid over-approximation in certain cases, an important goal for implicit flow tracking, in general, is to determine when the influence of a control construct has ended, i.e., to determine the correct *pc* label for every statement in the program. In other words, it is important to have a precise control scope analysis in place. For block-structured control flow limited to `if` and `while` commands, this is straightforward: The effect of a control construct ends with its lexical scope, e.g., in

```
if (h) {l = 1;} l = 2
```

`h` influences the control flow at `l = 1` but not at `l = 2`. This leads to a straightforward implementation of a *pc* upgrading and downgrading strategy: One maintains a *stack* of *pc* labels [104]; the effective *pc* is the top one. When entering a control flow construct like `if` or `while`, a new *pc* label, equal to the join of labels of all values on which the construct's guard depends with the previous effective *pc*, is pushed. When exiting the construct, the label is popped.

Unfortunately, it is unclear how to extend this simple strategy to non-block-structured control flow constructs such as `break`, `continue` and `return-in-the-middle` for functions, all of which occur in high-level languages. For example, consider the program

```
l = 1; while(1) {... if(h) {break;}; l = 0; break;}
```

The content of this chapter is based partly on the work published as part of the paper, "Information Flow Control in WebKit's JavaScript Bytecode" [22]

with h labeled H . This program leaks the value of h into l , but no assignment to l appears in a block-scope guarded by h . Indeed, the pc upgrading and downgrading strategy just described is ineffective for this program.

Implicit flow in the form of error handling is also a source of information leak as it helps the adversary to learn about the system [47]. For instance, an exception handler might print the stack trace of the error, from which an attacker can determine what sort of attacks the system is vulnerable to.

Tracking information flow in the presence of unstructured control flows is non-trivial as the control breaks out of block structures. Exceptions are much more difficult to handle as they allow for non-local control transfer. Much work on error handling has focused in the context of static analysis [11, 76] and the work on IFC for dynamic languages has mostly ignored exceptions and other unstructured control flow constructs [13, 14, 32, 44, 50]. Just et al. [60] present dynamic IFC for JavaScript bytecode with static analysis to determine implicit flows precisely but ignore implicit flows due to exceptions. Hedin and Sabelfeld propose a dynamic IFC approach for a language modeling the core features of JavaScript [54] but ignore unstructured control flow constructs like *break*, *continue* and *return-in-the-middle* for functions. For handling exceptions, they introduce annotations and an additional class of labels. An extension introduces similar annotations to deal with unstructured control flows [55]. These labels are more restrictive than needed, e.g., the code indicated by dots in the example above is executed irrespective of the condition h in the first iteration, and thus there is no need to raise the pc before checking that condition.

To solve this issue, this chapter presents a precise dynamic analysis approach using *post-dominator* analysis [42, 60].

5.1 Control Flow Graphs and Post-dominator Analysis

The approach presented in this chapter performs on-the-fly post-dominator analysis at runtime to handle implicit flows. A control flow graph (CFG), which is a directed graph, is constructed for every new function before it is executed with every instruction being represented as a node and whose edges represent the possible control flows. For every branch node, its immediate post-dominator (IPD) is computed [22, 42, 60]. A stack of pc labels is maintained. When executing a branch node, a new pc label is pushed on the stack *along with* the node's IPD. When the IPD is actually reached, the pc label along with the IPD is popped from the stack. In [70, 101], the authors prove that the IPD marks the end of the scope of an operation and hence the security context of the operation, so our strategy is sound. The IPD-based solution works for all forms of unstructured control flow like *break*, *continue*, *return-in-the-middle*, and exceptions. Multiple *return* statements in a function are represented by a single *return* node. Theorem 6 shows that the IPD of a node is the most precise node where the context of an operation can be removed. For proving Theorem 6, a few definitions are defined below.

Definition 9. (Control flow graph)

A control flow graph is a directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{E}, n_s, n_e, \mathcal{L})$. \mathcal{N} is the set of nodes. \mathcal{E} is the set of control flow edges $(n_1, n_2) \in \mathcal{E}$, where $n_i \in \mathcal{N}$. (n_1, n_2) represents n_2 may immediately execute after n_1 . The nodes

$n_s, n_e \in \mathcal{N}$ are special nodes representing the start and end point of \mathcal{G} . The function \mathcal{L} maps the edges in \mathcal{E} to labels.

Definition 10. (Path)

A path in a CFG \mathcal{G} is a sequence of nodes (n_1, n_2, \dots, n_m) such that $(n_i, n_{i+1}) \in \mathcal{E}$, written as $n_1 \rightarrow_p n_m$. A node n that lies on the path $n_1 \rightarrow_p n_m$ is written as $n \in n_1 \rightarrow_p n_m$. The notation $n_1 < n_2$ with respect to two nodes n_1 and n_2 in a CFG \mathcal{G} indicates that n_2 lies on a path $n_1 \rightarrow_p n_e$.

Definition 11. (Branch-point)

A branch-point b is a node in a CFG \mathcal{G} that has more than one successor, i.e., $\text{outdegree}(b) > 1$.

Definition 12. (Post-dominator)

In a CFG \mathcal{G} , a node n_d is said to be the post-dominator of a node n if all paths from n to the end-node pass through n_d , i.e., $\forall p. n \rightarrow_p n_e \implies n_d \in p$. The notation $n_d \text{ pd } n$ indicates that n_d is a post-dominator of n .

Definition 13. (Immediate post-dominator)

A node i is the immediate post-dominator of the node n , denoted as $\text{IPD}(n)$, iff:

1. $i \text{ pd } n$ and
2. $\nexists n_o \in \mathcal{N}. ((n_o \text{ pd } n) \wedge (n_o < i))$ or
 $\forall n_o \in \mathcal{N}. ((n_o \neq n) \implies (n_o \text{ pd } n \implies n_o \text{ pd } i))$.

Theorem 6 (Precision). *Choosing any node other than the IPD to lower the pc-label will either give unsound results or be less precise.*

Proof. The proof of the theorem is described in Appendix B. □

5.2 Exceptions and Synthetic Exit Nodes

Maintaining a precise CFG for post-dominator analysis at runtime in the presence of exceptions is expensive. The CFG of a function is constructed statically on-the-fly when compiling the function at runtime. An exception-throwing node in a function that does not catch that exception should have an outgoing control flow edge to the next related exception handler in the runtime call-stack. This means that the CFG is, in general, inter-procedural, and edges going out of a function depend on its calling context, so IPDs of nodes in the function must be computed *every time the function is called* (the IPDs change based on the earlier functions in the call-stack that called the particular function where the exception occurs). Moreover, in the case of recursive functions, the nodes must be replicated for every call. This is rather expensive. Ideally, the function's CFG should be built only once when *the function is compiled* and work intra-procedurally.

In the design presented in this chapter, every function that may throw an unhandled exception and has an exception handler present earlier in the call-stack (which is assumed to be known at runtime through additional data structures in the system) has a special *synthetic exit node* (SEN), which is placed after the regular return node of the function in the CFG. Every exception-throwing node, whose exception will not be caught within the function, has an outgoing edge to the SEN. In essence,

the SEN is treated as the IPD for nodes whose actual IPDs lie outside of the function. By doing this, all cross-function edges are eliminated and the CFGs become intra-procedural. This allows the computation of the CFGs just once as compared to the inter-procedural case. For every exception-throwing instruction that has an associated handler, its context is maintained during dynamic information flow analysis until the handler is reached. Thus, function calls and all potential exception-throwing instructions are represented as nodes with multiple edges (branches) and push a node on the *pc*-stack. However, a new node is *not* pushed on the *pc*-stack if the IPD of the current node is the same as the IPD on the top of the *pc*-stack or if the IPD of the current node is the SEN, as in this case the *real* IPD, which is outside of this method, is already on top of the *pc*-stack. In fact, the actual IPD of a node having SEN as its IPD is the node that is currently on the top of the stack. This result is shown in Theorem 7. The proof is shown in Appendix B.

Theorem 7. *The actual IPD of a node having SEN as its IPD is the node on the top of the *pc*-stack, which lies in a previously called function.*

In summary, these semantics emulate the effect of having cross-function edges. For illustration, consider the following two functions *f* and *g*. The \diamond at the end of *g* denotes its SEN. Note that there is an edge from `throw 9` to \diamond because `throw 9` is not handled within *g*. \square denotes the IPD of the function call *g*() and handler `catch(e)`.

```
function f() = {
  l = 0;
  try { g(); } catch(e) { l = 1; }
   $\square$  return 1;
}

function g() = {
  if (h) {throw 9;}
  return 7;
}  $\diamond$ 
```

In the absence of instrumentation, when *f* is invoked with *pc* = *L*, the two functions together leak the value of *h*, which is assumed to have a label *H*, into the return value of *f*. When calling *g*, the current *pc* and IPD (*L*, \square) are pushed on the *pc*-stack. When executing the condition `if (h)` a new node is not pushed again, but the top element is merely updated to (*H*, \square) as its IPD is the SEN \diamond . If *h* is `false`, control reaches the `return` statement but with *pc* = *H*. At \square , *pc* is lowered to *L*, so *f* ends with the return value 0 and public label *L*. If *h* is `true`, control reaches the handler, which is in *f* and invokes it with the same *pc* as at the point of exception, i.e., *H*. Consequently, permissive-upgrade marks the assignment in the `catch` block as partially-leaked and prevents the implicit information leak in this case.

5.3 Formal Model

This section formally models the semantics of the language with dynamic IFC instrumentation including unstructured control flow and exceptions. The language from Figure 3.1 is extended to include unstructured control flow constructs like `break`, `continue`, `return` and exceptions. Programs are considered as a collection of functions (without parameter-passing). The control flow analysis is performed on a function before it is executed and is abstractly represented as a CFG in the formal model. Thus, the program itself is modeled as a huge control flow graph (*G*). IPDs are computed using the algorithm by Lengauer and Tarjan [65] when the CFG is created. The CFG is statically

$$\begin{aligned}
e &:= \mathbf{n} \mid \mathbf{x} \mid e_1 \odot e_2 \\
\iota &:= \mathbf{end} \mid \mathbf{x} := e \mid \mathbf{branch} \ e \mid \mathbf{jmp} \mid \mathbf{return} \mid \mathbf{throw} \ e \mid \mathbf{catch} \ \mathbf{x} \mid \mathbf{SEN}
\end{aligned}$$

Figure 5.1: Language Syntax

constructed (only once) as new functions are called or discovered at runtime. For a non-branching node $\iota \in \mathcal{G}$, $Succ(\iota)$ denotes ι 's unique successor. For a conditional branching node ι , $Left(\iota)$ and $Right(\iota)$ denote successors when the condition is true and false, respectively.

The syntax of the language modeling the nodes in a CFG is shown in Figure 5.1. The command `if e then c_1 else c_2` is represented as the node `branch e` with $Left(\mathcal{G}, \iota) = c_1$ and $Right(\mathcal{G}, \iota) = c_2$. Similarly, `while e do c` is represented as `branch e` with $Left(\iota) = c$ and $Right(\iota)$ being the command following `while` in the program. A `jmp` node in the CFG corresponds to `break` and `continue` with $Succ(\iota)$ pointing to the next node in the CFG according to the operation. It is also assumed that a function always ends with a `return` statement and thus a CFG normally ends with the `return` node. Multiple `return` statements in a function are represented using a single `return` node. $Succ(\iota)$ points the node to return to in the previous CFG, while the return value is saved in a global variable that can be accessed by the program later on. When a new CFG is added for a function, the `return` node of that CFG points to the successor node of the function call in the previous CFG. However, the IPDs are computed when the CFGs are intra-procedural. Every node in the program's CFG is uniquely identifiable.

In general, every function has an associated exception table that maps each potentially exception-throwing instruction in the function to the appropriate exception handler within the function. This is represented by adding a *Right* edge in the CFG from the instruction's node to the handler's node; `throw` has only one outgoing edge. It is conservatively assumed that any unknown code may throw an exception, so function call is exception-throwing for this purpose. If a function contains unhandled exceptions, the corresponding edges in the CFG point to the `SEN` of the CFG. The `SEN` is only created if one of the previous functions in the call-stack has an appropriate exception handler for the unhandled exceptions in the current function. When an `SEN` node is created, an edge is added from the `SEN` of the CFG to a node in the previous CFG, which is either the `catch` node or the `SEN` of that CFG. $Succ$ denotes one of these edges. If there are no appropriate handlers in the call-stack, the exception-throwing nodes have an edge to the `end` node of the program CFG. For simplicity of exposition, it is assumed here that all exceptions belong to a single class — for different types of exceptions, the exception class would also be matched for determining the appropriate exception handler.

Program configurations for commands (nodes) are represented as $\langle \sigma, \iota, \rho \rangle$, where σ represents the memory store as before, ι represents the currently executing node, and ρ is the *pc*-stack. The configuration for expressions is the same as before: $\langle \sigma, e \rangle$.

Each entry of the *pc*-stack ρ is a pair (ℓ, ι) , where ℓ is a security label, and ι is a node in the CFG. When a new control context is entered, the new *pc*-label, which is a join of the current context label and the existing *pc*-label (the label on the top of the stack), is pushed together with the IPD ι of

$$\begin{array}{c}
\begin{array}{c}
\iota = (x := e) \quad \langle \sigma, e \rangle \Downarrow \mathbf{n}^m \\
\begin{array}{c}
l = \Gamma(\sigma(\mathbf{x})) \quad l = \ell_x \vee l = \ell_x^* \quad pc = \Gamma(!\rho) \quad k = \left\{ \begin{array}{ll} pc \sqcup m, & pc \sqsubseteq \ell_x \\ ((pc \sqcup m) \sqcap \ell_x)^*, & pc \not\sqsubseteq \ell_x \end{array} \right\} \\
\sigma' = \sigma[\mathbf{x} \mapsto \mathbf{n}^k] \quad \iota' = Succ(\iota) \quad \rho' = isIPD(\iota', \rho)
\end{array} \\
\text{ASSN} \frac{}{\langle \sigma, \iota, \rho \rangle \rightarrow \langle \sigma', \iota', \rho' \rangle}
\end{array} \\
\\
\begin{array}{c}
\iota = \text{branch } e \quad \langle \sigma, e \rangle \Downarrow \mathbf{b}^\ell \\
\iota' = \left\{ \begin{array}{ll} \text{Left}(\iota), & \text{if } \mathbf{b} = \text{true} \\ \text{Right}(\iota), & \text{otherwise} \end{array} \right\} \quad \rho'' = \rho.\text{push}(\ell, IPD(\iota)) \quad \rho' = isIPD(\iota', \rho'') \\
\text{BRANCH} \frac{}{\langle \sigma, \iota, \rho \rangle \rightarrow \langle \sigma, \iota', \rho' \rangle}
\end{array} \\
\\
\begin{array}{c}
\iota = \text{jmp or return or SEN} \quad \iota' = Succ(\iota) \quad \rho' = isIPD(\iota', \rho) \\
\text{JMP, RET, SEN} \frac{}{\langle \sigma, \iota, \rho \rangle \rightarrow \langle \sigma, \iota', \rho' \rangle}
\end{array} \\
\\
\begin{array}{c}
\iota = \text{throw } e \\
\langle \sigma, e \rangle \Downarrow \mathbf{n}^k \quad pc = \Gamma(!\rho) \quad excValue = \mathbf{n}^{(k \sqcup pc)} \quad \iota' = Succ(\iota) \quad \rho' = isIPD(\iota', \rho) \\
\text{THROW} \frac{}{\langle \sigma, \iota, \rho \rangle \rightarrow \langle \sigma, \iota', \rho' \rangle}
\end{array} \\
\\
\begin{array}{c}
\iota = \text{catch } x \quad excValue = \mathbf{n}^m \\
l = \Gamma(\sigma(\mathbf{x})) \quad l = \ell_x \vee l = \ell_x^* \quad pc = \Gamma(!\rho) \quad k = \left\{ \begin{array}{ll} pc \sqcup m, & pc \sqsubseteq \ell_x \\ ((pc \sqcup m) \sqcap \ell_x)^*, & pc \not\sqsubseteq \ell_x \end{array} \right\} \\
\sigma' = \sigma[\mathbf{x} \mapsto \mathbf{n}^k] \quad \iota' = Succ(\iota) \quad \rho' = isIPD(\iota', \rho) \\
\text{CATCH} \frac{}{\langle \sigma, \iota, \rho \rangle \rightarrow \langle \sigma', \iota', \rho' \rangle}
\end{array} \\
\\
\begin{array}{c}
\iota = \text{end} \\
\text{END} \frac{}{\langle \sigma, \iota, \rho \rangle \rightarrow _}
\end{array}
\end{array}$$

Figure 5.2: Semantics

the entry point of the control context. (ι) uniquely identifies where the control of the context ends. In the semantics, the meta-function $isIPD$ pops the stack. It takes the current instruction and the current pc -stack, and returns a new pc -stack. $!\rho$ returns the top frame of the pc -stack. $\Gamma(!\rho)$ returns the current context label, also represented as pc in the semantics.

$$isIPD(\iota, \rho) := \begin{cases} \rho.\text{pop}() & \text{if } !\rho = (_, \iota) \\ \rho & \text{otherwise} \end{cases}$$

As explained in Section 5.2, a new node (ℓ, ι) is pushed onto ρ only when ι (the IPD) differs from the corresponding entry on the top of the stack or it is not SEN (Theorem 7). Otherwise, ℓ is joined with the label on the top of the stack. This is formally represented using the function $\rho.\text{push}(\ell, \iota)$.

The derivation rules in Figure 5.2 define small-step semantics that define the judgment: $\langle \sigma, \iota, \rho \rangle \rightarrow \langle \sigma', \iota', \rho' \rangle$ (as big-step semantics cannot model unstructured control flow). The rules for expressions

are the same as in Figure 3.2. The rules are informally explained above. The soundness of the analysis including unstructured control flow and exceptions is proven below.

To state the theorem formally, the equivalence of different data structures with respect to the adversary needs to be defined. Value and memory equivalence is defined as before in Definitions 7 and 8. To prove termination-insensitive non-interference, some additional definitions and auxiliary lemmas are defined below. The detailed proofs can be found in Appendix C.

Definition 14 (*pc-stack equivalence*). For two pc-stacks ρ_1, ρ_2 , $\rho_1 \sim_\ell \rho_2$ iff the corresponding nodes of ρ_1 and ρ_2 having label less than or equal to ℓ are equal.

Definition 15 (*State equivalence*). Two states $s_1 = \langle \sigma_1, \iota_1, \rho_1 \rangle$ and $s_2 = \langle \sigma_2, \iota_2, \rho_2 \rangle$ are equivalent, written as $s_1 \sim_\ell s_2$, iff $\sigma_1 \sim_\ell \sigma_2$, $\iota_1 = \iota_2$, and $\rho_1 \sim_\ell \rho_2$.

Lemma 8 (*Confinement Lemma*). If $\langle \sigma, \iota, \rho \rangle \rightarrow \langle \sigma', \iota', \rho' \rangle$ and $\Gamma(!\rho) \not\sqsubseteq \ell$, then $\sigma \sim_\ell \sigma'$, and $\rho \sim_\ell \rho'$.

Theorem 8. *Suppose:*

1. $\langle \sigma_1, \iota_1, \rho_1 \rangle \sim_\ell \langle \sigma_2, \iota_2, \rho_2 \rangle$
2. $\langle \sigma_1, \iota_1, \rho_1 \rangle \rightarrow^* \langle \sigma'_1, \mathit{end}, [] \rangle$
3. $\langle \sigma_2, \iota_2, \rho_2 \rangle \rightarrow^* \langle \sigma'_2, \mathit{end}, [] \rangle$

Then, $\sigma'_1 \sim_\ell \sigma'_2$.

Part IV

Budget-based Limited Information Release

Chapter 6

Bounding Information Leaks Dynamically

Information flow control allows tracking of sensitive and private information while preventing leaks to unauthorized public channels. However, in some cases it is required that (some part of) the sensitive data be accessible over certain public channels. This violates the security property of non-interference but is needed for practical reasons. This is generally achieved by specifying special declassification policies, which allow for such release of information. For instance, the declassification policy `declassify(pwd == input)` would treat the result of the equality check (`pwd == input`) as public, thus allowing information to be released to the user. The enforcement generally requires `declassify` annotations to be added to the code to specify which operations need to release the value. Most of the approaches require policy specifications by the developer that define *what* information can be released by the program when the data to be declassified is part of the untrusted code; else the untrusted code can arbitrarily add `declassify` annotations and release information at will.

An alternative approach to declassification is to *quantify* the amount of information a program releases about a sensitive value and to determine an acceptable upper bound, which accepts or rejects the program. The information released is generally quantified by computing the difference in the entropy, a measure of uncertainty, of information contained in the sensitive value before and after the release with respect to an adversary [7, 34, 71, 92, 93]. Research in quantitative information flow has focused on statically analyzing the program and determining a worst-case bound on the amount of information that a program can release as a whole [16, 33, 34, 42, 64, 93]. These approaches rely on the assumption that the probability distribution of the output values is known upfront.

Quantitative information flow analysis methods, in general, yield leakage bounds over all possible executions of a program. A dynamic analysis, on the other hand, considers only the path being taken in the current execution, which in general consensus seems to be unsound for quantifying

```
1  age = getCurrentAge(birthday, birthyear);
2  if (age < 18)
3      preference = "child";
4  else
5      preference = "adult";
```

Listing 6.1: Age-based Advertisement

information leaks. However, such an analysis would prove useful in scenarios like the Web where the program to be analyzed might only be available at runtime and one needs to only bound the amount of information released by the program about a sensitive value to untrusted sources [25].

In the context of dynamic analysis, bounding information leaks helps control the information leakage at runtime. The motivation of this work is, thus, to provide an analysis that bounds information leaks dynamically at runtime and to prove that the approach is sound. McCamant and Ernst [73] provide the only dynamic quantitative information flow analysis but do not prove the soundness of their approach. They, however, provide a simulation-based proof [72] for quantifying the amount of information released by the program about a secret for a simple system with a two-level lattice. They track information flow more finely at the granularity of bits by restricting the language's variables to single bits. They maintain a shadow bit for every variable that indicates the secrecy level (public or secret) of the value in the variable. However, it is unclear how their technique scales for multiple security levels or multiple secrets in the system.

Moreover, the approach by McCamant and Ernst [73] is more suitable for quantification rather than for *bounding* information leaks; bounding information leaks introduces additional non-trivial complications and leakage channels as detailed in later sections. Other prior work by Besson *et al.* [20] used self-information to quantify the number of bits of information that could be leaked about a secret value. However, this technique requires the knowledge of the probability distribution or individual probabilities of the outputs that are based on secret values. In cases where it is difficult to determine the probability of the outputs upfront, the use of this measure for analysis does not help much.

This chapter explores an alternative approach, where the developer can specify a *budget* for every sensitive value in the system (defaulting to zero if not specified), which is basically an upper bound on the amount of information allowed to leak about that value. An underlying enforcement mechanism ensures that no more information about that value is leaked than what is specified by the budget.

Additionally, in scenarios like the Web where the untrusted third-party code changes quite often, it is difficult for the developer to keep up with the changes and modify the *what* policy specifications accordingly. For instance, consider the program snippet in Listing 6.1 where the displayed advertisements for some web-page depend on the viewer's age. The exact age is considered sensitive data, but in order to provide its functionality only an abstraction of the age is required, namely whether `age < 18`. The developer might specify `declassify(age < 18)` as a declassification policy. However, different jurisdictions might entail adaptations to this rule (requiring a different check on age), resulting in a cumbersome process where the developer needs to update the policy frequently.

If, however, the developer associates a budget of 3 with the initial value in the variable `age` (meaning that only 3 bits of information about `age` are allowed to leak), this example will be accepted in the envisioned model. If the advertiser changes the criteria to include another check whether the user is a teenager, this comparison can be included without requiring any change to the policy. The soundness of the underlying enforcement would ensure that no more information than the specified budget can ever leak.

Bounding information leaks has additional advantages — it provides a formal meaning for allowing certain operations like the one described in Listing 6.1 while establishing important security properties related to them. There are a number of real-world programs that allow a limited number of such operations to be performed, e.g., a password checker allows 3 tries before it locks the user's account.

Bounding information leaks in a purely dynamic setting is still largely an open problem. As the property of non-interference does not allow any information release, this chapter proposes an end-to-end guarantee that allows the flow of secret information to public channels in a controlled manner by declassifying fragments of the secret. The policy enforces that the only information leaked about the secrets by a program is bounded by their pre-specified budgets. A sound enforcement of the property is described and proven sound information-theoretically by coding the outputs generated through information releases and showing that the output codes are uniquely-decodable. We utilize an important result by Shannon [92], which states that the average length of uniquely-decodable codes upper bounds the Shannon entropy of the outputs.

6.1 Quantifying Information Leaks

The information released, or the *mutual information*, is generally quantified as the difference in the entropy, which is a measure of the adversary's uncertainty, of the sensitive information before and after the information is released, i.e.,

$$\text{mutual information} = \text{initial uncertainty} - \text{final uncertainty}$$

Roughly speaking, this amounts to the gain in knowledge of the adversary about a sensitive value. The adversary computes a probable initial set of values for the sensitive value. By observing the information released, the adversary can refine his/her knowledge set by removing the improbable values. Various information-theoretic measures have been proposed [7, 34, 71, 92, 93] for quantifying information leaks. Many of these measures determine an average worst-case bound on the amount of information that a program can leak. But for deterministic systems in which the prior probability distribution is not known, the maximum leakage is generally realized on uniform priors and most of these measures equal Shannon entropy.

Clark et al. [33] propose an approach to statically quantify the amount of information leaked in an imperative language. An important result that they prove in their work is that the information released by a program in a deterministic system is equal to the Shannon entropy of the outputs given the public inputs. Thus, for computing the information released by a program about a secret, it is enough to compute the Shannon entropy of the public outputs given the public inputs.

```
1 dbPwd = getActualPassword(user);
2 uPwd = readUserPassword();
3 login = (dbPwd == uPwd);
```

Listing 6.2: Password Checker

Another important result is the Shannon’s source coding theorem [92] that intuitively states that if one can associate variable-sized bit-codes with different outputs of a program such that the codes are uniquely-decodable, then the average of the code-lengths of these codes has been shown to be bounded by the Shannon entropy of the outputs of the program. The soundness of our approach builds on top of these two results. We associate bit-codes with the outputs of the program being analyzed and show that these codes are uniquely-decodable. In other words, the information release as computed by our approach averaged over all executions is bounded by the actual information released by the program.

6.2 Limited Information Release

Limited information release (LIR) is an information release policy that declassifies some (limited) parts of sensitive information. The motivation for LIR is that, in general, certain information flows leak only an insignificant part of a secret. As an example, the *comparison* of a secret with a constant value is largely considered acceptable and its rejection by standard IFC analyses is too restrictive in practice [18, 61]. Such information leaks are usually acceptable if one can guarantee that an adversary cannot widen the declassification to launder information.

In the real-world, such policies find a wide range of applications. For instance, in the advertisement example from Listing 6.1, the third-party ad service does not need to know the exact age of the viewer for determining the preferred advertisement. It is sufficient to know whether the viewer is a child or an adult (or a teenager). A similar example is that of a music app that hosts advertisements for music shows and concerts in a town. Based on the age and the preferences of the user, the advertisements might differ in each case.

Another common use-case of such policies is user authentication based on a secret password or a PIN. The password checker in Listing 6.2 compares the secret password to the password entered by the user. The public variable `login` reflects whether these match, indicating whether the log-in was successful. Strict non-interference would normally prohibit assignments to the low variable `login` as the value is derived from the secret password. However, releasing the log-in status to the user cannot be avoided. Normally, a user can log-in in a single try but the probability of an adversary guessing the correct password in one (or even a few) tries can be assumed to be negligible when the password is strong and chosen randomly [99]. Thus, the assignment to the `login` variable may be permitted for a few tries before the user’s account gets locked. Generally in these applications, there is a trade-off between some private information leak and better services.

In practice, certain operations involving Boolean-valued expressions, and implicit flows leak very little information, which might be required for providing proper functionality. King *et al.* [63] inves-

```
1 pub = 0; i = 1;
2 while (i ≤ 231) {
3   if ((sec & i) == i)
4     pub = pub | i;
5   i = i << 1;
6 }
```

Listing 6.3: Laundering attack via implicit flows

tigate the occurrence of implicit flows in some standard algorithms and discuss the pros and cons of handling implicit flows. Russo *et al.* [86] also observe that implicit flows cannot be exploited in non-malicious code to leak secrets efficiently. The LIR policy is, thus, guided by two key tenets:

- **Declassification of Boolean-valued expressions.** As has been observed previously [18, 61] and is exemplified above, comparison operations and other Boolean values often provide a low bandwidth channel for information leak. For instance, $h \neq l$ only reveals whether the two values in h and l are equal or not. Similarly, $h < l$ and $h > l$ only reveal that the value of h is lesser and greater than the value of l , respectively. With LIR, such comparison operations and Boolean-valued expressions are treated as potential points of information release. As all operations need not be declassified, information is released only for those operations that are explicitly annotated with `declassify`.
- **Bounding the information released.** Unfortunately, allowing information release when only an insignificant amount of information is leaked can be widened to leak the complete secret [88, 99]. The example in Listing 6.3 is a classical example of a laundering attack. It implicitly leaks the secret value in `sec` to the variable `pub` without any direct assignments. Every time the check on line 3 is performed, it leaks one bit of `sec`. The whole secret gets implicitly laundered into the variable `pub` as this comparison is performed in a loop of length equal to the size of the secret (in bits).

To limit the amount of information that can be leaked through laundering attacks, LIR introduces a notion of *budget* (a natural number) associated with a sensitive value, which is an upper bound on the amount of information that is allowed to leak about the sensitive value. A budget is associated with every secret in the program and defaults to 0, if left unspecified.

Intuitively, a program is said to satisfy limited information release, if the information released about the initial value of each secret in the system is limited by its pre-specified budget. If no information is released about any secret or if the budgets for all the secrets reduce to 0, LIR reduces to the standard non-interference policy. It is important to note that even when staying within the bounds of the budget one can leak the entire secret via comparison operations. For instance, if an equality comparison $h == l$ involving a secret value h and a public value l returns `true`, a public adversary knows that the value of h is the same as the value in l .

Although declassification of only Boolean-valued expressions is allowed by LIR, this can be extended to other expressions as well. In such cases, where the expression being declassified has a

```

1  pub = 0;
2  h = sec % 2;
3  sec = sec / 2;
4  if (declassify(h == 1))
5      pub = pub | 1;
6  i = sec % 2;
7  if (declassify(i == 1))
8      pub = pub | 2;

```

Listing 6.4: Leak due to dependent variables

data type having n possible values, the leak is accounted for by assuming that $\log_2 n$ bits are leaked. In this thesis, information release is restricted to Boolean expressions. Declassifying an integer value would effectively mean that all bits of information about the secret value have been released and the secret value could have been simply declassified.

6.3 LIR Enforcement

6.3.1 Language and Syntax

This section describes a runtime enforcement of LIR for the simple imperative language shown in Figure 3.1 extended with the `declassify` operator as shown in Figure 6.1. The comparison and arithmetic operators are separated as \oplus and \odot , respectively, for the rules. The language is sufficient for describing the key idea of LIR — appropriate deduction of budgets at declassification of Boolean expressions involving secrets.

$$\begin{aligned}
 e &::= n \mid x \mid e_1 \odot e_2 \mid e_1 \oplus e_2 \mid \text{declassify}(e_1 \oplus e_2) \\
 c &::= \text{skip} \mid x := e \mid c_1; c_2 \mid \text{if } e \text{ then } c_1 \text{ else } c_2 \mid \text{while } e \text{ do } c
 \end{aligned}$$

Figure 6.1: Syntax of the language

6.3.2 Key Aspects of the Enforcement

Normally, in dynamic information flow control mechanisms, every input to the program has a label associated with it, which is an element of a security lattice. Along with the label, every input is associated with a *budget* that represents an upper bound on the amount of information that is allowed to leak about that input.

To understand how the enforcement works with budgets, consider the example in Listing 6.4. Consider a security lattice $L \sqsubseteq H$. Assume that the initial values in `sec` and `h` are both secret with label H and initial budget of 1. The initial value in `pub` and `i` are both public having the value-label

L. Suppose that at the assignment on line 2, the label of `sec` is carried over to the label of `h` (like any flow-sensitive IFC monitor would). The declassification on line 4 would then release the value of `h` (the last bit of `sec`) to `pub` deducting the budget of `h` but not `sec` as it has no information of `h` being dependent on `sec`. As a result, the second bit of `sec` can also be leaked later (through the declassify on line 7), which should not have been allowed as the budget of `sec` is just 1.

6.3.2.1 Labels and dependencies

Instead of directly tracking the label on the values, variable dependencies are tracked. As labels are an abstraction of the dependencies, they carry lesser information than the dependencies. To track the confidentiality level of the variables, every variable has an initial immutable security label associated with it, referred to as its *value-label* (which is its initial label at the start of the program). An immutable map \mathbb{L} maps a variable to its initial value-label. Tracking dependencies enables individual access to the additional meta-data — value-label and budget — for every dependency. Thus, the enforcement can determine the right variables to deduct budget from, which is required at the point of declassification.

With dependency tracking, in the above example, initially `h` and `sec` will only have self-dependencies and value-labels H . The monitor would then update the dependence of `h` on `sec` on line 2. As a result, at the declassify instruction on line 4, it would deduct the budget from `sec` and not `h`. On line 7, as the budget of the variable `sec` has exhausted, the leak of another bit of `sec` is prevented.

The dependencies in the label are split into two parts and the label on a value is represented as a pair: (l, δ) . The variable is itself represented as $x = n^{(l, \delta)}$, where n is the value in the variable x and (l, δ) is the label on the value. The second part of the label, δ , also referred to as dependency-set, is a set of dependencies about whom some information can be released, i.e., those variables that have a non-zero budget. Once the budget of a variable expires (becomes 0) and no more information about that variable can be released, the variable is removed from the dependency-set of the current value and added to the first part of the label to indicate that the current value is still dependent on the variable but cannot release any information about it. For instance, if $x = n^{(\perp, \{x, y\})}$ such that $\mathbb{L}(y) = k$, then current value in x is dependent on its original value and y 's original value and the budgets of both x and y are non-zero. The initial value-label of y is k . When the budget of y expires, y is removed from the dependency-set of the value in x and the value-label of y is joined with \perp and added to the secrecy-level of the value, i.e., $x = n^{(k, \{x\})}$. It suffices to track the value-labels of the variables whose budget has expired as the individual dependencies are not required for any further monitoring. Thus, the first part of the label, l , is the join (upper-bound) of the value-labels of all the dependencies for which no more information is allowed to (or can) be released. The first part of the label, l , is referred to as the secrecy-level of the value and is an element of the security lattice. The secrecy-level l also represents the minimal level at which the value can be observed, i.e., the current value in the variable can never be released or declassified to a level below l . Initially, every variable x depends only on itself and is labeled $(\perp, \{x\})$, where \perp represents the least element in the lattice. The \mapsto notation is used to map a variable to its label in the rest of this section, e.g., $[x \mapsto (\perp, \{x\})]$.

The join operation on the labels returns the join of the secrecy-levels and the union of the dependency-

sets as the final label, i.e.,

$$(l, \delta) \sqcup (l', \delta') = (l \sqcup l', \delta \cup \delta')$$

Similarly, the ordering on the labels is defined as:

$$(l, \delta) \sqsubseteq (l', \delta') = l \sqsubseteq l' \wedge \delta \subseteq \delta'$$

The function $\Gamma(\mathfrak{n}^{(l, \delta)})$ returns the current confidentiality level of the value, i.e., $\Gamma(\mathfrak{n}^{(l, \delta)}) = l \bigsqcup_{x \in \delta} \mathbb{L}(x)$. Note that the value in x is visible to an ℓ -level adversary only if $\Gamma(x) \sqsubseteq \ell$. pc denotes the current program-context level.

With the dependency tracking and the new labeling mechanism in place, let us trace the enforcement of budget reduction in a more involved example in Listing 6.5. Assume a security lattice $L \sqsubset M \sqsubset H$, and two secrets a and b having value-labels $\mathbb{L}(a) = M$ and $\mathbb{L}(b) = H$ and tainted $[a \mapsto (M, \{\}); b \mapsto (L, \{b\})]$. The initial budgets of a and b are 0 and 1, respectively. Also assume that the taint on the value in x is $[x \mapsto (M, \{\})]$. Consider two executions of the program depending on whether the value of a is 0 or not.

```

1  if (a == 0)
2    x = b
3    z = declassify(x == 1)
4    y = declassify(b == 0)

```

Listing 6.5: Example to illustrate the need of budget-labels

When a is 0, x is assigned b on line 2 with a label $[x \mapsto (M, \{b\})]$ (the join of the label of a and b). On line 3, x can be declassified to level M (as the budget of b is 1). Thus on line 3, $[z \mapsto (M, \{\})]$ and the budget of b becomes 0. On line 4, as b 's budget has expired, $[b \mapsto (H, \{\}); y \mapsto (H, \{\})]$. In the other run when a is not 0, $[x \mapsto (M, \{\})]$ on line 3. As the dependency-set is empty, no declassification occurs on line 3. As the budget of b on line 4 is still 1, the value on line 4 is declassified to L . Thus, after the assignment $[y \mapsto (L, \{\})]$. Thus, in the first run an L -level adversary cannot observe any value while in the second run, the adversary can see the value as it is labeled L . As the source code of the program and the initial budgets are publicly visible, the adversary knows the budget of b and can conclude that in the first case the branch on line 1 was taken and in the second run it was not taken. This leaks an unaccounted additional bit of information about a although the budget of a doesn't allow any information release.

6.3.2.2 Budget-labels

The reason for the additional leak is that a purely dynamic information flow monitor cannot account for dependencies in alternate branches. To prevent this leak, an immutable *budget-label* on the budgets is introduced, which is also an element of the security lattice. Additionally, the budget-label indicates the confidentiality level to which information about the secret value can be released. An immutable budget-label map \mathbb{B} tracks the budget-label associated with each variable. The budgets are represented as \mathfrak{n}^l , where \mathfrak{n} is the budget and l is the budget-label. For meaningful declassifications, it is required that the budget-label is not higher than the value-label for a variable, i.e., $\forall x. \mathbb{L}(x) \not\sqsubseteq \mathbb{B}(x)$.

To prevent the leak described above, the budgets are only reduced for those variables in the dependency-set δ that have a budget-label that is at least as high as l when declassifying a value labeled (l, δ) , i.e., $\forall x \in \delta. l \sqsubseteq \mathbb{B}(x)$. By imposing this condition, in the first run of the above example, the value of x labeled $[x \mapsto (M, \{b\})]$ on line 3 is not declassified as the budget-label of b ($\mathbb{B}(b) = L$) is lower than the secrecy-level of the value (M). The budget of b on line 4 remains 1 and can release the value of $b == 0$ to y , which is similar to how the second run proceeds, thereby avoiding the additional leak.

6.3.2.3 Handling Implicit Flows with Budgets

Dynamic IFC approaches handle implicit flows via control constructs by employing a program-context (pc) label, which is a join of the taint on all the values on which the current statement is control dependent. Since variable dependencies are used in labels in our enforcement, those dependencies also need to be checked for during assignment operations under a pc . However, using dependencies in the pc could lead to additional information leaks. This happens because the assignments that happen under a particular pc would carry the dependencies of the operands too, while in an alternate run those dependencies might not be present in the label of the assigned variable. Such mismatch of dependencies can leak information as illustrated below.

```

1  if (x ≤ 10)
2    x = y
3  z = declassify(x ≤ 15)

```

Listing 6.6: Illustration of handling of implicit flows

For the program in Listing 6.6, assume that x , y and z have value-labels of H , H and L ($L \sqsubseteq H$), and budgets of 1^L , 0 and 0 assigned to them, respectively. Also assume that their current labels are $[x \mapsto (L, \{x\}); y = (H, \{\}); z = (L, \{\})]$. Suppose that the check for dependencies is included in the pc , i.e., the assignment succeeds (as per the NSU check presented earlier) only if $pc = (l, \delta) \sqsubseteq (k_o, \delta_o)$ where (k_o, δ_o) is the label of the variable x .

Consider two runs of the program with $x \leq 10$ and $x > 10$. When $x \leq 10$, the branch on line 1 is taken and $pc = (L, \{x\})$ on line 2. This allows the assignment as the label of x is not less sensitive than the pc . As a result, $[x \mapsto (H, \{x\})]$, which is the join of pc and label of y . On line 3, no information is released as the budget-label of x is less than the secrecy-level of x , i.e., $\mathbb{B}(x) \sqsubseteq H$ (as explained above). In the other run when $x > 10$, the branch is not taken and the value of $x \leq 15$ on line 3 is released to the level L as $[x \mapsto (L, \{x\})]$ on line 3 and its budget is 1^L . This leads to different sensitivities of z in the two runs which leads to leaking an additional bit about x in the first run without being accounted for in its budget.

The above leak can be prevented by restricting the pc to contain *only* the secrecy-level, represented as $(l, \{\})$, i.e., none of the dependencies in the pc are allowed to release any information through the variables being assigned within the branch. This ensures that an assignment is visible only to those levels that can *actually* see the assignment. Before entering a branch or a loop, the pc 's secrecy-level l is joined with the value-labels of all the variables in its dependency-set. The function $\Gamma(pc)$ returns the secrecy-level l of the pc . When the meaning is clear from the context, pc is used instead of $\Gamma(pc)$.

With this condition in place, in the above example the combined label of all the dependencies of the value in the predicate is computed before using it in the pc . In the first run, on line 2 $pc = (H, \{\})$ (because $\mathbb{L}(x) = H$) and the assignment fails due to the NSU check as $pc \not\sqsubseteq \Gamma(x)$, i.e., $(H, \{\}) \not\sqsubseteq (L, \{x\})$. In other words, to ensure soundness it is required that $pc \sqsubseteq k$, where k is the secrecy-level of the variable being assigned, i.e., the secrecy-level of the variable being assigned is at least as high as the secrecy-level of the pc .

6.3.2.4 Budget Reduction Constraints

Now suppose that the example in Listing 6.5 is modified to instead declassify inside the branch as shown in Listing 6.7. As before, assume that a and b are two secrets having value-labels $\mathbb{L}(a) = M$ and $\mathbb{L}(b) = H$ and current labels $[a \mapsto (M, \{\}); b = (L, \{b\})]$, such that $L \subseteq M \subseteq H$. The initial budgets of a and b are 0 and 1^L , respectively. Also assume that $[x \mapsto (M, \{\})]$. Again, consider two executions of the program depending on whether the value of a is 0 or not.

```

1  if (a == 0)
2    x = declassify(b == 0)
3    z = declassify(b == 1)

```

Listing 6.7: Example to illustrate budget reduction

When a is 0, the value of $b == 0$ is declassified on line 2 making b 's budget 0 and $x \mapsto (M, \{\})$ (the join of the label of a and the budget-label of b). On line 3, as b 's budget has expired, $[b \mapsto (H, \{\}); z \mapsto (H, \{\})]$. In the other run when a is not 0, $[x \mapsto (M, \{\})]$ on line 3 and the budget of b also remains unchanged. As the budget of b on line 3 is still 1^L , the value is declassified to the level L . Thus, in the second run $[z \mapsto (L, \{\})]$. While in the first run an L -level adversary cannot observe any value, in the second run the adversary can see the value as it is labeled L . As the adversary knows the budget of b , it can conclude that in the first case the branch on line 1 was taken and in the second case it was not taken. This leaks an additional bit of information about a although the initial budget of a is 0.

As the budget of a secret variable is publicly visible, budget reduction in a secret context leads to such additional leaks. Two additional conditions need to be imposed to prevent such leaks:

1. A value cannot be declassified to a level lower than the current program context pc , as this could leak information about the context itself
2. The budget of a variable can only be reduced in a pc lower than or equal to the budget-label of the variable.

In the current setting, this would mean that if the label of the value being declassified is (l, δ) where δ contains a variable x , then information about x is only declassified (the budget of x is only reduced) if $pc \sqsubseteq l \sqcup \mathbb{B}(x)$. With this condition in place, in the first run of the above program, the value of $b == 0$ on line 2 is not declassified as the current program context M is not lesser than or equal to L , which is the join of the secrecy-level and the budget-label of b . Instead, the declassification occurs on line 3 thereby preventing any additional information leaks.

6.3.2.5 Permissiveness

The overall confidentiality of the underlying value with a label (l, δ) is the join of l with the value-label of all the variables in the dependency-set. Thus, all those dependencies whose value-label is below the secrecy-level of the label can be removed as keeping those dependencies does not affect the confidentiality of the value. This improves the permissiveness of the technique by not marking the release of information for variables that have a label lesser than or equal to the current secrecy-level or the pc because no information can be released below the current secrecy-level of the variable or the current pc . For instance, consider a value $n^{(k, \{x\})}$. If $\mathbb{L}(x) \sqsubseteq k$, then x is removed from the dependency-set and the value is represented as $n^{(k, \{\})}$. As explained earlier, budgets are deducted only when $pc \sqsubseteq k$, thus, the budgets of any variable whose value-label is below the pc are not deducted from.

```

1  if (med ≤ 0)
2    x = declassify(sec == y);
3  if (declassify(y ≠ 100))
4    pub = 1;

```

Listing 6.8: Example to illustrate permissiveness

To illustrate the gain in permissiveness by enforcing the condition specified above, consider the program in Listing 6.8. Assume $[med \mapsto (M, \{\}); x \mapsto (M, \{\}); sec \mapsto (M, \{sec\}); y \mapsto (L, \{y\}); pub \mapsto (L, \{pub\})$, with the value-labels being $\mathbb{L}(sec) = H, \mathbb{L}(y) = M, \mathbb{L}(pub) = L$ such that $L \sqsubset M \sqsubset H$. Also, assume that the budgets of sec and y are 1^M and 1^L , respectively. The program context label (pc) on line 2 is $(M, \{\})$. The expression $sec == y$ on line 2 has the label $(M, \{sec, y\})$.

Without the check for the current secrecy-level, the expression evaluation on line 2 releases the value of sec and y into x . However as the secrecy-level is M , the final value of x would have the secrecy-level of at least M . Thus, information about y is released to the levels M and above, who already had access to the value of y . Subsequently, on line 3 as the budget of y has expired, no more information about y is released and the assignment on line 4 fails. Thus, no useful information about y is released at any point in the program, yet the program is terminated because of the NSU check. With the check for current secrecy-level and the context label, the budget of y is not deducted on line 2, which allows the check on line 3 to release information about y thereby allowing the assignment on line 4 to succeed.

6.3.3 Semantics

To formally define the information released by a program starting from some initial memory containing secret values, big-step semantics is defined for the language shown in Figure 6.1 that releases some information about initial secret values at declassification of Boolean expressions.

Program configurations are extended with ι representing the budget store, which is a map from the variables to the budget of these initial values (inputs to the program). The configurations for expressions (e) and commands (c) are, thus, represented by $\langle \sigma, \iota, e \rangle$ and $\langle \sigma, \iota, c \rangle$, respectively where σ represents the memory store as before. $\langle \sigma, \iota, e \rangle \Downarrow_{pc} n^{(k, \delta)}, \iota', \tau$ defines expression evaluation. It means

that under some pc , starting with memory σ and budget store ι , the expression e evaluates to a value n labeled (k, δ) resulting in a budget store ι' . Additionally, the evaluation generates a trace τ , which is a list of values of the form $n_1^{k_1} :: n_2^{k_2} :: \dots :: n_j^{k_j}$. Every declassified value is recorded on the trace along with the level to which it is declassified. The two immutable maps — initial label map (\mathbb{L}) and budget-label map (\mathbb{B}) are assumed to be available and omitted from the rules for clarity.

The evaluation rules for expressions are shown in Figure 6.2. The rule `CONST` evaluates constant values to themselves with the label $(\perp, \{\})$ as they do not have any dependencies. Rule `VAR` evaluates a variable and returns its value n along with the label (k_o, δ_o) . The function \mathbb{R} returns those variables in δ_o whose budget has expired or have a budget-label lower than the secrecy-level k_o . The current secrecy-level k_o is then joined with the value-labels of the variables returned by \mathbb{R} while removing them from the dependency-set δ_o . As no new information is released in both the cases, there is no change to the budget store ι and no trace is generated (indicated as an empty trace ϵ).

In the rules for evaluating arithmetic (`AOP`) and normal comparison (`COP`) operations, the join of the label of the values that the two sub-expressions evaluate to is set as the label of the expression. The individual sub-expressions may release some information as captured in τ_1 and τ_2 , thus updating ι to ι' .

The `DCOPR` rule corresponds to the comparison operation with `declassify`, and allows information release in certain cases. The separation of `declassify` from the normal comparison operations is to offer more control over what information is required to be released. The rule applies only when the current context pc is lower than or equal to the label of the value obtained by evaluation ($pc \sqsubseteq k_o$). This avoids the leaks specified above in Section 6.3.2.2 and Section 6.3.2.4. The function $\Delta(\delta, l, \mathbb{L})$ returns the variables in the dependency-set δ that have an value-label greater than l . Additionally, variables whose budget has expired during the evaluation of the sub-expressions are removed from the dependency-set (using the function \mathbb{R}). To prevent the leak described earlier in Section 6.3.2.2, if the budget-label of any of the remaining variables in the dependency-set is not at least as high as the original secrecy-level of the computed value ($k_o \not\sqsupseteq \mathbb{B}(x)$), then the variable is removed from the dependency-set. For all the variables removed from the dependency-set, their secrecy-level is joined with that of the actual label. The budget of all remaining variables in the dependency-set δ is deducted by 1, corresponding to the 1-bit Boolean value released on the trace. The function $\mathbb{I}(\iota, \delta)$ reduces the budget and returns a new budget store. Their budget-label is also joined with the actual secrecy-level, which gives the final label of the declassified value. Depending on whether the provenance set δ is empty or not, either:

- the budget of all remaining variables in the provenance set δ is deducted by 1, corresponding to the 1-bit Boolean value released on the trace. The function $\mathbb{I}(\iota, \delta)$ reduces the budget and returns a new budget store. Their budget label is also joined with the actual security level, which gives the final label of the declassified value *or*
- if the remaining provenance set is empty, no declassification occurs

If declassification occurs, the declassified value is appended to the trace.

If the current context (pc) is not lower than or equal to the secrecy-level of the value obtained by evaluation ($pc \not\sqsubseteq k$), no information is released as shown in the rule `DCOPN`. The final label is a

$$\begin{array}{c}
\text{CONST} \frac{}{\langle \sigma, \iota, \mathbf{n} \rangle \Downarrow_{pc} \mathbf{n}^{(\perp, \{\})}, \iota, \epsilon} \\
\\
\text{VAR} \frac{\sigma(\mathbf{x}) = \mathbf{n}^{(k_o, \delta_o)} \quad \delta' = \mathbb{R}(\delta_o, \iota, k_o, \mathbb{B}) \quad \delta = \delta_o \setminus \delta' \quad k = k_o \bigsqcup_{\mathbf{x} \in \delta'} \mathbb{L}(\mathbf{x})}{\langle \sigma, \iota, \mathbf{x} \rangle \Downarrow_{pc} \mathbf{n}^{(k, \delta)}, \iota, \epsilon} \\
\\
\text{AOP} \frac{\langle \sigma, \iota, e_1 \rangle \Downarrow_{pc} \mathbf{n}_1^{(k_1, \delta_1)}, \iota_1, \tau_1 \quad \langle \sigma, \iota_1, e_2 \rangle \Downarrow_{pc} \mathbf{n}_2^{(k_2, \delta_2)}, \iota', \tau_2 \quad \mathbf{n} = \mathbf{n}_1 \odot \mathbf{n}_2 \quad (k, \delta) = (k_1, \delta_1) \sqcup (k_2, \delta_2)}{\langle \sigma, \iota, (e_1 \odot e_2) \rangle \Downarrow_{pc} \mathbf{n}^{(k, \delta)}, \iota', (\tau_1 :: \tau_2)} \\
\\
\text{COP} \frac{\langle \sigma, \iota, e_1 \rangle \Downarrow_{pc} \mathbf{n}_1^{(k_1, \delta_1)}, \iota_1, \tau_1 \quad \langle \sigma, \iota_1, e_2 \rangle \Downarrow_{pc} \mathbf{n}_2^{(k_2, \delta_2)}, \iota', \tau_2 \quad \mathbf{n} = \mathbf{n}_1 \oplus \mathbf{n}_2 \quad (k, \delta) = (k_1, \delta_1) \sqcup (k_2, \delta_2)}{\langle \sigma, \iota, (e_1 \oplus e_2) \rangle \Downarrow_{pc} \mathbf{n}^{(k, \delta)}, \iota', (\tau_1 :: \tau_2)} \\
\\
\text{DCOPN} \frac{\langle \sigma, \iota, (e_1 \oplus e_2) \rangle \Downarrow_{pc} \mathbf{n}^{(k, \delta)}, \iota', \tau \quad pc \not\sqsubseteq k}{\langle \sigma, \iota, \text{declassify}(e_1 \oplus e_2) \rangle \Downarrow_{pc} \mathbf{n}^{(k, \delta)}, \iota', \tau} \\
\\
\text{DCOPR} \frac{\langle \sigma, \iota, (e_1 \oplus e_2) \rangle \Downarrow_{pc} \mathbf{n}^{(k_o, \delta_o)}, \iota_1, \tau_1 \quad pc \sqsubseteq k_o \quad \delta' = \Delta(\delta_o, k_o, \mathbb{L}) \quad \delta'' = \mathbb{R}(\delta', \iota_1, k_o, \mathbb{B}) \quad \delta = \delta' \setminus \delta'' \quad k = k_o \bigsqcup_{\mathbf{x} \in \delta''} \mathbb{L}(\mathbf{x}) \bigsqcup_{\mathbf{y} \in \delta} \mathbb{B}(\mathbf{y}) \quad \iota' = \mathbb{I}(\iota_2, \delta) \quad \tau = \begin{cases} \tau_1 :: \tau_2 :: \mathbf{n}^k, & \text{if } (\delta \neq \emptyset) \\ \tau_1 :: \tau_2, & \text{otherwise} \end{cases}}{\langle \sigma, \iota, \text{declassify}(e_1 \oplus e_2) \rangle \Downarrow_{pc} \mathbf{n}^{(k, \{\})}, \iota', \tau}
\end{array}$$

Auxiliary functions:

$$\begin{aligned}
\Delta(\delta, k, \mathbb{L}) &= \{ \mathbf{x} \mid (\mathbf{x} \in \delta) \wedge (\mathbb{L}(\mathbf{x}) \not\sqsubseteq k) \} \\
\mathbb{R}(\delta, \iota, k, \mathbb{B}) &= \{ \mathbf{x} \mid (\mathbf{x} \in \delta) \wedge (\iota(\mathbf{x}) = 0 \vee k \not\sqsubseteq \mathbb{B}(\mathbf{x})) \} \\
\mathbb{I}(\iota, \delta) &= \lambda \mathbf{x}. \begin{cases} \iota(\mathbf{x}) - 1, & \text{if } (\mathbf{x} \in \delta) \\ \iota(\mathbf{x}), & \text{otherwise} \end{cases}
\end{aligned}$$

Figure 6.2: LIR - Semantics of expressions

$$\begin{array}{c}
\langle \sigma, \iota, e \rangle \Downarrow_{pc} \mathbf{n}^{(k_o, \delta)}, \iota', \tau_1 \quad i = \left\{ \begin{array}{l} 1, \text{ if } (\mathbf{n} = \mathbf{true}) \\ 2, \text{ otherwise} \end{array} \right\} \\
k = k_o \bigsqcup_{x \in \delta} \mathbb{L}(x) \quad \langle \sigma, \iota', c_i \rangle \Downarrow_{pc \sqcup k} \langle \sigma', \iota'', \tau_2 \rangle \\
\text{SKIP} \frac{}{\langle \sigma, \iota, \mathbf{skip} \rangle \Downarrow_{pc} \langle \sigma, \iota, \epsilon \rangle} \quad \text{IF-ELSE} \frac{}{\langle \sigma, \iota, (\mathbf{if } e \mathbf{ then } c_1 \mathbf{ else } c_2) \rangle \Downarrow_{pc} \langle \sigma', \iota'', \tau_1 :: \tau_2 \rangle} \\
\langle \sigma, \iota, c_1 \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau_1 \rangle \quad \langle \sigma, \iota, x \rangle \Downarrow_{pc} _^{(l, -)}, \iota, \epsilon \quad pc \sqsubseteq l \\
\langle \sigma', \iota', c_2 \rangle \Downarrow_{pc} \langle \sigma'', \iota'', \tau_2 \rangle \quad \langle \sigma, \iota, e \rangle \Downarrow_{pc} \mathbf{n}^{(m, \delta')}, \iota', \tau \\
\text{SEQ} \frac{}{\langle \sigma, \iota, c_1; c_2 \rangle \Downarrow_{pc} \langle \sigma'', \iota'', (\tau_1 :: \tau_2) \rangle} \quad \text{ASSN} \frac{k = pc \sqcup m \quad \delta = \Delta(\delta', k, \mathbb{L})}{\langle \sigma, \iota, x := e \rangle \Downarrow_{pc} \langle \sigma[x \mapsto \mathbf{n}^{(k, \delta)}], \iota', \tau \rangle} \\
\langle \sigma, \iota, e \rangle \Downarrow_{pc} \mathbf{false}^{(k, \delta)}, \iota', \tau \quad \langle \sigma, \iota, e \rangle \Downarrow_{pc} \mathbf{true}^{(k_o, \delta)}, \iota_1, \tau_1 \quad k = k_o \bigsqcup_{x \in \delta} \mathbb{L}(x) \\
\text{WHILE-F} \frac{}{\langle \sigma, \iota, \mathbf{while } e \mathbf{ do } c \rangle \Downarrow_{pc} \langle \sigma, \iota', \tau \rangle} \quad \text{WHILE-T} \frac{}{\langle \sigma, \iota, \mathbf{while } e \mathbf{ do } c \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau_1 :: \tau_2 \rangle}
\end{array}$$

Figure 6.3: LIR - Semantics of commands

join of the secrecy-levels of all the dependencies. The individual sub-expressions may release some information τ_1 and τ_2 , thus updating ι to ι_2 .

The judgment for a command execution is given by $\langle \sigma, \iota, c \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau \rangle$ — under a program context pc , the execution of a command c starting with memory σ and budget store ι results in a final memory σ' and budget store ι' while generating a trace τ of released values. The semantics is shown in Figure 6.3 whose rules are standard except for some changes in the assignment and branching rules. The assignment rule `ASSN` does the standard NSU check that disallows assignment to a public variable in a secret context — $pc \sqsubseteq l$, where l is the secrecy level of the variable x to which the assignment is being made. The branching rules `IF-ELSE` and `WHILE-T` compute the context label by joining the value-labels (\mathbb{L}) of all the dependencies in δ .

6.4 Formalization of LIR

The following section formalizes the property of LIR and prove it for the semantics presented above. LIR is essentially a property stating that the total number of bits that can be leaked about a secret (by a program) is upper bounded by its pre-specified budget (in an average sense). The trace generated in the LIR semantics captures the data about the declassified value. An adversary at level ℓ can view those values in a memory store σ that have a label less than or equal to ℓ ($\forall x. \Gamma(\sigma(x)) \sqsubseteq \ell$). The adversary can also observe the projection of the trace generated as part of the semantics, defined formally in Definition 16. Similarly, the adversary can view a projection of the budget-map as

defined in Definition 17.

Definition 16 (Trace projection). *Given a trace, τ , the trace projection w.r.t. an adversary at level ℓ , written $\tau \uparrow_\ell$, is defined as:*

$$\begin{aligned} \emptyset \uparrow_\ell &= \emptyset \\ (\mathbf{n}^m :: \tau) \uparrow_\ell &= \begin{cases} \mathbf{n}^m :: \tau \uparrow_\ell & \text{if } m \sqsubseteq \ell, \\ \tau \uparrow_\ell & \text{else.} \end{cases} \end{aligned}$$

Definition 17 (Budget-map projection). *The projection of a budget-map, ι , w.r.t. an adversary at level ℓ , written $\iota \uparrow_\ell$, is defined as:*

$$\iota \uparrow_\ell = \lambda \mathbf{x}. \begin{cases} \iota(\mathbf{x}), & \text{if } (\mathbb{B}(\mathbf{x}) \sqsubseteq \ell) \\ 0, & \text{otherwise} \end{cases}$$

The observational equivalence of various data structures used in the semantics with respect to an adversary needs to be defined for formally defining LIR. Definition 18 and 19 define the observational equivalence of two values and two memory stores, respectively, with respect to an adversary at level ℓ . Definition 20 and 21 define the observational equivalence of the budget-maps and the generated traces.

Definition 18. *Two labeled values $v_1 = \mathbf{n}_1^{(l_1, \delta_1)}$ and $v_2 = \mathbf{n}_2^{(l_2, \delta_2)}$ are observationally equivalent at level ℓ , written $v_1 \sim_\ell v_2$ iff either:*

1. $\mathbf{n}_1 = \mathbf{n}_2$, $l_1 = l_2 \sqsubseteq \ell$ and $\delta_1 = \delta_2$ (or)
2. $\Gamma(v_1) \not\sqsubseteq \ell$ and $\Gamma(v_2) \not\sqsubseteq \ell$ and $l_1 = l_2 \sqsubseteq \ell$ and $\delta_1 = \delta_2$ (or)
3. $\Gamma(v_1) \not\sqsubseteq \ell$ and $\Gamma(v_2) \not\sqsubseteq \ell$ and $l_1 \not\sqsubseteq \ell \vee \exists \mathbf{x} \in \delta_1. \mathbb{B}(\mathbf{x}) \not\sqsubseteq \ell$ and $l_2 \not\sqsubseteq \ell \vee \exists \mathbf{x} \in \delta_2. \mathbb{B}(\mathbf{x}) \not\sqsubseteq \ell$

Definition 19. *Two memory stores σ_1 and σ_2 are observationally equivalent at level ℓ , written $\sigma_1 \sim_\ell \sigma_2$ iff $\forall \mathbf{x}. \sigma_1(\mathbf{x}) \sim_\ell \sigma_2(\mathbf{x})$.*

Definition 20. *Two budget-maps ι and ι' are equivalent at level ℓ , written $\iota \sim_\ell \iota'$, iff $\forall \mathbf{x}. \mathbb{B}(\mathbf{x}) \sqsubseteq \ell \implies \iota(\mathbf{x}) = \iota'(\mathbf{x})$.*

Definition 21. *Two traces τ and τ' are equivalent at level ℓ , written $\tau \sim_\ell \tau'$, iff $\tau \uparrow_\ell = \tau' \uparrow_\ell$.*

As per LIR, the length of the projected trace is bounded by the total budget deducted. The length of a trace τ is represented as $|\tau|$. This intuition of LIR is formalized in Definition 22.

Definition 22 (Limited information release). *A program c is said to satisfy limited information release w.r.t. an adversary at level ℓ if for any given memory store σ , ι , and pc , $\langle \sigma, \iota, c \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau \rangle$ then $|\tau \uparrow_\ell| \leq \iota \uparrow_\ell - \iota' \uparrow_\ell$*

Theorem 10 shows that the monitored semantics presented in Section 6.3.3 satisfy limited information release for every secret in the memory. Lemma 9 proves that in a secret context with respect to an adversary at level ℓ , no declassification visible to the adversary occurs. Theorem 9 proves that

in a store containing only one secret value with respect to ℓ -level adversary (the other secrets can be assumed to be visible to the adversary), the length of the trace is bounded by the budget reduced, i.e., the number of bits declassified is bounded by the deduction in budget of that secret value. Theorem 10 generalizes this result to all secrets in the system. The proofs are detailed in Appendix D.

Lemma 9 (Confinement). *If $\langle \sigma, \iota, c \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau \rangle$, and $\Gamma(pc) \not\sqsubseteq \ell$, then $\sigma \sim_\ell \sigma'$, $\iota \sim_\ell \iota'$ and $\tau \uparrow_\ell = \epsilon$*

Theorem 9 (Limited information release for a single secret). *If $\langle \sigma, \iota, c \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau \rangle$, and $\exists x \in \sigma. (\Gamma(\sigma(x)) \not\sqsubseteq \ell \wedge (\forall y \in \sigma. y \neq x \wedge \Gamma(\sigma(y)) \sqsubseteq \ell))$, then $|\tau \uparrow_\ell| \leq \iota \uparrow_\ell(x) - \iota' \uparrow_\ell(x)$.*

Theorem 10 (Limited information release). *For any memory store σ , budget store ι and program c if $\langle \sigma, \iota, c \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau \rangle$, then c satisfies limited information release for σ and ι*

6.5 Soundness and Decoding Semantics

The LIR policy defined earlier enforces that the information released by a program about every secret is bounded by its pre-determined budget. However, to prove that the actual information released by the approach is soundly accounted for, it needs to be shown that the the budget reduction performed as part of LIR is well-founded in an information theoretic sense. To this end, the result of Shannon's source coding theorem [92] is leveraged to utilize the trace generated by the semantics containing the declassified values.

Recall that Shannon's source coding theorem states that if different outputs of a program can be associated with uniquely-decodable codes, i.e., given a bit string there are no two different interpretations of the sequence of code-words, the information leaked by the program about the secret inputs through these outputs as computed using Shannon entropy is upper-bounded by the average length of these codes. As the prior probability distribution is assumed to be unknown in this setting, channel capacity is used instead of Shannon entropy, which is the maximum information leakage across all possible distributions of a random variable. For deterministic programs, the channel capacity is realized on uniform priors. The trace projected to an adversary contains all the bits the adversary can observe about the secret inputs to the program. The projected trace can, thus, be regarded as a coding of the information released about the secrets to the adversary. Hence, if it can be proven that the trace projected to an adversary is uniquely-decodable then its average length over all executions is an upper bound on the information leaked by the program as computed by the definition of Shannon entropy for uniform distributions.

In order to prove the soundness of the approach, this section presents a decoding semantics that simulates the adversary's approach to decode the information it obtains via the trace. The purpose of the decoding semantics is to show that the adversary-projected trace corresponds to a uniquely-decodable code, i.e., decoding the trace would result in a unique final memory store for the adversary, thereby showing that our enforcement is sound. The salient features of the semantics are:

1. Decoding semantics is specialized to a fixed adversary at level ℓ .
2. Decoding semantics operates on adversary-projected data structures, i.e., an ℓ -projected memory, where the secret inputs are replaced by \star that represents an unknown value (as shown in

Definition 23) and ℓ -projected budget store, where the budgets having a budget-label higher than the adversary's level are replaced by 0 (as in Definition 17).

3. At declassification points, the decoding semantics reads a value from the ℓ -projected trace.
4. Decoding semantics completely skips the code that is executed under a pc influenced by secrets (having the value \star).

Definition 23 (Memory store projection). *The projection of a memory store, σ , w.r.t. an adversary at level ℓ , written $\sigma \uparrow_\ell$, is defined as:*

$$\sigma \uparrow_\ell = \lambda \mathbf{x}. \begin{cases} \sigma(\mathbf{x}), & \text{if } (\Gamma(\sigma(\mathbf{x})) \sqsubseteq \ell) \\ \star(\perp, \{\mathbf{x}\}), & \text{otherwise} \end{cases}$$

Program configurations for expressions and commands are extended with the projected trace and projections of the memory and budget store, given by $\langle \sigma \uparrow_\ell, \iota \uparrow_\ell, e, \tau \uparrow_\ell \rangle$ and $\langle \sigma \uparrow_\ell, \iota \uparrow_\ell, c, \tau \uparrow_\ell \rangle$, respectively. $\sigma \uparrow_\ell$, $\iota \uparrow_\ell$, and $\tau \uparrow_\ell$ represent the ℓ -projected memory, budget-map and trace, respectively.

The semantics of expression and command evaluation is shown in Figures 6.4 and 6.5 and is mostly similar to LIR semantics presented earlier (Figures 6.2 and 6.3) except that the decoding semantics takes the trace τ as input. The rules generate a new trace τ' , which is a suffix of the original trace τ . The main difference in the semantics of expressions shows up in the `S-DCOPR` rule — when declassification occurs, the declassified value is read from the trace (as opposed to the one computed by local evaluation in the sub-expressions, which is generally \star). The main difference in the semantics of commands is in the branching rules. Branching or looping on \star values skips the execution as shown in rules `S-IF-ELSE-S` and `S-WHILE-FS`.

The proof for the projected-trace being uniquely-decodable is based on the decoding semantics progressing on the trace generated by the LIR semantics and resulting in a unique final memory with respect to the adversary. This requires defining equivalence between the memory stores in the two semantics. The memory equivalence (Definition 25) is defined based on the observational equivalence of the values in the original memory store and the projected memory store (Definition 24). The idea behind the definition of value equivalence is that if a value is visible to the adversary in the original memory store, then the value is also visible to the adversary in the projected store (Definition 24.1). If the value is not visible to the adversary in the original memory store, then either the value shall never be declassified to the adversary in either semantics (Definition 24.2(a) and 24.2(b)) or the value can be declassified to the adversary under both the semantics (Definition 24.2(c)).

Definition 24. *Two values v_o and v_s , where v_o is a value in the original memory store and v_s is a value in the projected memory store for simulation, are observationally equivalent at level ℓ , written $v_o \simeq_\ell v_s$ iff*

1. $v_o = \mathbf{n}_o^{(k, \delta)}$, $v_s = \mathbf{n}_s^{(k', \delta')}$, $\mathbf{n}_o = \mathbf{n}_s$ and $(k = k') \sqsubseteq \ell$ and $\delta = \delta'$ (or)
2. $v_o = \mathbf{n}_o^{(k, \delta)}$, $\Gamma(v_o) \not\sqsubseteq \ell$, $v_s = \star^{(k', \delta')}$ and either:
 - (a) $k \not\sqsubseteq \ell \wedge k' \not\sqsubseteq \ell$
 - (b) $k \sqsubseteq \ell \wedge \exists \mathbf{x} \in \delta. (\mathbb{B}(\mathbf{x}) \not\sqsubseteq \ell) \wedge k' \not\sqsubseteq \ell$
 - (c) $(k = k') \sqsubseteq \ell \wedge \delta = \delta'$

$$\begin{array}{c}
\text{S-CONST} \frac{}{\langle \sigma, \iota, \mathbf{n}, \tau \rangle \downarrow_{pc} \mathbf{n}^{(\perp, \{\})}, \iota, \tau} \\
\\
\text{S-VAR} \frac{\sigma(\mathbf{x}) = \mathbf{v}^{(k, \delta_o)} \quad \delta' = \mathbb{R}(\delta_o, \iota, k_o, \mathbb{B}) \quad \delta = \delta_o \setminus \delta' \quad k = k_o \bigsqcup_{\mathbf{x} \in \delta'} \mathbb{L}(\mathbf{x})}{\langle \sigma, \iota, \mathbf{x}, \tau \rangle \downarrow_{pc} \mathbf{v}^{(k, \delta)}, \iota, \tau} \\
\\
\text{S-AOP} \frac{\langle \sigma, \iota, e_1, \tau \rangle \downarrow_{pc} \mathbf{v}_1^{(k_1, \delta_1)}, \iota_1, \tau_1 \quad \langle \sigma, \iota_1, e_2, \tau_1 \rangle \downarrow_{pc} \mathbf{v}_2^{(k_2, \delta_2)}, \iota', \tau' \quad \mathbf{v} = \mathbf{v}_1 \odot \mathbf{v}_2 \quad (k, \delta) = (k_1, \delta_1) \sqcup (k_2, \delta_2)}{\langle \sigma, \iota, (e_1 \odot e_2), \tau \rangle \downarrow_{pc} \mathbf{v}^{(k, \delta)}, \iota', \tau'} \\
\\
\text{S-COP} \frac{\langle \sigma, \iota, e_1, \tau \rangle \downarrow_{pc} \mathbf{v}_1^{(k_1, \delta_1)}, \iota_1, \tau_1 \quad \langle \sigma, \iota_1, e_2, \tau_1 \rangle \downarrow_{pc} \mathbf{v}_2^{(k_2, \delta_2)}, \iota', \tau' \quad \mathbf{v} = \mathbf{v}_1 \oplus \mathbf{v}_2 \quad (k, \delta) = (k_1, \delta_1) \sqcup (k_2, \delta_2)}{\langle \sigma, \iota, (e_1 \oplus e_2), \tau \rangle \downarrow_{pc} \mathbf{v}^{(k, \delta)}, \iota', \tau'} \\
\\
\text{S-DCOPN} \frac{\langle \sigma, \iota, (e_1 \oplus e_2), \tau \rangle \downarrow_{pc} \mathbf{v}^{(k, \delta)}, \iota', \tau' \quad pc \not\sqsubseteq k}{\langle \sigma, \iota, \text{declassify}(e_1 \oplus e_2), \tau \rangle \downarrow_{pc} \mathbf{v}^{(k, \delta)}, \iota', \tau'} \\
\\
\text{S-DCOPR} \frac{\langle \sigma, \iota, (e_1 \oplus e_2), \tau \rangle \downarrow_{pc} \mathbf{v}_o^{(k_o, \delta_o)}, \iota_1, \tau_1 \quad pc \sqsubseteq k_o \quad \delta' = \Delta(\delta_o, k_o, \mathbb{L}) \quad \delta'' = \mathbb{R}(\delta', \iota_1, k_o, \mathbb{B}) \quad \delta = \delta' \setminus \delta'' \quad m = k_o \bigsqcup_{\mathbf{x} \in \delta''} \mathbb{L}(\mathbf{x}) \bigsqcup_{\mathbf{y} \in \delta} \mathbb{B}(\mathbf{y})}{(\mathbf{v}, k, \iota', \tau') = \left\{ \begin{array}{l} (\mathbf{v}_t, k_t, \mathbb{I}(\iota_1, \delta_t), \tau''), \text{ if } (\tau_1 = \mathbf{v}_t^{k_t} :: \tau'') \wedge (k_t = m) \wedge (\delta \neq \emptyset) \\ (\mathbf{v}_o, m, \iota_1, \tau_1), \text{ otherwise} \end{array} \right\}} \\
\langle \sigma, \iota, \text{declassify}(e_1 \oplus e_2), \tau \rangle \downarrow_{pc} \mathbf{v}^{(k, \{\})}, \iota', \tau'
\end{array}$$

Figure 6.4: Decoding semantics of expressions

$$\begin{array}{c}
\text{s-SKIP} \frac{}{\langle \sigma, \iota, \text{skip}, \tau \rangle \downarrow_{pc} \langle \sigma, \iota, \tau \rangle} \\
\text{s-ASSN} \frac{pc \sqsubseteq \Gamma(\sigma(\mathbf{x})) \quad \langle \sigma, \iota, e, \tau \rangle \downarrow_{pc} \mathbf{n}^{(m, \delta)}, \iota', \tau' \quad k = pc \sqcup m \quad \delta = \Delta(\delta', k, \mathbb{L})}{\langle \sigma, \iota, (\mathbf{x} := e), \tau \rangle \downarrow_{pc} \langle \sigma[\mathbf{x} \mapsto \mathbf{n}^{(k, \delta)}], \iota', \tau' \rangle} \\
\text{s-SEQ} \frac{\langle \sigma, \iota, c_1, \tau \rangle \downarrow_{pc} \langle \sigma', \iota', \tau_1 \rangle \quad \langle \sigma', \iota', c_2, \tau_1 \rangle \downarrow_{pc} \langle \sigma'', \iota'', \tau' \rangle}{\langle \sigma, \iota, c_1; c_2, \tau \rangle \downarrow_{pc} \langle \sigma'', \iota'', \tau' \rangle} \\
\text{s-IF-ELSE-N} \frac{\langle \sigma, \iota, e, \tau \rangle \downarrow_{pc} \mathbf{v}^{(k_o, \delta)}, \iota', \tau_1 \quad \mathbf{v} \neq \star \quad i = \begin{cases} 1, & \text{if } (\mathbf{v} = \text{true}) \\ 2, & \text{otherwise} \end{cases} \quad k = k_o \bigsqcup_{x \in \delta} \mathbb{L}(x) \quad \langle \sigma, \iota', c_i, \tau_1 \rangle \downarrow_{pc \sqcup k} \langle \sigma', \iota'', \tau' \rangle}{\langle \sigma, \iota, (\text{if } e \text{ then } c_1 \text{ else } c_2), \tau \rangle \downarrow_{pc} \langle \sigma', \iota'', \tau' \rangle} \\
\text{s-IF-ELSE-S} \frac{\langle \sigma, \iota, e, \tau \rangle \downarrow_{pc} \star^{(-)}, \iota', \tau'}{\langle \sigma, \iota, (\text{if } e \text{ then } c_1 \text{ else } c_2), \tau \rangle \downarrow_{pc} \langle \sigma, \iota', \tau' \rangle} \\
\text{s-WHILE-FS} \frac{\langle \sigma, \iota, e, \tau \rangle \downarrow_{pc} \mathbf{v}^{(-)}, \iota', \tau' \quad (\mathbf{v} = \star) \vee (\mathbf{v} = \text{false})}{\langle \sigma, \iota, \text{while } e \text{ do } c, \tau \rangle \downarrow_{pc} \langle \sigma, \iota', \tau' \rangle} \\
\text{s-WHILE-T} \frac{\langle \sigma, \iota, e, \tau \rangle \downarrow_{pc} \text{true}^{(k_o, \delta)}, \iota_1, \tau_1 \quad k = k_o \bigsqcup_{x \in \delta} \mathbb{L}(x) \quad \langle \sigma, \iota_1, c; \text{while } e \text{ do } c, \tau_1 \rangle \downarrow_{pc \sqcup k} \langle \sigma', \iota', \tau_2 \rangle}{\langle \sigma, \iota, \text{while } e \text{ do } c, \tau \rangle \downarrow_{pc} \langle \sigma'', \iota'', \tau'' \rangle}
\end{array}$$

Figure 6.5: Decoding semantics of commands

Definition 25. Given a memory store σ and a projected memory store $\sigma' \uparrow_\ell$, their equivalence \simeq_ℓ at level ℓ is defined as: $\forall x. \sigma(x) \simeq_\ell \sigma' \uparrow_\ell(x)$

Under a given memory σ and budget store ι , if a program executes to completion under the LIR semantics and generates a trace τ , then for any chosen adversary ℓ the same program executing under ℓ -projected memory $(\sigma \uparrow_\ell)$ and budget store $(\iota \uparrow_\ell)$ along with the projected trace $\tau \uparrow_\ell$ results in an observationally equivalent final memory stores obtained at the end of both the executions with respect to the adversary ℓ (Theorem 11).

Theorem 11 (Simulation Theorem). If $\langle \sigma, \iota, c \rangle \downarrow_{pc} \langle \sigma', \iota', \tau \rangle$, then $\langle \sigma \uparrow_\ell, \iota \uparrow_\ell, c, \tau \uparrow_\ell \rangle \downarrow_{pc} \langle \sigma'', \iota'', \tau'' \rangle$ such that

$$\sigma' \simeq_{\ell} \sigma'', l' \sim_{\ell} l'' \text{ and } \tau'' = \epsilon$$

The proof of the theorem and the various auxiliary lemmas required to prove the theorem are described in Appendix D. As a sanity check, when the budgets of all the secrets in the store are 0, the LIR semantics satisfy non-interference (Corollary 3).

Corollary 3 (Non-interference). *If $\langle \sigma_1, \iota_1, c \rangle \Downarrow_{pc} \langle \sigma'_1, \iota'_1, \tau_1 \rangle$, $\langle \sigma_2, \iota_2, c \rangle \Downarrow_{pc} \langle \sigma'_2, \iota'_2, \tau_2 \rangle$, $\sigma_1 \sim_{\ell} \sigma_2$, and $\iota_1 \uparrow_{\ell} = \iota_2 \uparrow_{\ell} = 0$, then $\sigma'_1 \sim_{\ell} \sigma'_2$.*

Part V

Application to an IFC-enabled Web Browser

Chapter 7

Information Flow Policies for Web Browsers

Previously proposed approaches [22, 31, 55, 81], for enforcing dynamic information flow control in web browsers lack adequate support for specifying policies conveniently. Flowfox [97] provides a rich policy framework but all websites are subject to the same policy, and the underlying IFC technique, secure multi-execution [43], does not handle shared state soundly. COWL [94] uses coarse-grained isolation, allowing scripts' access to either remote domains or the shared state, but not both. This requires significant code changes when both are needed simultaneously (see Chapter 9 for more details).

This chapter presents WebPol, a policy framework that allows a web-page developer to release data selectively to third-party scripts (to obtain useful functionality), yet control what the scripts can do with the data. WebPol policies label sensitive content (page elements and user-generated events) at source, and selectively declassify them by specifying where (to which domains) the content and its derivatives can flow. Host page developers specify WebPol policies in JavaScript, a language already familiar to them. WebPol is integrated with the instrumentation of the dynamic IFC framework for WebKit presented thus far (WebPol integrates with any taint-based IFC solution to overcome the shortcomings listed above). The expressiveness of WebPol policies is demonstrated through examples and by applying WebPol to two real websites.

7.1 Policy Component for Web Browsers

IFC is a broad term for techniques that control the flow of sensitive information in accordance with pre-defined policies. Sensitive information is information derived from sources that are confidential

The content of this chapter is based on the work published as part of the paper, "WebPol: Fine-grained Information Flow Policies for Web Browsers" [24]

or private. Any IFC system has two components—the *policy component* and the *enforcement component*. The policy component allows labeling of private information sources. The label on a source specifies how private information from that source can be used and where it can flow. The collection of rules for labeling is called the policy. The enforcement component enforces policies. This could, for example, be the dynamic taint tracking approach described earlier. WebPol contributes a policy component to complement existing work on enforcement components in web browsers.

In the context of web-pages, data sources are objects generated in response to user events like the content of a password box generated due to key presses or a mouse click on a sensitive button, and data obtained in a network receive event. In WebPol, data sources can be labeled with three kinds of labels, in increasing order of confidentiality: 1) the label `public` represents non-sensitive data, 2) for each domain `domain`, the label `domain` represents data private to the domain; such data's flow should be limited only to the browser and servers belonging to `domain` and its sub-domains, and 3) the label `local` represents very confidential data that must never leave the browser. These labels are ordered `public < domaini < local`. These labels are fairly expressive. For example, labeling a data source with the domain of the hosting page restricts its transfer to only the host, and prevents exfiltration to third-parties. Labeling a data source with the domain of a service provider such as an analytics provider allows transfer to only that service.

Since most data on a web-page is not sensitive, it is reasonable to label data sources `public` by default and only selectively assign a different label. WebPol uses this blacklisting approach. Two nuances of source labeling are noteworthy. The first is its fine granularity. Not all objects generated by the same class of events have the same label. For instance, characters entered in a password field may have the domain label of the hosting page, limiting their flow only to the host, but characters entered in other fields may be accessible to third-party advertising or analytics scripts unrestricted. This leads to the following requirement on the policy component.

Requirement 1: The policy component must allow associating different policies with different elements of the page.

The second is that the label of an object can be dynamic, i.e., history-dependent. Consider a policy that hides from an analytics script how many times a user clicked within an interactive panel, but wants to share whether or not the user clicked at least once. The label of a click event on the panel is `public` the first time the user clicks on it and `private` afterwards and, hence, it depends on the history of user interaction. This yields the following requirement on the policy component.

Requirement 2: Labels may be determined dynamically. This requirement means that labels must be set by *trusted policy code* that is executed on-the-fly and that has local state.

7.2 WebPol policy model

WebPol works on a browser that has already been augmented with IFC enforcement based on taint tracking. It provides a framework that allows setting labels at fine-granularity, thus expressing and enforcing rich policies. This section explains the design of WebPol. WebPol prevents under-the-hood

exfiltration of sensitive data that has been provided to third-party scripts for legitimate reasons. So, third-party scripts are not trusted but code from the host domain is trusted. WebPol's policies are agnostic to specific channels of leak. However, current IFC enforcements in browsers track only explicit and implicit flows. Consequently, leaks over other channels such as timing and memory-usage are currently out of scope.

7.2.1 Policies as event handlers

The first question in the design of WebPol is who should specify policies. Since the goal here is to prevent exfiltration of data by third-party scripts and it is the developer of the host page who bootstraps the inclusion of scripts and best understands how data on the page should be used, it is natural and pragmatic to have the developer specify policies, possibly as part of the page itself.

The next question is how the developer specifies policies. To answer this, recall the two requirements identified in Section 7.1 — it should be possible to specify different policies on different page elements and policies should be allowed to include code that is executed on-the-fly to generate labels. Considering the fact that sensitive data is usually generated by input events, it is clear that policies should *be* page element-specific (trusted) code that is executed after events have occurred (this code labels event-generated data). Fortunately, web browsers provide exactly this abstraction in the form of event handlers! So, the event-handling logic in web browsers is extended to express WebPol policies. This allows leveraging a lot of the existing browser logic for event handler installation, parsing and event dispatch in interpreting policies. The rest of this section explains how this is done.

In WebPol, policies are special event handlers, specified using a special marker in the HTML source of the hosting page. These special handlers, called *policy handlers*, follow standard JavaScript syntax, can be attached to any page element, listening for any event and, like other handlers, are triggered every time the event is dispatched on the element or any of its descendants. However, unlike other handlers, the sole goal of policy handlers is to assign labels to other sensitive objects, including the event being dispatched. To allow the policy handlers to do this, the browser is slightly modified to afford these handlers two special privileges:

- Policy handlers can execute two new JavaScript API functions that set labels on other objects. No other JavaScript code can execute these two functions. These functions are described later.
- During event dispatch all applicable policy handlers are executed before ordinary handlers. This ensures that labels are set before ordinary handlers (including those of third-party scripts) execute.

To maintain the integrity of the policies, policy handlers must be included in the HTML source of the page directly. They *cannot* be installed dynamically by JavaScript code. Otherwise, third-party scripts could install policy handlers that set very permissive labels. Also, if a DOM element has a policy handler, third-party scripts are disallowed from detaching that element or moving it elsewhere, as that can change the interpretation of the policy. Similarly, changing the attributes of such an element is restricted.

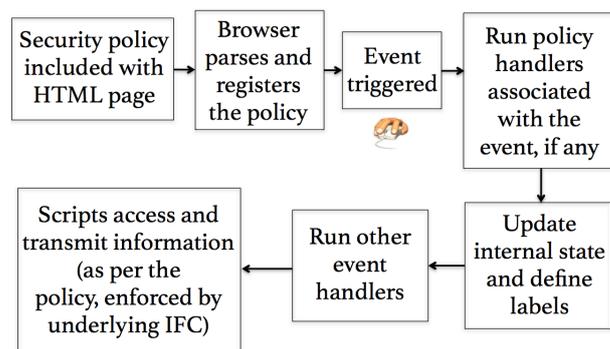


Figure 7.1: Work-flow of the WebPol policy model

Since different policy handlers can be associated with different elements, Requirement 1 is satisfied. Moreover, policy handlers are ordinary JavaScript code, so they can also maintain local state in private variables, thus satisfying Requirement 2.

The work-flow of policy interpretation in WebPol is shown in Figure 7.1. Briefly, the steps are:

1. The web page developer specifies the policy in the host HTML page in the form of special event handlers.
2. The browser parses the policy and registers its handlers (mostly like usual handlers, but with the two special privileges mentioned above).
3. When an event dispatches, listening policy handlers are executed first.
4. These policy handlers set labels on objects affected by the event, including the event object itself. They may also update any local state they maintain.
5. The remaining event handlers are dispatched as usual. The IFC enforcement in the browser enforces all labels that have been set by the policy handlers (during any prior event's dispatch), thus preventing any data leak in contravention of the labels.

7.2.2 Integration with the web browser

WebPol needs minor modifications to the browser to parse and interpret policies and to expose additional JavaScript API functions to set labels.

HTML and event dispatch changes. WebPol adds an HTML extension to differentiate policy code from other JavaScript code. Concretely, the browser's parser is changed to interpret any script file with the extension `.policy` included directly in the host page as a policy. If such a *policy script* installs a handler, it is treated as a policy handler. Additionally, a policy script can set labels on the page's global variables and DOM elements (like password fields). If a script does this, it should be included in the host page before third-party scripts that use those variables. WebPol also requires a small change to the browser's event dispatch mechanism to execute policy handlers before other handlers.

```
1 var p = document.getElementById("pwd");
2 p.addEventListener("keypress", function (e){
3     var score = checkPwdStrength(p.value);
4     document.getElementById("pwdStrength").innerText = score;
5     new Image().src = "http://stealer.com/pwd.jsp?pwd="+p +score;
6 });
```

Listing 7.1: Password strength checking script that leaks the password

Label-setting APIs. WebPol exposes two new JavaScript API functions to set labels. These functions can be called only by the policy code in `.policy` files and handlers installed by such files (the browser is modified to enforce this).

The function `setLabel(label)` sets the label of the object on which it is called to `label`. As explained earlier, `label` can be `public`, a domain name, or `local` (the default is `public`). Once an object's label is set, it is enforced by the underlying IFC enforcement. The special label `HOST` is a proxy for the domain of the host page.

The function `setContext(label)` can be called only on an event object. It restricts the *visibility* of the event to label `label` and higher. In simple terms, if `label` is a domain, then only that domain can ever learn that this event occurred, whereas if `label` is `local`, then no domain can ever learn that this event occurred. Technically, this is accomplished by setting the *pc* of the event handlers running during the dispatch to `label`, which ensures that their side-effects (writes to DOM and network communication) are labeled `label` or higher.

As opposed to `setLabel`, which makes individual data objects (like password fields) private, `setContext` makes the *existence* of an event private. This is useful. For instance, clicking on the "politics" section of a news feed might indicate that the user is interested in politics, which may be private information, so the page may want to hide even the existence of click events from third-party scripts. (The distinction between the privacy of event content and event occurrence has been previously described by Rafnsson and Sabelfeld [80].)

7.3 Expressiveness of WebPol

The expressiveness of WebPol policies is illustrated through the following examples.

7.3.1 Example 1: Password strength checker

Many websites deploy *password strength checkers* on pages where users set new passwords. A password strength checker is an event handler from a third-party library that is triggered each time the user enters a character in the new password field. The handler provides visual feedback to the user about the strength of the password entered so far. Strength checkers usually check the length of the password and the diversity of characters used. Consequently, they do not require any network communication. However, standard browser policies cannot enforce this and the password strength

```

1 function currencyConverter() {
2     var toCur = document.getElementById("to").value;
3     var xh = new XMLHttpRequest();
4     xh.onreadystatechange = function() {
5         if (xh.readyState == 4) {
6             currencyRate = eval(xhttp.responseText);
7             var aAmt = document.getElementById("amt").value;
8             var convAmt = aAmt * currencyRate;
9             document.getElementById("camt").innerHTML = convAmt;
10            xh.open("GET", "http://currConv.com/amount.jsp?atc=" + aAmt);
11            xh.send(); }}
12    xh.open("GET", "http://currConv.com/conv.jsp?toCur=" + toCur, true);
13    xh.send(); }

```

Listing 7.2: Currency converter script that leaks a private amount

checker can easily leak the password if it wants to. Listing 7.1 shows such a “leaky” password checker. The checker installs a listener for keypresses in the password field (line 2). In response to every keypress, the listener delivers its expected functionality by checking the strength of the password and indicating this to the user (lines 3, 4), but then it leaks out the password to `stealer.com` by requesting an image at a URL that includes the password (lines 5, 6).

With WebPol, the developer of the host web-page can prevent any exfiltration of the password by including the policy script:

```
document.getElementById("pwd").setLabel("HOST");
```

This policy sets the label of the password field to the host’s own domain using the function `setLabel()`. Subsequently, the IFC enforcement restricts all outgoing communication that depends on the password field to the host.

Conceptually, this example is simple because it does not really leverage the fine-granularity of WebPol policies and fine-grained dynamic IFC. Here, the third-party script does not need any network communication for its intended functionality and, hence, simpler confinement mechanisms that prohibit a third-party script from communicating with remote servers would also suffice. The next example is a scenario where the third-party script legitimately needs remote communication, which leverages the fine-granularity of WebPol policies and fine-grained dynamic IFC.

7.3.2 Example 2: Currency conversion

Consider a web-page from an e-commerce website which displays the cost of an item that the user intends to buy. The amount is listed in the site’s native currency, say US dollars (USD), but for the user’s convenience, the site also allows the user to see the amount converted to a currency of his/her choice. For this, the user selects a currency from a drop-down list. A third-party JavaScript library reads both the USD amount and the second currency, converts the amount to the second currency and inserts it into the web-page, next to the USD amount. The third-party script fetches the current

```
1 var p = document.getElementById("sect_name");
2 p.addEventListener("click",function(event){
3   event.setLabel("HOST"); });
```

Listing 7.3: Policy that allows counting clicks but hides details of the clicks

```
1 clickCount = 0;
2 var p = document.getElementById("sect_name");
3 p.addEventListener("click",function (e){ clickCount += 1; });
```

Listing 7.4: Analytics script that counts clicks

conversion rate from its back-end service at `currConv.com`. Consequently, it must send the *name* of the second currency to its back-end service, but must not send the amount being converted (that is private information). The web browser's same-origin policy has been relaxed (using, say, CORS [2]) to allow the script to talk to its back-end service at `currConv.com`. The risk is that the script can now exfiltrate the private amount. Listing 7.2 shows a leaky script that does this. On line 11, the script makes a request to its back-end service passing to it the two currencies. The callback handler (lines 4–11) reads the amount from the page element `amt`, converts it and inserts the result into the page (lines 6–9). Later, it leaks out the amount to the back-end service on line 10, in contravention of the intended policy.

With WebPol, this leak can be prevented with the following policy that sets the label of the amount to the host only:

```
document.getElementById("amt").setLabel("HOST")
```

This policy will prevent exfiltration of the amount and will not interfere with the requirement to exfiltrate the second currency. Importantly, no modifications are required to a script that does not try to leak data (e.g., the script obtained by dropping the leaky line 10 of Listing 7.2).

7.3.3 Example 3: Web analytics

To better understand how users interact with their websites, web developers often include third-party analytics scripts that track user clicks and keypresses to generate useful artifacts like page heat-maps (which part of the page did the user interact with most?). Although a web developer might be interested in tracking only certain aspects of their users' interaction, the inclusion of the third-party scripts comes with the risk that the scripts will also record and exfiltrate other private user behavior (possibly for monetizing it later). Using WebPol, the web developer can write precise policies on which user events an analytics script can access and when. Several examples of this are shown.

To allow a script to only count the number of occurrences of a class of events (e.g., mouse clicks) on a section of the page, but to hide the details of the individual events (e.g., the coordinates of every individual click), the web developer can add a policy handler on the top-most element of the section

```
1 var alreadyClicked = false;
2 var p = document.getElementById("sect_name");
3 p.addEventListener("click",function (event){
4   if (alreadyClicked = true)
5     event.setContext("HOST");
6   else {
7     alreadyClicked = true;
8     event.setLabel("HOST");
9   }});
```

Listing 7.5: Policy that only tracks whether a click happened or not

```
1 document.body.addEventListener("keypress", function (event){
2   var o = window.getComputedStyle(event.target).getPropertyValue("opacity");
3   if (o < 0.5)
4     event.setLabel("HOST");
5 });
```

Listing 7.6: Example policy to prevent overlay-based stealing of keystrokes

to set the label of the individual event objects to HOST. This prevents the analytics script's listening handler from examining the details of individual events, but since the handler is still invoked at each event, it can count their total number. Listings 7.3 and 7.4 show the policy handler and the corresponding analytics script that counts clicks in a page section named `sect_name`.

Next, consider a restriction of this policy, which allows the analytics script to learn only whether or not *at least one* click happened in the page section, completely hiding clicks beyond the first. This policy can be represented in WebPol using a local state variable in the policy to track whether or not a click has happened and the function `setContext()`. Listing 7.5 shows the policy. The policy uses a variable `alreadyClicked` to track whether or not the user has clicked in the section. Upon the user's first click, the policy handler sets the event's label to the host's domain (line 8). This makes the event object private but allows the analytics handler to trigger and record the occurrence of the event. On every subsequent click, the policy handler sets the event's *context* to the host domain using `setContext()` (line 5). This prevents the analytics script from exfiltrating any information about the event, including the fact that it occurred.

Finally, note that a developer can subject different page sections to different policies by attaching different policy handlers to them. The most sensitive sections may have a policy that unconditionally sets the event context to the host's, effectively hiding all user events in those sections. Less sensitive sections may have policies like those of Listings 7.5 and 7.3. Non-sensitive sections may have no policies at all, allowing analytics scripts to see all events in them.

7.3.4 Example 4: Defending against overlay-based attacks

To bypass WebPol policies, an adversarial script may “trick” a user using transparent overlays. For example, suppose a script wants to exfiltrate the contents of a password field that is correctly protected by a WebPol policy. The script can create a transparent overlay on top of the password field. Any password the user enters will go into the overlay, which *isn't* protected by any policy and, hence, the script can leak the password.

Such attacks can be prevented easily in WebPol using a *single* policy, attached to the top element of the page, that labels data entered into all significantly transparent overlays as HOST. Listing 7.6 shows such a policy. This particular policy labels all keypress events on elements of opacity below 0.5 as HOST, thus preventing their exfiltration. The threshold value 0.5 can be changed, and the policy can be easily extended to other user events like mouse clicks.

7.3.5 Summary of WebPol expressiveness

The security community has extensively studied several aspects of policy labeling, colloquially called the *dimensions of declassification* (see [91] for a survey). Broadly speaking, WebPol policies cover three of these dimensions—the policies specify what data is declassified (dimension: what), to which domains (dimension: to whom) and under what state (dimension: when). “What data is declassified” is specified by selectively attaching policies to elements of the page. “Which domains get access” is determined directly by the labels that the policy sets. Finally, labels generated by policy handlers can depend on state, as illustrated in Listing 7.5.

There are two other common dimensions of labeling—who can label the data (dimension: who) and where in the code can the labels change (dimension: where). These dimensions are fixed in WebPol due to the specifics of the problem: All policies are specified by the host page in statically defined policies.

Chapter 8

Policy Implementation and Evaluation in an IFC-enabled Browser

This chapter describes the implementation of WebPol (Chapter 7) and the LIR policy (Chapter 6) on top of an existing IFC-instrumentation [22, 60, 81] (referred to as WebIFC in the remaining of this chapter). WebIFC targets WebKit, the browser engine used in Safari, for enforcing dynamic IFC in the three main components of the engine — JavaScript bytecode interpreter, the document object model (DOM) engine, and the event handling mechanism.

Labels in WebIFC are word size bit-sets (currently 64 bits); each bit in the bit-set represents label from a distinct domain (like google.com). Join on labels is simply bitwise or. WebIFC adds labels to all data structures, including registers, objects, object properties and scope chain pointers, and attach security labels to every node in the DOM graph and all its properties, including pointer to other nodes. In WebIFC, the code is instrumented to propagate explicit and implicit labels and implement the permissive-upgrade check. Specifically, WebKit’s JavaScript bytecode interpreter (JavaScriptCore) for enforcing dynamic IFC for JavaScript [22, 60]. In WebKit, bytecode is generated by a source-code compiler and organized into code blocks. Each code block is a sequence of bytecodes with line numbers and corresponds to the instructions for a function or an `eval` statement. A code block is generated when a function is created or an `eval` is executed. WebIFC performs control flow analysis on a code block when it is created and generates a CFG for it before it starts executing. The IPDs of its nodes are calculated by static analysis of its bytecode; they are computed using an algorithm by Lengauer and Tarjan with CFG as an input to the algorithm [65]. The formalization of the bytecodes, the semantics of its bytecode interpreter with the instrumentation of dynamic IFC, and the proof of correctness of instrumentation of the bytecodes with the IFC semantics is shown in [23]. Additionally, all native JavaScript methods in the Array, RegExp, and String objects are instrumented in WebIFC and appropriate IFC checks in the native C code implementing all DOM APIs

The content of this chapter is based partly on the work published as part of the paper, “WebPol: Fine-grained Information Flow Policies for Web Browsers” [24]

up to Level 3 are added [81]. WebIFC also modifies the event handling loop, labeling every event and event handler based on their formalization of the event handling loop with the IFC checks [81].

WebIFC enforces termination-insensitive non-interference in all the major components of WebKit while tracking labels at a fine-granularity. This thesis adds additional instrumentation on top of WebIFC for enforcing the ideas presented earlier. For exceptions, a synthetic exit node is added to the CFG along with the edges as described earlier in Section 5.2. Additional changes are required to the compiler to make it compliant with the instrumentation. The modification is to emit a slightly different, but functionally equivalent bytecode sequence for `finally` blocks; this is needed for accurate computation of IPDs.

8.1 Implementation and Evaluation of WebPol

8.1.1 Implementation of WebPol

WebPol is prototyped in WebKit on top of WebIFC. To implement WebPol, the HTML parser was modified to distinguish policy files (extension `.policy`) from other JavaScript files and to give policy code extra privileges. Two new JavaScript API functions — `setLabel()` and `setContext()` — were added. Finally, the event dispatch logic was modified to trigger policy handlers before other handlers. In all, 25 lines in the code of the parser were modified, 60 lines for the two new API functions were added and 110 lines in the event dispatch logic were modified. Thus, implementing WebPol has low overhead, and can be ported to other browsers easily.

8.1.2 Evaluation of WebPol

The goal of the evaluation is two-fold. The first goal is to measure the overhead of the system with IFC enforcement and WebPol, both in parsing and installing policies during page load and for executing policy handlers later. This is done by running a few benchmarks, by measuring the overhead for the examples presented in Section 7.3, and for two real-world websites. Second, to understand whether WebPol can be used easily, WebPol policies are applied to two real-world websites. All the experiments were performed on a 3.2GHz Quad-core Intel Xeon processor with 8GB RAM, running Mac OS X version 10.7.4 using Safari 6.0. The implementation and evaluation is done on WebKit nightly build `#r122160`. As WebIFC does not cover JIT, JIT support was disabled in all the experiments.

8.1.2.1 Performance Overheads on Synthetic Examples

To measure the instrumentation's runtime overhead, four examples from Chapter 7.3 (Examples 1, 2 and the two sub-examples of Example 3) were tested in three different configurations: **Base**—uninstrumented browser, no enforcement; **WebIFC**—existing instrumented browser with IFC checks, but no policy handlers (everything is labeled public); **WebPol**—instrumented browser running policy handlers.

Example #	JavaScript Execution Time			Page Load Time		
	Base	WebIFC	WebPol	Base	WebIFC	WebPol
Example 1	2430	2918 (+20.1%)	2989 (+1.9%)	16	17 (+6.3%)	19 (+12.5%)
Example 2	3443	4361 (+26.7%)	5368 (+29.2%)	41	43 (+4.9%)	46 (+7.2%)
Example 3 (count)	1504	1737 (+15.5%)	1911 (+11.6%)	24	25 (+4.2%)	31 (+25.0%)
Example 3 (presence)	1780	2095 (+17.7%)	2414 (+18.9%)	26	28 (+7.7%)	30 (+7.7%)

Table 8.1: Performance of examples from Section 7.3. All time in ms. The numbers in parenthesis are additional overheads relative to **Base**.

JavaScript execution time: The overheads of executing policy handler code were measured by interacting with all four programs manually by entering relevant data and performing clicks a fixed number of times. For each of these configurations, the total time spent *only in executing JavaScript* was measured, including scripts and policies loaded initially with the page and the scripts and policies executed in response to events. The difference between **WebIFC** and **Base** run times is the overhead of dynamic IFC, while the difference between the **WebPol** and **WebIFC** run times is the overhead of evaluating policy handlers. Since only JavaScript execution time is measured and there are no time-triggered handlers in these examples, variability in the inter-event gap introduced by the human actor does not affect the measurements.

The left half of Table 8.1 shows these observations. All numbers are averages of 5 runs and the standard deviations are all below 7%. **WebIFC** adds overheads ranging from 15.5% to 26.7% over **Base**. To this, the policy handlers (**WebPol**) add overheads ranging from 1.9% to 29.2%. The **WebPol** overheads are already modest, but this is also a very challenging (conservative) experiment for WebPol. The scripts in both sub-examples of Example 3 do almost nothing. The scripts in Examples 1 and Example 2 are slightly longer, but are still much simpler than real scripts. On real and longer scripts, the relative overheads of evaluating the policy handlers is significantly lower as shown later. Moreover, the baseline in this experiment does not include other browser costs, such as the cost of page parsing and rendering, and network delays. Compared to those, both **WebIFC** and **WebPol** overheads are negligible.

Page load time: The time taken for loading the initial page (up to the DOMContentLoaded event) was measured separately. The difference between **WebPol** and **WebIFC** is the overhead for parsing and loading policies. The right half of Table 8.1 shows these observations. All numbers are the average of 20 runs and standard deviations were below 8%. WebPol overheads due to policy parsing and loading range from 7.2% to 25% (last column). When the overheads due to taint tracking (column **WebIFC**) are added, the numbers increase to 12.1% to 29.2%. Note that page-load overheads are incurred only once on every page (re-)load.

```

1 document.getElementById("passwordPwd").setLabel("secret");
2 document.getElementById("passwordTxt").setLabel("secret");
3 var x = document.getElementsByTagName("div");
4 var i = 0;
5 for (i = 0; i < x.length; i++)
6     x[i].setLabel("secret");

```

Listing 8.1: Policy code for password strength checking website

```

1 var x = document.getElementsByClassName("user"); // username
2 var y = document.getElementsByClassName("pwd"); // password
3 for (i = 0; i < x.length; i++) {
4     x[i].addEventListener("keypress", function(event){
5         event.setLabel("HOST");
6     });}
7 for (i = 0; i < y.length; i++) {
8     y[i].addEventListener("keypress", function(event){
9         event.setLabel("HOST");
10    });}

```

Listing 8.2: Policy code for bank log-in website with an analytics script

8.1.2.2 Policies on Real-world Websites.

Further, WebPol was evaluated by writing policies for two real-world applications—a website that deploys a password-strength checker (similar to Example 1) and a bank log-in page that includes third-party analytics scripts (similar to Example 3). The WebPol policies specified for the password strength checking website and the bank website with an analytics script are shown in Listings 8.1 and 8.2, respectively. The code on lines 3–6 of Listing 8.1 allows the strength-checking script to write back the visual indicator of password strength to the host page’s DOM.

Experience writing policies: In both cases, meaningful policies could be specified easily after understanding the code, suggesting that WebPol policies can be (and should be) written by website developers. The policy for the password-strength checker is similar to Listing 7.1 and prevents the password from being leaked to third-parties. Additional policy code of 4 lines was needed to allow the script to write the results of the password strength check (which depends on the password) into the host page. The analytics script on the bank website communicates all user-behavior to its server. The policy specified disallows exfiltration of keypresses on the username and the password text-boxes to third-parties.

Performance overheads: The performance overheads on the two websites were also measured, in the same configurations as for the synthetic examples. Table 8.2 shows the results. On real-world websites, where actual computation is long, the overheads of WebPol are rather small. The overheads of executing policy handlers, even relative only to **Base**’s JavaScript execution time, are 4.0% and 13.2%, while the overheads of parsing and loading policies are no more than 4.0%. Even the to-

Website	JavaScript Execution Time			Page Load Time		
	Base	WebIFC	WebPol	Base	WebIFC	WebPol
Password	79.5	115.5 (+45.3%)	126 (+13.2%)	303	429 (+41.6%)	441 (+4.0%)
Analytics	273.4	375.1 (+37.2%)	386.1 (+4.0%)	2151	2422 (+12.6%)	2499 (+3.6%)

Table 8.2: Performance on two real-world websites. All time in ms. The numbers in parenthesis are additional overheads relative to **Base**.

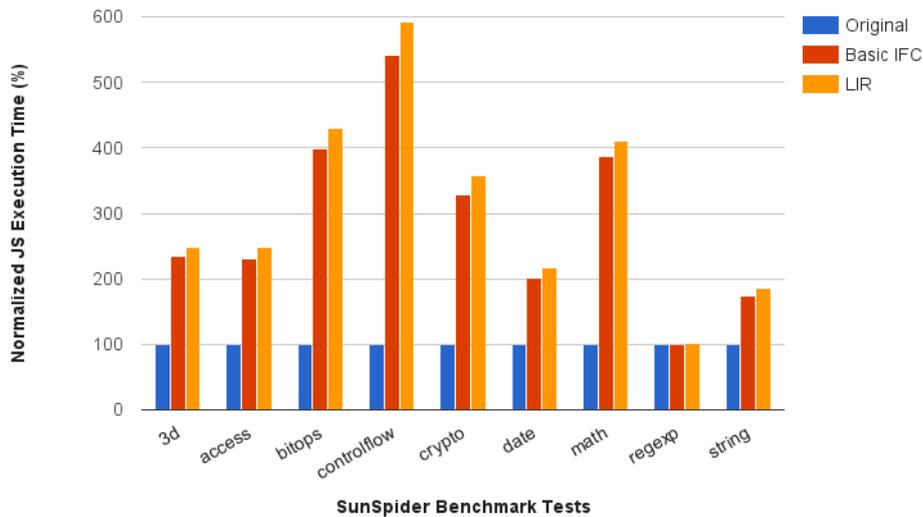


Figure 8.1: Overheads of LIR on SunSpider Benchmark Tests

tal overhead of **WebIFC** and **WebPol** does not adversely affect the user experience in any significant way.

8.2 Implementation and Evaluation of LIR

To implement the LIR semantics described in Chapter 6, the security label attached to the JavaScript objects and the DOM nodes was modified to carry a *provenance-label*, representing the dependency set. The provenance-label is basically a bit-vector where each bit represents a distinct JavaScript object or DOM node. Each bit in the provenance-label is mapped to a budget, a budget label and an actual label of the object. The `setLabel()` API is extended to include the budget, and the budget label for a JavaScript object or a DOM node. If the budget is not specified, it is assumed to be 0. Similarly, if a budget label is not specified, it is assumed to be \perp . The checks are performed as per the semantics at every comparison operation assuming an implicit declassification at each comparison operation.

The performance overhead added as part of the LIR instrumentation in the IFC-enabled browser is evaluated and compared against the uninstrumented browser and WebIFC (Basic IFC). The per-

formance evaluation was done for the standard SunSpider 1.0.2 JavaScript benchmark suite on the uninstrumented browser, on the IFC-instrumented browser and on the LIR instrumented browser. The average overhead for LIR instrumentation over the original uninstrumented browser is 170% and adds only 17% to the overhead of WebPol.

Chapter 9

Related Work

Browser security is a very widely-studied topic. Here, only closely related work on browser security policies and policy enforcement techniques is described.

Information flow control and script isolation. With the widespread use of JavaScript, research in dynamic techniques for IFC has regained momentum. Nonetheless, static analyses are not completely futile. Guarnieri *et al.* [50] present a static abstract interpretation for tracking taints in JavaScript. However, the omnipresent `eval` construct is not supported and this approach does not take implicit flows into account. Chugh *et al.* propose a staged information flow approach for JavaScript [32]. They perform static server-side policy checks on statically available code and generate residual policy-checks that are applied to dynamically loaded code. This approach is limited to certain JavaScript constructs excluding dynamic features like dynamic field access or the `with` construct.

Austin and Flanagan [13] propose purely dynamic IFC for dynamically-typed languages like JavaScript. They use the no-sensitive-upgrade (NSU) check [104] to handle implicit flows. Their permissive-upgrade strategy [14] is more permissive than NSU but retains termination-insensitive non-interference. This work builds on the permissive-upgrade strategy. They also present faceted evaluation [15], which is the most permissive technique of the three. However, given the performance considerations, the technique is not suitable for enforcing information flow control in browsers. Just *et al.* [60] present dynamic IFC for JavaScript bytecode with static analysis to determine implicit flows precisely even in the presence of semi-unstructured control flow like `break` and `continue`. Again, NSU is leveraged to prevent implicit flows. This thesis builds on top of their work to enforce information flow control in browsers.

JSFlow [54, 55] is a stand-alone implementation of a JavaScript interpreter with fine-grained taint tracking. Many seminal ideas for labeling and tracking flows in JavaScript owe their lineage to JSFlow, but since JSFlow is written from scratch it has very high overheads and introduces annotations to deal with semi-structured control flow. It detects security violations due to branches that have not been executed and injects annotations to prevent these in subsequent runs. The approach presented

in this thesis relies on analyzing CFGs and does not require annotations. To improve permissiveness, their subsequent work [27] uses testing.

Kerschbaumer *et al.* [62] build an implementation of an information flow monitor for WebKit but do not handle all implicit flows. A black-box approach to enforcing non-interference is based on secure multi-execution (SME) [43]. Bielova *et al.* [26] and De Groef *et al.* [39] implement SME for web browsers. These systems do not attach labels to specific fields in the DOM. Instead, labels are attached to individual DOM APIs.

Chudnov and Naumann [31] present another approach to fine-grained IFC for JavaScript. They rewrite source programs to add shadow variables that hold labels and additional code that tracks taints. This approach is inherently more portable than that of JSFlow or the work presented in this thesis, both of which are tied to specific, instrumented browsers. However, it is unclear how this approach could be extended with a policy framework like WebPol that assigns state-dependent labels at runtime.

The work most closely related to WebPol is that of Vanhoef *et al.* [97] on stateful declassification policies in reactive systems, including web browsers. Their policies are similar to the ones presented here, but there are significant differences. First, their policies are attached to the browser and they are managed by the browser user rather than website developers. Second, the policies have coarse-granularity: They apply uniformly to all events of a certain type. Hence, it is impossible to specify a policy that makes keypresses in a password field secret, but makes other keypresses public. Third, the enforcement is based on secure multi-execution [43], which is, so far, not compatible with shared state like the DOM.

COWL [94] enforces mandatory access control at coarse-granularity. In COWL, third-party scripts are sandboxed. Each script gets access to either remote servers or the host's DOM, but not both. Scripts that need both must be re-factored to pass DOM elements over a message-passing API (`postMessage`). This can be both difficult and have high overhead. For scripts that do not need this factorization, COWL is more efficient than solutions based on fine-grained information flow tracking.

Mash-IF [67] uses static analysis to enforce IFC policies. Mash-IF's model is different from WebPol's model. Mash-IF policies are attached only to DOM nodes and there is no support for adding policies to new objects or events. Also, in Mash-IF, the browser user (not the website developer) decides what declassification is allowed. Mash-IF is limited to a JavaScript subset that excludes commonly used features such as `eval` and dynamic property access.

JSand [6] uses server-side changes to the host page to introduce wrappers around sensitive objects, in the style of object capabilities [75]. These wrappers mediate every access by third-party scripts and can enforce rich access policies. Through secure multi-execution, coarse-grained information flow policies are also supported. However, as mentioned earlier, it is unclear how secure multi-execution can be used with scripts that share state with the host page.

WebPol policies are enforced using an underlying IFC component. Although, in principle, any IFC technique such as fine-grained taint tracking [22, 55, 59], coarse-grained taint tracking [94] or secure multi-execution [43] can be used with WebPol, to leverage the full expressiveness of WebPol's finely-granular policies, a fine-grained IFC technique is needed.

Access control. The traditional browser security model is based on restricting scripts' access to data, not on tracking how scripts use data. In the traditional model, it is impossible to allow scripts access to data they need for legitimate purposes and, simultaneously, to prevent them from leaking the data on the side, which is the goal of IFC and WebPol. More broadly, no mechanism based only on access control can solve this problem. Nonetheless, some closely related work on access control in web browsers is discussed below.

All browsers today implement the same-origin policy (SOP) [19], which prevents a page and scripts included in it from making requests to domains other than the page's host. However, for pragmatic reasons, image requests are exempt, which is sufficient to leak information. Consequently, the SOP is not effective against malicious or buggy scripts. Cross-origin resource sharing (CORS) [2] relaxes the SOP further to allow some cross-origin requests. Content security policies (CSPs) [3] allow white- and black-listing scripts. Unlike WebPol, CSPs offer no protection against scripts that have been included by the developer without realizing that they leak information.

Conscript [74] allows the specification of fine-grained access policies on individual scripts, limiting what actions every script can perform. Similarly, AdJail [68] limits the execution of third-party scripts to a shadow page and restricts communication between the script and the host page.

Zhou and Evans [105] take a dual approach, where fine-grained access control rules are attached to DOM elements. The rules specify which scripts can and cannot access individual elements. Along similar lines, Dong *et al.* [45] present a technique to isolate sensitive data using authenticated encryption. Their goal is to reduce the size of the trusted computing base.

ADsafe [38] and FBJS [4] restrict third-party code to subsets of JavaScript, and use static analysis to check for illegitimate access. Caja [1] uses object capabilities to mediate all access by third-party scripts. Webjail [95] supports least privilege integration of third-party scripts by restricting script access based on high-level policies specified by the developer.

All these techniques enforce only access policies and cannot control what a script does with data it has access to.

Part VI

Conclusion and Outlook

Chapter 10

Conclusion

Security and privacy have been of paramount importance since computer systems and applications have started handling private, sensitive and confidential information. The need of the hour is to carefully handle such data as it interacts with various untrusted third-party applications. Although helpful in many scenarios, static information flow analyses are mostly ineffective when working with dynamic languages like JavaScript, which is an indispensable part of the modern web. The dynamic nature of JavaScript makes sound static analysis difficult. Dynamic information flow control is a promising step forward, though the permissiveness of dynamic analyses presents a major challenge to the practicality of these techniques. This thesis focuses on improving the usability of dynamic information flow techniques by:

- **Developing mechanisms that enhance the precision and permissiveness of the analyses:**
To improve the precision of dynamic information flow analysis, this thesis develops a sound approach for handling complex language features like unstructured control flow and exceptions. The existing approaches to handle these features are too conservative generating a lot of false-positives and often require additional annotations by the developer. In contrast, the methodology presented in this thesis performs a sound and precise dynamic information flow analysis using post-dominator analysis at runtime to handle these features without requiring any additional annotations from the developer. The approach is also shown to be the most precise approach for handling such complex features dynamically.

To further improve the permissiveness of dynamic information flow analysis, this thesis presents the design of a sound improvement and enhancement of the permissive-upgrade strategy. The development improves the original strategy's permissiveness by relaxing the rules for handling partially-leaked data while retaining soundness. The original strategy's enforcement was limited to a two-point security lattice, and lacked generalization to an arbitrary lattice (as is required for real-world scenarios). To this end, this thesis presents a non-trivial approach to generalize the applicability of the approach to an arbitrary security lattice.

- **Proposing a technique to bound the release of sensitive information in realistic applications:** Although dynamic quantification has been studied earlier, bounding information leaks dynamically remains an open problem. To this end, this thesis develops a sound approach to bound information leaks dynamically by allowing information release in accordance to a pre-specified budget that specifies the amount of information that can be released about the secret. The thesis proposes the property of limited information release to capture this security condition and proves its enforcement sound, information-theoretically.
- **Describing a comprehensive policy mechanism for easy specification of security policies:** To complement the work on enforcement components in web browsers, the thesis also explores a policy specification mechanism to specify flexible and useful information flow policies for web applications.

Chapter 11

Future Directions

11.1 Evaluating Information Flow Policies on Real-World Websites

The dissertation focused on bridging the gap between the theory of dynamic information flow control and its practicality in the real-world. As part of understanding the practicality of such approaches in real-world websites, it would be interesting to investigate the types of security policies developers would be interested in. The idea would be to come up with a set of generic policies that could help evaluate different real-world websites for privacy and confidentiality violations. Another interesting direction for evaluation of information flow policies on real-world websites is to explore integrity violations by third-party scripts. Broadly, these policies could be related to cookies, web storage, and other user-data used by different websites. For example, a policy related to cookies could be that third-party scripts should not be able to extract information from first-party cookies, which is currently allowed if the cookie is not marked *HttpOnly*. Similarly, a third-party script (loaded from the same domain) can modify the `localStorage` object of the host without any restrictions, which can be restricted or prevented by ensuring that such information does not leak via the third-party script.

11.2 Exploring Alternative Granularities for Enforcing Information Flow Control

The current enforcement of dynamic information flow control tracks the flow of information between all data structures in the system at a fine-granularity. Although this increases the precision of the analysis significantly, it affects the performance adversely. Alternatively, one could explore approaches that are a little coarser than the current approach without giving up much on the pre-

cision. One such approach would be to only track information flow globally and at the level of a function-call, rather than tracking it through all local data-structures. The challenge here would be to retain the precision of fine-grained techniques, which can probably be achieved by code-rewriting. The high-level idea would be to execute the part of the code not dependent on sensitive data before executing that part of the code that operates on sensitive data.

11.3 Handling Timing Leaks on the Web

Primarily, the work done in the area of dynamic information flow control has focused on handling two types of leaks — leaks due to explicit and implicit flows. However, there are various other covert channels like timing, resource-usage etc. that can leak sensitive information. Among these, web-based timing attacks have been known to be a non-trivial source of information leaks with no feasible client-side solutions till date [96]. For instance, the time taken for performing sensitive computation can be measured by sending out public requests before and after the computation. Although the two requests do not carry any sensitive information, the time between the two requests can leak information about the data on which the computation was performed [96]. As part of future work, it would be interesting to explore a language-based solution to handling timing leaks, which would mostly be based on balancing the time taken for performing sensitive operations. The challenging part would be to handle loops and the different browser features that result in such leaks.

Appendix

A Proofs for Improved and Generalized Permissive Upgrade

A.1 Proofs for Improved Permissive Upgrade Strategy

Lemma 1 (Expression Evaluation). *If $\langle \sigma_1, e \rangle \Downarrow n_1^{k_1}$ and $\langle \sigma_2, e \rangle \Downarrow n_2^{k_2}$ and $\sigma_1 \sim \sigma_2$, then $n_1^{k_1} \sim n_2^{k_2}$.*

Proof. Induction on the derivation and case analysis on the last expression rule.

1. CONST: $n_1 = n_2 = n$ and $k_1 = k_2 = \perp$.
2. VAR: AS $\sigma_1 \sim \sigma_2, \forall x. \sigma_1(x) = n_1^{k_1} \sim \sigma_2(x) = n_2^{k_2}$.
3. OPER: IH1: If $\langle \sigma_1, e_1 \rangle \Downarrow n_1^{k_1}, \langle \sigma_2, e_1 \rangle \Downarrow n_2^{k_2}, \sigma_1 \sim \sigma_2$, then $n_1^{k_1} \sim n_2^{k_2}$.
 IH2: If $\langle \sigma_1, e_2 \rangle \Downarrow n_1^{k_1}, \langle \sigma_2, e_2 \rangle \Downarrow n_2^{k_2}, \sigma_1 \sim \sigma_2$, then $n_1^{k_1} \sim n_2^{k_2}$.
 T.S. $n_1^{k_1} \sim n_2^{k_2}$, where $n_1 = n_1' \odot n_1'', n_2 = n_2' \odot n_2''$ and $k_1 = k_1' \sqcup k_1'', k_2 = k_2' \sqcup k_2''$.
 As $\sigma_1 \sim \sigma_2$, from IH1 and IH2, $n_1^{k_1} \sim n_2^{k_2}$ and $n_1^{k_1} \sim n_2^{k_2}$.
 Proof by case analysis on low-equivalence definition (Definition 3) for $n_1^{k_1} \sim n_2^{k_2}$ followed by case analysis on low-equivalence definition for $n_1^{k_1} \sim n_2^{k_2}$.

- $n_1' = n_2'$ and $k_1' = k_2' = L$:
 - $n_1'' = n_2''$ and $k_1'' = k_2'' = L$: $n_1 = n_2$ and $k_1 = k_2 = L$
 - $k_1'' = k_2'' = H$: $k_1 = k_2 = H$
 - $k_1'' = P$ or $k_2'' = P$: $k_1 = P$ or $k_2 = P$
- $k_1' = k_2' = H$:
 - $n_1'' = n_2''$ and $k_1'' = k_2'' = L$: $k_1 = k_2 = H$
 - $k_1'' = k_2'' = H$: $k_1 = k_2 = H$
 - $k_1'' = P$ or $k_2'' = P$: $k_1 = H$ and $k_2 = H$
- $k_1' = P$ or $k_2' = P$:
 - $n_1'' = n_2''$ and $k_1'' = k_2'' = L$: $k_1 = P$ or $k_2 = P$
 - $k_1'' = k_2'' = H$: $k_1 = k_2 = H$
 - $k_1'' = P$ or $k_2'' = P$: $k_1 = P$ and/or $k_2 = P$

□

Lemma 2 (Evolution). *If $pc = H$ and $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$, then $\forall x. \Gamma(\sigma(x)) = P \implies \Gamma(\sigma'(x)) = P$.*

Proof. Proof by induction on the derivation rules and case analysis on the last rule.

- SKIP, WHILE-F: $\sigma = \sigma'$
- ASSN-PUS: If $pc = H$ and $l = P$, then $k = P$. All other $\sigma(x)$ remain unchanged.
- SEQ:
 - IH1: If $pc = H$ and $\langle \sigma, c_1 \rangle \Downarrow_{pc} \sigma''$, then $\forall x. \Gamma(\sigma(x)) = P \implies \Gamma(\sigma''(x)) = P$.
 - IH2: If $pc = H$ and $\langle \sigma'', c_2 \rangle \Downarrow_{pc} \sigma'$, then $\forall x. \Gamma(\sigma''(x)) = P \implies \Gamma(\sigma'(x)) = P$.
 - From IH1 and IH2, if $pc = H$ and $\langle \sigma, c_1; c_2 \rangle \Downarrow_{pc} \sigma'$, then $\forall x. \Gamma(\sigma(x)) = P \implies \Gamma(\sigma'(x)) = P$.

- IF-ELSE:

IH: If $pc = H$ and $\langle \sigma, c_i \rangle \Downarrow_{pc \sqcup \ell} \sigma'$, then $\forall x. \Gamma(\sigma(x)) = P \implies \Gamma(\sigma'(x)) = P$.

As $H \sqcup \ell = H$, from IH.

- WHILE-T: Similar to SEQ and IF-ELSE

□

Lemma 3 (Confinement for improved permissive-upgrade with a two-point lattice). *If $pc = H$ and $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$, then $\sigma \sim \sigma'$.*

Proof. Proof by induction on the derivation rules and case analysis on the last step.

- SKIP,WHILE-F: $\sigma = \sigma'$
- ASSN-PUS: If $l = L$, then $k = P$ else if $l = H$, then $k = H$, else if $l = P$, then $k = P$. Thus, $\sigma \sim \sigma'$
- SEQ: IH1: If $pc = H$ and $\langle \sigma, c_1 \rangle \Downarrow_{pc} \sigma''$, then $\sigma \sim \sigma''$ and
IH2: if $pc = H$ and $\langle \sigma'', c_2 \rangle \Downarrow_{pc} \sigma'$, then $\sigma'' \sim \sigma'$.
From Lemma 2, $\forall x. \Gamma(\sigma(x)) = P \implies \Gamma(\sigma''(x)) = P$ and $\forall x. \Gamma(\sigma''(x)) = P \implies \Gamma(\sigma'(x)) = P$.
From definition, $\forall x$ either:
 $\sigma(x) = \sigma''(x)$ and $\Gamma(\sigma(x)) = \Gamma(\sigma''(x)) = L$: From IH2, either $\sigma''(x) = \sigma'(x)$ and $\Gamma(\sigma''(x)) = \Gamma(\sigma'(x)) = L$ or $\Gamma(\sigma'(x)) = P$
or $\Gamma(\sigma(x)) = \Gamma(\sigma''(x)) = H$: From IH2, $\Gamma(\sigma''(x)) = \Gamma(\sigma'(x)) = H$
or either $\Gamma(\sigma(x)) = P$ or $\Gamma(\sigma''(x)) = P$: If $\Gamma(\sigma(x)) = P$, then from Lemma 2, $\Gamma(\sigma''(x)) = P$.
Hence, $\Gamma(\sigma'(x)) = P$.
- IF-ELSE: IH: If $pc = H$ and $\langle \sigma, c_i \rangle \Downarrow_{pc \sqcup \ell} \sigma'$, then $\sigma \sim \sigma'$. If $pc = H$, then $H \sqcup \ell = H$. Thus, from IH.
- WHILE-T: Similar to IF-ELSE and SEQ.

□

Theorem 1 (TINI for improved permissive-upgrade with a two-point lattice). *With the assignment rule ASSN-PUS and the modified syntax of Figure 3.4, if $\sigma_1 \sim \sigma_2$ and $\langle \sigma_1, c \rangle \Downarrow_{pc} \sigma'_1$ and $\langle \sigma_2, c \rangle \Downarrow_{pc} \sigma'_2$, then $\sigma'_1 \sim \sigma'_2$.*

Proof. Proof by induction on the derivation rules and case analysis on the last step.

- SKIP,WHILE-F: $\sigma'_1 = \sigma_1 \sim \sigma_2 = \sigma'_2$
- ASSN-PUS: From Lemma 1, $n_1^{m_1} \sim n_2^{m_2}$. If $pc = L$, then $k = m$. If $pc = H$ and $l = H$, then $k_1 = k_2 = H$. If $pc = H$ and $l = L$, then $k_1 = k_2 = P$. Hence, $\sigma'_1 \sim \sigma'_2$.
- SEQ: IH1: If $\sigma_1 \sim \sigma_2$ and $\langle \sigma_1, c_1 \rangle \Downarrow_{pc} \sigma'_1$, and $\langle \sigma_2, c_1 \rangle \Downarrow_{pc} \sigma'_2$, then $\sigma'_1 \sim \sigma'_2$ and
IH2: If $\sigma'_1 \sim \sigma'_2$ and $\langle \sigma'_1, c_2 \rangle \Downarrow_{pc} \sigma''_1$, and $\langle \sigma'_2, c_2 \rangle \Downarrow_{pc} \sigma''_2$, then $\sigma''_1 \sim \sigma''_2$.
From IH1 and IH2, $\sigma''_1 \sim \sigma''_2$

- **IF-ELSE:** IH: If $\sigma_1 \sim \sigma_2$ and $\langle \sigma_1, c_i \rangle \Downarrow_{pc \sqcup \ell_1} \sigma'_1$, and $\langle \sigma_2, c_j \rangle \Downarrow_{pc \sqcup \ell_2} \sigma'_2$, and $\ell_1 = \ell_2$, and $c_i = c_j$ then $\sigma'_1 \sim \sigma'_2$. From Lemma 1, $\mathbf{n}_1^{k_1} \sim \mathbf{n}_2^{k_2}$. Thus, either $\ell_1 = \ell_2 = L$ or $\ell_1 = \ell_2 = H$. If $\ell_1 = \ell_2 = L$, then $\mathbf{n}_1 = \mathbf{n}_2$. Thus, $c_i = c_j$ and hence, from IH $\sigma'_1 \sim \sigma'_2$.

If $\ell_1 = \ell_2 = H$, then $pc \sqcup H = H$. From Lemma 3, $\sigma_1 \sim \sigma'_1$ and $\sigma_2 \sim \sigma'_2$, and $\sigma_1 \sim \sigma_2$.

T.S. $\sigma'_1 \sim \sigma'_2$, i.e., $\forall \mathbf{x}. \sigma'_1(\mathbf{x}) \sim \sigma'_2(\mathbf{x})$.

Let $\sigma_1(\mathbf{x}) = \mathbf{n}_1^{k_1}$ and $\sigma_2(\mathbf{x}) = \mathbf{n}_2^{k_2}$ and $\sigma'_1(\mathbf{x}) = \mathbf{n}'_1^{k'_1}$ and $\sigma'_2(\mathbf{x}) = \mathbf{n}'_2^{k'_2}$. Case analysis on the definition of equivalence:

- $\mathbf{n}_1 = \mathbf{n}_2$ and $k_1 = k_2 = L$: Either $\mathbf{n}'_1 = \mathbf{n}_1$ and $k'_1 = k_1 = L$ and $\mathbf{n}'_2 = \mathbf{n}_2$ and $k'_2 = k_2 = L$ or $k'_1 = P$ or $k'_2 = P$
- $k_1 = k_2 = H$: $k'_1 = k_1 = H$ and $k'_2 = k_2 = H$
- $k_1 = P$ or $k_2 = P$: From Lemma 2, $k'_1 = P$ or $k'_2 = P$

- **WHILE-T:** Similar to IF-ELSE and SEQ.

□

A.2 Examples for Equivalence Definition

Consider the following notations for the examples:

l, m, h, l^* represent any variable with label L, M, H, L^* , respectively, such that $L \sqsubseteq M \sqsubseteq H$.

An ℓ -level adversary is assumed. ℓ represents the labels that are above the level of the attacker.

Table A.1 shows example programs for the transition from low-equivalent values to low-equivalent values. First column and first row of the table represents all the possible ways in which two values can be low-equivalent (from definition 7).

A.3 Proofs and Results for Generalized Permissive Upgrade for Arbitrary Lattices

Lemma 4. Expression Evaluation Lemma

If $\sigma_1 \sim_\ell \sigma_2$,
 $\langle \sigma_1, e \rangle \Downarrow \mathbf{n}_1^{k_1}$,
 $\langle \sigma_2, e \rangle \Downarrow \mathbf{n}_2^{k_2}$,
 then $\mathbf{n}_1^{k_1} \sim_\ell \mathbf{n}_2^{k_2}$.

Proof. Proof by induction on the derivation and case analysis on the last expression rule.

1. **CONST:** $\mathbf{n}_1 = \mathbf{n}_2 = \mathbf{n}$ and $k_1 = k_2 = \perp$.
2. **VAR:** AS $\sigma_1 \sim_\ell \sigma_2$, $\forall \mathbf{x}. \sigma_1(\mathbf{x}) = \mathbf{n}_1^{k_1} \sim_\ell \sigma_2(\mathbf{x}) = \mathbf{n}_2^{k_2}$.
3. **OPER:** IH1: If $\langle \sigma_1, e_1 \rangle \Downarrow \mathbf{n}_1^{k'_1}$, $\langle \sigma_2, e_1 \rangle \Downarrow \mathbf{n}_2^{k'_2}$, $\sigma_1 \sim_\ell \sigma_2$, then $\mathbf{n}_1^{k'_1} \sim_\ell \mathbf{n}_2^{k'_2}$.
 IH2: If $\langle \sigma_1, e_2 \rangle \Downarrow \mathbf{n}_1^{k''_1}$, $\langle \sigma_2, e_2 \rangle \Downarrow \mathbf{n}_2^{k''_2}$, $\sigma_1 \sim_\ell \sigma_2$, then $\mathbf{n}_1^{k''_1} \sim_\ell \mathbf{n}_2^{k''_2}$.
 T.S. $\mathbf{n}_1^{k_1} \sim_\ell \mathbf{n}_2^{k_2}$, where $\mathbf{n}_1 = \mathbf{n}'_1 \odot \mathbf{n}''_1$, $\mathbf{n}_2 = \mathbf{n}'_2 \odot \mathbf{n}''_2$ and $k_1 = k'_1 \sqcup k''_1$, $k_2 = k'_2 \sqcup k''_2$.

	l, l	l_1^*, l_2	l_1, l_2^*	l_1^*, l_2^*	l_1, l_2	l_1^*, l_2	l_1, l_2^*
l, l	-	if(h) x1 = l	if(h) x1 = l	if(h) x1 = l else x1 = l	x1 = h	x1 = m if(h) x1 = 4 if(m) x1 = l*	x1 = m if(h) x1 = 4 if(m) x1 = l*
l_1^*, l_2	x1 = l	-	x1 = l if(h) x1 = l	if(h) x1 = l	x1 = h	x1 = m if (h) x1 = l if(m) x1 = l*	x1 = m if (h) x1 = l if(m) x1 = l*
l_1, l_2^*	x1 = l	x1 = l if(h) x1 = l	-	if(h) x1 = l	x1 = h	x1 = m if (h) x1 = l if(m) x1 = l*	x1 = m if (h) x1 = l if(m) x1 = l*
l_1^*, l_2^*	x1 = l	x1 = l if (h) x1 = l	x1 = l if (h) x1 = l	-	x1 = h	x1 = m if (h) x1 = l if (m) x1 = l*	x1 = m if (h) x1 = l if (m) x1 = l*
l_1, l_2	x1 = l	x1 = l if (h) x1 = l	x1 = l if (h) x1 = l	if (h) x1 = l else x1 = l	-	x1 = m if (h) x1 = l if(m) x1 = l*	x1 = m if (h) x1 = l if (m) x1 = l*
l_1^*, l_2	x1 = l	x1 = l if (h) x1 = l	x1 = l if (h) x1 = l	x1 = l if (h) x1 = l else x1 = l	x1 = h	-	x1 = m if (h) x1 = l if (m) x1 = l*
l_1, l_2^*	x1 = l	x1 = l if (h) x1 = l	x1 = l if (h) x1 = l	x1 = l if (h) x1 = l else x1 = l	x1 = h	x1 = m if (h) x1 = l if(m) x1 = l*	-

Table A.1: Examples for all possible transitions of low-equivalent to low-equivalent values

As $\sigma_1 \sim_\ell \sigma_2$, from IH1 and IH2, $n_1^{k'_1} \sim_\ell n_2^{k'_2}$ and $n_1^{k''_1} \sim_\ell n_2^{k''_2}$.

Proof by case analysis on low-equivalence definition for $n_1^{k'_1} \sim_\ell n_2^{k'_2}$ followed by case analysis on low-equivalence definition for $n_1^{k''_1} \sim_\ell n_2^{k''_2}$.

□

Lemma 5. \star -preservation Lemma

$\forall x.$ If $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$, $\Gamma(\sigma(x)) = \ell^*$ and $pc \not\sqsubseteq \ell$, then $\Gamma(\sigma'(x)) = \ell'^* \wedge \ell' \sqsubseteq \ell$

Proof. Proof by induction on the derivation and case analysis on the last rule.

1. SKIP : $\sigma = \sigma'$.
2. ASSN-N: As $pc \not\sqsubseteq \ell$, these cases do not apply.
3. ASSN-S: From the premises, for x in statement c , $\Gamma(\sigma'(x)) = ((pc \sqcup m) \sqcap \ell)^* = \ell'$. Thus, $\ell' \sqsubseteq \ell$.
For any other y , $\sigma(y) = \sigma'(y)$. Thus, $\ell' = \ell$.
4. SEQ : IH1 : $\forall x.$ If $\langle \sigma, c \rangle \Downarrow_{pc} \sigma''$, $\Gamma(\sigma(x)) = \ell^*$ and $pc \not\sqsubseteq \ell$, then $\Gamma(\sigma''(x)) = \ell''^* \wedge \ell'' \sqsubseteq \ell$
IH2 : $\forall x.$ If $\langle \sigma', c \rangle \Downarrow_{pc} \sigma'$, $\Gamma(\sigma''(x)) = \ell'^*$ and $pc \not\sqsubseteq \ell'$, then $\Gamma(\sigma'(x)) = \ell'^* \wedge \ell' \sqsubseteq \ell''$
Thus, from IH1 and IH2, $\Gamma(\sigma'(x)) = \ell'^* \wedge \ell' \sqsubseteq \ell$.
5. IF-ELSE: Let $k = \ell''$.
IH: $\forall x.$ If $\langle \sigma, c \rangle \Downarrow_{pc \sqcup \ell''} \sigma'$, $\Gamma(\sigma(x)) = \ell^*$ and $pc \sqcup \ell'' \not\sqsubseteq \ell$, then $\Gamma(\sigma'(x)) = \ell'^* \wedge \ell' \sqsubseteq \ell$
As $pc \not\sqsubseteq \ell$, so $pc \sqcup \ell'' \not\sqsubseteq \ell$.
Thus from IH, $\Gamma(\sigma'(x)) = \ell'^* \wedge \ell' \sqsubseteq \ell$
6. WHILE-T: Let $k = \ell_e$.
IH1: $\forall x.$ If $\langle \sigma, c \rangle \Downarrow_{pc \sqcup \ell_e} \sigma''$, $\Gamma(\sigma(x)) = \ell^*$ and $pc \sqcup \ell_e \not\sqsubseteq \ell$, then $\Gamma(\sigma''(x)) = \ell''^* \wedge \ell'' \sqsubseteq \ell$
IH2: $\forall x.$ If $\langle \sigma', c \rangle \Downarrow_{pc \sqcup \ell_e} \sigma'$, $\Gamma(\sigma''(x)) = \ell'^*$ and $pc \sqcup \ell_e \not\sqsubseteq \ell$, then $\Gamma(\sigma'(x)) = \ell'^* \wedge \ell' \sqsubseteq \ell$
As $pc \not\sqsubseteq \ell$, so $pc \sqcup \ell_e \not\sqsubseteq \ell$.
Thus from IH1 and IH2, $\Gamma(\sigma'(x)) = \ell'^* \wedge \ell' \sqsubseteq \ell$
7. WHILE-F : $\sigma = \sigma'$.

□

Corollary 1. If $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$ and $\Gamma(\sigma(x)) = \ell^*$ and $\Gamma(\sigma'(x)) = \ell'$, then $pc \sqsubseteq \ell$.

Proof. Immediate from Lemma 5.

□

Lemma 6. pc Lemma

If $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$, then $\forall x. \Gamma(\sigma'(x)) = \ell \implies (\sigma(x) = \sigma'(x)) \vee pc \sqsubseteq \ell$.

Proof. Proof by induction on the derivation and case analysis on the last rule.

- SKIP: $\sigma(x) = \sigma'(x)$.

- **ASSN-N**: For \mathbf{x} in the statement c , by premises, $\ell = pc \sqcup \ell_e$. Thus, $pc \sqsubseteq \ell$.
For any other y s.t. $\Gamma(\sigma'(y)) = \ell'$, $\sigma(y) = \sigma'(y)$. For **ASSN-S**, case does not apply.
- **SEQ**: **IH1**: If $\langle \sigma, c_1 \rangle \Downarrow_{pc} \sigma''$, then $\forall \mathbf{x}. \Gamma(\sigma''(\mathbf{x})) = \ell'' \implies (\sigma(\mathbf{x}) = \sigma''(\mathbf{x})) \vee pc \sqsubseteq \ell''$.
IH2: If $\langle \sigma', c_2 \rangle \Downarrow_{pc} \sigma'$, then $\forall \mathbf{x}. \Gamma(\sigma'(\mathbf{x})) = \ell \implies (\sigma''(\mathbf{x}) = \sigma'(\mathbf{x})) \vee pc \sqsubseteq \ell$.
From **IH2**, if $\sigma''(\mathbf{x}) \neq \sigma'(\mathbf{x})$, then $pc \sqsubseteq \ell$.
If $\sigma''(\mathbf{x}) = \sigma'(\mathbf{x})$, then from **IH1**:
 - If $\sigma(\mathbf{x}) = \sigma''(\mathbf{x})$: $\sigma(\mathbf{x}) = \sigma'(\mathbf{x})$.
 - If $\sigma(\mathbf{x}) \neq \sigma''(\mathbf{x})$: $pc \sqsubseteq \ell''$, where $\ell'' = \Gamma(\sigma''(\mathbf{x}))$. As $\sigma''(\mathbf{x}) = \sigma'(\mathbf{x})$, $\ell'' = \Gamma(\sigma'(\mathbf{x})) = \ell$.
Thus, $pc \sqsubseteq \ell$.
- **IF-ELSE**: **IH**: If $\langle \sigma, c \rangle \Downarrow_{pc \sqcup \ell_e} \sigma'$, then $\forall \mathbf{x}. \Gamma(\sigma'(\mathbf{x})) = \ell \implies (\sigma(\mathbf{x}) = \sigma'(\mathbf{x})) \vee pc \sqcup \ell_e \sqsubseteq \ell$.
From **IH**, either $(\sigma(\mathbf{x}) = \sigma'(\mathbf{x}))$ or $pc \sqcup \ell_e \sqsubseteq \ell$. Thus, $(\sigma(\mathbf{x}) = \sigma'(\mathbf{x})) \vee pc \sqsubseteq \ell$.
- **WHILE-T**: **IH1**: If $\langle \sigma, c \rangle \Downarrow_{pc \sqcup \ell_e} \sigma''$, then $\forall \mathbf{x}. \Gamma(\sigma''(\mathbf{x})) = \ell'' \implies (\sigma(\mathbf{x}) = \sigma''(\mathbf{x})) \vee pc \sqsubseteq \ell''$.
IH2: If $\langle \sigma', c \rangle \Downarrow_{pc \sqcup \ell_e} \sigma'$, then $\forall \mathbf{x}. \Gamma(\sigma'(\mathbf{x})) = \ell \implies (\sigma''(\mathbf{x}) = \sigma'(\mathbf{x})) \vee pc \sqsubseteq \ell$.
From similar reasoning as in “**SEQ**”, either $\sigma(\mathbf{x}) = \sigma'(\mathbf{x})$ or $pc \sqcup \ell_e \sqsubseteq \ell$. Thus, $\sigma(\mathbf{x}) = \sigma'(\mathbf{x}) \vee pc \sqsubseteq \ell$.
- **WHILE-F**: $\sigma(\mathbf{x}) = \sigma'(\mathbf{x})$.

□

Corollary 2. If $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$ and $\Gamma(\sigma(\mathbf{x})) = \ell^*$ and $\Gamma(\sigma'(\mathbf{x})) = \ell'$, then $pc \sqsubseteq \ell'$.

Proof. Immediate from Lemma 6. □

Lemma 7. Confinement Lemma If $pc \not\sqsubseteq \ell$, $\langle \sigma, c \rangle \Downarrow_{pc} \sigma'$, then $\sigma \sim_\ell \sigma'$.

Proof. Proof by induction on the derivation and case analysis on the last rule.

1. **SKIP** : $\sigma = \sigma'$.
2. **ASSN-N**: Let $\mathbf{x}_i = \mathbf{v}_i^{k_i}$ and $\mathbf{x}_f = \mathbf{v}_f^{k_f}$, s.t. $k_i = \ell_i \vee k_i = \ell_i^*$ and $pc \sqsubseteq \ell_i$: As $pc \not\sqsubseteq \ell$, $\ell_i \not\sqsubseteq \ell$. By premises of **ASSN-N**, $k_f = \ell_f \vee k_f = \ell_f^*$, where $\ell_f = pc \sqcup \ell_e$. As $pc \not\sqsubseteq \ell$, $\ell_f \not\sqsubseteq \ell$. Thus, by definition 7.2, 7.3, 7.4 or 7.5, $\mathbf{x}_i \sim_\ell \mathbf{x}_f$.
3. **ASSN-S**: Let $\mathbf{x}_i = \mathbf{v}_i^{k_i}$ and $\mathbf{x}_f = \mathbf{v}_f^{k_f}$, s.t. $k_i = \ell_i \vee k_i = \ell_i^*$ and $pc \not\sqsubseteq \ell_i$: By premise, $k_f = ((pc \sqcup m) \sqcap \ell_i)^*$. Thus, $\ell_f \sqsubseteq \ell_i$ and by definition 7.3 or 7.5 $\mathbf{x}_i \sim_\ell \mathbf{x}_f$.
4. **SEQ** : **IH1**: $\sigma \sim_\ell \sigma''$ and **IH2**: $\sigma'' \sim_\ell \sigma'$. **T.S** : $\sigma \sim_\ell \sigma'$.
For all $\mathbf{x} \in \text{dom}(\sigma)$, respective $\mathbf{x}'' \in \text{dom}(\sigma'')$ and respective $\mathbf{x}' \in \text{dom}(\sigma')$, $\mathbf{x} \sim_\ell \mathbf{x}''$ and $\mathbf{x}'' \sim_\ell \mathbf{x}'$.
To show: $\mathbf{x} \sim_\ell \mathbf{x}'$.
Let $\mathbf{x} = \mathbf{v}_1^{k_1}$, $\mathbf{x}'' = \mathbf{v}_2^{k_2}$, $\mathbf{x}' = \mathbf{v}_3^{k_3}$, where $k_1 = \ell_1 \vee k_1 = \ell_1^*$, $k_2 = \ell_2 \vee k_2 = \ell_2^*$ and $k_3 = \ell_3 \vee k_3 = \ell_3^*$.

Case-analysis on definition 7 for **IH1**.

- $(k_1 = k_2) = \ell' \sqsubseteq \ell \wedge \mathbf{v}_1 = \mathbf{v}_2$: By **IH2** and definition 7,

-
- (a) $(k_2 = k_3) = \ell' \sqsubseteq \ell \wedge v_2 = v_3$ (case 1): Transitivity of equality, $(k_1 = k_3) = \ell' \sqsubseteq \ell \wedge v_1 = v_3$. Thus, $x \sim_\ell x'$.
 - (b) $k_2 = \ell'$ and $k_3 = \ell_3^* \wedge \ell_3 \sqsubseteq \ell' \sqsubseteq \ell$ (case 5): By definition 7.5 $x \sim_\ell x'$.
 - $k_1 = \ell_1 \not\sqsubseteq \ell \wedge k_2 = \ell_2 \not\sqsubseteq \ell$: By IH2, either
 - (a) $k_2 = \ell_2 \not\sqsubseteq \ell \wedge k_3 = \ell_3 \not\sqsubseteq \ell$. By definition 7.2, $x \sim_\ell x'$.
 - (b) $k_2 = \ell_2 \not\sqsubseteq \ell \wedge k_3 = \ell_3^* : \ell_1 \not\sqsubseteq \ell$. Thus, by definition 7.5, $x \sim_\ell x'$.
 - $k_1 = \ell_1^* \wedge k_2 = \ell_2^*$: By IH2,
 - (a) $k_2 = \ell_2^* \wedge k_3 = \ell_3^*$ (case 3): By definition 7.3, $x \sim_\ell x'$.
 - (b) $k_2 = \ell_2^* \wedge k_3 = \ell_3 \wedge (\ell_3 \not\sqsubseteq \ell)$ (case 4): By definition 7.4, $x \sim_\ell x'$.
 - (c) $k_2 = \ell_2^* \wedge k_3 = \ell_3 \wedge (\ell_2 \sqsubseteq \ell_3)$ (case 4): By corollary 1, $pc \sqsubseteq \ell_2$. As $pc \not\sqsubseteq \ell$ and $\ell_2 \sqsubseteq \ell_3$, so $\ell_3 \not\sqsubseteq \ell$. By definition 7.4, $x \sim_\ell x'$.
 - $k_1 = \ell_1^* \wedge k_2 = \ell_2$ s.t. $(\ell_2 \not\sqsubseteq \ell)$ (case 4): Either
 - $k_2 = \ell_2 \not\sqsubseteq \ell \wedge k_3 = \ell_3 \not\sqsubseteq \ell$: By definition 7.4, $x \sim_\ell x'$.
 - $k_2 = \ell_2 \not\sqsubseteq \ell \wedge k_3 = \ell_3^*$: By definition 7.3, $x \sim_\ell x'$.
 - $k_1 = \ell_1^* \wedge k_2 = \ell_2$ s.t. $(\ell_1 \sqsubseteq \ell_2)$ (case 4):
 - $k_2 = k_3 = \ell_2$: By definition 7.4, $x \sim_\ell x'$.
 - $k_2 = \ell_2 \not\sqsubseteq \ell \wedge k_3 = \ell_3 \not\sqsubseteq \ell$: By definition 7.4, $x \sim_\ell x'$.
 - $k_2 = \ell_2 \not\sqsubseteq \ell \wedge k_3 = \ell_3^*$: By definition 7.3, $x \sim_\ell x'$.
 - $k_1 = \ell_1 \wedge k_2 = \ell_2^*$ s.t. $(\ell_1 \not\sqsubseteq \ell)$: By IH2,
 - (a) $k_2 = \ell_2^* \wedge k_3 = \ell_3^*$ (case 3): By definition 7.5, $x \sim_\ell x'$.
 - (b) $k_2 = \ell_2^* \wedge k_3 = \ell_3$ s.t. $(\ell_3 \not\sqsubseteq \ell)$ (case 4): By definition 7.2, $x \sim_\ell x'$.
 - (c) $k_2 = \ell_2^* \wedge k_3 = \ell_3$ s.t. $(\ell_2 \sqsubseteq \ell_3)$ (case 4): By corollary 1, $pc \sqsubseteq \ell_2$. As $pc \not\sqsubseteq \ell$ and $\ell_2 \sqsubseteq \ell_3$, so $\ell_3 \not\sqsubseteq \ell$. By definition 7.2, $x \sim_\ell x'$.
 - $k_1 = \ell_1 \wedge k_2 = \ell_2^*$ s.t. $(\ell_2 \sqsubseteq \ell_1)$: Also, $(\ell_2 \sqsubseteq \ell_1 \sqsubseteq \ell)$. By IH2,
 - (a) $k_2 = \ell_2^* \wedge k_3 = \ell_3^*$ (case 3): As $\ell_2 \sqsubseteq \ell$ and $pc \not\sqsubseteq \ell$, $pc \not\sqsubseteq \ell_2$. By lemma 5, $\ell_3 \sqsubseteq \ell_2$. Thus, $\ell_3 \sqsubseteq \ell_2 \sqsubseteq \ell_1$. By definition 7.5, $x \sim_\ell x'$.
 - (b) $k_2 = \ell_2^* \wedge k_3 = \ell_3$ (case 4): As $\ell_2 \sqsubseteq \ell$ and $pc \not\sqsubseteq \ell$, $pc \not\sqsubseteq \ell_2$. But, by corollary 1, $pc \sqsubseteq \ell_2$. By contradiction, this case does not hold.
5. IF-ELSE : IH : $k = \ell'$. If $(pc \sqcup \ell') \not\sqsubseteq \ell$, then $\sigma \sim_\ell \sigma'$.
As $pc \not\sqsubseteq \ell$, $pc \sqcup \ell' \not\sqsubseteq \ell$. Thus, by IH, $\sigma \sim_\ell \sigma'$.
6. WHILE-T: IH1 : $k = \ell'$. If $(pc \sqcup \ell') \not\sqsubseteq \ell$, then $\sigma \sim_\ell \sigma'$.
As $pc \not\sqsubseteq \ell$, $pc \sqcup \ell' \not\sqsubseteq \ell$. Thus, by IH1, $\sigma \sim_\ell \sigma''$.
IH2 : $k = \ell'$. If $(pc \sqcup \ell') \not\sqsubseteq \ell$, then $\sigma' \sim_\ell \sigma''$.
As $pc \not\sqsubseteq \ell$, $pc \sqcup \ell' \not\sqsubseteq \ell$. Thus, by IH, $\sigma'' \sim_\ell \sigma'$.
Therefore, $\sigma \sim_\ell \sigma''$ and $\sigma'' \sim_\ell \sigma'$.
(Reasoning similar to SEQ.)
7. WHILE-F : $\sigma = \sigma'$

□

Theorem 2. Termination-insensitive non-interference

If $\sigma_1 \sim_\ell \sigma_2$, $\langle \sigma_1, c \rangle \Downarrow_{pc} \sigma'_1$, $\langle \sigma_2, c \rangle \Downarrow_{pc} \sigma'_2$, then $\sigma'_1 \sim_\ell \sigma'_2$.

Proof. By induction on the derivation and case analysis on the last step

1. **SKIP:** $\sigma'_1 = \sigma_1 \sim_\ell \sigma_2 = \sigma'_2$
2. **ASSN-N and ASSN-S:** As $\sigma_1 \sim_\ell \sigma_2$, $\forall x. \sigma_1(x) \sim_\ell \sigma_2(x)$. Let $\sigma_1(x) = v_1^{k_1}$, $\sigma_2(x) = v_2^{k_2}$ and $\sigma'_1(x) = v_1^{k'_1}$, $\sigma'_2(x) = v_2^{k'_2}$
 - s. t. $k_i = \ell_i \vee k_i = \ell_i^*$ and $k'_i = \ell'_i \vee k'_i = \ell_i^*$ for $i = 1, 2$.

Let $\langle e_1, \sigma_1 \rangle \Downarrow w_1^{k_1^e} \wedge \langle e_2, \sigma_2 \rangle \Downarrow w_2^{k_2^e}$

 - s. t. $k_i^e = \ell_i^e \vee k_i^e = \ell_i^{e*}$ for $i = 1, 2$. For low-equivalence of e_1 and e_2 , the following cases arise:
 - (a) $k_i^e = \ell_i^e$, s.t. $(\ell_1^e = \ell_2^e) = \ell^e \sqsubseteq \ell \wedge w_1 = w_2$:
 - i. $pc \not\sqsubseteq \ell_1 \wedge pc \not\sqsubseteq \ell_2$: By premise of ASSN-S rules, $k'_i = ((pc \sqcup \ell^e) \sqcap \ell_i)^*$. By definition 7.3, $\sigma'_1 \sim_\ell \sigma'_2$.
 - ii. $pc \not\sqsubseteq \ell_1 \wedge pc \sqsubseteq \ell_2$: $k'_1 = ((pc \sqcup \ell^e) \sqcap \ell_1)^*$ and $k'_2 = pc \sqcup \ell^e$. As $\ell_1^e \sqsubseteq \ell_2^e$, by definition 7.4, $\sigma'_1 \sim_\ell \sigma'_2$.
 - iii. $pc \sqsubseteq \ell_1 \wedge pc \not\sqsubseteq \ell_2$: $k'_2 = ((pc \sqcup \ell^e) \sqcap \ell_2)^*$ and $k'_1 = pc \sqcup \ell^e$. As $\ell_2^e \sqsubseteq \ell_1^e$, by definition 7.5, $\sigma'_1 \sim_\ell \sigma'_2$.
 - iv. $pc \sqsubseteq \ell_1 \wedge pc \sqsubseteq \ell_2$: $k'_1 = pc \sqcup \ell^e$ and $k'_2 = pc \sqcup \ell^e$. If $pc \sqsubseteq \ell$ and $\ell^e \sqsubseteq \ell$ and $w_1 = w_2$, by definition 7.1, $\sigma'_1 \sim_\ell \sigma'_2$. If $pc \not\sqsubseteq \ell$, $pc \sqcup \ell^e \not\sqsubseteq \ell$. By definition 7.2, $\sigma'_1 \sim_\ell \sigma'_2$.
 - (b) $\ell_1^e \not\sqsubseteq \ell \wedge \ell_2^e \not\sqsubseteq \ell$: From premise of assignment rules, $k'_1 = pc \sqcup \ell_1^e \vee k'_1 = (pc \sqcup \ell_1^e)^* \vee k'_1 = ((pc \sqcup \ell_1^e) \sqcap \ell_1)^*$. Similarly, $k'_2 = pc \sqcup \ell_2^e \vee k'_2 = (pc \sqcup \ell_2^e)^* \vee k'_2 = ((pc \sqcup \ell_2^e) \sqcap \ell_2)^*$. Since $\ell_1^e \not\sqsubseteq \ell$ and $\ell_2^e \not\sqsubseteq \ell$, $pc \sqcup \ell_1^e \not\sqsubseteq \ell$ and $pc \sqcup \ell_2^e \not\sqsubseteq \ell$. Therefore, from Definition 7.2, 7.3, 7.4 or 7.5 $\sigma'_1 \sim_\ell \sigma'_2$.
 - (c) $k_i^e = \ell_i^{e*}$: By premise of ASSN-S rules, $k'_i = ((pc \sqcup \ell_i^e) \sqcap \ell_i)^*$ or $k'_i = (pc \sqcup \ell_i^e)^*$. By definition 7.3, $\sigma'_1 \sim_\ell \sigma'_2$.
 - (d) $k_1^e = \ell_1^{e*} \wedge k_2^e = \ell_2^e$:
 - i. $pc \not\sqsubseteq \ell_1 \wedge pc \not\sqsubseteq \ell_2$: By premise of ASSN-S rules, $k'_i = ((pc \sqcup \ell_i^e) \sqcap \ell_i)^*$. By definition 7.3, $\sigma'_1 \sim_\ell \sigma'_2$.
 - ii. $pc \not\sqsubseteq \ell_1 \wedge pc \sqsubseteq \ell_2$: $k'_1 = ((pc \sqcup \ell_1^e) \sqcap \ell_1)^*$ and $k'_2 = pc \sqcup \ell_2^e$. From definition 7.4, $\ell_1^e \sqsubseteq \ell_2^e$, so $(pc \sqcup \ell_1^e) \sqcap \ell_1 \sqsubseteq pc \sqcup \ell_2^e$. By definition 7.4, $\sigma'_1 \sim_\ell \sigma'_2$.
 - iii. $pc \sqsubseteq \ell_1 \wedge pc \not\sqsubseteq \ell_2$: $k'_2 = ((pc \sqcup \ell_2^e) \sqcap \ell_2)^*$ and $k'_1 = (pc \sqcup \ell_1^e)^*$. By definition 7.3, $\sigma'_1 \sim_\ell \sigma'_2$.
 - iv. $pc \sqsubseteq \ell_1 \wedge pc \sqsubseteq \ell_2$: $k'_1 = (pc \sqcup \ell_1^e)^*$ and $k'_2 = pc \sqcup \ell_2^e$. If $\ell_2^e \not\sqsubseteq \ell$, so $pc \sqcup \ell_2^e \not\sqsubseteq \ell$. Else if $\ell_1^e \sqsubseteq \ell_2^e$, then $pc \sqcup \ell_1^e \sqsubseteq pc \sqcup \ell_2^e$. By definition 7.4, $\sigma'_1 \sim_\ell \sigma'_2$.
 - (e) $k_1^e = \ell_1^e \wedge k_2^e = \ell_2^{e*}$:
 - i. $pc \not\sqsubseteq \ell_1 \wedge pc \not\sqsubseteq \ell_2$: By premise of ASSN-S rules, $k'_i = ((pc \sqcup \ell_i^e) \sqcap \ell_i)^*$. By definition 7.3, $\sigma'_1 \sim_\ell \sigma'_2$.
 - ii. $pc \not\sqsubseteq \ell_1 \wedge pc \sqsubseteq \ell_2$: $k'_1 = ((pc \sqcup \ell_1^e) \sqcap \ell_1)^*$ and $k'_2 = (pc \sqcup \ell_2^e)^*$. By definition 7.3, $\sigma'_1 \sim_\ell \sigma'_2$.

- iii. $pc \sqsubseteq \ell_1 \wedge pc \not\sqsubseteq \ell_2$: $k'_1 = pc \sqcup \ell_1^e$ and $k'_2 = ((pc \sqcup \ell_2^e) \sqcap \ell_2)^*$. $(pc \sqcup \ell_2^e) \sqcap \ell_2 \sqsubseteq pc \sqcup \ell_1^e$. By definition 7.5, $\sigma'_1 \sim_\ell \sigma'_2$.
- iv. $pc \sqsubseteq \ell_1 \wedge pc \sqsubseteq \ell_2$: $k'_1 = (pc \sqcup \ell_1^e)^*$ and $k'_2 = pc \sqcup \ell_2^e$. If $\ell_1^e \not\sqsubseteq \ell$, so $pc \sqcup \ell_1^e \not\sqsubseteq \ell$. Else if $\ell_2^e \sqsubseteq \ell_1^e$, then $pc \sqcup \ell_2^e \sqsubseteq pc \sqcup \ell_1^e$. By definition 7.5, $\sigma'_1 \sim_\ell \sigma'_2$.
3. SEQ: IH1: If $\sigma_1 \sim_\ell \sigma_2$ then $\sigma'_1 \sim_\ell \sigma'_2$
 IH2: If $\sigma'_1 \sim_\ell \sigma'_2$ then $\sigma_1 \sim_\ell \sigma_2$
 Since $\sigma_1 \sim_\ell \sigma_2$, therefore, from IH1 and IH2 $\sigma'_1 \sim_\ell \sigma'_2$.
4. IF-ELSE: IH: If $\sigma_1 \sim_\ell \sigma_2$, $\langle \sigma_1, c \rangle \Downarrow_{pc \sqcup \ell_1^e} \sigma'_1$, $\langle \sigma_2, c \rangle \Downarrow_{pc \sqcup \ell_2^e} \sigma'_2$ and $pc \sqcup \ell_1^e = pc \sqcup \ell_2^e$ then $\sigma'_1 \sim_\ell \sigma'_2$.
- If $\ell_1^e \sqsubseteq \ell$, $\ell_1^e = \ell_2^e$ and $n_1 = n_2$. By IH, $\sigma'_1 \sim_\ell \sigma'_2$.
 - If $\ell_1^e \not\sqsubseteq \ell$, then $\ell_2^e \not\sqsubseteq \ell$, $pc \sqcup \ell_i^e \not\sqsubseteq \ell$ for $i = 1, 2$. By Lemma 7, $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$. T.S. $\sigma'_1 \sim_\ell \sigma'_2$, i.e., $(\forall x. \sigma'_1(x) \sim_\ell \sigma'_2(x))$
- Case analysis on the definition of low-equivalence of values, x , in σ_1 and σ_2 . Let $\sigma_1(x) = v_1^{k_1}$ and $\sigma_2(x) = v_2^{k_2}$ and $\sigma'_1(x) = v_1^{k'_1}$ and $\sigma'_2(x) = v_2^{k'_2}$
- (a) $(k_1 = k_2) = \ell' \sqsubseteq \ell \wedge v_1 = v_2 = v$:
- If $k'_1 = \ell'_1 \wedge k'_2 = \ell'_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by definition 7.1, $\ell' = \ell'_1 \wedge v = v'_1$ and $\ell' = \ell'_2 \wedge v = v'_2$. Thus, $\ell'_1 = \ell'_2 \wedge v'_1 = v'_2$, so $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
 - If $k'_1 = \ell'_1^* \wedge k'_2 = \ell'_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by definition 7.5 $\ell'_1 \sqsubseteq \ell_1 = \ell'$ and by definition 7.1 $k'_2 = \ell'_2 = \ell_2 = \ell'$. So, $\ell'_1 \sqsubseteq \ell'_2$. By definition 7.4, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
 - If $k'_1 = \ell'_1 \wedge k'_2 = \ell'_2^*$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by definition 7.1 $k'_1 = \ell'_1 = \ell_1 = \ell'$ and by definition 7.5 $\ell'_2 \sqsubseteq \ell_2 = \ell'$. So, $\ell'_2 \sqsubseteq \ell'_1$. By definition 7.5, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
 - If $k'_1 = \ell'_1^* \wedge k'_2 = \ell'_2^*$, then by definition 7.3, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
- (b) $(k_1 = \ell_1 \not\sqsubseteq \ell) \wedge (k_2 = \ell_2 \not\sqsubseteq \ell)$:
- If $k'_1 = \ell'_1 \wedge k'_2 = \ell'_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by definition 7.2, $(k'_1 = \ell'_1 \not\sqsubseteq \ell) \wedge (k'_2 = \ell'_2 \not\sqsubseteq \ell)$. So, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
 - If $k'_1 = \ell'_1^* \wedge k'_2 = \ell'_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by definition 7.2 $k'_2 = \ell'_2 \not\sqsubseteq \ell$. By definition 7.4, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
 - If $k'_1 = \ell'_1 \wedge k'_2 = \ell'_2^*$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by definition 7.2 $k'_1 = \ell'_1 \not\sqsubseteq \ell$. By definition 7.5, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$. If $k'_1 = \ell'_1^* \wedge k'_2 = \ell'_2^*$, then by definition 7.3, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
- (c) $(k_1 = \ell_1^* \wedge k_2 = \ell_2^*)$:
- If $k'_1 = \ell'_1^* \wedge k'_2 = \ell'_2^*$, by definition 7.3, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
 - If $k'_1 = \ell'_1 \wedge k'_2 = \ell'_2^*$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by corollary 2, $pc \sqcup \ell_1^e \sqsubseteq \ell'_1$. As $pc \sqcup \ell_1^e \not\sqsubseteq \ell$ and by definition 7.2, $\ell'_1 \not\sqsubseteq \ell$. By definition 7.5, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
 - If $k'_1 = \ell'_1^* \wedge k'_2 = \ell'_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by corollary 2, $pc \sqcup \ell_2^e \sqsubseteq \ell'_2$. As $pc \sqcup \ell_2^e \not\sqsubseteq \ell$ and by definition 7.2, $\ell'_2 \not\sqsubseteq \ell$. By definition 7.4, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
 - If $k'_1 = \ell'_1 \wedge k'_2 = \ell'_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by corollary 2, $pc \sqcup \ell_1^e \sqsubseteq \ell'_1$ and $pc \sqcup \ell_2^e \sqsubseteq \ell'_2$. As $pc \sqcup \ell_i^e \not\sqsubseteq \ell$ and by definition 7.2, $\ell'_1 \not\sqsubseteq \ell$ and $\ell'_2 \not\sqsubseteq \ell$. By definition 7.2, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.

(d) ($k_1 = \ell_1^* \wedge k_2 = \ell_2$):

– $\ell_2 \not\sqsubseteq \ell$:

- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2^*$, by definition 7.3, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2^*$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by corollary 2, $pc \sqcup \ell_1^e \sqsubseteq \ell_1^*$. As $pc \sqcup \ell_1^e \not\sqsubseteq \ell$ and by definition 7.2, $\ell_1^* \not\sqsubseteq \ell$. By definition 7.5, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by definition 7.2, $\ell_2 \not\sqsubseteq \ell$. By definition 7.4, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by corollary 2, $pc \sqcup \ell_1^e \sqsubseteq \ell_1^*$. As $pc \sqcup \ell_1^e \not\sqsubseteq \ell$ and by definition 7.2, $\ell_1^* \not\sqsubseteq \ell$. By definition 7.2, $\ell_2 \not\sqsubseteq \ell$. By definition 7.2, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.

– $\ell_1 \sqsubseteq \ell_2 \sqsubseteq \ell$:

- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2^*$, by definition 7.3, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2^*$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by corollary 2, $pc \sqcup \ell_1^e \sqsubseteq \ell_1^*$. As $pc \sqcup \ell_1^e \not\sqsubseteq \ell$, and by definition 7.2, $\ell_1^* \not\sqsubseteq \ell$. By definition 7.5, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, $\ell_1^* \sqsubseteq (pc \sqcup \ell_1^e) \sqcap \ell_1$ as $pc \sqcup \ell_1^e \not\sqsubseteq \ell_1$ and $\ell_2 = \ell_2$ by corollary 1 and definition 7.1. Thus, $\ell_1^* \sqsubseteq \ell_2$. By definition 7.4, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by corollary 1, $pc \sqcup \ell_1^e \sqsubseteq \ell_1$. As $pc \sqcup \ell_1^e \not\sqsubseteq \ell$, by contradiction the case does not hold.

(e) ($k_1 = \ell_1 \wedge k_2 = \ell_2^*$):

– $\ell_1 \not\sqsubseteq \ell$:

- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2^*$, by definition 7.3, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by corollary 2, $pc \sqcup \ell_2^e \sqsubseteq \ell_2$. As $pc \sqcup \ell_2^e \not\sqsubseteq \ell$ and by definition 7.2, $\ell_2 \not\sqsubseteq \ell$. By definition 7.5, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2^*$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by definition 7.2, $\ell_1^* \not\sqsubseteq \ell$. By definition 7.4, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by corollary 2, $pc \sqcup \ell_2^e \sqsubseteq \ell_2$. As $pc \sqcup \ell_2^e \not\sqsubseteq \ell$ and by definition 7.2, $\ell_2 \not\sqsubseteq \ell$. By definition 7.2, $\ell_1^* \not\sqsubseteq \ell$. By definition 7.2, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.

– $\ell_2 \sqsubseteq \ell_1$:

- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2^*$, by definition 7.3, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2^*$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, $\ell_2^* \sqsubseteq (pc \sqcup \ell_2^e) \sqcap \ell_2$ as $pc \sqcup \ell_2^e \not\sqsubseteq \ell_2$ and $\ell_1^* = \ell_1$ by corollary 1 and definition 7.1. Thus, $\ell_2^* \sqsubseteq \ell_1^*$. By definition 7.5, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by corollary 2, $pc \sqcup \ell_2^e \sqsubseteq \ell_2$. As $pc \sqcup \ell_2^e \not\sqsubseteq \ell$, and by definition 7.2, $\ell_2 \not\sqsubseteq \ell$. By definition 7.4, $\sigma'_1(\mathbf{x}) \sim_\ell \sigma'_2(\mathbf{x})$.
- * If $k'_1 = \ell_1^* \wedge k'_2 = \ell_2$, then as $\sigma_1 \sim_\ell \sigma'_1$ and $\sigma_2 \sim_\ell \sigma'_2$, by corollary 1, $pc \sqcup \ell_2^e \sqsubseteq \ell_2$. As $pc \sqcup \ell_2^e \not\sqsubseteq \ell$, by contradiction the case does not hold.

5. WHILE-T: IH1: If $\sigma_1 \sim_\ell \sigma_2$, $\langle \sigma_1, c \rangle \Downarrow_{pc \sqcup \ell_1^e} \sigma'_1$, $\langle \sigma_2, c \rangle \Downarrow_{pc \sqcup \ell_2^e} \sigma'_2$ and $pc \sqcup \ell_1^e = pc \sqcup \ell_2^e$ then $\sigma'_1 \sim_\ell \sigma'_2$.

IH2: If $\sigma'_1 \sim_\ell \sigma'_2$, $\langle \sigma'_1, c \rangle \Downarrow_{pc \sqcup \ell_1^e} \sigma_1$, $\langle \sigma'_2, c \rangle \Downarrow_{pc \sqcup \ell_2^e} \sigma_2$ and $pc \sqcup \ell_1^e = pc \sqcup \ell_2^e$ then $\sigma_1 \sim_\ell \sigma_2$.

- If $\ell_1^e \sqsubseteq \ell$, $\ell_1^e = \ell_2^e$ and $n_1 = n_2$. By IH1 and IH2, $\sigma_1' \sim_\ell \sigma_2'$.
- If $\ell_1^e \not\sqsubseteq \ell$, then $\ell_2^e \not\sqsubseteq \ell$, $pc \sqcup \ell_i^e \not\sqsubseteq \ell$ for $i = 1, 2$. By Lemma 7, $\sigma_1 \sim_\ell \sigma_1''$ and $\sigma_2 \sim_\ell \sigma_2''$.
T.S. $\sigma_1'' \sim_\ell \sigma_2''$: By similar reasoning as IF-ELSE.
As $\sigma_1'' \sim_\ell \sigma_2''$, and by Lemma 7, $\sigma_1' \sim_\ell \sigma_1''$ and $\sigma_2' \sim_\ell \sigma_2''$.
T.S. $\sigma_1' \sim_\ell \sigma_2'$: By similar reasoning as IF-ELSE.

6. WHILE-F: $\sigma_1' = \sigma_1 \sim_\ell \sigma_2 = \sigma_2'$

□

B Proofs of Precision for Dynamic IFC with Unstructured Control Flow and Exceptions

Theorem 3 (Precision). *Choosing any node other than the IPD to lower the pc-label will either give unsound results or be less precise.*

Proof. Consider a branch-point $b \in \mathcal{N}$ with IPD $IPD(b) = i \in \mathcal{N}$.

Assume that $(n \in \mathcal{N}) \neq i$ is the node where the context of the predicate expression in b is removed.

Thus, either:

- $b < n < i$: Then, $\exists p.n \notin b \rightarrow_p n_e$. Thus, if n performs an action that should not have been performed in the context of the predicate expression in b , it might leak information about the predicate expression in b .
- $b < i < n$: Then, for any $n' \in \mathcal{N}$ such that $i < n' < n$ performing an operation that should not be performed in the context of b would be reported illicit as n' would be executed in the context of b .
 - If $n' \text{ pd } b$, then $\forall p.n' \in b \rightarrow_p n_e$. Hence, the statement n' executes irrespective of whether the branch at b is taken or not and hence, does not depend on the predicate expression in b , i.e., there is no implicit flow from the predicate expression in b to n' , but still the program might be rejected.
 - If n' is a statement executing under the context of another branch-point b' , such that $b' \text{ pd } b$, then as b' does not have any implicit flow from the predicate expression in b , any statement executing under the context of the predicate expression in b' should not be influenced by the context of the predicate expression in b . Hence, the program might be rejected even though there is no information leak.
- $i \langle \rangle n$: $\forall p.n \notin i \rightarrow_p n_e$ or $\forall p.n \notin b \rightarrow_p n_e$, n will never be reached. Thus, the context of b shall not be removed until n_e such that $b < i < n_e$. Similar reasoning as in the second case with $n = n_e$.

Hence, the most precise node where one can safely remove the context of b is $n = IPD(b) = i$. \square

Theorem 4. *The actual IPD of a node having SEN as its IPD is the node on the top of the pc-stack, which lies in a previously called function.*

Proof. Assume two functions F and G given by the CFGs $\mathcal{G} = (\mathcal{N}, \mathcal{E}, n_s, n_e, \mathcal{L})$ and $\mathcal{G}' = (\mathcal{N}', \mathcal{E}', n'_s, n'_e, \mathcal{L}')$, respectively. The program's start and exit node are given by N_s and N_e , respectively. Consider a branch-point $b' \in \mathcal{N}'$ having SEN as its IPD. Assume a branch-point $b \in \mathcal{N}$ such that $b < b' < (i = IPD(b))$, $(i \in \mathcal{N}) \neq \text{SEN}$, b is the last executed branch-point and top of the pc-stack contains i .

$\forall p.i \in b \rightarrow_p n_e$. Thus, $i \text{ pd } b'$ such that $b < b' < i$.

T.S. $\nexists n \in b' \rightarrow_p i \mid (n \text{ pd } b') \wedge (b' < n < i)$.

Proof by contradiction: Assume $\exists n.n \text{ pd } b' \mid b' < n < i$. Then the node n either lies in the function F

or G or in another function H given by the CFG $\mathcal{G}'' = (\mathcal{N}'', \mathcal{E}'', n_s'', n_e'', \mathcal{L}'')$, such that F calls H and H calls G .

- $n \in \mathcal{N}'$: As $IPD(b') = \text{SEN}$ and SEN is the last node in a function(\mathcal{G}'), $\exists p.n \notin b \rightarrow_p n_e$.
- $n \in \mathcal{N}''$: As $\forall p.n \in b' \rightarrow_p N_e$ and $G() < b'$, thus, $\forall p.n \in G() \rightarrow_p N_e$, which means $IPD(G()) \neq \text{SEN}$. Hence, the top of the pc-stack then would have $IPD(G()) = (n'' \in \mathcal{N}'') \leq n$ and not i .
- $n \in \mathcal{N}$: When the call to G or any other function H is made, it would push i , IPD of the branch-point on the top of the pc-stack.

Thus, $\nexists n \in b' \rightarrow_p i \mid (n \text{ pd } b') \wedge (b' < n < i)$. Hence, the top of the pc-stack, i is the actual IPD of any node b' having SEN as its intra-procedural IPD. \square

C Proofs for IFC with Unstructured Control Flow and Exceptions

Lemma 8 (Confinement Lemma). *If $\langle \sigma, \iota, \rho \rangle \rightarrow \langle \sigma', \iota', \rho' \rangle$ and $\Gamma(!\rho) \not\sqsubseteq \ell$, then $\sigma \sim_\ell \sigma'$, and $\rho \sim_\ell \rho'$.*

Proof. $\Gamma(!\rho) = pc$ in the proof that follows.

As $pc \not\sqsubseteq \ell$, the nodes in the pc-stack that have label less than or equal to ℓ will remain unchanged. Branching instructions pushing a new node would have label of at least pc due to monotonicity of pc-stack. Even if ι' is the IPD corresponding to the $!\rho.ipd$, it would only pop the top node. Thus, all the nodes that have label less than or equal to ℓ will remain unchanged. Hence, $\rho \sim_\ell \rho'$.

To show: $\sigma \sim_\ell \sigma'$.

By case analysis on the instruction type:

- ASSN, CATCH: Similar to cases ASSN-N and ASSN-S of Lemma 7.
- BRANCH, JMP, RET, SEN, THROW: $\sigma = \sigma'$

□

Lemma 9. *If $\langle \sigma_0, \iota_0, \rho_0 \rangle \rightarrow^n \langle \sigma_n, \iota_n, \rho_n \rangle$ and $\forall (0 \leq i \leq n). \Gamma(!\rho_i) \not\sqsubseteq \ell$, then $\rho_0 \sim_\ell \rho_n$.*

Proof. Proof by induction on n .

Basis: $\rho_0 \sim_\ell \rho_0$

IH: $\rho_0 \sim_\ell \rho_{n-1}$

From Definition 14, all nodes labeled less than or equal to ℓ of ρ_0 and ρ_{n-1} are equal. From Lemma 8, $\rho_{n-1} \sim \rho_n$ so, all nodes labeled less than or equal to ℓ of ρ_{n-1} and ρ_n are equal. Thus, all nodes labeled less than or equal to ℓ of ρ_0 and ρ_n are equal and by Definition 14, $\rho_0 \sim \rho_n$. □

Lemma 10. \star -preservation Lemma

If $\langle \sigma, \iota, \rho \rangle \rightarrow \langle \sigma', \iota', \rho' \rangle$,

then $\forall \mathbf{x}. \Gamma(\sigma(\mathbf{x})) = \ell^ \wedge (\Gamma(!\rho) \not\sqsubseteq \ell) \implies \Gamma(\sigma'(\mathbf{x})) = \ell'^* \wedge \ell' \sqsubseteq \ell$*

Proof. Proof by case analysis on the instruction type:

- ASSN, CATCH: from the premise
- BRANCH, JMP, RET, SEN, THROW: $\sigma = \sigma'$

□

Corollary 3. *If $\langle \sigma, \iota, \rho \rangle \rightarrow \langle \sigma', \iota', \rho' \rangle$, and $\Gamma(\sigma(\mathbf{x})) = \ell^*$ and $\Gamma(\sigma'(\mathbf{x})) = \ell'$, then $\Gamma(!\rho) \sqsubseteq \ell$.*

Proof. Immediate from Lemma 10. □

Lemma 11. *If $\langle \sigma_0, \iota_0, \rho_0 \rangle \rightarrow^* \langle \sigma_n, \iota_n, \rho_n \rangle$ and $\forall (0 \leq i \leq n). \Gamma(!\rho_i) \not\sqsubseteq \ell$, then $\sigma_0 \sim_\ell \sigma_n$.*

Proof. By induction on n .

Basis: $\sigma_0 \sim_\ell \sigma_0$ by Definition 8.

IH: $\sigma_0 \sim_\ell \sigma_{n-1}$.

From IH and Definition 8, $\forall \mathbf{x}. (\sigma_0(\mathbf{x}) \sim_\ell \sigma_{n-1}(\mathbf{x}))$. From Lemma 8, $\sigma_{n-1} \sim_\ell \sigma_n$. Thus, $\forall \mathbf{x}. (\sigma_{n-1}(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x}))$

Assume that $\sigma_0(\mathbf{x}) = \mathbf{v}_0^k$, $\sigma_{n-1}(\mathbf{x}) = \mathbf{v}_{n-1}^{k'}$, and $\sigma_n(\mathbf{x}) = \mathbf{v}_n^{k''}$ either:

- $(k = k') = \ell' \sqsubseteq \ell \wedge v_0 = v_{n-1}$:
 1. $(k' = k'') = \ell' \sqsubseteq \ell \wedge v_{n-1} = v_n$: $(k = k'') = \ell' \sqsubseteq \ell \wedge v_0 = v_n$. Thus, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$.
 2. $k' = \ell'$ and $k'' = \ell''^* \wedge \ell'' \sqsubseteq \ell' \sqsubseteq \ell$: By definition 7.5. Thus, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$.
- $k = \ell_1 \not\sqsubseteq \ell \wedge k' = \ell_2 \not\sqsubseteq \ell$:
 1. $k' = \ell_2 \not\sqsubseteq \ell \wedge k'' = \ell_3 \not\sqsubseteq \ell$. By definition 7.2, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$.
 2. $k' = \ell_2 \not\sqsubseteq \ell \wedge k'' = \ell_3^*$: $\ell_1 \not\sqsubseteq \ell$. Thus, by definition 7.5, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$.
- $k = \ell_1^* \wedge k' = \ell_2^*$:
 1. $k' = \ell_2^* \wedge k'' = \ell_3^*$: By definition 7.3, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$.
 2. $k' = \ell_2^* \wedge k'' = \ell_3 \wedge (\ell_3 \not\sqsubseteq \ell)$: By definition 7.4, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$.
 3. $k' = \ell_2^* \wedge k'' = \ell_3 \wedge (\ell_2 \sqsubseteq \ell_3)$: By corollary 3, $\Gamma(!\rho_{n-1}) \sqsubseteq \ell_2$. As $\Gamma(!\rho_{n-1}) \not\sqsubseteq \ell$ and $\ell_2 \sqsubseteq \ell_3$, so $\ell_3 \not\sqsubseteq \ell$. By definition 7.4, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$.
- $k = \ell_1^* \wedge k' = \ell_2$ s.t. $(\ell_2 \not\sqsubseteq \ell)$: Either
 - $k' = \ell_2 \not\sqsubseteq \ell \wedge k'' = \ell_3 \not\sqsubseteq \ell$: By definition 7.4, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$
 - $k' = \ell_2 \not\sqsubseteq \ell \wedge k'' = \ell_3^*$: By definition 7.3, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$
- $k = \ell_1^* \wedge k' = \ell_2$ s.t. $(\ell_1 \sqsubseteq \ell_2)$:
 - $k' = k'' = \ell_2$: By definition 7.4, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$
 - $k' = \ell_2 \not\sqsubseteq \ell \wedge k'' = \ell_3 \not\sqsubseteq \ell$: By definition 7.4, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$
 - $k' = \ell_2 \not\sqsubseteq \ell \wedge k'' = \ell_3^*$: By definition 7.3, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$
- $k = \ell_1 \wedge k' = \ell_2^*$ s.t. $(\ell_1 \not\sqsubseteq \ell)$:
 1. $k' = \ell_2^* \wedge k'' = \ell_3^*$: By definition 7.5, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$
 2. $k' = \ell_2^* \wedge k'' = \ell_3$ s.t. $(\ell_3 \not\sqsubseteq \ell)$: By definition 7.2, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$
 3. $k' = \ell_2^* \wedge k'' = \ell_3$ s.t. $(\ell_2 \sqsubseteq \ell_3)$: By corollary 3, $\Gamma(!\rho_n) \sqsubseteq \ell_2$. As $\Gamma(!\rho_n) \not\sqsubseteq \ell$ and $\ell_2 \sqsubseteq \ell_3$, so $\ell_3 \not\sqsubseteq \ell$. By definition 7.2, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$.
- $k = \ell_1 \wedge k' = \ell_2^*$ s.t. $(\ell_2 \sqsubseteq \ell_1)$: Also, $(\ell_2 \sqsubseteq \ell_1 \sqsubseteq \ell)$.
 1. $k' = \ell_2^* \wedge k'' = \ell_3^*$: As $\ell_2 \sqsubseteq \ell$ and $\Gamma(!\rho_n) \not\sqsubseteq \ell$, $\Gamma(!\rho_n) \not\sqsubseteq \ell_2$. By lemma 10, $\ell_3 \sqsubseteq \ell_2$. Thus, $\ell_3 \sqsubseteq \ell_2 \sqsubseteq \ell_1$. By definition 7.5, $\sigma_0(\mathbf{x}) \sim_\ell \sigma_n(\mathbf{x})$
 2. $k' = \ell_2^* \wedge k'' = \ell_3$: As $\ell_2 \sqsubseteq \ell$ and $\Gamma(!\rho_n) \not\sqsubseteq \ell$, $\Gamma(!\rho_n) \not\sqsubseteq \ell_2$. But, by Lemma 9, $\Gamma(!\rho_n) \sqsubseteq \ell_2$. By contradiction, this case does not hold.

□

Lemma 12. *Suppose*

$$\langle \sigma_1, \iota, \rho_1 \rangle \rightarrow \langle \sigma'_1, \iota'_1, \rho'_1 \rangle,$$

$$\langle \sigma_2, \iota, \rho_2 \rangle \rightarrow \langle \sigma'_2, \iota'_2, \rho'_2 \rangle,$$

$\sigma_1 \sim_\ell \sigma_2, \rho_1 \sim_\ell \rho_2, \Gamma(!\rho_1) = \Gamma(!\rho_2) \sqsubseteq \ell$, and either $\Gamma(!\rho'_1) = \Gamma(!\rho'_2) \sqsubseteq \ell$ or $\Gamma(!\rho'_1) \not\sqsubseteq \ell \wedge \Gamma(!\rho'_2) \not\sqsubseteq \ell$)
then $\sigma'_1 \sim_\ell \sigma'_2$ and $\rho'_1 \sim_\ell \rho'_2$.

Proof. Every instruction executes *isIPD* at the end of the operation. If ι'_i is the IPD corresponding to the $!\rho_i$. *ipd*, then it pops the first node on the pc-stack. As $\rho_1 \sim \rho_2$ and $\Gamma(!\rho_1) = \Gamma(!\rho_2)$, ι'_i would either pop in both the runs or in none. Thus, $\rho'_1 \text{ sim } \rho'_2$ (BRANCH rule is explained below).

Assume $\sigma_1(x) = v_1^{k_1}, \sigma_2(x) = v_2^{k_2}, \sigma'_1(x) = v_1^{k'_1}$ and $\sigma'_2(x) = v_2^{k'_2}$.

Proof by case analysis on the instruction type:

- ASSN, CATCH: $\Gamma(!\rho_1) = \Gamma(!\rho_2) = pc \sqsubseteq \ell$
 - $pc \sqsubseteq \ell_1 \wedge pc \sqsubseteq \ell_2$: As \mathbf{n}^m is equivalent, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
 - $pc \not\sqsubseteq \ell_1 \wedge pc \not\sqsubseteq \ell_2$: By Definition 7.3, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
 - $pc \sqsubseteq \ell_1 \wedge pc \not\sqsubseteq \ell_2$: $k'_2 \sqsubseteq pc$ and $pc \sqsubseteq k'_1$. By Definition 7.3 and 7.4, $\sigma'_1(x) \sim_\ell \sigma'_2(x)$.
Similarly for the analogous case.
- BRANCH: As \mathbf{b}^{ℓ_i} is equivalent in the two runs, either $\ell_1 = \ell_2 \sqsubseteq \ell$ or $\ell_1 \not\sqsubseteq \ell \wedge \ell_2 \not\sqsubseteq \ell$ (ℓ_i does not have *). The IPD of ι would be the same in both the cases. If the IPD is SEN, then the label of $!\rho_i$ is joined with the label obtained above, which is either less than or equal to ℓ and same in both the runs (or) not less than or equal to ℓ in both the runs. Thus, either $\Gamma(!\rho'_1) = \Gamma(!\rho'_2)$ or $\Gamma(!\rho'_1) \not\sqsubseteq \ell \wedge \Gamma(!\rho'_2) \not\sqsubseteq \ell$. Because $\rho_1 \sim \rho_2, \rho'_1 \sim_\ell \rho'_2$.
If the IPD is not SEN, then it is some other node, which makes the *ipd* field the same. Thus, the pushed node is the same in both the cases or has label not less than or equal to ℓ and hence, $\rho'_1 \sim_\ell \rho'_2, \sigma'_1 = \sigma_1 \sim_\ell \sigma_2 = \sigma'_2$.
- Other rules: $\sigma'_1 = \sigma_1 \sim_\ell \sigma_2 = \sigma'_2$.

□

Lemma 13. *Suppose*

$$1. \langle \sigma'_0, \iota_0, \rho'_0 \rangle \rightarrow \langle \sigma'_1, \iota'_1, \rho'_1 \rangle \rightarrow^{n-1} \langle \sigma'_n, \iota'_n, \rho'_n \rangle,$$

$$2. \langle \sigma''_0, \iota_0, \rho''_0 \rangle \rightarrow \langle \sigma''_1, \iota'_1, \rho''_1 \rangle \rightarrow^{m-1} \langle \sigma''_m, \iota'_m, \rho''_m \rangle,$$

$$3. (\rho'_0 \sim_\ell \rho''_0), (\sigma'_0 \sim_\ell \sigma''_0),$$

$$4. (\Gamma(!\rho'_0) = \Gamma(!\rho''_0) \sqsubseteq \ell), (\Gamma(!\rho'_n) = \Gamma(!\rho''_m) \sqsubseteq \ell),$$

$$5. \forall(0 < i < n).(\Gamma(!\rho'_i) \not\sqsubseteq \ell) \wedge \forall(0 < j < m).(\Gamma(!\rho''_j) \not\sqsubseteq \ell),$$

then $(\iota'_n = \iota'_m), (\rho'_n \sim_\ell \rho''_m)$, and $(\sigma'_n \sim_\ell \sigma''_m)$.

Proof. Starting with the same instruction and high context in both the runs can result in two different instructions, ι'_1 and ι''_1 . This is only possible if ι was some branching instruction in the first place and this divergence happened in a high context.

1. To prove $\iota'_n = \iota''_m$:

From the property of the IPDs, if ι_0 pushes a node with label $\not\sqsubseteq \ell$ on top of pc -stack which was originally $\sqsubseteq \ell$, $IPD(\iota_0)$ pops that node. Since the runs start from the same instruction ι_0 , $\iota'_n = \iota''_m = IPD(\iota)$, where $\Gamma(!\rho) \sqsubseteq \ell$.

2. To prove $\rho'_n \sim_\ell \rho''_m$:

- $n > 1$ and $m > 1$: $\Gamma(!\rho'_1) \not\sqsubseteq \ell \wedge \Gamma(!\rho''_1) \not\sqsubseteq \ell$, because ι_0 has the same IPD and ι'_1, ι''_1 are not the IPDs. As $\rho'_0 \sim \rho''_0$ and $\Gamma(!\rho'_1) \not\sqsubseteq \ell \wedge \Gamma(!\rho''_1) \not\sqsubseteq \ell$, from Lemma 12, $\rho'_1 \sim \rho''_1$ and $!\rho'_1.ipd = !\rho''_1.ipd = IPD(\iota_0)$, if $\iota'_1 \neq IPD(\iota_0)$ and $\iota''_1 \neq IPD(\iota_0)$. As $\iota'_n = \iota''_m = IPD(\iota_0)$, it pops the $!\rho'_1$ and $!\rho''_1$, which correspond to ρ'_n and ρ''_m in the n th and m th step. Because $\rho'_1 \sim \rho''_1$ and from Lemma 9, $\rho'_n \sim \rho''_m$.
- $n = 1$ and $m > 1$: If $\iota'_1 = IPD(\iota_0)$, and $\Gamma(!\rho'_1) \sqsubseteq \ell$. It pops the node pushed by ι_0 , i.e., $\Gamma(!\rho'_n) \sqsubseteq \ell$. In the other run as $\Gamma(!\rho''_1) \not\sqsubseteq \ell$ and $\Gamma(!\rho''_m) \sqsubseteq \ell$, by the property of IPD $\iota''_m = IPD(\iota_0)$, which would pop from the pc -stack $!\rho''_m$, the first frame labelled $\not\sqsubseteq \ell$ on the pc -stack. Thus, $\rho'_n \sim \rho''_m$.
- $n > 1$ and $m = 1$: Similar to the above case.

3. To prove $\sigma'_n \sim_\ell \sigma''_m$:

- (a) $n > 1$ and $m > 1$: From Lemma 12, $\sigma'_1 \sim_\ell \sigma''_1$. From Lemma 11, $\sigma'_1 \sim_\ell \sigma'_{n-1}$ and $\sigma''_1 \sim_\ell \sigma''_{m-1}$. And from Lemma 8 $\sigma'_{n-1} \sim_\ell \sigma'_n$ and $\sigma''_{m-1} \sim_\ell \sigma''_m$. Similar case analysis as above for different cases of equivalence.
- (b) $n = 1$ and $m > 1$: In case of BRANCH, $\sigma'_0 = \sigma'_1$ and $\sigma''_0 = \sigma''_1$. Thus, $\sigma'_1 \sim_\ell \sigma''_1$. From the above case, if $\sigma'_1 \sim_\ell \sigma''_1$, then $\sigma'_n \sim_\ell \sigma''_m$.
- (c) $n > 1$ and $m = 1$: Symmetric case of the above.

□

Definition 1 (Trace). A trace is defined as a sequence of configurations or states resulting from a program evaluation, i.e., for a program evaluation $\mathcal{P} = s_1 \rightarrow s_2 \rightarrow \dots \rightarrow s_n$ where $s_i = \langle \sigma_i, \iota_i, \rho_i \rangle$, the corresponding trace is given as $\mathcal{T}(\mathcal{P}) := s_1 :: s_2 :: \dots :: s_n$.

Definition 2 (Epoch-trace). An epoch-trace for an adversary at level ℓ , (\mathcal{E}_ℓ) over a trace $\mathcal{T} = s_1 :: s_2 :: \dots :: s_n$ where $s_i = \langle \sigma_i, \iota_i, \rho_i \rangle$ is defined inductively as:

$$\begin{aligned} \mathcal{E}_\ell(\text{nil}) &:= \text{nil} \\ \mathcal{E}_\ell(s_i :: \mathcal{T}) &:= \begin{cases} s_i :: \mathcal{E}(\mathcal{T}) & \text{if } \Gamma(!\rho_i) \sqsubseteq \ell, \\ \mathcal{E}(\mathcal{T}) & \text{else if } \Gamma(!\rho_i) \not\sqsubseteq \ell. \end{cases} \end{aligned}$$

Theorem 5 (Termination-Insensitive Non-interference). Suppose \mathcal{P} and \mathcal{P}' are two program evaluations. Then for their respective epoch-traces with respect to an adversary at level ℓ given by:

$\mathcal{E}_\ell(\mathcal{T}(\mathcal{P})) = s_1 :: s_2 :: \dots :: s_n,$
 $\mathcal{E}_\ell(\mathcal{T}(\mathcal{P}')) = s'_1 :: s'_2 :: \dots :: s'_m,$
 if $s_1 \sim_\ell s'_1$ and $n \leq m,$
 then
 $s_n \sim_\ell s'_n$

Proof. Proof by induction on n .

Basis: $s_1 \sim_\ell s'_1$, by assumption.

IH: $s_k \sim_\ell s'_k$

To prove: $s_{k+1} \sim_\ell s'_{k+1}$.

Let $s_k \rightarrow_i s_{k+1}$ and $s_k \rightarrow_{i'} s'_{k+1}$, then:

- $i = i' = 1$: From Lemma 12, $s_{k+1} \sim_\ell s'_{k+1}$.
- $i > 1$ or $i' > 1$: From Lemma 13, $s_{k+1} \sim_\ell s'_{k+1}$.

□

Corollary 4. *Suppose:*

1. $\langle \sigma_1, \iota_1, \rho_1 \rangle \sim_\ell \langle \sigma_2, \iota_2, \rho_2 \rangle$
2. $\langle \sigma_1, \iota_1, \rho_1 \rangle \rightarrow^* \langle \sigma'_1, \mathbf{end}, [] \rangle$
3. $\langle \sigma_2, \iota_2, \rho_2 \rangle \rightarrow^* \langle \sigma'_2, \mathbf{end}, [] \rangle$

Then, $\sigma'_1 \sim_\ell \sigma'_2$.

Proof. σ_1, σ_2 and ρ_1, ρ_2 are empty at the end of $*$ steps. From the semantics, in context $\sqsubseteq \ell$ both runs would push and pop the same number of nodes. Thus, both take same number of steps in the epoch-trace. Assume it to be k . Then in Theorem 5, $n = m = k$. Thus, $s_k \sim_\ell s'_k$, where $s_k = \langle \sigma'_1, \mathbf{end}, [] \rangle$ and $s'_k = \langle \sigma'_2, \mathbf{end}, [] \rangle$. By Definition 15, $\sigma'_1 \sim_\ell \sigma'_2$. □

D Proofs for Limited Information Release

Lemma 14 (Trace-projection). $(\tau_1 :: \tau_2)\uparrow_\ell = \tau_1\uparrow_\ell :: \tau_2\uparrow_\ell$

Proof. By induction on τ_1 .

Base case: If $\tau_1 = []$, then $([] :: \tau_2)\uparrow_\ell = \tau_2\uparrow_\ell = [] :: \tau_2\uparrow_\ell = []\uparrow_\ell :: \tau_2\uparrow_\ell$.

Inductive case: IH: $(\tau_1 :: \tau_2)\uparrow_\ell = \tau_1\uparrow_\ell :: \tau_2\uparrow_\ell$

TS: $((\mathbf{n}^m :: \tau_1) :: \tau_2)\uparrow_\ell = (\mathbf{n}^m :: \tau_1)\uparrow_\ell :: \tau_2\uparrow_\ell$

$$(\mathbf{n}^m :: \tau_1)\uparrow_\ell = \begin{cases} \mathbf{n}^m :: \tau_1\uparrow_\ell & \text{if } m \sqsubseteq \ell, \\ \tau_1\uparrow_\ell & \text{else.} \end{cases} \quad (1)$$

Also by associativity of list concatenation, $((\mathbf{n}^m :: \tau_1) :: \tau_2) = (\mathbf{n}^m :: (\tau_1 :: \tau_2))$. Again,

$$(\mathbf{n}^m :: (\tau_1 :: \tau_2))\uparrow_\ell = \begin{cases} \mathbf{n}^m :: (\tau_1 :: \tau_2)\uparrow_\ell & \text{if } m \sqsubseteq \ell, \\ (\tau_1 :: \tau_2)\uparrow_\ell & \text{else.} \end{cases} \quad (2)$$

If $m \sqsubseteq \ell$: from (2), LHS = $\mathbf{n}^m :: (\tau_1 :: \tau_2)\uparrow_\ell$ and from (1), RHS = $\mathbf{n}^m :: \tau_1\uparrow_\ell :: \tau_2\uparrow_\ell$. From IH, LHS = RHS. Else: from (2), LHS = $(\tau_1 :: \tau_2)\uparrow_\ell$ and from (1), RHS = $\tau_1\uparrow_\ell :: \tau_2\uparrow_\ell$. From IH, LHS = RHS. \square

Lemma 15 (Confinement of expressions). *If*

$\langle \sigma, \iota, e \rangle \Downarrow_{pc} \mathbf{n}^{(m, \delta)}$, ι', τ , and $\Gamma(pc) \not\sqsubseteq \ell$, then $\iota \sim_\ell \iota'$ and $\tau\uparrow_\ell = \epsilon$

Proof. Proof by induction on the derivation for expressions and case analysis on the last rule.

- CONST and VAR: $\iota = \iota'$ and $\tau = \epsilon$
- AOP and COP:
 - IH1: If $\langle \sigma, \iota, e_1 \rangle \Downarrow_{pc} \mathbf{n}_1^{(m_1, \delta_1)}$, ι_1, τ_1 , and $pc \not\sqsubseteq \ell$, then $\iota \sim_\ell \iota_1$ and $\tau_1\uparrow_\ell = \epsilon$
 - IH2: If $\langle \sigma, \iota_1, e_2 \rangle \Downarrow_{pc} \mathbf{n}_2^{(m_2, \delta_2)}$, ι', τ_2 , and $pc \not\sqsubseteq \ell$, then $\iota_1 \sim_\ell \iota'$ and $\tau_2\uparrow_\ell = \epsilon$

As $pc \not\sqsubseteq \ell$, $\iota \sim_\ell \iota_1$ and $\tau_1\uparrow_\ell = \epsilon$ and $\iota_1 \sim_\ell \iota'$ and $\tau_2\uparrow_\ell = \epsilon$. Thus, $(\tau_1\uparrow_\ell :: \tau_2\uparrow_\ell) = (\tau_1 :: \tau_2)\uparrow_\ell = \epsilon$.

As $\iota \sim_\ell \iota_1$, $\forall \mathbf{x} \in \iota.\mathbb{B}(\mathbf{x}) \sqsubseteq \ell \implies \iota(\mathbf{x}) = \iota_1(\mathbf{x})$ and $\iota_1 \sim_\ell \iota'$, $\forall \mathbf{x} \in \iota_1.\mathbb{B}(\mathbf{x}) \sqsubseteq \ell \implies \iota_1(\mathbf{x}) = \iota'(\mathbf{x})$. Thus, $\forall \mathbf{x} \in \iota.\mathbb{B}(\mathbf{x}) \sqsubseteq \ell \implies \iota(\mathbf{x}) = \iota'(\mathbf{x})$.
- DCOPR: From above reasoning, $\tau_1\uparrow_\ell :: \tau_2\uparrow_\ell = (\tau_1 :: \tau_2)\uparrow_\ell = \epsilon$ and $\forall \mathbf{x} \in \iota.\mathbb{B}(\mathbf{x}) \sqsubseteq \ell \implies \iota(\mathbf{x}) = \iota_2(\mathbf{x})$. As $pc \sqsubseteq k_o$ and $pc \not\sqsubseteq \ell$, $k_o \not\sqsubseteq \ell$. ι_2 changes for only those x that have a budget label $\mathbb{B}(x) \supseteq k_o$. Thus, $\forall \mathbf{x} \in \iota.\mathbb{B}(\mathbf{x}) \sqsubseteq \ell \implies \iota(\mathbf{x}) = \iota'(\mathbf{x})$. As $l \not\sqsubseteq \ell$, $\mathbf{n}^l\uparrow_\ell = \epsilon$. Thus, $\tau\uparrow_\ell = (\tau_1 :: \tau_2 :: \mathbf{n}^l)\uparrow_\ell = \epsilon$
- DCOPN: From above reasoning, $\tau_1\uparrow_\ell :: \tau_2\uparrow_\ell = (\tau_1 :: \tau_2)\uparrow_\ell = \epsilon$ and $\forall \mathbf{x} \in \iota.\mathbb{B}(\mathbf{x}) \sqsubseteq \ell \implies \iota(\mathbf{x}) = \iota_2(\mathbf{x})$. $\tau\uparrow_\ell = \epsilon$.

\square

Lemma 16 (Confinement). *If* $\langle \sigma, \iota, c \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau \rangle$, and $\Gamma(pc) \not\sqsubseteq \ell$, then $\sigma \sim_\ell \sigma'$, $\iota \sim_\ell \iota'$ and $\tau\uparrow_\ell = \epsilon$

Proof. Proof by induction on the derivation for commands and case analysis on the last rule.

- SKIP: $\sigma = \sigma', \iota = \iota'$ and $\tau = \epsilon$
- ASSN: As $pc \not\sqsubseteq \ell$ and $pc \sqsubseteq \Gamma_f(\sigma(\mathbf{x})), \Gamma_f(\sigma(\mathbf{x})) \not\sqsubseteq \ell$. Also, $(\Gamma_f(\sigma'(\mathbf{x})) = pc \sqcup m) \not\sqsubseteq \ell$ and all other $\sigma(\mathbf{x})$ remains unchanged. Thus, $\sigma \sim_\ell \sigma'$. From Lemma 15, $\iota \sim_\ell \iota'$ and $\tau \uparrow_\ell = \epsilon$.
- SEQ: IH1: If $\langle \sigma, \iota, c_1 \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau_1 \rangle$, and $pc \not\sqsubseteq \ell$, then $\sigma \sim_\ell \sigma', \iota \sim_\ell \iota'$ and $\tau_1 \uparrow_\ell = \epsilon$
 IH2: If $\langle \sigma', \iota', c_2 \rangle \Downarrow_{pc} \langle \sigma'', \iota'', \tau_2 \rangle$, and $pc \not\sqsubseteq \ell$, then $\sigma' \sim_\ell \sigma'', \iota' \sim_\ell \iota''$ and $\tau_2 \uparrow_\ell = \epsilon$
 As \sim_ℓ is an equivalence relation, from IH1 and IH2, $\sigma \sim_\ell \sigma'', \iota \sim_\ell \iota''$. From Lemma 14, $(\tau_1 :: \tau_2) \uparrow_\ell = \epsilon$.
- IF-ELSE: IH: If $\langle \sigma, \iota', c_i \rangle \Downarrow_{pc \sqcup k} \langle \sigma', \iota'', \tau_2 \rangle$, and $pc \sqcup (k, \delta) \not\sqsubseteq \ell$, then $\sigma \sim_\ell \sigma', \iota' \sim_\ell \iota''$ and $\tau_2 \uparrow_\ell = \epsilon$
 From Lemma 15, $\iota \sim_\ell \iota'$ and $\tau_1 \uparrow_\ell = \epsilon$. Also, as $pc \not\sqsubseteq \ell, pc \sqcup l \not\sqsubseteq \ell$ and from IH, $\sigma \sim_\ell \sigma', \iota' \sim_\ell \iota''$ and $\tau_2 \uparrow_\ell = \epsilon$. Thus, $\iota \sim_\ell \iota''$ and $(\tau_1 :: \tau_2) \uparrow_\ell = \epsilon$
- WHILE-T: Similar to IF-ELSE and SEQ
- WHILE-F: $\sigma = \sigma'$. From Lemma 15, $\iota \sim_\ell \iota'$ and $\tau \uparrow_\ell = \epsilon$

□

Lemma 17. *If*

$\langle \sigma, \iota, e \rangle \Downarrow_{pc} \mathbf{n}^{(\iota, \delta)}, \iota', \tau$ and $\exists \mathbf{x} \in \sigma. \left(\Gamma(\sigma(\mathbf{x})) \not\sqsubseteq \ell \wedge (\forall \mathbf{y} \in \sigma. \mathbf{y} \neq \mathbf{x} \wedge \Gamma(\sigma(\mathbf{y})) \sqsubseteq \ell) \right)$,
then $|\tau \uparrow_\ell| \leq \iota \uparrow_\ell(\mathbf{x}) - \iota' \uparrow_\ell(\mathbf{x})$.

Proof. Induction on the derivation for expressions and case analysis on the last rule.

- CONST: $\iota = \iota'$ and $|\tau \uparrow_\ell| = 0$.
- VAR: $\iota = \iota'$ and $|\tau \uparrow_\ell| = 0$.
- AOP and COP: IH1: If $\langle \sigma, \iota, e_1 \rangle \Downarrow_{pc} \mathbf{n}_1^{(k_1, \delta_1)}, \iota_1, \tau_1$, and
 $\exists \mathbf{x} \in \sigma. \left(\Gamma(\sigma(\mathbf{x})) \not\sqsubseteq \ell \wedge (\forall \mathbf{y} \in \sigma. \mathbf{y} \neq \mathbf{x} \wedge \Gamma(\sigma(\mathbf{y})) \sqsubseteq \ell) \right)$ then $|\tau_1 \uparrow_\ell| \leq \iota_1 \uparrow_\ell(\mathbf{x}) - \iota' \uparrow_\ell(\mathbf{x})$.
 IH2: If $\langle \sigma, \iota_1, e_2 \rangle \Downarrow_{pc} \mathbf{n}_2^{(k_2, \delta_2)}, \iota', \tau_2$, and
 $\exists \mathbf{x} \in \sigma. \left(\Gamma(\sigma(\mathbf{x})) \not\sqsubseteq \ell \wedge (\forall \mathbf{y} \in \sigma. \mathbf{y} \neq \mathbf{x} \wedge \Gamma(\sigma(\mathbf{y})) \sqsubseteq \ell) \right)$ then $|\tau_2 \uparrow_\ell| \leq \iota_1 \uparrow_\ell(\mathbf{x}) - \iota' \uparrow_\ell(\mathbf{x})$.
 From IH1 and IH2, $|\tau_1 \uparrow_\ell| + |\tau_2 \uparrow_\ell| = \iota \uparrow_\ell(\mathbf{x}) - \iota' \uparrow_\ell(\mathbf{x})$. Thus, $|(\tau_1 :: \tau_2) \uparrow_\ell| \leq \iota \uparrow_\ell(\mathbf{x}) - \iota' \uparrow_\ell(\mathbf{x})$.
- DCOPR: Similar to COP, $|(\tau_1 :: \tau_2) \uparrow_\ell| \leq \iota \uparrow_\ell(\mathbf{x}) - \iota_2 \uparrow_\ell(\mathbf{x})$. Either:
 - $k_o \sqsubseteq \ell \wedge \mathbf{x} \in \delta_o. \mathbb{L}(\mathbf{x}) \not\sqsubseteq \ell$: If $\mathbf{x} \in \delta. (\mathbb{B}(\mathbf{x}) \not\sqsubseteq \ell) \implies l \not\sqsubseteq \ell$. Thus, $|\tau \uparrow_\ell| = |\tau_1 \uparrow_\ell| + |\tau_2 \uparrow_\ell|$.
 As $\iota_2 = \iota'$, $|\tau \uparrow_\ell| \leq \iota \uparrow_\ell(\mathbf{x}) - \iota' \uparrow_\ell(\mathbf{x})$.
 If $\mathbf{x} \in \delta. (\mathbb{B}(\mathbf{x}) \sqsubseteq \ell) \implies l \sqsubseteq \ell$ and $\iota'(\mathbf{x}) = \iota_2(\mathbf{x}) - 1$. $\tau = \tau_1 :: \tau_2 :: \mathbf{n}^l$. $|\tau \uparrow_\ell| = |\tau_1 \uparrow_\ell| + |\tau_2 \uparrow_\ell| + 1$. Thus, $|\tau \uparrow_\ell| \leq \iota \uparrow_\ell(\mathbf{x}) - \iota' \uparrow_\ell(\mathbf{x})$.
 - If $\delta = \emptyset$ then $\tau \uparrow_\ell = \tau_1 \uparrow_\ell :: \tau_2 \uparrow_\ell$ and $\iota' = \iota_2$. Thus, $|\tau \uparrow_\ell| \leq \iota \uparrow_\ell(\mathbf{x}) - \iota' \uparrow_\ell(\mathbf{x})$
 - $k_o \not\sqsubseteq \ell$: $l \not\sqsubseteq \ell$. Thus, $|\tau \uparrow_\ell| = |\tau_1 \uparrow_\ell| + |\tau_2 \uparrow_\ell|$. As $\iota_2 = \iota'$, $|\tau \uparrow_\ell| \leq \iota \uparrow_\ell(\mathbf{x}) - \iota' \uparrow_\ell(\mathbf{x})$.

- **DCOPN:** Similar to **COP**, $|(\tau_1 :: \tau_2)\uparrow\ell| \leq \iota\uparrow\ell - \iota_2\uparrow\ell$. Thus, $|\tau\uparrow\ell| \leq \iota\uparrow\ell(\mathbf{x}) - \iota'\uparrow\ell(\mathbf{x})$

□

Theorem 6 (LIR for a single secret). *If $\langle \sigma, \iota, c \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau \rangle$, and*

$\exists \mathbf{x} \in \sigma. \left(\Gamma(\sigma(\mathbf{x})) \not\sqsubseteq \ell \wedge (\forall \mathbf{y} \in \sigma. \mathbf{y} \neq \mathbf{x} \wedge \Gamma(\sigma(\mathbf{y})) \sqsubseteq \ell) \right)$,
then $|\tau\uparrow\ell| \leq \iota\uparrow\ell(\mathbf{x}) - \iota'\uparrow\ell(\mathbf{x})$.

Proof. Induction on the derivation for commands and case analysis on the last rule.

- **SKIP:** $\iota = \iota'$ and $|\tau\uparrow\ell| = 0$.
- **ASSN:** From Lemma 17.
- **SEQ:** IH1: If $\langle \sigma, \iota, c_1 \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau_1 \rangle$, and
 $\exists \mathbf{x} \in \sigma. \left(\Gamma(\sigma(\mathbf{x})) \not\sqsubseteq \ell \wedge (\forall \mathbf{y} \in \sigma. \mathbf{y} \neq \mathbf{x} \wedge \Gamma(\sigma(\mathbf{y})) \sqsubseteq \ell) \right)$ then $|\tau_1\uparrow\ell| \leq \iota\uparrow\ell(\mathbf{x}) - \iota'\uparrow\ell(\mathbf{x})$.
 IH2: If $\langle \sigma', \iota', c_2 \rangle \Downarrow_{pc} \langle \sigma'', \iota'', \tau_2 \rangle$, and
 $\exists \mathbf{x} \in \sigma. \left(\Gamma(\sigma(\mathbf{x})) \not\sqsubseteq \ell \wedge (\forall \mathbf{y} \in \sigma. \mathbf{y} \neq \mathbf{x} \wedge \Gamma(\sigma(\mathbf{y})) \sqsubseteq \ell) \right)$ then $|\tau_2\uparrow\ell| \leq \iota'\uparrow\ell(\mathbf{x}) - \iota''\uparrow\ell(\mathbf{x})$.
 $x \in \sigma \implies \mathbf{x} \in \sigma_1$.
 From IH1 and IH2, $|\tau_1\uparrow\ell| + |\tau_2\uparrow\ell| \leq \iota\uparrow\ell(\mathbf{x}) - \iota''\uparrow\ell(\mathbf{x})$. Thus, $|\tau\uparrow\ell| \leq \iota\uparrow\ell(\mathbf{x}) - \iota''\uparrow\ell(\mathbf{x})$.
- **IF-ELSE:** From Lemma 17, $|\tau_1\uparrow\ell| \leq \iota\uparrow\ell(\mathbf{x}) - \iota'\uparrow\ell(\mathbf{x})$.
 IH: If $\langle \sigma, \iota', c_i \rangle \Downarrow_{pc \sqcup k} \langle \sigma', \iota'', \tau_2 \rangle$, and
 $\exists \mathbf{x} \in \sigma. \left(\Gamma(\sigma(\mathbf{x})) \not\sqsubseteq \ell \wedge (\forall \mathbf{y} \in \sigma. \mathbf{y} \neq \mathbf{x} \wedge \Gamma(\sigma(\mathbf{y})) \sqsubseteq \ell) \right)$ then $|\tau_2\uparrow\ell| \leq \iota'\uparrow\ell(\mathbf{x}) - \iota''\uparrow\ell(\mathbf{x})$.
 From IH and Lemma 17, $|\tau\uparrow\ell| \leq \iota\uparrow\ell(\mathbf{x}) - \iota''\uparrow\ell(\mathbf{x})$.
- **WHILE-T:** Similar to **IF-ELSE**
- **WHILE-F:** From Lemma 17.

□

Lemma 18. *If $\langle \sigma, \iota, e \rangle \Downarrow_{pc} \mathbf{n}^{(\iota, \delta)}, \iota', \tau$, then $|\tau\uparrow\ell| \leq \iota\uparrow\ell - \iota'\uparrow\ell$.*

Proof. Induction on the derivation for expressions and case analysis on the last rule.

- **CONST:** $\iota = \iota'$ and $|\tau\uparrow\ell| = 0$.
- **VAR:** $\iota = \iota'$ and $|\tau\uparrow\ell| = 0$.
- **AOP and COP:**
 IH1: If $\langle \sigma, \iota, e_1 \rangle \Downarrow_{pc} \mathbf{n}_1^{(k_1, \delta_1)}, \iota_1, \tau_1$, then $|\tau_1\uparrow\ell| \leq \iota\uparrow\ell - \iota_1\uparrow\ell$.
 IH2: If $\langle \sigma, \iota_1, e_2 \rangle \Downarrow_{pc} \mathbf{n}_2^{(k_2, \delta_2)}, \iota', \tau_2$, then $|\tau_2\uparrow\ell| \leq \iota_1\uparrow\ell - \iota'\uparrow\ell$.
 From IH1 and IH2, $|\tau_1\uparrow\ell| + |\tau_2\uparrow\ell| = \iota\uparrow\ell - \iota'\uparrow\ell$. Thus, $|(\tau_1 :: \tau_2)\uparrow\ell| \leq \iota\uparrow\ell - \iota'\uparrow\ell$.
- **DCOPR:** Similar to **COP**, $|(\tau_1 :: \tau_2)\uparrow\ell| \leq \iota\uparrow\ell - \iota_2\uparrow\ell$. Either:

- $k_o \sqsubseteq \ell \wedge \forall \mathbf{x} \in \delta_o. \mathbb{L}(\mathbf{x}) \sqsubseteq \ell$: If $\delta \neq \emptyset$ then $\exists x. \iota_2(\mathbf{x}) > 0$ and $\iota_2(\mathbf{x}) - \iota'(\mathbf{x}) = 1$. For more than one x , $\iota_2 \uparrow \ell - \iota' \uparrow \ell \geq 1$. Also, $\tau = \tau_1 :: \tau_2 :: \mathbf{n}^l$. $|\tau \uparrow \ell| = |\tau_1 \uparrow \ell| + |\tau_2 \uparrow \ell| + 1$. Thus, $|\tau \uparrow \ell| \leq \iota \uparrow \ell - \iota' \uparrow \ell$
- $k_o \sqsubseteq \ell \wedge \exists \mathbf{x} \in \delta_o. \mathbb{L}(\mathbf{x}) \not\sqsubseteq \ell$: If $\exists \mathbf{x} \in \delta. (\mathbb{B}(\mathbf{x}) \not\sqsubseteq \ell) \implies l \not\sqsubseteq \ell$. Thus, $|\tau \uparrow \ell| = |\tau_1 \uparrow \ell| + |\tau_2 \uparrow \ell|$. As $\iota_2 = \iota'$, $|\tau \uparrow \ell| \leq \iota \uparrow \ell - \iota' \uparrow \ell$.
If $\forall \mathbf{x} \in \delta. (\mathbb{B}(\mathbf{x}) \sqsubseteq \ell) \implies l \sqsubseteq \ell$ and $\iota'(\mathbf{x}) = \iota_2(\mathbf{x}) - 1$. $\tau = \tau_1 :: \tau_2 :: \mathbf{n}^l$. $|\tau \uparrow \ell| = |\tau_1 \uparrow \ell| + |\tau_2 \uparrow \ell| + 1$. Thus, $|\tau \uparrow \ell| \leq \iota \uparrow \ell - \iota' \uparrow \ell$.
If $\delta = \emptyset$ then $\tau \uparrow \ell = \tau_1 \uparrow \ell :: \tau_2 \uparrow \ell$ and $\iota' = \iota_2$. Thus, $|\tau \uparrow \ell| \leq \iota \uparrow \ell - \iota' \uparrow \ell$
- $k_o \not\sqsubseteq \ell$: $l \not\sqsubseteq \ell$. Thus, $|\tau \uparrow \ell| = |\tau_1 \uparrow \ell| + |\tau_2 \uparrow \ell|$. As $\iota_2 = \iota'$, $|\tau \uparrow \ell| \leq \iota \uparrow \ell - \iota' \uparrow \ell$. If $\delta = \emptyset$ then $\tau \uparrow \ell = \tau_1 \uparrow \ell :: \tau_2 \uparrow \ell$ and $\iota' = \iota_2$. Thus, $|\tau \uparrow \ell| \leq \iota \uparrow \ell - \iota' \uparrow \ell$
- DCOPN: Similar to COP, $|(\tau_1 :: \tau_2) \uparrow \ell| \leq \iota \uparrow \ell - \iota_2 \uparrow \ell$. Thus, $|\tau \uparrow \ell| \leq \iota \uparrow \ell - \iota' \uparrow \ell$

□

Theorem 7 (LIR). If $\langle \sigma, \iota, c \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau \rangle$, then $|\tau \uparrow \ell| \leq \iota \uparrow \ell - \iota' \uparrow \ell$.

Proof. Induction on the derivation for commands and case analysis on the last rule.

- SKIP: $\iota = \iota'$ and $|\tau \uparrow \ell| = 0$.
- ASSN: From Lemma 18.
- SEQ: IH1: If $\langle \sigma, \iota, c_1 \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau_1 \rangle$, then $|\tau_1 \uparrow \ell| \leq \iota \uparrow \ell - \iota' \uparrow \ell$.
IH2: If $\langle \sigma', \iota', c_2 \rangle \Downarrow_{pc} \langle \sigma'', \iota'', \tau_2 \rangle$, then $|\tau_2 \uparrow \ell| \leq \iota' \uparrow \ell - \iota'' \uparrow \ell$.
From IH1 and IH2, $|\tau_1 \uparrow \ell| + |\tau_2 \uparrow \ell| \leq \iota \uparrow \ell - \iota'' \uparrow \ell$. Thus, $|\tau \uparrow \ell| \leq \iota \uparrow \ell - \iota'' \uparrow \ell$.
- IF-ELSE: From Lemma 18, $|\tau_1 \uparrow \ell| \leq \iota \uparrow \ell - \iota' \uparrow \ell$.
IH: If $\langle \sigma, \iota', c_i \rangle \Downarrow_{pc \sqcup k} \langle \sigma', \iota'', \tau_2 \rangle$, then $|\tau_2 \uparrow \ell| \leq \iota' \uparrow \ell - \iota'' \uparrow \ell$.
From IH and Lemma 18, $|\tau \uparrow \ell| \leq \iota \uparrow \ell - \iota'' \uparrow \ell$.
- WHILE-T: Similar to IF-ELSE
- WHILE-F: From Lemma 18.

□

Lemma 19. If $\langle \sigma, \iota, e, \tau \rangle \Downarrow_{pc} \mathbf{v}, \iota', \tau'$, then $\langle \sigma, \iota, e, \tau :: \tau^r \rangle \Downarrow_{pc} \mathbf{v}, \iota', \tau' :: \tau^r$

Proof. Proof by induction on the derivation for expressions and case analysis on the last rule.

- S-CONST and S-VAR: $\tau = \tau' = \epsilon$, thus $\tau :: \tau^r = \tau' :: \tau^r$.
- S-AOP and S-COP:
 - IH1: If $\langle \sigma, \iota, e_1, \tau \rangle \Downarrow_{pc} \mathbf{v}, \iota_1, \tau'$, then $\langle \sigma, \iota, e_1, \tau :: \tau^r \rangle \Downarrow_{pc} \mathbf{v}, \iota_1, \tau' :: \tau^r$
 - IH2: If $\langle \sigma, \iota, e_2, \tau' \rangle \Downarrow_{pc} \mathbf{v}, \iota_1, \tau''$, then $\langle \sigma, \iota, e, \tau' :: \tau^r \rangle \Downarrow_{pc} \mathbf{v}, \iota', \tau'' :: \tau^r$
 - T.S. If $\langle \sigma, \iota, e, \tau \rangle \Downarrow_{pc} \mathbf{v}, \iota', \tau''$, then $\langle \sigma, \iota, e, \tau :: \tau^r \rangle \Downarrow_{pc} \mathbf{v}, \iota', \tau'' :: \tau^r$. From IH1 and IH2.

- S-DCOPR: As above, if $\langle \sigma, \iota, e, \tau \rangle \downarrow_{pc} v, \iota', \tau''$, then $\langle \sigma, \iota, e, \tau :: \tau^r \rangle \downarrow_{pc} v, \iota', \tau'' :: \tau^r$. As v is the same, $\tau'' = \mathbf{n}^l :: \tau_1$ and $\tau'' :: \tau^r = \mathbf{n}^l :: \tau_1 :: \tau^r$.

Thus, if $\langle \sigma, \iota, e, \tau \rangle \downarrow_{pc} v, \iota', \tau_f$, then $\langle \sigma, \iota, e, \tau :: \tau^r \rangle \downarrow_{pc} v, \iota', \tau_f :: \tau^r$

- S-DCOPN: Similar to above case

□

Lemma 20. If $\langle \sigma, \iota, c, \tau \rangle \downarrow_{pc} \sigma', \iota', \tau'$, then $\langle \sigma, \iota, c, \tau :: \tau^r \rangle \downarrow_{pc} \sigma', \iota', \tau' :: \tau^r$

Proof. Proof by induction on the derivation for commands and case analysis on the last rule.

- S-SKIP: $\tau = \tau'$, thus $\tau :: \tau^r = \tau' :: \tau^r$.
- S-ASSN: From Lemma 19.
- S-SEQ: IH1: If $\langle \sigma, \iota, c, \tau \rangle \downarrow_{pc} \sigma', \iota', \tau_1$, then $\langle \sigma, \iota, c, \tau :: \tau^r \rangle \downarrow_{pc} \sigma', \iota', \tau_1 :: \tau^r$
IH2: If $\langle \sigma, \iota, c, \tau_1 \rangle \downarrow_{pc} \sigma', \iota', \tau'$, then $\langle \sigma, \iota, c, \tau_1 :: \tau^r \rangle \downarrow_{pc} \sigma', \iota', \tau' :: \tau^r$.
From IH1 and IH2.

- S-IF-ELSE-N and S-WHILE-T: Similar to above case
- S-IF-ELSE-S and S-WHILE-FS: From Lemma 19.

□

Corollary 5 (Trace Reduction). If $\langle \sigma, \iota, c, \tau \rangle \downarrow_{pc} \sigma', \iota', \epsilon$, then $\langle \sigma, \iota, c, \tau :: \tau^r \rangle \downarrow_{pc} \sigma', \iota', \tau^r$

Proof. From Lemma 20

□

Lemma 21 (Equivalence). If $\sigma \sim_\ell \sigma'$ and $\sigma \simeq_\ell \sigma''$, then $\sigma' \simeq_\ell \sigma''$

Proof. Consider $\mathbf{x} = \mathbf{n}_1^{(l, \delta)}$ in σ and $\mathbf{x} = \mathbf{n}_2^{(l', \delta')}$ in σ' and $\mathbf{x} = \mathbf{v}_s$ such that $\sigma(\mathbf{x}) \sim_\ell \sigma'(\mathbf{x})$. Thus, either:

- $\mathbf{n}_1 = \mathbf{n}_2 \wedge l = l' \wedge \delta = \delta'$: As $\sigma \simeq_\ell \sigma''$, $\mathbf{v}_s = \mathbf{n}_s^{(l_s, \delta_s)}$ and $\mathbf{n}_1 = \mathbf{n}_s \wedge l = l_s \wedge \delta = \delta_s$. Thus, $\mathbf{n}_2 = \mathbf{n}_s \wedge l' = l_s \wedge \delta' = \delta_s$. Hence, $\sigma' \simeq_\ell \sigma''$ from Definition 24.1
- $l = l' \wedge \delta = \delta' \wedge \exists y^m \in \delta.m \not\sqsubseteq l$: As $\sigma \simeq_\ell \sigma''$, $\mathbf{v}_s = \star^{(l_s, \delta_s)}$ and $l \sqsubseteq l$ and $l_s = \top$ and $\exists y^m \in \delta.\mathbb{B}(\mathbf{x}) \not\sqsubseteq l$. Thus, $l' \sqsubseteq l$ and $l_s = \top$ and $\exists y^m \in \delta'.\mathbb{B}(\mathbf{x}) \not\sqsubseteq l$. Else, $l = l_s \wedge \delta = \delta_s$. Thus, $l' = l_s \wedge \delta' = \delta_s$. Hence, $\sigma' \simeq_\ell \sigma''$ from Definition 24.2(b), 24.2(c)
- $l \not\sqsubseteq l \wedge l' \not\sqsubseteq l$: As $\sigma \simeq_\ell \sigma''$, $\mathbf{v}_s = \star^{(l_s, \delta_s)}$ and $l_s = \top$. Hence, $\sigma' \simeq_\ell \sigma''$ from Definition 24.2(a).

□

Lemma 22 (Expression Simulation). If $\langle \sigma, \iota, e \rangle \downarrow_{pc} v, \iota', \tau$, then

$\langle \sigma \uparrow_\ell, \iota \uparrow_\ell, e, \tau \uparrow_\ell \rangle \downarrow_{pc} v', \iota'', \tau'$ such that $\text{TTv} \simeq_\ell v', \iota'' \uparrow_\ell = \iota''$, and $\tau' = \epsilon$

Proof. Assume $\sigma^s = \sigma \uparrow_\ell$. Proof by induction on the derivation for expressions and case analysis on the last rule.

- **CONST:** $\mathbf{n}^\perp = \mathbf{n}^\perp$ and $\delta = \delta' = \{\}$. Thus, $\mathbf{v} \simeq_\ell \mathbf{v}'$. $\iota = \iota'$ and $\iota \uparrow_\ell = \iota''$, thus, $\iota' \uparrow_\ell = \iota''$. $\tau = \epsilon$, thus, $\tau \uparrow_\ell = \tau' = \epsilon$.
- **VAR:** As $\sigma \simeq_\ell \sigma^s$, $\sigma(\mathbf{x}) \simeq_\ell \sigma^s(\mathbf{x})$, thus $\mathbf{v} \simeq_\ell \mathbf{v}'$. $\iota = \iota'$ and $\iota \uparrow_\ell = \iota''$, thus, $\iota' \uparrow_\ell = \iota''$. $\tau = \epsilon$, thus, $\tau \uparrow_\ell = \tau' = \epsilon$.
- **AOP and COP:**
 IH1: If $\langle \sigma, \iota, e_1 \rangle \Downarrow_{pc} \mathbf{v}_1, \iota_1, \tau_1$ and $\sigma \simeq_\ell \sigma^s$, then $\langle \sigma^s, \iota \uparrow_\ell, e_1, \tau_1 \uparrow_\ell \rangle \Downarrow_{pc} \mathbf{v}'_1, \iota'_1, \tau'_1$ such that $\mathbf{v}_1 \simeq_\ell \mathbf{v}'_1$, $\iota_1 \uparrow_\ell = \iota'_1$ and $\tau'_1 = \epsilon$.
 IH2: If $\langle \sigma, \iota_1, e_2 \rangle \Downarrow_{pc} \mathbf{v}_2, \iota_2, \tau_2$ and $\sigma \simeq_\ell \sigma^s$, then $\langle \sigma^s, \iota_1 \uparrow_\ell, e_2, \tau_2 \uparrow_\ell \rangle \Downarrow_{pc} \mathbf{v}'_2, \iota'_2, \tau'_2$ such that $\mathbf{v}_2 \simeq_\ell \mathbf{v}'_2$, $\iota_2 \uparrow_\ell = \iota'_2$ and $\tau'_2 = \epsilon$.
 As $\sigma \simeq_\ell \sigma^s$, from IH1 and IH2 $\mathbf{v}_1 \simeq_\ell \mathbf{v}'_1$ and $\mathbf{v}_2 \simeq_\ell \mathbf{v}'_2$, where $\mathbf{v}_1 = \mathbf{n}_1^{(k_1, \delta_1)}$, $\mathbf{v}_2 = \mathbf{n}_2^{(k_2, \delta_2)}$, $\mathbf{v}'_1 = (\mathbf{n}_{s_1}^{(k'_1, \delta'_1)} \vee \star^{(k'_1, \delta'_1)})$, $\mathbf{v}'_2 = (\mathbf{n}_{s_2}^{(k'_2, \delta'_2)} \vee \star^{(k'_2, \delta'_2)})$.
 $\iota' = \iota_2$ and $\iota'' = \iota'_2$, thus from IH1 and IH2, $\iota' \uparrow_\ell = \iota''$.
 T.S: $\mathbf{v} \simeq_\ell \mathbf{v}'$, i.e, $\mathbf{v}_1 \odot \mathbf{v}_2 \simeq_\ell \mathbf{v}'_1 \odot \mathbf{v}'_2$.
 As $\mathbf{v}_1 \simeq_\ell \mathbf{v}'_1$, either:

1. $\mathbf{v}_1 = \mathbf{n}_1^{(k_1, \delta_1)}$, $\mathbf{v}'_1 = \mathbf{n}_{s_1}^{(k'_1, \delta'_1)}$, $\mathbf{n}_1 = \mathbf{n}_{s_1}$ and $k_1 = k'_1 \sqsubseteq \ell$ and $\delta_1 = \delta'_1$:

As $\mathbf{v}_2 \simeq_\ell \mathbf{v}'_2$, either:

- (a) $\mathbf{v}_2 = \mathbf{n}_2^{(k_2, \delta_2)}$, $\mathbf{v}'_2 = \mathbf{n}_{s_2}^{(k'_2, \delta'_2)}$, $\mathbf{n}_2 = \mathbf{n}_{s_2}$ and $k_2 = k'_2 \sqsubseteq \ell$ and $\delta_2 = \delta'_2$:
 $\mathbf{n}_1 \odot \mathbf{n}_2 = \mathbf{n}_{s_1} \odot \mathbf{n}_{s_2}$ and $(k_1 \sqcup k_2 = k'_1 \sqcup k'_2) \sqsubseteq \ell$ and $\delta_1 \cup \delta_2 = \delta'_1 \cup \delta'_2$. From Definition 24.1, $\mathbf{v} \simeq_\ell \mathbf{v}'$.
- (b) $\mathbf{v}_2 = \mathbf{n}_2^{(k_2, \delta_2)}$ and $\Gamma(\mathbf{v}_2) \not\sqsubseteq \ell$ and $\mathbf{v}'_2 = \star^{(k'_2, \delta'_2)}$:
 - $k'_2 = \top$ and $k_2 \not\sqsubseteq \ell$:
 $(l = k_1 \sqcup k_2) \not\sqsubseteq \ell$, thus $\Gamma(\mathbf{v}) \not\sqsubseteq \ell$, $\mathbf{n}_{s_1} \odot \star = \star$ and $l' = k'_1 \sqcup \top = \top$. From Definition 24.2(a), $\mathbf{v} \simeq_\ell \mathbf{v}'$.
 - $k'_2 = \top$, $k_2 \sqsubseteq \ell$ and $\exists \mathbf{x} \in \delta_2. (\mathbb{B}(\mathbf{x}) \not\sqsubseteq \ell)$:
 $(l = k_1 \sqcup k_2) \sqsubseteq \ell$, $(\Gamma(\mathbf{v}) = \Gamma(\mathbf{v}_1) \sqcup \Gamma(\mathbf{v}_2)) \not\sqsubseteq \ell$ and $l' = \top$. $(\mathbf{n}_{s_1} \odot \star = \star)$. $\delta = \delta_1 \cup \delta_2$, thus $\mathbf{x} \in \delta_2 \implies \mathbf{x} \in \delta$, i.e., $\exists \mathbf{x} \in \delta. (\mathbb{B}(\mathbf{x}) \not\sqsubseteq \ell)$. From Definition 24.2(b), $\mathbf{v} \simeq_\ell \mathbf{v}'$.
 - $(k_2 = k'_2) \sqsubseteq \ell \wedge \delta_2 = \delta'_2$:
 $(l = k_1 \sqcup k_2) \sqsubseteq \ell$ and $(\Gamma(\mathbf{v}) = \Gamma(\mathbf{v}_1) \sqcup \Gamma(\mathbf{v}_2)) \not\sqsubseteq \ell$. As $\delta_1 = \delta'_1$, $\delta = \delta_1 \cup \delta_2$, and $\delta' = \delta'_1 \cup \delta'_2$. $\delta = \delta'$. $k_1 = k'_1$ and $k_2 = k'_2$, so, $l = l'$. From Definition 24.2(c), $\mathbf{v} \simeq_\ell \mathbf{v}'$.

2. $\mathbf{v}_1 = \mathbf{n}_1^{(k_1, \delta_1)}$, $\Gamma(\mathbf{v}_1) \not\sqsubseteq \ell$, $\mathbf{v}'_1 = \star^{(k'_1, \delta'_1)}$:

As $\mathbf{v}_2 \simeq_\ell \mathbf{v}'_2$, either:

- $\mathbf{v}_2 = \mathbf{n}_2^{(k_2, \delta_2)}$, $\mathbf{v}'_2 = \mathbf{n}_{s_2}^{(k'_2, \delta'_2)}$, $\mathbf{n}_2 = \mathbf{n}_{s_2}$, $k_2 = k'_2 \sqsubseteq \ell$ and $\delta_2 = \delta'_2$:

Similar to the case (1.b) above

- $\mathbf{v}_2 = \mathbf{n}_2^{(k_2, \delta_2)}$, $\Gamma(\mathbf{v}_2) \not\sqsubseteq \ell$ and $\mathbf{v}'_2 = \star^{(k'_2, \delta'_2)}$:

* $k'_1 = \top$ and $k_1 \not\sqsubseteq \ell$:

$l' = k'_1 \sqcup k'_2 = \top$ and $(l = k_1 \sqcup k_2) \not\sqsubseteq \ell$. From Definition 24.2(a) $\mathbf{v} \simeq_\ell \mathbf{v}'$

- * $v'_1 = \star^{(k'_1, \delta'_1)}$, $k'_1 = \top$, $k_1 \sqsubseteq \ell$, and $\exists x \in \delta_1. (\mathbb{B}(x) \not\sqsubseteq \ell)$:
 $\delta = \delta_1 \cup \delta_2$ and $l' = k'_1 \sqcup k'_2 = \top$. As $l' = \top$ either:
 $(l = k_1 \sqcup k_2) \sqsubseteq \ell$. As $x \in \delta_1 \implies x \in \delta$. Thus, $\exists x \in \delta. (\mathbb{B}(x) \not\sqsubseteq \ell)$. From Definition 24.2(b), $v \simeq_\ell v'$ (or)
 $k_2 \not\sqsubseteq \ell$ and $(l = k_1 \sqcup k_2) \not\sqsubseteq \ell$. From Definition 24.2(a), $v \simeq_\ell v'$
- * $(k_1 = k'_1) \sqsubseteq \ell$ and $\delta_1 = \delta'_1$:
As $v_2 \simeq_\ell v'_2$, either:
 $k'_2 = \top$ and $k_2 \not\sqsubseteq \ell$. $l = k_1 \sqcup k_2$ and $l' = k'_1 \sqcup k'_2$. From Definition 24.2(a), $v \simeq_\ell v'$. (or)
 $k_2 \sqsubseteq \ell$, $k'_2 = \top$ and $\exists x \in \delta_2. (\mathbb{B}(x) \not\sqsubseteq \ell)$. $k_1 \sqcup k_2 \sqsubseteq \ell$, $k'_1 \sqcup k'_2 = \top$, $\delta = \delta_1 \cup \delta_2$ and $\exists x \in \delta. (\mathbb{B}(x) \not\sqsubseteq \ell)$. From Definition 24.2(b), $v \simeq_\ell v'$ (or)
 $(k_2 = k'_2) \sqsubseteq \ell$, $\delta_2 = \delta'_2$: $l = k_1 \sqcup k_2$ and $l' = k'_1 \sqcup k'_2$, so $l = l'$. $\delta = \delta_1 \cup \delta_2$ and $\delta' = \delta'_1 \cup \delta'_2$. Thus, $\delta = \delta'$. From Definition 24.2(c), $v \simeq_\ell v'$.

T.S.: If $\langle \sigma, \iota, e \rangle \Downarrow_{pc} \mathbf{n}^{(l, \delta)}$, ι', τ , and $\langle \sigma^s, \iota, e, \tau \uparrow \ell \rangle \Downarrow_{pc} n', \delta', \iota'', \tau'$, and $\sigma \simeq_\ell \sigma^s$, then $\tau' = \epsilon$. $\tau = \tau_1 :: \tau_2$.

$\tau \uparrow \ell = \tau_1 \uparrow \ell :: \tau_2 \uparrow \ell$. From AOP, COP and IH1, $\langle \sigma^s, \iota, e_1, \tau_1 \uparrow \ell :: \tau_2 \uparrow \ell \rangle \Downarrow_{pc} n'_1, \delta'_1, \iota'_1, \tau_2 \uparrow \ell$ because $\tau'_1 = \epsilon$. From AOP, COP and IH2, $\tau' = \epsilon$

- DCOPR: IH1: If $\langle \sigma, \iota, e_1 \rangle \Downarrow_{pc} v_1, \iota_1, \tau_1$ and $\sigma \simeq_\ell \sigma^s$, then $\langle \sigma^s, \iota \uparrow \ell, e_1, \tau_1 \uparrow \ell \rangle \Downarrow_{pc} v'_1, \iota'_1, \tau'_1$ such that $v_1 \simeq_\ell v'_1$, $\iota_1 \uparrow \ell = \iota'_1$ and $\tau'_1 = \epsilon$.
IH2: If $\langle \sigma, \iota_1, e_2 \rangle \Downarrow_{pc} v_2, \iota_2, \tau_2$ and $\sigma \simeq_\ell \sigma^s$, then $\langle \sigma^s, \iota_1 \uparrow \ell, e_2, \tau_2 \uparrow \ell \rangle \Downarrow_{pc} v'_2, \iota'_2, \tau'_2$ such that $v_2 \simeq_\ell v'_2$, $\iota_2 \uparrow \ell = \iota'_2$ and $\tau'_2 = \epsilon$.
As $\sigma \simeq_\ell \sigma^s$, from IH1 and IH2 $v_1 \simeq_\ell v'_1$, $v_2 \simeq_\ell v'_2$ and $\tau'_1 = \tau'_2 = \epsilon$.
Also from IH1 and IH2, $\iota_2 \uparrow \ell = \iota'_2$.
T.S: $v \simeq_\ell v'$, $\iota' \uparrow \ell = \iota''$ and $\tau' = \epsilon$.
 $v = v_1 \oplus v_2$ and $v' = v'_1 \oplus v'_2$.
As $v_1 \simeq_\ell v'_1$, either:

1. $v_1 = \mathbf{n}_1^{(k_1, \delta_1)}$, $v'_1 = \mathbf{n}_{s_1}^{(k'_1, \delta'_1)}$, $\mathbf{n}_1 = \mathbf{n}_{s_1}$, $k_1 = k'_1 \sqsubseteq \ell$ and $\delta_1 = \delta'_1$:

As $v_2 \simeq_\ell v'_2$, either:

- (a) $v_2 = \mathbf{n}_2^{(k_2, \delta_2)}$, $v'_2 = \mathbf{n}_{s_2}^{(k'_2, \delta'_2)}$, $\mathbf{n}_2 = \mathbf{n}_{s_2}$, $k_2 = k'_2 \sqsubseteq \ell$ and $\delta_2 = \delta'_2$:

Assume $\delta_i = \delta_1 \cup \delta_2$ and $\delta'_i = \delta'_1 \cup \delta'_2$, thus, $\delta_i = \delta'_i$. As $(l = k_1 \sqcup k_2 = k'_1 \sqcup k'_2 = l')$, $\delta_i = \delta'_i$ and $\iota_2 \uparrow \ell = \iota'_2$, $\delta_i = \delta'_i$. Either:

- $\delta_i = \delta'_i = \emptyset$: $n = \mathbf{n}_1 \oplus \mathbf{n}_2 = \mathbf{n}_{s_1} \oplus \mathbf{n}_{s_2} = \mathbf{n}_s$ and $(l = k_1 \sqcup k_2 = k'_1 \sqcup k'_2 = l') \sqsubseteq \ell$, thus, $v \simeq_\ell v'$.

Also, $\iota' = \iota_2$ and $\iota'' = \iota'_2$. Thus, $\iota' \uparrow \ell = \iota''$.

$\tau = \tau_1 :: \tau_2$. From s-DCOPR and IH1, $\langle \sigma^s, \iota \uparrow \ell, e_1, \tau_1 \uparrow \ell :: \tau_2 \uparrow \ell \rangle \Downarrow_{pc} v'_1, \iota'_1, \tau_2 \uparrow \ell$ because $\tau'_1 = \epsilon$. From s-DCOPR and IH2, $\tau' = \tau'_2 = \epsilon$.

- $(\delta_i = \delta'_i) \neq \emptyset$: $\tau = \tau_1 :: \tau_2 :: \mathbf{n}^l$. From s-DCOPR and IH1, $\langle \sigma^s, \iota \uparrow \ell, e_1, \tau_1 :: \tau_2 :: n \rangle \Downarrow_{pc} v'_1, \iota'_1, \tau_2$ because $\tau'_1 = \epsilon$. From s-DCOPR and IH2, $\langle \sigma^s, \iota_1 \uparrow \ell, e_2, \tau_2 :: n \rangle \Downarrow_{pc} v'_2, \iota'_2, n$ because $\tau'_2 = \epsilon$. From s-DCOPR, $\tau' = \epsilon$ and $\tau_2 = \mathbf{n}^l$. Thus, $v \simeq_\ell v'$.
As $\delta_i = \delta'_i$ and $\iota_2 \uparrow \ell = \iota'_2$, $\iota' \uparrow \ell = \iota''$.

(b) $v_2 = n_2^{(k_2, \delta_2)}$, $\Gamma(v_2) \not\sqsubseteq \ell$, $v'_2 = \star^{(k'_2, \delta'_2)}$ and either:

- $k_2 \not\sqsubseteq \ell$ and $k'_2 = \top$: ($k_o = k_1 \sqcup k_2$) $\not\sqsubseteq \ell$, thus, $l \not\sqsubseteq \ell$. As $l' = \top$ from s-DCOPR, $\delta'_i = \emptyset$ and $v' = \star^{(\top, \{\})}$. From Definition 24.2(a), $v \simeq_\ell v'$. From s-DCOPR, IH1, and IH2, $\iota_2 \uparrow \ell = \iota'_2$ and $\tau'_2 = \epsilon$. Thus, $\iota' \uparrow \ell = \iota''$ and $\tau' = \epsilon$.
- $k'_2 = \top$ and $k_2 \sqsubseteq \ell$ and $\exists x \in \delta_2. (\mathbb{B}(x) \not\sqsubseteq \ell)$:
 $v' = \star^{(\top, \delta')}$. ($k_o = k_1 \sqcup k_2$) $\sqsubseteq \ell$ and $\delta = \delta_1 \cup \delta_2$. As $\exists x \in \delta. (\mathbb{B}(x) \not\sqsubseteq \ell)$, $l \not\sqsubseteq \ell$. If $\iota_2(x) = 0$, then $l \sqsupseteq k_o \sqcup m$, else $l \sqsupseteq k_o \sqcup \mathbb{B}(x)$. In either case, $l \not\sqsubseteq \ell$. Also, $l = k'_1 \sqcup k'_2 = \top$ so $\delta'_i = \emptyset$. From Definition 24.2(a), $v \simeq_\ell v'$. From s-DCOPR, IH1, and IH2, $\iota_2 \uparrow \ell = \iota'_2$ and $\tau'_2 = \epsilon$. Thus, $\iota' \uparrow \ell = \iota''$ and $\tau' = \epsilon$.
- $k_2 = k'_2 \sqsubseteq \ell$ and $\delta_2 = \delta'_2$:
 $\delta = \delta_1 \cup \delta_2$ and $\delta' = \delta'_1 \cup \delta'_2$. As $\delta_1 = \delta'_1$, $\delta = \delta'$. From IH1, and IH2, $\iota_2 \uparrow \ell = \iota'_2$ and $\tau'_2 = \epsilon$. If $\exists x. \iota_2 \uparrow \ell(x) = \iota_2(x) = 0 \wedge \mathbb{L}(x) \not\sqsubseteq \ell$, then $l \not\sqsubseteq \ell$ and $l' = \top$. Else $l = l' \sqsubseteq \ell$. Thus, $\delta_i = \delta'_i$. Thus, if $\delta_i = \delta'_i = \emptyset$, $l' = \top$ and $l \not\sqsubseteq \ell$, $\iota' = \iota_2 \uparrow \ell = \iota'_2 = \iota''$ and $\tau' = \tau'_2 = \epsilon$. Else, $\tau = \tau_1 :: \tau_2 :: n^l$ and $\tau_2 \uparrow \ell = n^l :: \tau'$ and $v' = n^l$. Hence, $v \simeq_\ell v'$ and $\tau' = \epsilon$. As $\delta = \delta'$ and $\iota_2 \uparrow \ell = \iota'_2$, $\iota' \uparrow \ell = \iota''$

2. $v_1 = n_1^{(k_1, \delta_1)}$ and $\Gamma(v_1) \not\sqsubseteq \ell$ and $v'_1 = \star^{(k'_1, \delta'_1)}$:

As $v_2 \simeq_\ell v'_2$, either:

- (a) $v_2 = n_2^{(k_2, \delta_2)}$, $v'_2 = n_{s_2}^{(k'_2, \delta'_2)}$, $n_2 = n_{s_2}$, $k_2 = k'_2 \sqsubseteq \ell$ and $\delta_2 = \delta'_2$: Similar to case (1.b)
- (b) $v_2 = n_2^{(k_2, \delta_2)}$ and $\Gamma(v_2) \not\sqsubseteq \ell$ and $n'_2 = \star^{(k'_2, \delta'_2)}$:

Different cases of Definition 24.2:

- $k'_1 = \top$ and $k_1 \not\sqsubseteq \ell$: From DCOPR, $l \not\sqsubseteq \ell$ and from s-DCOPR $l' = \top$. $\forall x \in \delta. \mathbb{L}(x) \not\sqsubseteq \ell$, so $\iota_2 \uparrow \ell = \iota' \uparrow \ell$ and $\tau \uparrow \ell = \tau_1 \uparrow \ell :: \tau_2 \uparrow \ell$. As $\delta = \emptyset$, $\iota'_2 = \iota''$ and $\tau' = \tau'_2 = \epsilon$. Thus, $\iota'' = \iota' \uparrow \ell$ and from Definition 24.2(a), $v \simeq_\ell v'$.
- $k_1 \sqsubseteq \ell$ and $\exists x \in \delta_1. (\mathbb{B}(x) \not\sqsubseteq \ell)$ and $k'_1 = \top$: $\exists x \in \delta. (\mathbb{B}(x) \not\sqsubseteq \ell)$ as $\delta = \delta_1 \cup \delta_2$. From DCOPR, $l \not\sqsubseteq \ell$ and $l' = \top$. Similar to above case.
- $k_1 \sqsubseteq \ell$ and $\delta_1 = \delta'_1$:
 If $k'_2 = \top$ and either $k_2 \not\sqsubseteq \ell$ or $k_2 \sqsubseteq \ell$ and $\exists x \in \delta_2. (\mathbb{B}(x) \not\sqsubseteq \ell)$ - similar to above two cases.
 If $k_2 \sqsubseteq \ell$ and $\delta_2 = \delta'_2$: $\delta = \delta_1 \cup \delta_2$ and $\delta' = \delta'_1 \cup \delta'_2$ - similar to last case of 1.b

- DCOPN: Similar to AOP, COP.

□

Theorem 8 (Simulation Theorem). *If $\langle \sigma, \iota, c \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau \rangle$, then*

$\langle \sigma \uparrow \ell, \iota \uparrow \ell, c, \tau \uparrow \ell \rangle \Downarrow_{pc} \langle \sigma'', \iota'', \tau'' \rangle$ such that $\sigma' \simeq_\ell \sigma''$, $\iota' \uparrow \ell = \iota''$ and $\tau'' = \epsilon$

Proof. Assume $\sigma \uparrow \ell = \sigma^s$. Thus, $\sigma \simeq_\ell \sigma^s$. As ℓ simulates the run, initial $pc \sqsubseteq \ell$. Proof by induction on the derivation for commands and case analysis on the last rule.

- SKIP: $\sigma' = \sigma \simeq_\ell \sigma^s = \sigma''$. $\tau = \epsilon$. Thus $\tau \uparrow \ell = \tau'' = \epsilon$. $\iota = \iota'$ and $\iota \uparrow \ell = \iota''$. So, $\iota' \uparrow \ell = \iota''$

- ASSN: As $\sigma \simeq_\ell \sigma^s$, $\sigma(\mathbf{x}) \simeq_\ell \sigma^s(\mathbf{x})$. Thus, either $(\Gamma_f(\sigma(\mathbf{x})) = \Gamma_f(\sigma^s(\mathbf{x})))$ (or) $\Gamma_f(\sigma(\mathbf{x})) \not\sqsubseteq \ell$ and $\Gamma_f(\sigma^s(\mathbf{x})) = \top$. In either case, $pc \sqsubseteq \Gamma_f(\sigma(\mathbf{x}))$ and $pc \sqsubseteq \Gamma_f(\sigma^s(\mathbf{x}))$. From Lemma 22, $\mathbf{n}^{(m, \delta')} \simeq_\ell \mathbf{n}_s^{(m', \delta'')}$, $\iota' \uparrow_\ell = \iota''$ and $\tau'' = \epsilon$. From Definition 24, $\mathbf{n}^{(l, \delta)} \simeq_\ell \mathbf{n}_s^{(l_s, \delta_s)}$. Thus, $\sigma' \simeq_\ell \sigma''$, $\iota' \uparrow_\ell = \iota''$ and $\tau'' = \epsilon$.
- SEQ: IH1: If $\langle \sigma, \iota, c_1 \rangle \Downarrow_{pc} \langle \sigma_1, \iota_1, \tau_1 \rangle$, then $\langle \sigma \uparrow_\ell, \iota \uparrow_\ell, c_1, \tau_1 \uparrow_\ell \rangle \Downarrow_{pc} \langle \sigma'_1, \iota'_1, \tau'_1 \rangle$ such that $\sigma_1 \simeq_\ell \sigma'_1$, $\iota_1 \uparrow_\ell = \iota'_1$ and $\tau'_1 = \epsilon$
 IH2: If $\langle \sigma_1, \iota_1, c_2 \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau_2 \rangle$, then $\langle \sigma_1 \uparrow_\ell, \iota_1 \uparrow_\ell, c_2, \tau_2 \uparrow_\ell \rangle \Downarrow_{pc} \langle \sigma'', \iota'', \tau'' \rangle$ such that $\sigma' \simeq_\ell \sigma''$, $\iota' \uparrow_\ell = \iota''$ and $\tau'' = \epsilon$.
 T.S: If $\langle \sigma, \iota, c_1; c_2 \rangle \Downarrow_{pc} \langle \sigma', \iota', \tau_1 :: \tau_2 \rangle$, then $\langle \sigma \uparrow_\ell, \iota \uparrow_\ell, c_1; c_2, \tau_1 \uparrow_\ell :: \tau_2 \uparrow_\ell \rangle \Downarrow_{pc} \langle \sigma'', \iota'', \tau'' \rangle$ such that $\sigma' \simeq_\ell \sigma''$, $\iota' \uparrow_\ell = \iota''$ and $\tau'' = \epsilon$.
 From IH1, $\sigma'_1 = \sigma_1 \uparrow_\ell$ and $\iota'_1 \uparrow_\ell = \iota'_1$, thus, from IH2 $\sigma' \simeq_\ell \sigma''$ and $\iota' \uparrow_\ell = \iota''$.
 From IH1 $\tau'_1 = \epsilon$, and Corollary 5, $\langle \sigma \uparrow_\ell, \iota \uparrow_\ell, c_1, \tau_1 \uparrow_\ell :: \tau_2 \uparrow_\ell \rangle \Downarrow_{pc} \langle \sigma'', \iota'', \tau_2 \uparrow_\ell \rangle$. From IH2, $\tau'' = \epsilon$.
- IF-ELSE: From Lemma 22, $\mathbf{v} \simeq_\ell \mathbf{v}'$, $\iota_1 \uparrow_\ell = \iota'_1$, and $\tau'_1 = \epsilon$.
 IH1: If $\langle \sigma, \iota_1, c_i \rangle \Downarrow_{pc \sqcup k} \langle \sigma', \iota', \tau_2 \rangle$, then $\langle \sigma \uparrow_\ell, \iota_1 \uparrow_\ell, c_i, \tau_2 \uparrow_\ell \rangle \Downarrow_{pc \sqcup k} \langle \sigma'', \iota'', \tau'' \rangle$ such that $\sigma' \simeq_\ell \sigma''$, $\iota' \uparrow_\ell = \iota''$ and $\tau'' = \epsilon$.
 If $\mathbf{v}' = \mathbf{n}_s^{(l', \delta')}$, then from IH1 $\sigma' \simeq_\ell \sigma''$, $\iota' \uparrow_\ell = \iota''$ and $\tau'' = \epsilon$.
 Else from confinement lemma 16 and Lemma 21, $\sigma' \simeq_\ell \sigma''$, $\iota' \uparrow_\ell = \iota''$ and $\tau'' = \epsilon$.
- WHILE-T and WHILE-F: Similar to if and seq above

□

Lemma 23. *If $\sigma_1 \sim_\ell \sigma_2$ then $\sigma_1 \uparrow_\ell = \sigma_2 \uparrow_\ell$*

Proof. From Definition 19, $\sigma_1 \sim_\ell \sigma_2 \implies \forall \mathbf{x}. \sigma_1(\mathbf{x}) \sim_\ell \sigma_2(\mathbf{x})$. Assume $\sigma_1(\mathbf{x}) = \mathbf{n}_1^{(k_1, \delta_1)}$ and $\sigma_2(\mathbf{x}) = \mathbf{n}_2^{(k_2, \delta_2)}$. From Definition 18, either:

- $\mathbf{n}_1 = \mathbf{n}_2$, $k_1 = k_2 \sqsubseteq \ell$ and $\delta_1 = \delta_2$: By definition of $\sigma \uparrow_\ell$, $\sigma_1 \uparrow_\ell(\mathbf{x}) = \mathbf{n}_1^{(k_1, \delta_1)}$ and $\sigma_2 \uparrow_\ell(\mathbf{x}) = \mathbf{n}_2^{(k_2, \delta_2)}$. Hence, $\sigma_1 \uparrow_\ell(\mathbf{x}) = \sigma_2 \uparrow_\ell(\mathbf{x})$
- $\Gamma(\mathbf{v}_1) \not\sqsubseteq \ell$ and $\Gamma(\mathbf{v}_2) \not\sqsubseteq \ell$ such that $k_1 = k_2 \sqsubseteq \ell$ and $\delta_1 = \delta_2$: By definition of $\sigma \uparrow_\ell$, $\sigma_1 \uparrow_\ell(\mathbf{x}) = \star^{(k_1, \delta_1)}$ and $\sigma_2 \uparrow_\ell(\mathbf{x}) = \star^{(k_2, \delta_2)}$. Hence, $\sigma_1 \uparrow_\ell(\mathbf{x}) = \sigma_2 \uparrow_\ell(\mathbf{x})$
- $\Gamma(\mathbf{v}_1) \not\sqsubseteq \ell$ and $\Gamma(\mathbf{v}_2) \not\sqsubseteq \ell$ such that $k_1 \not\sqsubseteq \ell$ and $k_2 \not\sqsubseteq \ell$: By definition of $\sigma \uparrow_\ell$, $\sigma_1 \uparrow_\ell(\mathbf{x}) = \star^{(\top, \{\})}$ and $\sigma_2 \uparrow_\ell(\mathbf{x}) = \star^{(\top, \{\})}$. Hence, $\sigma_1 \uparrow_\ell(\mathbf{x}) = \sigma_2 \uparrow_\ell(\mathbf{x})$

□

Lemma 24. *If $\iota_1 \sim_\ell \iota_2$ then $\iota_1 \uparrow_\ell = \iota_2 \uparrow_\ell$*

Proof. From Definition 20, $\iota_1 \sim_\ell \iota_2 \implies \forall \mathbf{x}. \mathbb{B}(\mathbf{x}) \sqsubseteq \ell \implies \iota_1(\mathbf{x}) = \iota_2(\mathbf{x})$. By definition of $\iota \uparrow_\ell$, if $(\mathbb{B}(\mathbf{x}) \not\sqsubseteq \ell$ then $\iota_1 \uparrow_\ell(\mathbf{x}) = 0$ else $\iota_1 \uparrow_\ell(\mathbf{x}) = \iota_1(\mathbf{x}))$ and if $(\mathbb{B}(\mathbf{x}) \not\sqsubseteq \ell$ then $\iota_2 \uparrow_\ell(\mathbf{x}) = 0$ else $\iota_2 \uparrow_\ell(\mathbf{x}) = \iota_2(\mathbf{x}))$. Thus, if $\mathbb{B}(\mathbf{x}) \sqsubseteq \ell$, then $\iota_1 \uparrow_\ell(\mathbf{x}) = \iota_2 \uparrow_\ell(\mathbf{x})$ else $\iota_1 \uparrow_\ell(\mathbf{x}) = \iota_2 \uparrow_\ell(\mathbf{x}) = 0$. Thus, $\iota_1 \uparrow_\ell = \iota_2 \uparrow_\ell$. □

Lemma 25. *If $\sigma_1 \simeq_\ell \sigma'$ and $\sigma_2 \simeq_\ell \sigma'$, then $\sigma_1 \sim_\ell \sigma_2$*

Proof. From Definition 25, $\sigma_1 \simeq_\ell \sigma' \implies \forall \mathbf{x}. \sigma_1(\mathbf{x}) \simeq_\ell \sigma'(\mathbf{x})$ and $\sigma_2 \simeq_\ell \sigma' \implies \forall \mathbf{x}. \sigma_2(\mathbf{x}) \simeq_\ell \sigma'(\mathbf{x})$. Assume $\sigma_1(\mathbf{x}) = \mathbf{n}_1^{(k_1, \delta_1)}$, $\sigma_2(\mathbf{x}) = \mathbf{n}_2^{(k_2, \delta_2)}$ and $\sigma'(\mathbf{x}) = \mathbf{v}^{(l', \delta')}$.

T.S.: $\forall \mathbf{x}. \sigma_1(\mathbf{x}) \sim_\ell \sigma_2(\mathbf{x})$ From Definition 24 for $\sigma_1(\mathbf{x}) \simeq_\ell \sigma'(\mathbf{x})$, either:

- $\mathbf{v} = \mathbf{n}_s$ and $\mathbf{n}_1 = \mathbf{n}_s$ and $(k_1 = l') \sqsubseteq \ell$ and $\delta_1 = \delta'$: As $\sigma_2(\mathbf{x}) \simeq_\ell \sigma'(\mathbf{x})$, by Definition 24.1 $\mathbf{n}_2 = \mathbf{n}_s$ and $(k_2, \delta_2) = (l', \delta')$. Thus, $\sigma_1(\mathbf{x}) = \sigma_2(\mathbf{x})$.
- $\Gamma(\sigma_1(\mathbf{x})) \not\sqsubseteq \ell$, $\mathbf{v} = \star$ and either:
 1. $k_1 \not\sqsubseteq \ell \wedge l' = \top$: Either $k_2 \not\sqsubseteq \ell$ or $k_2 \sqsubseteq \ell \wedge \exists \mathbf{y} \in \delta. (\mathbb{B}(\mathbf{y}) \not\sqsubseteq \ell)$: From Definition 18.
 2. $k_1 \sqsubseteq \ell \wedge \exists \mathbf{y} \in \delta. (\mathbb{B}(\mathbf{y}) \not\sqsubseteq \ell) \wedge l' = \top$: Similar to above case
 3. $(k_1 = l') \sqsubseteq \ell \wedge \delta_1 = \delta'$: As $l' \sqsubseteq \ell$, $k_2 \sqsubseteq \ell \wedge \delta_2 = \delta'$. From Definition 18.(2).

□

Lemma 26. *If $\langle \sigma_1, \iota_1, c \rangle \Downarrow_{pc} \langle \sigma'_1, \iota'_1, \tau_1 \rangle$, $\langle \sigma_2, \iota_2, c \rangle \Downarrow_{pc} \langle \sigma'_2, \iota'_2, \tau_2 \rangle$, $\sigma_1 \sim_\ell \sigma_2$, $\iota_1 \sim_\ell \iota_2$, and $\tau_1 \sim_\ell \tau_2$, then $\sigma'_1 \sim_\ell \sigma'_2$.*

Proof. From Theorem 8, if $\langle \sigma_1, \iota_1, c \rangle \Downarrow_{pc} \langle \sigma'_1, \iota'_1, \tau_1 \rangle$,

then $\langle \sigma_1 \uparrow \ell, \iota_1 \uparrow \ell, c, \tau_1 \uparrow \ell \rangle \Downarrow_{pc} \langle \sigma''_1, \iota''_1, \tau''_1 \rangle$ such that $\sigma'_1 \simeq_\ell \sigma''_1$, $\iota'_1 \uparrow \ell = \iota''_1$ and $\tau''_1 = \epsilon$ and

if $\langle \sigma_2, \iota_2, c \rangle \Downarrow_{pc} \langle \sigma'_2, \iota'_2, \tau_2 \rangle$, then

$\langle \sigma_2 \uparrow \ell, \iota_2 \uparrow \ell, c, \tau_2 \uparrow \ell \rangle \Downarrow_{pc} \langle \sigma''_2, \iota''_2, \tau''_2 \rangle$ such that $\sigma'_2 \simeq_\ell \sigma''_2$, $\iota'_2 \uparrow \ell = \iota''_2$ and $\tau''_2 = \epsilon$.

$\sigma_1 \sim_\ell \sigma_2 \implies \sigma_1 \uparrow \ell = \sigma_2 \uparrow \ell$. $\iota_1 \sim_\ell \iota_2 \implies \iota_1 \uparrow \ell = \iota_2 \uparrow \ell$.

As $\tau_1 \sim_\ell \tau_2$, from Definition 21 $\tau_1 \uparrow \ell = \tau_2 \uparrow \ell$.

Thus, $\sigma''_1 = \sigma''_2$, $\iota''_1 = \iota''_2$ and $\tau''_1 = \tau''_2$.

$\sigma'_1 \simeq_\ell \sigma''_1$, and $\sigma'_2 \simeq_\ell \sigma''_2$, i.e., $\sigma'_2 \simeq_\ell \sigma''_1$. From Lemma 25. □

Corollary 6 (Non-interference). *If $\langle \sigma_1, \iota_1, c \rangle \Downarrow_{pc} \langle \sigma'_1, \iota'_1, \tau_1 \rangle$, $\langle \sigma_2, \iota_2, c \rangle \Downarrow_{pc} \langle \sigma'_2, \iota'_2, \tau_2 \rangle$, $\sigma_1 \sim_\ell \sigma_2$, and $\iota_1 \uparrow \ell = \iota_2 \uparrow \ell = 0$, then $\sigma'_1 \sim_\ell \sigma'_2$.*

Proof. From Theorem 7 and Lemma 26 □

List of Figures

3.1	Syntax of the Language	16
3.2	Semantics	16
3.3	Assignment rule for NSU	17
3.4	Syntax of labels including the partially-leaked label P	19
4.1	Labels and label operations	24
4.2	Caption	25
4.3	Lattice explaining rule $ASSN-s$	26
4.4	A powerset/product lattice	29
5.1	Language Syntax	37
5.2	Semantics	38
6.1	Syntax of the language	48
6.2	LIR - Semantics of expressions	55
6.3	LIR - Semantics of commands	56
6.4	Decoding semantics of expressions	60
6.5	Decoding semantics of commands	61
7.1	Work-flow of the WebPol policy model	68
8.1	Overheads of LIR on SunSpider Benchmark Tests	79

List of Tables

4.1	Execution steps in two runs of the program from Listing 4.1, with two variants of the rule <code>ASSN-S</code>	27
8.1	Performance of examples from Section 7.3. All time in ms. The numbers in parenthesis are additional overheads relative to Base	77
8.2	Performance on two real-world websites. All time in ms. The numbers in parenthesis are additional overheads relative to Base	79
A.1	Examples for all possible transitions of low-equivalent to low-equivalent values . . .	96

Bibliography

- [1] Google Caja - A source-to-source translator for securing JavaScript-based web content. <https://developers.google.com/caja/>. [Online; accessed 25-Apr-2017].
- [2] Cross-Origin Resource Sharing. <http://www.w3.org/TR/cors/>. [Online; accessed 25-Apr-2017].
- [3] Content Security Policy Level 3. <https://www.w3.org/TR/CSP3/>. [Online; accessed 25-Apr-2017].
- [4] Facebook. FBJS. <https://developers.facebook.com/docs/javascript>. [Online; accessed 25-Apr-2017].
- [5] World's Biggest Data Breaches. <http://www.informationisbeautiful.net/visualizations/worlds-biggest-data-breaches-hacks/>, 2017. [Online; accessed 25-Apr-2017].
- [6] P. Agten, S. Van Acker, Y. Brondsema, P. H. Phung, L. Desmet, and F. Piessens. JSand: Complete client-side sandboxing of third-party JavaScript without browser modifications. In *Proc. Annual Computer Security Applications Conference*, pages 1–10, 2012.
- [7] M. S. Alvim, K. Chatzikokolakis, C. Palamidessi, and G. Smith. Measuring information leakage using generalized gain functions. In *Proc. 2012 IEEE 25th Computer Security Foundations Symposium, CSF '12*, pages 265–279, 2012.
- [8] A. Askarov and A. Sabelfeld. Gradual release: Unifying declassification, encryption and key release policies. In *Proc. 2007 IEEE Symposium on Security and Privacy (SP '07)*, pages 207–221, 2007.
- [9] A. Askarov and A. Sabelfeld. Localized delimited release: Combining the what and where dimensions of information release. In *Proc. 2007 Workshop on Programming Languages and Analysis for Security, PLAS '07*, pages 53–60, 2007.
- [10] A. Askarov and A. Sabelfeld. Tight enforcement of information-release policies for dynamic languages. In *Proc. IEEE Computer Security Foundations Symposium*, pages 43–59, 2009.

- [11] A. Askarov and A. Sabelfeld. Catch me if you can: permissive yet secure error handling. In *Proc. ACM SIGPLAN Fourth Workshop on Programming Languages and Analysis for Security*, pages 45–57, 2009.
- [12] A. Askarov, S. Hunt, A. Sabelfeld, and D. Sands. Termination-insensitive noninterference leaks more than just a bit. In *Proc. European Symposium on Research in Computer Security*, pages 333–348, 2008.
- [13] T. H. Austin and C. Flanagan. Efficient purely-dynamic information flow analysis. In *Proc. ACM SIGPLAN Workshop on Programming Languages and Analysis for Security*, pages 113–124, 2009.
- [14] T. H. Austin and C. Flanagan. Permissive dynamic information flow analysis. In *Proc. ACM SIGPLAN Workshop on Programming Languages and Analysis for Security*, pages 1–12, 2010.
- [15] T. H. Austin and C. Flanagan. Multiple facets for dynamic information flow. In *Proc. ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 165–178, 2012.
- [16] M. Backes, B. Kopf, and A. Rybalchenko. Automatic discovery and quantification of information leaks. In *Proc. 2009 30th IEEE Symposium on Security and Privacy*, pages 141–153, 2009.
- [17] A. Banerjee, D. A. Naumann, and S. Rosenberg. Expressive declassification policies and modular static enforcement. In *Proc. 2008 IEEE Symposium on Security and Privacy*, SP '08, pages 339–353, 2008.
- [18] T. Bao, Y. Zheng, Z. Lin, X. Zhang, and D. Xu. Strict control dependence and its effect on dynamic information flow analyses. In *Proc. 19th International Symposium on Software Testing and Analysis*, ISSTA '10, pages 13–24, 2010.
- [19] A. Barth. The web origin concept. <http://tools.ietf.org/html/rfc6454>. [Online; accessed 25-Apr-2017].
- [20] F. Besson, N. Bielova, and T. Jensen. Hybrid information flow monitoring against web tracking. In *Proc. IEEE 26th Computer Security Foundations Symposium (CSF '13)*, pages 240–254, 2013.
- [21] A. Bichhawat, V. Rajani, D. Garg, and C. Hammer. Generalizing permissive-upgrade in dynamic information flow analysis. In *Proc. 9th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security*, pages 15:15–15:24, 2014.
- [22] A. Bichhawat, V. Rajani, D. Garg, and C. Hammer. Information flow control in WebKit's JavaScript bytecode. In *Proc. Principles of Security and Trust*, pages 159–178. Springer Berlin Heidelberg, 2014.
- [23] A. Bichhawat, V. Rajani, D. Garg, and C. Hammer. Information Flow Control in WebKit's JavaScript Bytecode. *CoRR*, abs/1401.4339, 2014. URL <http://arxiv.org/abs/1401.4339>.
- [24] A. Bichhawat, V. Rajani, J. Jain, D. Garg, and C. Hammer. WebPol: Fine-grained Information Flow Policies for Web Browsers. In *Proc. 22nd European Symposium on Research in Computer Security (ESORICS '17)*, 2017.

- [25] N. Bielova. Dynamic leakage - a need for a new quantitative information flow measure. In *Proc. 11th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security*, 2016.
- [26] N. Bielova, D. Devriese, F. Massacci, and F. Piessens. Reactive non-interference for a browser model. In *Proc. International Conference on Network and System Security*, pages 97–104, 2011.
- [27] A. Birgisson, D. Hedin, and A. Sabelfeld. Boosting the permissiveness of dynamic information-flow tracking by testing. In *Proc. European Symposium on Research in Computer Security (ESORICS)*, pages 55–72, 2012.
- [28] P. Buiras, D. Stefan, and A. Russo. On dynamic flow-sensitive floating-label systems. In *Proc. 2014 IEEE 27th Computer Security Foundations Symposium, CSF '14*, pages 65–79. IEEE Computer Society, 2014.
- [29] S. Chong. Required information release. In *Proc. 2010 23rd IEEE Computer Security Foundations Symposium*, pages 215–227, July 2010.
- [30] S. Chong and A. C. Myers. Security policies for downgrading. In *Proc. 11th ACM Conference on Computer and Communications Security, CCS '04*, pages 198–209, 2004.
- [31] A. Chudnov and D. A. Naumann. Inlined information flow monitoring for javascript. In *Proc. 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, pages 629–643, 2015.
- [32] R. Chugh, J. A. Meister, R. Jhala, and S. Lerner. Staged information flow for JavaScript. In *Proc. ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 50–62, 2009.
- [33] D. Clark, S. Hunt, and P. Malacaria. A static analysis for quantifying information flow in a simple imperative language. *J. Comput. Secur.*, 15(3):321–371, 2007.
- [34] M. R. Clarkson, A. C. Myers, and F. B. Schneider. Quantifying information flow with beliefs. *J. Comput. Secur.*, 17(5):655–701, Oct. 2009.
- [35] E. Cohen. Information transmission in computational systems. In *Proc. Sixth ACM Symposium on Operating Systems Principles, SOSP '77*, pages 133–139, 1977.
- [36] E. Cohen. *Information Transmission in Sequential Programs*. Foundations of Secure Computation. Academic Press, 1978.
- [37] T. M. Cover and J. A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006.
- [38] D. Crockford. ADsafe. <http://adsafe.org/>. [Online; accessed 25-Apr-2017].
- [39] W. De Groef, D. Devriese, N. Nikiforakis, and F. Piessens. Flowfox: a web browser with flexible and precise information flow control. In *Proc. ACM Conference on Computer and Communications Security*, pages 748–759, 2012.
- [40] D. E. Denning. A lattice model of secure information flow. *Commun. ACM*, 19(5):236–243, May 1976.

- [41] D. E. Denning and P. J. Denning. Certification of programs for secure information flow. *Communications of the ACM*, 20(7):504–513, July 1977.
- [42] D. E. R. Denning. *Cryptography and Data Security*. Addison-Wesley Longman Publishing Co., Inc., 1982.
- [43] D. Devriese and F. Piessens. Noninterference through secure multi-execution. In *Proc. IEEE Symposium on Security and Privacy*, pages 109–124, 2010.
- [44] M. Dhawan and V. Ganapathy. Analyzing information flow in JavaScript-based browser extensions. In *Proc. Annual Computer Security Applications Conference*, pages 382–391, 2009.
- [45] X. Dong, Z. Chen, H. Siadati, S. Tople, P. Saxena, and Z. Liang. Protecting sensitive web content from client-side vulnerabilities with cryptons. In *Proc. 2013 ACM SIGSAC Conference on Computer and Communications Security*, pages 1311–1324, 2013.
- [46] J. S. Fenton. Memoryless subsystems. *The Computer Journal*, 17(2):143, 1974.
- [47] O. Foundation. Information leakage. https://www.owasp.org/index.php/Information_Leakage, June 2013.
- [48] R. Giacobazzi and I. Mastroeni. Abstract non-interference: Parameterizing non-interference by abstract interpretation. In *Proc. 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '04*, pages 186–197, 2004.
- [49] J. A. Goguen and J. Meseguer. Security policies and security models. In *Proc. IEEE Symposium on Security and Privacy*, pages 11–20, 1982.
- [50] S. Guarnieri, M. Pistoia, O. Tripp, J. Dolby, S. Teilhet, and R. Berg. Saving the world wide web from vulnerable javascript. In *Proc. 2011 International Symposium on Software Testing and Analysis*, pages 177–187, 2011.
- [51] G. L. Guernic. Automaton-based confidentiality monitoring of concurrent programs. In *Proc. IEEE Computer Security Foundations Symposium*, pages 218–232, 2007.
- [52] G. L. Guernic, A. Banerjee, T. Jensen, and D. A. Schmidt. Automata-based confidentiality monitoring. In *Proc. Asian Computing Science Conference on Secure Software*, pages 75–89, 2006.
- [53] C. Hammer and G. Snelting. Flow-sensitive, context-sensitive, and object-sensitive information flow control based on program dependence graphs. *International Journal of Information Security*, 8(6):399–422, Dec. 2009.
- [54] D. Hedin and A. Sabelfeld. Information-flow security for a core of JavaScript. In *Proc. IEEE Computer Security Foundations Symposium*, pages 3–18, 2012.
- [55] D. Hedin, A. Birgisson, L. Bello, and A. Sabelfeld. JSFlow: Tracking information flow in JavaScript and its APIs. In *Proc. ACM Symposium on Applied Computing*, pages 1663–1671, 2014.
- [56] D. Hedin, L. Bello, and A. Sabelfeld. Value-sensitive hybrid information flow control for a javascript-like language. In *Proc. 2015 IEEE 28th Computer Security Foundations Symposium, CSF '15*, pages 351–365, 2015.

- [57] C. Hritcu, M. Greenberg, B. Karel, B. C. Pierce, and G. Morrisett. All your IFCEException are belong to us. In *Proc. IEEE Symposium on Security and Privacy (Oakland)*, pages 3–17, 2013.
- [58] S. Hunt and D. Sands. On flow-sensitive security types. In *Proc. ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 79–90, 2006.
- [59] D. Jang, R. Jhala, S. Lerner, and H. Shacham. An empirical study of privacy-violating information flows in JavaScript web applications. In *Proc. ACM Conference on Computer and Communications Security*, pages 270–283, 2010.
- [60] S. Just, A. Cleary, B. Shirley, and C. Hammer. Information flow analysis for JavaScript. In *Proc. ACM SIGPLAN International Workshop on Programming Language and Systems Technologies for Internet Clients*, pages 9–18, 2011.
- [61] M. G. Kang, S. McCamant, P. Poosankam, , and D. Song. Dta++: Dynamic taint analysis with targeted control-flow propagation. In *Proc. 18th Network and Distributed System Security Symposium (NDSS '11)*, 2011.
- [62] C. Kerschbaumer, E. Hennigan, P. Larsen, S. Brunthaler, and M. Franz. Towards precise and efficient information flow control in web browsers. In *Proc. Trust and Trustworthy Computing*, pages 187–195. 2013.
- [63] D. King, B. Hicks, M. Hicks, and T. Jaeger. Implicit flows: Can't live with 'em, can't live without 'em. In R. Sekar and A. K. Pujari, editors, *Proc. International Conference on Information Systems Security (ICISS)*, volume 5352 of *Lecture Notes in Computer Science*, pages 56–70. Springer, Dec. 2008.
- [64] B. Kopf and A. Rybalchenko. Approximation and randomization for quantitative information-flow analysis. In *Computer Security Foundations Symposium (CSF), 2010 23rd IEEE*, pages 3–14, July 2010.
- [65] T. Lengauer and R. E. Tarjan. A fast algorithm for finding dominators in a flowgraph. *ACM Trans. Program. Lang. Syst.*, 1(1):121–141, Jan. 1979.
- [66] P. Li and S. Zdancewic. Downgrading policies and relaxed noninterference. In *Proc. 32Nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '05*, pages 158–170, 2005.
- [67] Z. Li, K. Zhang, and X. Wang. Mash-if: Practical information-flow control within client-side mashups. In *Proc. 2010 IEEE/IFIP International Conference on Dependable Systems Networks (DSN)*, pages 251–260, June 2010.
- [68] M. T. Louw, K. T. Ganesh, and V. N. Venkatakrishnan. Adjail: Practical enforcement of confidentiality and integrity policies on web advertisements. In *Proc. 19th USENIX Conference on Security*, pages 24–24, 2010.
- [69] A. Lux and H. Mantel. Declassification with explicit reference points. In *Proc. 14th European Conference on Research in Computer Security, ESORICS'09*, pages 69–85, 2009.

- [70] W. Masri and A. Podgurski. Algorithms and tool support for dynamic information flow analysis. *Information & Software Technology*, 51(2):385–404, 2009.
- [71] J. L. Massey. Guessing and entropy. In *In Proc. 1994 IEEE International Symposium on Information Theory*, page 204, 1994.
- [72] S. McCamant and M. D. Ernst. A simulation-based proof technique for dynamic information flow. In *ACM SIGPLAN Workshop on Programming Languages and Analysis for Security, PLAS '07*, pages 41–46, June 2007.
- [73] S. McCamant and M. D. Ernst. Quantitative information flow as network flow capacity. In *Proc. 29th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '08*, pages 193–205, 2008.
- [74] L. A. Meyerovich and B. Livshits. Conscript: Specifying and enforcing fine-grained security policies for javascript in the browser. In *Proc. 2010 IEEE Symposium on Security and Privacy*, pages 481–496, 2010.
- [75] M. Miller. *Robust composition: Towards a unified approach to access control and concurrency control*. PhD thesis, Johns Hopkins University, 2006.
- [76] A. C. Myers. JFlow: practical mostly-static information flow control. In *Proc. 26th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, POPL '99*, pages 228–241, 1999.
- [77] F. Nentwich, N. Jovanovic, E. Kirda, C. Kruegel, and G. Vigna. Cross-site scripting prevention with dynamic data tainting and static analysis. In *Proc. Network and Distributed System Security Symposium*, 2007.
- [78] N. Nikiforakis, L. Invernizzi, A. Kapravelos, S. Van Acker, W. Joosen, C. Kruegel, F. Piessens, and G. Vigna. You are what you include: Large-scale evaluation of remote javascript inclusions. In *Proc. 2012 ACM Conference on Computer and Communications Security, CCS '12*, pages 736–747, 2012.
- [79] F. Pottier and V. Simonet. Information flow inference for ml. *ACM Trans. Program. Lang. Syst.*, 25(1):117–158, Jan. 2003.
- [80] W. Rafnsson and A. Sabelfeld. Secure multi-execution: Fine-grained, declassification-aware, and transparent. In *Proc. 2013 IEEE 26th Computer Security Foundations Symposium*, pages 33–48, 2013.
- [81] V. Rajani, A. Bichhawat, D. Garg, and C. Hammer. Information flow control for event handling and the DOM in web browsers. In *Proc. IEEE 28th Computer Security Foundations Symposium (CSF '15)*, pages 366–379, 2015.
- [82] G. Richards, C. Hammer, B. Burg, and J. Vitek. The eval that men do – a large-scale study of the use of eval in JavaScript applications. In *Proc. European Conference on Object-Oriented Programming*, pages 52–78, 2011.

- [83] G. Richards, C. Hammer, F. Zappa Nardelli, S. Jagannathan, and J. Vitek. Flexible access control for javascript. In *Proc. 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA '13*, pages 305–322, 2013.
- [84] B. P. S. Rocha, S. Bandhakavi, J. d. Hartog, W. H. Winsborough, and S. Etalle. Towards static flow-based declassification for legacy and untrusted programs. In *Proc. 2010 IEEE Symposium on Security and Privacy, SP '10*, pages 93–108, 2010.
- [85] A. Russo and A. Sabelfeld. Dynamic vs. static flow-sensitive security analysis. In *Proc. 2010 IEEE 23rd Computer Security Foundations Symposium*, pages 186–199, 2010.
- [86] A. Russo, A. Sabelfeld, and K. Li. Implicit flows in malicious and nonmalicious code. In *Proc. 2009 Marktoberdorf Summer School, 2009*.
- [87] A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1), Jan 2003.
- [88] A. Sabelfeld and A. C. Myers. A model for delimited information release. In K. Futatsugi, F. Mizoguchi, and N. Yonezaki, editors, *Software Security - Theories and Systems*, volume 3233 of *LNCS*, pages 174–191. Springer Berlin Heidelberg, 2004.
- [89] A. Sabelfeld and A. Russo. From dynamic to static and back: Riding the roller coaster of information-flow control research. In *Proc. Perspectives of Systems Informatics*, pages 352–365, 2010.
- [90] A. Sabelfeld and D. Sands. A per model of secure information flow in sequential programs. *Higher-Order and Symbolic Computation*, 14(1):59–91, 2001.
- [91] A. Sabelfeld and D. Sands. Declassification: Dimensions and principles. *J. Comput. Secur.*, 17(5):517–548, Oct. 2009.
- [92] C. E. Shannon. A mathematical theory of communication. *SIGMOBILE Mob. Comput. Commun. Rev.*, 5(1):3–55, Jan. 2001.
- [93] G. Smith. *Foundations of Software Science and Computational Structures: 12th International Conference, FOSSACS 2009*, chapter On the Foundations of Quantitative Information Flow, pages 288–302. Springer, 2009.
- [94] D. Stefan, E. Z. Yang, P. Marchenko, A. Russo, D. Herman, B. Karp, and D. Mazières. Protecting users by confining JavaScript with COWL. In *Proc. USENIX Symposium on Operating Systems Design and Implementation*, pages 131–146, 2014.
- [95] S. Van Acker, P. De Ryck, L. Desmet, F. Piessens, and W. Joosen. Webjail: Least-privilege integration of third-party components in web mashups. In *Proc. 27th Annual Computer Security Applications Conference*, pages 307–316, 2011.
- [96] T. Van Goethem, W. Joosen, and N. Nikiforakis. The clock is still ticking: Timing attacks in the modern web. In *Proc. 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, pages 1382–1393, 2015.

- [97] M. Vanhoef, W. De Groef, D. Devriese, F. Piessens, and T. Rezk. Stateful declassification policies for event-driven programs. In *Proc. 2014 IEEE 27th Computer Security Foundations Symposium*, pages 293–307, 2014.
- [98] P. Vogt, F. Nentwich, N. Jovanovic, E. Kirda, C. Krügel, and G. Vigna. Cross site scripting prevention with dynamic data tainting and static analysis. In *Proc. Network and Distributed System Security Symposium*, 2007.
- [99] D. Volpano and G. Smith. Verifying secrets and relative secrecy. In *Proc. 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '00, pages 268–276, 2000.
- [100] D. Volpano, C. Irvine, and G. Smith. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(2-3):167–187, Jan 1996.
- [101] B. Xin and X. Zhang. Efficient online detection of dynamic control dependence. In *Proc. 2007 International Symposium on Software Testing and Analysis*, pages 185–195, 2007.
- [102] J. Yang, K. Yessenov, and A. Solar-Lezama. A language for automatically enforcing privacy policies. In *Proc. 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '12, pages 85–96, 2012.
- [103] S. Zdancewic and A. C. Myers. Robust declassification. In *Proc. 14th IEEE Workshop on Computer Security Foundations*, CSFW '01, pages 15–23, 2001.
- [104] S. A. Zdancewic. *Programming Languages for Information Security*. PhD thesis, Cornell University, August 2002.
- [105] Y. Zhou and D. Evans. Protecting private web content from embedded scripts. In *Proc. 16th European Conference on Research in Computer Security*, pages 60–79, 2011.