

CERIAS Tech Report 2024-3
Language-Based Techniques for Policy-Agnostic Oblivious Computation
by Qianchuan Ye
Center for Education and Research
Information Assurance and Security
Purdue University, West Lafayette, IN 47907-2086

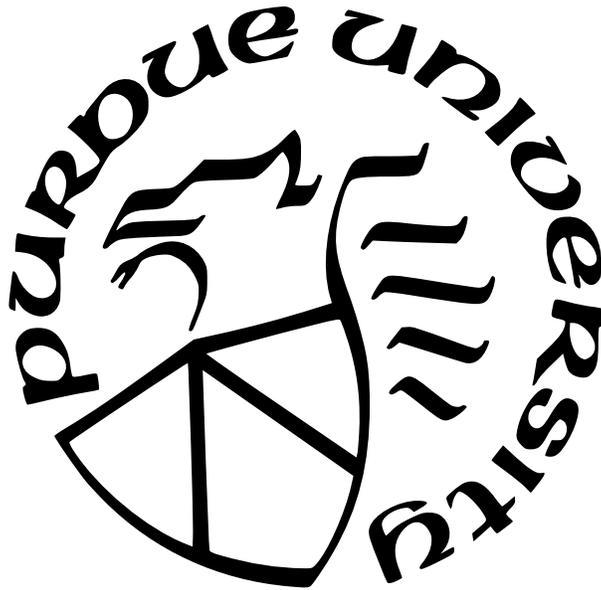
**LANGUAGE-BASED TECHNIQUES FOR POLICY-AGNOSTIC
OBLIVIOUS COMPUTATION**

by
Qianchuan Ye

A Dissertation

*Submitted to the Faculty of Purdue University
In Partial Fulfillment of the Requirements for the degree of*

Doctor of Philosophy



Department of Computer Science

West Lafayette, Indiana

May 2024

**THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL**

Dr. Benjamin Delaware, Chair

Department of Computer Science

Dr. Roopsha Samanta

Department of Computer Science

Dr. Milind Kulkarni

Elmore Family School of Electrical and Computer Engineering

Dr. David Darais

Galois, Inc.

Dr. Tianyi Zhang

Department of Computer Science

Approved by:

Dr. Kihong Park

To my loving mother

ACKNOWLEDGMENTS

I would not have survived my PhD career without the immense help and support from many humans and cats.

First of all, I thank my advisor Benjamin Delaware. I started my PhD program with zero background in any of the topics that I enjoy doing now. Ben taught me everything from doing research in general to solving a specific Coq error. I still remember the good old days when we sat down together and he patiently guided me through my Coq proofs. He is always willing to answer my dumb questions and listen to my rants. Ben provided me more than enough support and encouragement while giving me freedom to pursue whatever I was interested in at the time. His insights, questions, advice and sense of humor made me a better researcher, writer, presenter and a better person. I am forever in debt to him.

I also thank my PhD advisory committee. The “policy-agnostic” direction, a key component of this dissertation, is due to Roopsha Samanta’s suggestions when we worked on the HACCLE project together. Her “Reasoning about Programs” course is also my favorite class during my PhD. Milind Kulkarni’s sharp ideas, questions and comments shaped this dissertation and my presentations, as well as pointed out promising future directions. David Darais gave a lot of insightful feedback to this dissertation, and his work on oblivious languages has been a source of inspiration. The helpful feedback from my examining committee, Christina Garman and Tianyi Zhang, also improved this dissertation. I would like to thank other professors whom I have worked with or received valuable advice from: Suresh Jagannathan, Xiaokang Qiu, Jenna DiVincenzo, and Vasileios Zikas.

I have learned a lot from my now and former colleagues, Pedro Abreu, Rob Dickerson, Patrick LaFontaine, Prasita Mukherjee, Zhe Zhou, Kia Rahmani, Eric Bond and Paul Krogmeier. I have had many stimulating discussions with Raghav Malik and Kirshanthan Sundararajah. Their ideas and suggestions have been made into this dissertation. It was also delightful to study category theory with Raghav, Rob and Zhe.

Many thanks to my friends at PurPL group, including Guannan Wei, Chris Wagner, Nouraldin Jaber, Yongwei Yuan, Ashish Mishra, Pratyush Das, Adhitha Dias, Artem Pelenitsyn, Julia Belyakova, Oliver Bračevac, Durga Keerthi Mandarapu, Vidhush Singhal and

Vani Nagarajan. I have enjoyed the talks and conversations in our PurPL seminars (and free lunch). I thank all my friends at Purdue, including Guanhong Tao, Jianliang Wu, Le Yu, Yiran Hu, Yuyan Bao, Shengwei An, Zikang Xiong, Yilin Zheng, Chuyang Ke, Habiba Farrukh, Seunghoon Lee and Basavesh Ammanaghatta Shivakumar. They have made my PhD life at Purdue less boring.

I am grateful to my dear friend Jianfei Gao. He is a wonderful workout and show-watching buddy and an amazing chef, although I would prefer less garlic. His cats, Π and Ξ , deserve my special thanks for making me less productive and disrupting my sleep during the two years we spent together. Π is the sweetest lady in the world who loves sleeping on my lap when I need to get work done. She single-handedly converted me into a cat person. Ξ is, well, a little asshole.

I have met some wonderful people outside Purdue in various summer schools and conferences. Kuen-Bang Hou (Favonia) and Yao Li gave me much helpful advice. Natarajan Shankar has many fun stories to tell, and the conversations with him are always inspiring. I feel fortunate to have had the opportunity to work with Clément Pit-Claudel and Adam Chlipala. Ian Sweet and his work have inspired this dissertation, especially the implementations. Ashish Kundu has given valuable feedback to this work and helped me along the way. My friends from summer schools, Matthew Yacavone, Ramana Nagasamudram, Ziyang Li, Yiyun Liu, Jacob Prinz and Mako Bates, made these summer schools a much more enjoyable experience. Vikraman Choudhury helped me understand adjunction better. Loïc Pujet recommended me some materials about cubical type theory that I still have not read; one day I will go down this rabbit hole.

Finally, I want to thank my beloved mom and brother. They have been extremely supportive not just during my PhD, but throughout my life. Without them I would not even make it to a PhD program.

TABLE OF CONTENTS

LIST OF FIGURES	10
ABSTRACT	14
1 INTRODUCTION	15
1.1 Problem Description	16
1.2 Contributions and Outline	19
2 BACKGROUND	22
2.1 Threat Models	22
2.2 Security Specifications	24
2.3 Security-Type Systems	28
2.4 Dependent Type Systems	34
3 OBLIVIOUS ALGEBRAIC DATA TYPES	43
3.1 Overview	43
3.1.1 Encoding Private Data and Policies	45
3.1.2 Enforcing Policies	48
3.1.3 Performance Implication of the Threat Model	51
3.2 λ_{OADT} , Formally	51
3.2.1 Syntax	52
3.2.2 Semantics	54
3.2.3 Type System	56
3.2.4 Type Safety and Obliviousness	62
Obliviousness	63
3.3 Conclusion	65
4 TAPE SEMANTICS: DYNAMIC ENFORCEMENT OF POLICIES	66
4.1 Overview	67
4.2 λ_{OADT^+} , Formally	70

4.2.1	Syntax	70
4.2.2	Semantics	71
4.2.3	Type System	73
4.2.4	Type Safety and Obliviousness	77
4.3	Extending $\lambda_{\text{OADT}^{\oplus}}$	78
4.4	$\lambda_{\text{OADT}^{\oplus}}$ in Action	80
4.5	An Unsafe Reference Semantics	81
4.5.1	Metatheory of Reveal Semantics	83
4.6	Deriving Secure Implementations	85
4.6.1	Derivation Algorithm	87
4.6.2	Metatheory of the Derivation Algorithm	88
4.7	Conclusion	90
5	TAYPE: A POLICY-AGNOSTIC OBLIVIOUS LANGUAGE	91
5.1	Overview	93
5.1.1	Type Checking and Core TAYPE	93
5.1.2	Translating to OIL	95
	Translating Oblivious Type Definitions	96
	Translating Oblivious Types and Operations	96
	Translating Tape Semantics	97
5.2	TAYPE, Formally	98
5.2.1	Syntax	98
5.2.2	Semantics	101
5.2.3	Type System	103
5.2.4	Type Safety and Obliviousness	105
5.2.5	Surface Language and Bidirectional Type Checker	106
5.3	OIL and Translation	110
5.3.1	Syntax, Semantics and Type System	110
5.3.2	Translating from TAYPE to OIL	112
	Translating Types	112

	Translating Expressions	115
5.3.3	Translation for Conceal and Reveal Phases	120
5.4	Implementation	121
5.4.1	Optimizations	122
5.5	Experiments	124
5.5.1	Case Study: Medical Records	124
5.5.2	Case Study: Dating Application	124
5.5.3	Case Study: Secure Calculator	126
5.5.4	Discussion	127
5.5.5	Microbenchmarks	127
5.6	Conclusion	130
6	TAYPSI: STATIC ENFORCEMENT OF POLICIES	132
6.1	Overview	133
6.2	TAYPSI, Formally	142
6.2.1	Syntax	142
6.2.2	Semantics	144
6.2.3	Type System	146
6.2.4	Metatheory	149
6.3	Ψ -structures and Declarative Lifting	149
6.3.1	OADT Structures	150
6.3.2	Join Structures	152
6.3.3	Introduction and Elimination Structures	153
6.3.4	Coercion Structures	155
6.3.5	Declarative Lifting	156
6.3.6	Logical Refinement	158
6.3.7	Metatheory of Lifting	160
6.4	Algorithmic Lifting	161
6.4.1	Constraint Solving	165
6.4.2	Metatheory of Algorithmic Lifting	168

6.5	Implementation	169
6.5.1	Optimizations	169
6.6	Evaluation	171
6.6.1	Microbenchmark Performance	171
6.6.2	Impact of Optimization	173
6.6.3	Compilation Overhead	174
6.7	Conclusion	176
7	RELATED WORK	178
8	CONCLUSIONS AND FUTURE WORK	183
	REFERENCES	185
A	OMITTED RULES AND PROOFS FOR TAYPE	199
A.1	Semantics	199
A.2	Type System	202
A.3	Translation Algorithm	205
A.4	Totality of the Translation Algorithm	208
B	OMITTED RULES AND PROOFS FOR TAYPSI	211
B.1	Declarative Lifting	211
B.2	Algorithmic Lifting	213
B.3	Metatheory of Lifting	214

LIST OF FIGURES

2.1	Syntax of a simple security-typed calculus	28
2.2	Semantics of a simple security-typed calculus	28
2.3	Typing rules of a simple security-typed calculus	29
2.4	Indistinguishability of expressions in a simple security-typed language	30
2.5	Syntax of a simple dependently typed calculus	35
2.6	Parallel reduction of a simple dependently typed calculus	35
2.7	Typing rules of a simple dependently typed calculus	36
2.8	Kinding rules of a simple dependently typed calculus	36
3.1	Lookup element in a search tree	43
3.2	Execution trace of a lookup computation	44
3.3	Oblivious trees	46
3.4	Oblivious tree with a maximum depth of two	46
3.5	Either-or policy as an OADT	47
3.6	Public and oblivious types	48
3.7	Oblivious lookup function in λ_{OADT}	50
3.8	λ_{OADT} syntax	52
3.9	λ_{OADT} semantics	55
3.10	Semi-lattice on λ_{OADT} kinds	57
3.11	λ_{OADT} kinding rules	58
3.12	λ_{OADT} typing rules	59
3.13	Subset of λ_{OADT} parallel reduction rules	61
3.14	λ_{OADT} global definition typing rules	61
4.1	Sketch of a secure lookup function	67
4.2	Example λ_{OADT^+} execution traces	68
4.3	Oblivious lookup function in λ_{OADT^+}	69
4.4	λ_{OADT^+} syntax	71
4.5	λ_{OADT^+} semantics	72
4.6	Selected λ_{OADT^+} typing rules	74

4.7	Selected λ_{OADT^+} kinding rules	76
4.8	Selected λ_{OADT^+} parallel reduction rules	77
4.9	A subset of extended language for fixed-width integers	79
4.10	λ_{OADT^+} reveal semantics	82
4.11	Generalized section and retraction	87
4.12	Commuting diagram chasing	88
5.1	Compilation pipeline	92
5.2	List membership predicate	93
5.3	An oblivious implementation of $\widehat{\text{elem}}$	94
5.4	A fully annotated implementation of <code>elem</code> in core TAYPE and OIL	95
5.5	Selected leaky types and functions in OIL	96
5.6	Core TAYPE syntax	99
5.7	Selected small-step semantics rules of core TAYPE	102
5.8	Selected core TAYPE typing rules	104
5.9	Selected surface TAYPE bidirectional typing rules	107
5.10	Selected inference rules for dependent contexts	108
5.11	OIL source syntax	111
5.12	Rules for translating core TAYPE types to OIL types	113
5.13	Leaky structures for function types	114
5.14	Leaky structures for <code>lists</code>	114
5.15	Generating leaky ADT definitions	115
5.16	Selected rules for translating core TAYPE standard expressions to OIL expressions	116
5.17	Selected rules for translating core TAYPE leaky and oblivious expressions to OIL expressions	117
5.18	Selected rules for translating core TAYPE oblivious types to sizes in OIL	118
5.19	Oblivious injection	118
5.20	Selected rules for translating core TAYPE definitions to OIL definitions	119
5.21	Resolving leaky instances	119
5.22	Summary of programs used in our case studies	125
5.23	Workflow of the secure calculator	126

5.24	Definitions of oblivious decision trees with different public views	128
5.25	Decision tree	128
5.26	Microbenchmarks	129
6.1	Filtering a list	133
6.2	Oblivious lists with maximum and exact length public views	134
6.3	Naive translation of <code>filter</code> to secure versions	138
6.4	Ψ -structures of $\widehat{\text{list}}_{\leq}$	140
6.5	$\lambda_{\text{OADT}\Psi}$ syntax	143
6.6	Selected small-step semantics rules of $\lambda_{\text{OADT}\Psi}$	145
6.7	Selected typing rules of $\lambda_{\text{OADT}\Psi}$	146
6.8	Selected kinding rules of $\lambda_{\text{OADT}\Psi}$	147
6.9	$\lambda_{\text{OADT}\Psi}$ global definitions typing	147
6.10	Simple types, specification types and erasure	150
6.11	Mergeability	153
6.12	Coercion	155
6.13	Selected declarative lifting rules	157
6.14	A logical relation for refinement	159
6.15	Translation pipeline	161
6.16	Typed macros	163
6.17	Constraints	163
6.18	Selected algorithmic lifting rules	164
6.19	Constraint solving algorithm	167
6.20	Running times for each benchmark	172
6.21	Impact of turning off optimizations	174
6.22	Impact of turning off public view memoization	175
6.23	Impact of constraint solving on compilation speed	176
A.1	Core TAYPE evaluation context	199
A.2	Core TAYPE semantics rules	200
A.3	Core TAYPE semantics rules (cont.)	201
A.4	Core TAYPE kinding rules	202

A.5	Core TAYPE typing rules	203
A.6	Core TAYPE typing rules (cont.)	204
A.7	Translating core TAYPE oblivious types to OIL sizes	205
A.8	Translating core TAYPE expressions to OIL expressions	206
A.9	Translating core TAYPE expressions to OIL expressions (cont.)	207
B.1	Refinement as logical relation	211
B.2	Declarative lifting rules	212
B.3	Algorithmic lifting rules	213
B.4	Algorithmic lifting rules (cont.)	214

ABSTRACT

Protecting personal information is growing increasingly important to the general public, to the point that major tech companies now advertise the privacy features of their products. Despite this, it remains challenging to implement applications that do not leak private information either directly or indirectly, through timing behavior, memory access patterns, or control flow side channels. Existing security and cryptographic techniques such as secure multiparty computation (MPC) provide solutions to privacy-preserving computation, but they can be difficult to use for non-experts and even experts.

This dissertation develops the design, theory and implementation of various language-based techniques that help programmers write privacy-critical applications under a strong threat model. The proposed languages support private structured data, such as trees, that may hide their structural information and complex policies that go beyond whether a particular field of a record is private. More crucially, the approaches described in this dissertation decouple privacy and programmatic concerns, allowing programmers to implement privacy-preserving applications modularly, i.e., to independently develop application logic and independently update and audit privacy policies. Secure-by-construction applications are derived automatically by combining a standard program with a separately specified security policy.

1. INTRODUCTION

It is often the case that the owners of some private data want to compute some joint function of their data: a group of hospitals, for example, may want to calculate some statistics about their patients. In the case that this data is sensitive, the parties may not want (or be legally allowed) to simply pool their data and compute the result. *Secure computation* provides a solution in such scenarios, allowing multiple parties to perform a joint computation while keeping their sensitive data secure. Secure computation was formally introduced in the early 1980s by Yao [1], along with one of the first examples, Yao’s Millionaires’ Problem. In this problem, Alice and Bob are millionaires, and they wish to know who is richer without disclosing their wealth. In other words, this is a two-party secure computation that calculates a boolean result $x \leq y$ from private integer inputs x and y provided by the two parties, without revealing these two integers. Starting from this simple problem, secure computation has since found many privacy-focused applications, including secure auctions, voting, and privacy-preserving machine learning [2–4].

There are two major paradigms for secure computation: *secure multiparty computation* (MPC), wherein the computation is performed jointly by all parties involved; and *outsourced computation*, where a computationally powerful entity such as an untrusted cloud provider carries out the computation [2]. MPC is typically implemented using cryptography-based protocols, such as Yao’s Garbled Circuits [1] or secret-sharing [5, 6], while outsourced computation can be implemented using a variety of mechanisms, including cryptography-based fully homomorphic encryption [7, 8], virtualization [9, 10] and secure processors [11].

Writing secure applications that directly use these techniques can be quite challenging and error-prone, however, even if the author has the requisite cryptographic expertise. Thus, several high-level programming languages and compilers have been created to help programmers write secure applications, starting with Fairplay, the first publicly available MPC compiler [12]. For example, Obliv-C [13] is a C-like language for MPC applications which compiles down to Yao’s Garbled Circuits. Other notable languages include PICCO [14], OblivM [15], Wysteria/Wys* [16, 17], λ_{obliv} [18], Viaduct [19], and Symphony [20]. [Chapter 7](#) discusses these languages in greater detail.

1.1 Problem Description

While these languages raise the level of abstraction, it remains challenging to develop complex secure applications in them, due to their lack of or limited support for rich data structures, complex policies on data, and the separation of programmatic and privacy concerns.

Consider a private decision tree classification problem. Suppose Alice owns a medical record, represented as a *record* data type with several fields, such as her ID, age, height and weight. On the other hand, Bob provides a decision tree, each of whose nodes compares a feature (i.e., field) from Alice’s record against a threshold, and whose leaves are decisions. The computation is a standard decision tree classification algorithm that traverses Bob’s decision tree according to the comparison result of each node, until it reaches a decision. While it is straightforward to write this simple application in a conventional programming language, it is challenging to implement a secure version in the aforementioned languages.

To begin with, many of these languages have limited support for rich recursive data structures, like trees. When such data structures are supported, they typically require leaking information about the structure of the data: in Obliv-C, for example, users can define trees with secure nodes using pointers, but the “shape” of the underlying tree will always be visible to adversaries, as Obliv-C pointers are public data. What if Bob wants to hide also the structure of his decision tree, so that observing this data structure itself or how it is used does not disclose whether the tree is left-heavy or right-heavy? Unfortunately, in general it is impossible to hide everything about a recursive data: some public information has to be disclosed to bound the data’s in-memory representation and the computation over this data, e.g., recursion depth; the secure computation may simply not terminate otherwise. If Bob wishes to hide his tree’s structure, he at a minimum needs to reveal its maximum depth, for example. An important question for supporting secure data structures is thus what information about a private data may be publicly shared. Each choice of publicly disclosed information, or *public view*, defines a *privacy policy* on the data.

One major obstacle to securely implementing applications that respect their policies is the possibility of *timing channels*: as one example, the run time of any terminating computation reveals some approximate information about the “size” of the data structures it uses. Authors

of secure computations must be careful to not inadvertently reveal more information through such timing side-channels. As an example, consider the following Obliv-C program, which traverses an oblivious array `a`:

```
for (i = 0; i < MAX_BOUND; i++) {  
    // some secure computation on a[i]  
}
```

Here, the author has chosen to avoid timing channels by using an upper bound, `MAX_BOUND`, on the length of `a`. In effect, `MAX_BOUND` provides a public view on the structure of `a`, which is then used to ensure a consistent running time for the `for` loop. Of course, the author also must ensure `a` has been padded out to this maximum bound and that there are no `break` statements that depend on the contents of `a` in the body of the loop. In order to be secure, a computation over structured data types must be carried out without revealing any information outside this public view, including the structure of the private data.

While this trick to make programs *constant-time* [9] is easy to implement for computations over simple data types like arrays, it becomes more complicated for richer data structures, like the decision tree in our motivating example. First, users have to decide how to manually “pad” data structures, so that they are consistent with a particular public view. Second, programmers also have to track the public view throughout the program, making sure it remains consistent throughout. These manual efforts essentially force the programmers to explicitly enforce the privacy policies within the logic of the application itself; as a consequence, the entire application must be examined in order to audit its privacy policy. Lastly, richer data structures can have multiple public views, i.e., privacy policies, representing different trade-offs between privacy and performance. An application must be rewritten for each of these policies, due to the intermixing of policy enforcement and application logic.

This is particularly true for applications with the sorts of complex requirements that can occur in practice. Within the United States, for example, the Health Insurance Portability and Accountability Act (HIPAA) governs how patient data may be used. HIPAA allows *either* the personally identifiable information (PII) *or* medical data to be shared, but not both. Notably, this policy does not simply specify whether some particular field of a patient’s

medical record is private or public; rather it is a *relation* that dictates how a program can access and manipulate different parts of every individual record. To conform to this policy, a secure application must either pay the (considerable) cryptographic overhead of conservatively securing all accesses to the fields of a record, or adopt a more sophisticated strategy for monitoring how data is accessed. These challenges become more acute when dealing with recursive data, e.g., lists or trees, whose policies are necessarily more complex. In the private decision tree classification problem, if Bob, the owner of the tree, stipulates that only its depth may be disclosed, the classification function must use secure operations to ensure that no other information about the tree is leaked, e.g., its spine or the attributes it uses. If Bob is willing to share the latter bits of information, however, this function must either be rewritten to take advantage of the new, more permissive policy, or continue to pay the cost of providing stricter privacy guarantees. Thus, the intermixing of privacy and programmatic concerns in current languages require users to write different implementations of essentially the same program for each distinct privacy policy, and makes it difficult to read, write, and reason about secure applications. Ideally, these concerns should be separated, allowing programmers to write the functionality of their program once and for all, and then select the right public view for their security and performance requirements.

This dissertation considers the following research questions.

- Can we design a secure language that supports rich data structures that may hide their own structures? Can this language support complex policies on data, e.g., the either-or policy for medical records?
- How can we guarantee that no sensitive information is leaked when executing programs in this language, even if an attacker can observe the structured data itself and how it is manipulated? To faithfully model the domain of secure multiparty computation, we need to consider a strong threat model where a powerful attacker can observe every intermediate state of program executions, which also naturally covers timing side-channels.
- How do we decouple privacy policies from application logic? This form of modularity allows users to implement applications and specify policies independently.

1.2 Contributions and Outline

This dissertation develops the design, theory and implementation of various novel programming language techniques for *oblivious computation*. For generality, this dissertation uses the term oblivious computation to mean computation that does not leak private information directly or indirectly, e.g., through side-channels, under a strong threat model. The solutions described in this dissertation can be applied to any oblivious computation, including secure multiparty computation, fully homomorphic encryption, and other secure computation techniques that are not based on cryptography.

These solutions provide a foundation to a functional programming language that is rich, safe, and accessible.

- **Rich.** Using our language, programmers can implement oblivious applications that involve rich data structures with complex policies, and use other high-level functional programming features such as higher-order functions.
- **Safe.** No attacker can infer any private information beyond their own private input and the publicly shared information, even when they are able to observe the data representation and every single program state in an execution of a program written in this language.
- **Accessible.** Programmers can write the application logic in the standard way, without knowing the underlying secure computation model and the particular privacy policies used for the applications. In other words, the language is *policy-agnostic*. Privacy policies are separately defined, specified and audited. The functionality and privacy policies can then be composed to derive the secure implementation. This modular design allows functionality and policies to be reusable and makes it easy to switch policies for, e.g., making tradeoffs between privacy and performance.

The remainder of this dissertation is structured as follows: [Chapter 2](#) introduces the important concepts and ideas that underline or inspire the work in this dissertation, e.g., *noninterference* [21]. This chapter develops a simple security-typed language and a simple dependently typed language, to illustrate the formalization and proof techniques for these

systems, which are the key inspiration for the dissertation. The example core calculi in [Chapter 2](#) are designed to be simple while similar in style to the novel systems described in subsequent chapters.

The next four chapters discuss the main technical contributions of this dissertation.

- [Chapter 3](#) describes *oblivious algebraic data types* (OADTs), a form of dependent types that can encode complex privacy policies for structured data. This chapter also develops λ_{OADT} , a core calculus for writing oblivious programs using OADTs, and describes the strong security guarantees provided by its type system.
- [Chapter 4](#) tackles the problem of decoupling privacy and programmatic concerns. This chapter presents a dynamic approach for enforcing privacy policies automatically, using a novel operational semantics called *tape semantics*. This semantics allows programs to include unsafe computations, and then repairs these unsafe computations at runtime, which is the key to modularizing these concerns. A core calculus extending λ_{OADT} , dubbed λ_{OADT^+} , is presented, and various theoretic guarantees are established.
- [Chapter 5](#) discusses the implementation of a policy-agnostic programming language based on λ_{OADT^+} , called TAYPE. This chapter presents a bidirectional type checker and a compilation pipeline that addresses various challenges in implementing OADTs and tape semantics. The chapter concludes with an evaluation of this language.
- [Chapter 6](#) proposes a static approach for automatic policy enforcement. This approach transforms source programs into secure target programs that respect the given privacy policies, enabling significant performance improvements over previous dynamic approach, i.e., tape semantics, and better policy specifications. The TAYPSI language described in this chapter, and its underlying core calculus $\lambda_{\text{OADT}^\Psi}$, still provide the same security guarantee and separation of privacy concerns and application logic.

Finally, related work is discussed in [Chapter 7](#), and [Chapter 8](#) summarizes the dissertation and proposes future directions.

The material in this dissertation is mainly based on Ye and Delaware [22–24]. The formalization of all the core calculi presented in this dissertation and the proofs of their

metatheory have been mechanized in the Coq proof assistant. [Chapter 3](#), [Chapter 4](#), [Chapter 5](#) and [Chapter 6](#) include references to the publicly available artifacts of each of these formal developments, as well as an implementation of the languages.

2. BACKGROUND

Before describing the technical contributions of this dissertation, we begin with an introduction to several important concepts and established techniques.

2.1 Threat Models

This dissertation adopts a semi-honest setting [25]. In other words, the attackers are passive, in that they follow the computation or protocols faithfully, but may try to infer secrets from the information they can gather while doing so. Although they do not act maliciously, e.g., executing the programs in a deliberately wrong way, private information can still be leaked inadvertently. An attacker may be able to observe some extra information besides the input and output that they are permitted to see, e.g., how long a program runs, and infer sensitive information from such *side-channels* indirectly. One major challenge for oblivious computation is to protect against leakage through these side-channels.

The strength of a threat model in the semi-honest setting is determined by the capability of an attacker, i.e., what information they can observe. Consider a standard imperative programming model, where computations can be thought of as state transformations. Each state is a store, or memory, containing data, indexed by variable names. Some portion of this memory may be protected, and can only be accessed by trusted entities. A variable pointing to values in this secret section of memory is tagged with a *security label* \top (high-security), indicating that the contents of this variable are not visible to any attackers, the low-security observers. On the other hand, variables labeled with \perp (low-security) have values that are visible to everyone, including attackers. A secure program in this programming model should ensure no high-security data values can be inferred from low-security data or other information available via side-channels. For example, assigning a high-security variable \mathbf{h} to a low-security variable \mathbf{l} , $\mathbf{l} := \mathbf{h}$, breaks the security guarantee, as an attacker can obtain the secret \mathbf{h} by observing the value of \mathbf{l} after the assignment. A secure programming language should reject this program as insecure.

Undesirable information flow from high-security to low-security can happen implicitly as well. The following example modifies the low-security variable `l` according to whether the secret `h` is greater than 0.

```
if h > 0 then l := 1 else l := 0
```

By observing `l`, an attacker can infer some information about the secret `h`, which should be prohibited by the system.

An attacker may be able to observe more than just the low-security portion of the memory. For example, they may observe the running time of a program. This sort of timing behavior can also reveal secrets, as illustrated by the following example.

```
if h > 0 then h := 1 else (delay 10; h := 0)
```

Even though this program does not modify any low-security variables, an attacker can still infer some information about the initial value of `h` and thus obtain the final assignment to `h`, by measuring how much time this program takes to finish. A secure system needs to make sure high-security information does not influence the timing behavior of a program, if an attacker can gather information from such *timing channels*.

Perhaps surprisingly, the following innocent-looking program is also insecure if executed under a standard semantics.

```
if h > 0 then h := 1 else h := 0
```

While this program indeed reveals no private information through output and timing channels, a powerful attacker is still able to peek into `h` by observing which branch this program takes via *control flow channels*.

Our final strong threat model reflects both those of standard MPC protocols based on simultaneous execution of the programs, e.g., secret-sharing [6, 26], and those based on outsourced computation where untrusted evaluators perform the execution, e.g., fully homomorphic encryption [7]. In these protocols, *any* party involved in the computation could be an attacker, including the ones executing the programs, forcing us to protect against a powerful attacker that can observe the whole execution, including every intermediate program state. As a result, we have to obscure which branches the program takes, for example. This threat model also naturally covers weaker adversaries, including those who can only observe

the timing behavior. The language solutions developed in this thesis are designed to be secure under this strong threat model.

2.2 Security Specifications

One standard notion for specifying information security is *noninterference* [21], which roughly states that protected data can not influence observable information, e.g., low-security data. Hence, an observer cannot infer any protected information by observing low-security information. There are many variants of noninterference [27], reflecting different semantics and threat models. This property is usually defined relationally: for example, given *indistinguishable* inputs that differ only in their protected components, a program always produces outputs that are also *indistinguishable*. Intuitively, this means if the high-security values in the input change, the observable information in the execution remains unaffected, even when the high-security portion of the output can vary.

Indistinguishability specifies what information an observer can see, and is usually defined using *security labels*. For example, a high-security label (\top) may correspond to confidentiality, which describes private information, while a low-security label (\perp) describes public information. In the multiparty setting, principal identifiers such as `Alice` and `Bob` may be used as security labels. Alice may only have access to her own data, i.e., information labeled with `Alice`, for instance. Indistinguishability is then an equivalence relation (\approx) indexed by observers, stating that the observable information, according to its security labels, in the data is the same. In a system with high and low labels, for example, data are indistinguishable to an attacker if their low-security information is the same.

Another main component in a definition of noninterference are the outputs of a computation, i.e., what information is available to an observer in an execution. If we are only concerned about the computed result, then its output is simply the result. However, a program's output can include information in side-channels as well. For example, we may define the output of a program as the result it computes to *and* its running time, if we wish to account for timing behavior.

The general recipe for defining noninterference is thus to specify an indistinguishability relation on a program’s data and the output of a program, according to its semantics and threat model. Noninterference can then be formalized as when indistinguishable program states (configurations) produce indistinguishable outputs.

Example 2.2.1 (Termination-insensitive noninterference). Under a standard imperative semantics, a program configuration is a pair (P, σ) , where P is a program and σ is a map from variable names to their values, i.e., machine state, tagged by a \top or \perp label. A program’s behavior can be described by a big-step operational semantics $P, \sigma \Downarrow \sigma'$, meaning that the program P , when run in initial state σ , evaluates to final state σ' .

Consider a threat model where attackers can only see the values of variables labeled with \perp . In this threat model, we can define indistinguishability on a pair of states as those in which all the variables labelled with \perp have the same values:

$$\sigma \approx \sigma' \triangleq \sigma_{\perp} = \sigma'_{\perp}$$

where σ_{\perp} means the \perp projection of σ , consisting of all the variables that have the \perp label. The output is simply the final state. We can then define *termination-insensitive* noninterference as follows: given any initial states σ_1 and σ_2 such that $\sigma_1 \approx \sigma_2$, if a program P with these states evaluates to final states σ'_1 and σ'_2 , i.e., $P, \sigma_1 \Downarrow \sigma'_1$ and $P, \sigma_2 \Downarrow \sigma'_2$, then $\sigma'_1 \approx \sigma'_2$. Note that this definition is termination-insensitive because we assume P terminates with σ and σ' , otherwise this property is vacuously satisfied.

Alternatively, we can fit this definition into the previous framing more directly by lifting the indistinguishability notion to program configurations: $P, \sigma \approx P', \sigma'$ if and only if $P = P'$ and $\sigma \approx \sigma'$. We can also extend the notion of output to include a divergence symbol which indicates (P, σ) does not terminate (assuming this language does not have other diverging behaviors such as exceptions). We can then say two outputs are indistinguishable if one of them is the divergence symbol or they are indistinguishable states. With these generalized definitions, noninterference is simply that indistinguishable program configurations generate

indistinguishable outputs. Choosing to instead only relate the divergence symbol to itself leads to *termination-sensitive* noninterference:

Example 2.2.2 (Termination-sensitive noninterference). Consider an attacker who can also observe termination behavior. Noninterference is defined as follows: given any initial states σ_1 and σ_2 such that $\sigma_1 \approx \sigma_2$, a program P either diverges under both initial states, or evaluates to a pair of final states σ'_1 and σ'_2 , i.e., $P, \sigma_1 \Downarrow \sigma'_1$ and $P, \sigma_2 \Downarrow \sigma'_2$, such that $\sigma'_1 \approx \sigma'_2$.

Example 2.2.3 (Probabilistic noninterference). Under a nondeterministic and probabilistic computation model, program semantics may be defined as $P, \sigma \Downarrow D$, where D is a probability distribution on states. To protect against attackers who can sample multiple runs, we may consider two outputs indistinguishable if the generated distributions are statistically indistinguishable.

Example 2.2.4 (Timing channels). As illustrated in [Section 2.1](#), timing behavior can be used to infer private information. The big-step semantics we used previously is too coarse for this threat model. While it is possible to define a big-step semantics that also calculates the running time compositionally, it is more straightforward to define a small-step semantics $(P, \sigma) \longrightarrow (P', \sigma')$, which means a program P with state σ takes one computation step to the configuration (P', σ') . The program P in a configuration can be thought of as a program counter. For example, the program $(x := x + 1; y := 2)$ with state $\{x \mapsto 0\}$ steps to $(y := 2)$ with state $\{x \mapsto 1\}$.

With this small-step semantics, noninterference can be defined as follows: given initial states σ_1 and σ_2 , if $(P, \sigma_1) \longrightarrow^{n_1} (\epsilon, \sigma'_1)$ and $(P, \sigma_2) \longrightarrow^{n_2} (\epsilon, \sigma'_2)$, then $\sigma'_1 \approx \sigma'_2$ and $n_1 = n_2$, where ϵ is the empty program. The judgment $(P, \sigma) \longrightarrow^n (P', \sigma')$ generalizes the previous single-step version, meaning that (P, σ) takes n steps to (P', σ') . This noninterference definition counts how many steps it takes for a program to finish in order to ensure the number of steps taken is independent of the secrets. In this formulation, we assume each step costs the same amount of time for simplicity, but this can be easily adapted to capture more precise timing behavior, e.g., by annotating the small-step semantics with a more precise cost for each command. This version of the definition also does not consider termination channels, but it is easy to state a variant similar to termination-sensitive noninterference.

Example 2.2.5 (Access patterns). Attackers in some threat models can observe memory access patterns. To capture this capability, we may extend the previous small-step semantics to also record memory access events $(P, \sigma) \xrightarrow{\text{ev}} (P', \sigma')$. For example, the program $(x := \text{read } 1; \dots)$ may step to (\dots) with the new state $\{x \mapsto 0\}$ and the generated event $\text{READ}(1, \top, 0)$. This read event records the memory location 1, its security label and the read result 0.

The output of an execution is a list of the generated events under this small-step semantics. Indistinguishability of events and the noninterference definition can be straightforwardly formulated similarly to the previous attempts. One variant of this kind of noninterference is *memory trace obliviousness* [28].

Example 2.2.6 (Control flow channels). In a strong threat model, an attacker may be able to inspect the program counter in the CPU and the executing instructions. This is the case in many oblivious computations, including MPC, as discussed in [Section 2.1](#); as any party executing the programs can be thought of as an attacker who has full control over their own computing device. To model a powerful attacker who can observe every program state, we consider output to be the (possibly infinite) traces of configurations under a small-step semantics. In other words, the output is the (possibly infinite) list $(P_1, \sigma_1); (P_2, \sigma_2); \dots$ if the execution is $(P_1, \sigma_1) \longrightarrow (P_2, \sigma_2) \longrightarrow \dots$. Indistinguishability of traces is then simply pair-wise indistinguishability of program configurations.

Going back to the last example in [Section 2.1](#), the following two traces are distinguishable, despite the indistinguishable initial program configurations.

`(if h > 0 then h := 1 else h := 0, {h ↦ 1})` \longrightarrow `(h := 1, {h ↦ 1})` \longrightarrow `(ϵ, {h ↦ 1})`

`(if h > 0 then h := 1 else h := 0, {h ↦ 0})` \longrightarrow `(h := 0, {h ↦ 0})` \longrightarrow `(ϵ, {h ↦ 0})`

Observe that an attacker can infer the value of `h` by observing the second configuration in these traces, i.e., the branch it takes, and this program does not satisfy noninterference.

$e ::= b_l \mid x \mid \lambda x:\tau \Rightarrow e \mid e e \mid \text{if } e \text{ then } e \text{ else } e$	EXPRESSIONS
$\tau ::= \mathbb{B}_l \mid \tau \rightarrow \tau$	TYPES
$l ::= \top \mid \perp$	LABELS
$v ::= b_l \mid x \mid \lambda x:\tau \Rightarrow e$	VALUES

Figure 2.1. Syntax of a simple security-typed calculus

$e \Downarrow v$

$\frac{\text{E-VAL}}{v \Downarrow v}$	$\frac{\text{E-APP} \quad e_1 \Downarrow \lambda x:\tau \Rightarrow e'_1 \quad e_2 \Downarrow v_2 \quad [v_2/x]e'_1 \Downarrow v}{e_1 e_2 \Downarrow v}$
$\frac{\text{E-IFTRUE} \quad e_0 \Downarrow \text{true}_l \quad e_1 \Downarrow v}{\text{if } e_0 \text{ then } e_1 \text{ else } e_2 \Downarrow v}$	$\frac{\text{E-IFFALSE} \quad e_0 \Downarrow \text{false}_l \quad e_2 \Downarrow v}{\text{if } e_0 \text{ then } e_1 \text{ else } e_2 \Downarrow v}$

Figure 2.2. Semantics of a simple security-typed calculus

2.3 Security-Type Systems

Various language properties can be ensured statically by a well-designed type system. Security-typed languages [27, 29] have been extensively studied to enforce information-flow properties such as noninterference. This section presents a functional language as a small extension to simply typed lambda calculus (STLC) with high- and low-clearance booleans. This core calculus is designed to be minimal in order to exhibit the key idea of using a type system to restrict how certain resources can be used, while remaining similar in style to the proposed calculi in this dissertation.

Figure 2.1 shows the syntax of this core calculus. Terms and types in this language are mostly standard. Boolean literals and boolean types are annotated with a security label, indicating a boolean value is only visible by observers with enough permission according to its label.

This language has a standard (big-step) operational semantics, as shown in Figure 2.2. Note that the labels are ignored at runtime (and hence can be erased), and the language relies solely on its type system for the security guarantees. This is different from variants

$$\begin{array}{c}
\boxed{\Gamma \vdash e : \tau} \\
\\
\begin{array}{c}
\text{T-LIT} \\
\hline
\Gamma \vdash b_l : \mathbb{B}_l
\end{array}
\quad
\begin{array}{c}
\text{T-VAR} \\
x : \tau \in \Gamma \\
\hline
\Gamma \vdash x : \tau
\end{array}
\quad
\begin{array}{c}
\text{T-ABS} \\
x : \tau_1, \Gamma \vdash e : \tau_2 \\
\hline
\Gamma \vdash \lambda x : \tau_1 \Rightarrow e : \tau_1 \rightarrow \tau_2
\end{array}
\quad
\begin{array}{c}
\text{T-APP} \\
\Gamma \vdash e_2 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash e_1 : \tau_1 \\
\hline
\Gamma \vdash e_2 \ e_1 : \tau_2
\end{array}
\\
\\
\begin{array}{c}
\text{T-IF} \\
\Gamma \vdash e_0 : \mathbb{B}_l \quad \Gamma \vdash e_1 : \tau \quad \Gamma \vdash e_2 : \tau \quad l \sqsubseteq \text{label}(\tau) \\
\hline
\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : \tau
\end{array}
\end{array}$$

Figure 2.3. Typing rules of a simple security-typed calculus

of security-typed languages that may implicitly *promote* labels to ones of higher security, effectively “moving” low-security values to a more protected memory region.

The key idea of a security-type system is to rule out the programs with undesirable information flow by analyzing the security level of each component and imposing restrictions accordingly. A type system that enforces these restrictions for our simple security-typed language is presented in [Figure 2.3](#). Most of the typing rules are standard. Boolean literals are ascribed a boolean type with the same label in T-LIT. T-IF is the most interesting rule. This rule uses the expected label comparison operator \sqsubseteq (e.g., $\perp \sqsubseteq \top$), and the meta-function `label` which obtains a type’s security label: `label`(\mathbb{B}_l) = l and `label`($\tau_1 \rightarrow \tau_2$) = \perp . Note that for simplicity we assume functions are always visible to attackers (hence the \perp label). While it is possible to assign a high-security label to functions [29], functions are generally public information in secure multiparty computation and many other oblivious computations. The side condition in T-IF enforces a crucial policy that an `if`-conditional’s branches can only be “more secure” than its condition, otherwise we risk revealing the value of the condition from the result of a conditional. For example, the expression `if true $_{\top}$ then true $_{\perp}$ else false $_{\perp}$` is ill-typed, since a low-observer can see the computed result `true $_{\perp}$` and infer that the condition is also `true`.

The goal of this type system is to guarantee noninterference. Unlike imperative languages, this pure functional language is stateless, and its program configurations are simply expressions. However, establishing noninterference is not necessarily easier, due to the existence of higher-

$$\begin{array}{c}
\boxed{e \approx e'} \\
\hline
\frac{}{b_{\perp} \approx b_{\perp}} \quad \frac{}{b_{\top} \approx b'_{\top}} \quad \frac{}{x \approx x} \quad \frac{e \approx e'}{\lambda x:\tau \Rightarrow e \approx \lambda x:\tau \Rightarrow e'} \quad \frac{e_1 \approx e'_1 \quad e_2 \approx e'_2}{e_1 \ e_2 \approx e'_1 \ e'_2} \\
\hline
\frac{e_0 \approx e'_0 \quad e_1 \approx e'_1 \quad e_2 \approx e'_2}{\text{if } e_0 \text{ then } e_1 \text{ else } e_2 \approx \text{if } e'_0 \text{ then } e'_1 \text{ else } e'_2}
\end{array}$$

Figure 2.4. Indistinguishability of expressions in a simple security-typed language

order functions, i.e., lambda abstractions, even if they are low-security. Formally, we want to prove the following noninterference theorem.

Theorem 2.3.1 (Noninterference). *If $x : \mathbb{B}_{\top} \vdash e : \mathbb{B}_{\perp}$ and $\cdot \vdash v_1, v_2 : \mathbb{B}_{\top}$, then $[v_1/x]e \Downarrow v \iff [v_2/x]e \Downarrow v$.*

This definition is similar to Zdancewic [29, Theorem 3.1.1]. It says the (low-clearance) result can not be affected by a high-clearance value (i.e., secure boolean) in an expression.

This theorem is simply a special case of [Lemma 2.3.5](#) below. It can also be straightforwardly generalized to indistinguishable expressions and indistinguishable substitutions, but these stronger noninterference theorems are also corollaries of [Lemma 2.3.5](#), or can be achieved using the same proof techniques. A direct proof of these noninterference theorems (including [Theorem 2.3.1](#)) will not succeed, however, similar to the challenge in proving normalization of STLC: reduction of a lambda application may result in bigger terms that “escape” the induction hypotheses.

To formally establish noninterference, we first define indistinguishability of expressions as inference rules in [Figure 2.4](#). The first two rules establish that two high-security booleans are indistinguishable, and all other rules are simply congruence rules. Intuitively, two expressions are indistinguishable if they only differ in their unobservable, secure boolean values. Note that attackers are allowed to peek under the binder of a lambda abstraction, as it is the case under the strong threat model used later in this dissertation. This definition of indistinguishability for lambda abstractions is not strong enough to prove noninterference directly, because substituting indistinguishable arguments in indistinguishable function bodies may result in

bigger terms that are not obviously indistinguishable after evaluation. To strengthen this equivalence relation, we extend indistinguishability to the stronger *logical relation* [30, 31] below.

Definition 2.3.1. This relation is mutually defined via a pair of set-valued type denotations: a value interpretation $\mathcal{V}[\tau]$ and an expression interpretation $\mathcal{E}[\tau]$. We say closed and well-typed terms e and e' are equivalent at type τ if $(e, e') \in \mathcal{E}[\tau]$. In other words, they evaluate to equivalent values, per $\mathcal{V}[\tau]$. Importantly, the desirable property of substitution is directly encoded in the value interpretation of function type. In addition, the logical relation is naturally extended to typing contexts $\mathcal{G}[\Gamma]$, in order to relate equivalent substitutions. We use the notation $\sigma \vdash \Gamma$ to mean that $\cdot \vdash \sigma(x) : \tau$ for every $x : \tau \in \Gamma$, similar to the well-typedness side conditions in other interpretations.

$$\begin{aligned} \mathcal{V}[\mathbb{B}_\perp] &= \{ (b_\perp, b_\perp) \} & \mathcal{V}[\mathbb{B}_\top] &= \{ (b_\top, b'_\top) \} \\ \mathcal{V}[\tau_1 \rightarrow \tau_2] &= \left\{ \left(\lambda x : \tau_1 \Rightarrow e, \lambda x : \tau_1 \Rightarrow e' \right) \left| \begin{array}{l} \cdot \vdash \lambda x : \tau_1 \Rightarrow e : \tau_1 \rightarrow \tau_2 \wedge \cdot \vdash \lambda x : \tau_1 \Rightarrow e' : \tau_1 \rightarrow \tau_2 \wedge \\ e \approx e' \wedge \\ \forall (v, v') \in \mathcal{V}[\tau_1]. ([v/x]e, [v'/x]e') \in \mathcal{E}[\tau_2] \end{array} \right. \right\} \\ \mathcal{E}[\tau] &= \{ (e, e') \mid \cdot \vdash e : \tau \wedge \cdot \vdash e' : \tau \wedge e \Downarrow v \wedge e' \Downarrow v' \wedge (v, v') \in \mathcal{V}[\tau] \} \\ \mathcal{G}[\Gamma] &= \{ (\sigma, \sigma') \mid \sigma \vdash \Gamma \wedge \sigma' \vdash \Gamma \wedge \forall x : \tau \in \Gamma. (\sigma(x), \sigma'(x)) \in \mathcal{V}[\tau] \} \end{aligned}$$

The base cases of the denotation of values can also be defined directly using our indistinguishability relation. For example:

$$\mathcal{V}[\mathbb{B}_\perp] = \{ (v, v') \mid \cdot \vdash v : \mathbb{B}_\perp \wedge \cdot \vdash v' : \mathbb{B}_\perp \wedge v \approx v' \}$$

This style is equivalent because a well-typed value has a canonical form that is determined by its type.

Lemma 2.3.2 (Canonical forms).

- If $\cdot \vdash v : \mathbb{B}_l$, then $v = \text{true}_l$ or $v = \text{false}_l$.

- If $\cdot \vdash v : \tau_1 \rightarrow \tau_2$, then $v = \lambda x : \tau_1 \Rightarrow e$ for some e .

Proof. Straightforward case analysis on the type derivation. \square

A standard substitution lemma is needed to discharge the well-typedness side conditions in the logical relations.

Lemma 2.3.3 (Substitution preserves typing relation). *If $\Gamma \vdash e : \tau$ and $\sigma \vdash \Gamma$, then $\cdot \vdash \sigma(e) : \tau$.*

Proof. Routine induction on the typing derivation. \square

We also need another substitution lemma about indistinguishability. Indistinguishability is naturally extended to substitutions pointwise: $\sigma \approx \sigma'$ if $\sigma(x) \approx \sigma'(x)$ for each x in their domain.

Lemma 2.3.4 (Substitution preserves indistinguishability). *If $\sigma \approx \sigma'$, then $\sigma(e) \approx \sigma'(e)$.*

Proof. Routine induction on the structure of e . \square

Now we can state and prove a general lemma that substitution preserves the logical relations.

Lemma 2.3.5 (Substitution preserves logical relation). *If $\Gamma \vdash e : \tau$ and $(\sigma, \sigma') \in \mathcal{G}[\Gamma]$, then $(\sigma(e), \sigma'(e)) \in \mathcal{E}[\tau]$.*

Proof. By induction on the derivation of the typing judgment. The cases T-LIT and T-VAR are trivial. In the rest of the proof, we do not explicitly prove the well-typedness side conditions, since they are simply consequences of [Lemma 2.3.3](#).

Case T-ABS: We need to show $(\lambda x : \tau_1 \Rightarrow \sigma(e), \lambda x : \tau_1 \Rightarrow \sigma'(e)) \in \mathcal{V}[\tau_1 \rightarrow \tau_2]$. First, $\sigma(e) \approx \sigma'(e)$ follows from [Lemma 2.3.4](#) and the fact that $\sigma \approx \sigma'$ when $(\sigma, \sigma') \in \mathcal{G}[\Gamma]$. Next, suppose $(v, v') \in \mathcal{V}[\tau_1]$, we have $(\sigma[x \mapsto v], \sigma'[x \mapsto v']) \in \mathcal{G}[x : \tau_1, \Gamma]$. It then follows from the induction hypothesis that $(\sigma[x \mapsto v](e), \sigma'[x \mapsto v'](e)) \in \mathcal{E}[\tau_2]$. That is $([v/x]\sigma(e), [v'/x]\sigma'(e)) \in \mathcal{E}[\tau_2]$, as required.

Case T-APP: We have by the induction hypotheses:

- $(\sigma(\mathbf{e}_2), \sigma'(\mathbf{e}_2)) \in \mathcal{E}[\tau_1 \rightarrow \tau_2]$
- $(\sigma(\mathbf{e}_1), \sigma'(\mathbf{e}_1)) \in \mathcal{E}[\tau_1]$

Therefore:

- $\sigma(\mathbf{e}_2) \Downarrow \lambda \mathbf{x} : \tau_1 \Rightarrow \mathbf{u}_2$ for some \mathbf{u}_2
- $\sigma'(\mathbf{e}_2) \Downarrow \lambda \mathbf{x} : \tau_1 \Rightarrow \mathbf{u}'_2$ for some \mathbf{u}'_2
- $(\lambda \mathbf{x} : \tau_1 \Rightarrow \mathbf{u}_2, \lambda \mathbf{x} : \tau_1 \Rightarrow \mathbf{u}'_2) \in \mathcal{V}[\tau_1 \rightarrow \tau_2]$
- $\sigma(\mathbf{e}_1) \Downarrow \mathbf{v}_1$ for some \mathbf{v}_1
- $\sigma'(\mathbf{e}_1) \Downarrow \mathbf{v}'_1$ for some \mathbf{v}'_1
- $(\mathbf{v}_1, \mathbf{v}'_1) \in \mathcal{V}[\tau_1]$

The value interpretation of $\tau_1 \rightarrow \tau_2$ allows us to derive that $([\mathbf{v}_1/\mathbf{x}] \mathbf{u}_2, [\mathbf{v}'_1/\mathbf{x}] \mathbf{u}'_2) \in \mathcal{E}[\tau_2]$, which again gives us:

- $[\mathbf{v}_1/\mathbf{x}] \mathbf{u}_2 \Downarrow \mathbf{v}$ for some \mathbf{v}
- $[\mathbf{v}'_1/\mathbf{x}] \mathbf{u}'_2 \Downarrow \mathbf{v}'$ for some \mathbf{v}'
- $(\mathbf{v}, \mathbf{v}') \in \mathcal{V}[\tau_2]$

We then have $\sigma(\mathbf{e}_2) \sigma(\mathbf{e}_1) \Downarrow \mathbf{v}$ and $\sigma'(\mathbf{e}_2) \sigma'(\mathbf{e}_1) \Downarrow \mathbf{v}'$ from E-APP with the desired property $(\mathbf{v}, \mathbf{v}') \in \mathcal{V}[\tau_2]$. That is $(\sigma(\mathbf{e}_2) \sigma(\mathbf{e}_1), \sigma'(\mathbf{e}_2) \sigma'(\mathbf{e}_1)) \in \mathcal{E}[\tau_2]$.

Case T-IF: We want to show:

$$(\text{if } \sigma(\mathbf{e}_0) \text{ then } \sigma(\mathbf{e}_1) \text{ else } \sigma(\mathbf{e}_2), \text{if } \sigma'(\mathbf{e}_0) \text{ then } \sigma'(\mathbf{e}_1) \text{ else } \sigma'(\mathbf{e}_2)) \in \mathcal{E}[\tau]$$

By the induction hypothesis, we know $\sigma(\mathbf{e}_0) \Downarrow \mathbf{b}_l$ and $\sigma'(\mathbf{e}_0) \Downarrow \mathbf{b}'_l$ for some \mathbf{b} and \mathbf{b}' , such that $(\mathbf{b}_l, \mathbf{b}'_l) \in \mathcal{V}[\mathbb{B}_l]$. We consider two cases.

If $l = \perp$, then $\mathbf{b} = \mathbf{b}'$. Without loss of generality, assume they equal **true** (the case of **false** is the similar). It suffices to show $\sigma(\mathbf{e}_1) \Downarrow \mathbf{v}$ and $\sigma'(\mathbf{e}_1) \Downarrow \mathbf{v}'$ for some \mathbf{v} and \mathbf{v}' such that $(\mathbf{v}, \mathbf{v}') \in \mathcal{V}[\tau]$. But it follows immediately from the induction hypothesis.

If $l = \top$, then $\text{label}(\tau) = \top$, which means τ is \mathbb{B}_\top . By the induction hypotheses, we know all 4 branches evaluate to some values, and they are (secure) boolean literals by [Lemma 2.3.2](#). Since $(b_\top, b'_\top) \in \mathcal{E}[\mathbb{B}_\top]$ for any b and b' , the proof is concluded trivially. \square

Finally, [Theorem 2.3.1](#), i.e., noninterference, is a direct consequence of [Lemma 2.3.5](#).

2.4 Dependent Type Systems

Dependently typed languages [\[32\]](#) allow types to depend on terms, providing strong static guarantees about behavior of programs. For example, the type of lists may depend on a natural number which specifies the length of a list. Full-spectrum dependent type systems form the foundation of many theorem provers, such as Coq [\[33\]](#), Lean [\[34\]](#), and Agda [\[35\]](#). These languages support *large elimination*, or computing types from data. Another form of dependent types is refinement types [\[36–39\]](#). These systems usually do not support large elimination, but rather augment the types with predicates that restrict the values to a “subset” of the type. Refinement types are designed to be lighter-weight, and many program specifications can be expressed through these refinement types. Unlike Coq or Agda, where proofs are constructed inside the language based on the *Curry-Howard Correspondence*, refinement type systems rely heavily on automated theorem provers, e.g., Z3 [\[40\]](#), to discharge proof obligations, trading expressiveness for more automation. There are also works that aim to bring dependent types to general purpose languages [\[41–43\]](#).

Dependently typed languages, especially those with large elimination, directly inspired the novel type systems introduced in this dissertation. This section presents a simple dependently typed language that extends the simply typed lambda calculus with dependent functions and dependent conditionals. While this core calculus is similar in spirit to pure type systems (PTS) [\[32\]](#) in the lambda cube [\[44, 45\]](#), it is closer in style to λLF [\[46, Chapter 2.2\]](#) and the type systems in this dissertation.

[Figure 2.5](#) shows the syntax of this simple dependently typed language. Similar to most dependently typed languages, types and terms belong to the same syntactic class. By convention, we try to use the metavariable τ for types and e to refer to terms. The key extension to simply typed lambda calculus is the inclusion of dependent function types

EXPRESSIONS

$e, \tau ::= b \mid x \mid \lambda x:\tau \Rightarrow e \mid e e \mid \text{if } e \text{ then } e \text{ else } e \mid \mathbb{B} \mid \Pi x:\tau, \tau$

Figure 2.5. Syntax of a simple dependently typed calculus

$e \Rightarrow e'$

$\frac{}{e \Rightarrow e} \text{R-REFL}$	$\frac{e_1 \Rightarrow e'_1 \quad e_2 \Rightarrow e'_2}{(\lambda x:\tau \Rightarrow e_2) e_1 \Rightarrow [e'_1/x]e'_2} \text{R-APP}$	$\frac{e_1 \Rightarrow e'_1}{\text{if true then } e_1 \text{ else } e_2 \Rightarrow e'_1} \text{R-IFTRUE}$	
$\frac{e_2 \Rightarrow e'_2}{\text{if false then } e_1 \text{ else } e_2 \Rightarrow e'_2} \text{R-IFFALSE}$	$\frac{\tau \Rightarrow \tau' \quad e \Rightarrow e'}{\lambda x:\tau \Rightarrow e \Rightarrow \lambda x:\tau' \Rightarrow e'} \text{R-ABSCGR}$	$\frac{e_1 \Rightarrow e'_1 \quad e_2 \Rightarrow e'_2}{e_1 e_2 \Rightarrow e'_1 e'_2} \text{R-APPCGR}$	
$\frac{e_0 \Rightarrow e'_0 \quad e_1 \Rightarrow e'_1 \quad e_2 \Rightarrow e'_2}{\text{if } e_0 \text{ then } e_1 \text{ else } e_2 \Rightarrow \text{if } e'_0 \text{ then } e'_1 \text{ else } e'_2} \text{R-IFCGR}$	$\frac{\tau_1 \Rightarrow \tau'_1 \quad \tau_2 \Rightarrow \tau'_2}{\Pi x:\tau_1, \tau_2 \Rightarrow \Pi \tau'_1, \tau'_2} \text{R-PICGR}$		

Figure 2.6. Parallel reduction of a simple dependently typed calculus

(Π) and dependent conditionals (**if**) which allow for type-level computation. For example, $\Pi x:\mathbb{B}, \text{if } x \text{ then } \mathbb{B} \text{ else } \mathbb{B} \rightarrow \mathbb{B}$ is a valid type: the return type of this function is dictated by its argument x . We write $\tau_1 \rightarrow \tau_2$ for $\Pi x:\tau_1, \tau_2$ when x does not appear in τ_2 , indicating that there is no dependency in this type.

While we can give this language a standard call-by-value or call-by-name small-step operational semantics, we need a set of more permissive reduction rules to establish the desired metatheoretic properties. A standard technique is to define an alternative semantics called *parallel reduction*, shown in [Figure 2.6](#). These semantics rules allow reductions under the binder of a lambda abstraction (R-ABSCGR) and reductions in **if** branches (R-IFCGR). All the subcomponents of an expression are reduced simultaneously in the congruence rules and even in the β -reduction rules (R-APP). As type-level computation is supported, reductions can also happen in types (e.g., R-PICGR). Note that the standard call-by-value or call-by-name reduction relations are included in this “bigger” parallel reduction relation, so properties such as preservation (i.e., subject reduction) also hold for these more restricted semantics.

$$\begin{array}{c}
\boxed{\Gamma \vdash e : \tau} \\
\\
\text{T-LIT} \quad \frac{}{\Gamma \vdash b : \mathbb{B}} \quad \text{T-VAR} \quad \frac{x : \tau \in \Gamma \quad \Gamma \vdash \tau :: *}{\Gamma \vdash x : \tau} \quad \text{T-ABS} \quad \frac{x : \tau_1, \Gamma \vdash e : \tau_2 \quad \Gamma \vdash \tau_1 :: *}{\Gamma \vdash \lambda x : \tau_1 \Rightarrow e : \prod x : \tau_1, \tau_2} \\
\\
\text{T-APP} \quad \frac{\Gamma \vdash e_2 : \prod x : \tau_1, \tau_2 \quad \Gamma \vdash e_1 : \tau_1}{\Gamma \vdash e_2 \ e_1 : [e_1/x] \tau_2} \\
\\
\text{T-IF} \quad \frac{\Gamma \vdash e_0 : \mathbb{B} \quad \Gamma \vdash e_1 : [\text{true}/z] \tau \quad \Gamma \vdash e_2 : [\text{false}/z] \tau \quad \Gamma \vdash [e_0/z] \tau :: *}{\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : [e_0/z] \tau} \\
\\
\text{T-CONV} \quad \frac{\Gamma \vdash e : \tau \quad \tau \equiv \tau' \quad \Gamma \vdash \tau' :: *}{\Gamma \vdash e : \tau'}
\end{array}$$

Figure 2.7. Typing rules of a simple dependently typed calculus

$$\begin{array}{c}
\boxed{\Gamma \vdash \tau :: *} \\
\\
\text{K-BOOL} \quad \frac{}{\Gamma \vdash \mathbb{B} :: *} \quad \text{K-PI} \quad \frac{\Gamma \vdash \tau_1 :: * \quad x : \tau_1, \Gamma \vdash \tau_2 :: *}{\Gamma \vdash \prod x : \tau_1, \tau_2 :: *} \quad \text{K-IF} \quad \frac{\Gamma \vdash e_0 : \mathbb{B} \quad \Gamma \vdash \tau_1 :: * \quad \Gamma \vdash \tau_2 :: *}{\Gamma \vdash \text{if } e_0 \text{ then } \tau_1 \text{ else } \tau_2 :: *}
\end{array}$$

Figure 2.8. Kinding rules of a simple dependently typed calculus

Programs in this language are typed using a pair of typing and kinding judgments, $\Gamma \vdash e : \tau$ and $\Gamma \vdash \tau :: *$ respectively. Hence this is a 2-layer type system: terms have types, and types have kinds, although there is only one kind $*$ to keep this language simple. [Figure 2.7](#) and [Figure 2.8](#) show the (mutually defined) typing and kinding rules.

The kinding rules are mostly straightforward, with the most interesting rule K-IF allowing for large elimination on booleans. Some typing rules (T-VAR, T-ABS and T-IF) have side conditions about kinding to ensure the types used in these rules are well-formed. T-ABS is similar to the standard typing rule for lambda abstraction, although τ_2 is allowed to refer to the variable being bound. T-APP is also similar to its standard counterpart, but the return type of the function is specialized with its argument after the application. Typing a

dependent conditional (T-IF) relies on an implicit *motive* that is specialized when typing its branches, because the overall type of the expression may depend on its condition. This motive (τ) contains a special free variable (z) which stands in for the result of the condition. This variable is concretized with `true` ($[\text{true}/z]\tau$) when typing the `then` branch, for example. Finally, T-CONV allows any well-typed term to be typed with an equivalent type, using a type equivalence relation $\tau \equiv \tau'$. This equivalence is defined directly in terms of the parallel reduction relation. Two terms are said to be equivalent if they can parallelly reduce to the same term in zero or more steps:

$$\tau_1 \equiv \tau_2 \triangleq \exists \tau. \tau_1 \Rightarrow^* \tau \wedge \tau_2 \Rightarrow^* \tau$$

It is also possible to define this equivalence inductively using parallel reduction, or define it as the transitive and symmetric closure of parallel reduction. These definitions are equivalent to each other, although we may need the confluence property (Lemma 2.4.5) established below to prove that.

As an example, we can type the function `$\lambda x:\mathbb{B} \Rightarrow \text{if } x \text{ then false else } \lambda y:\mathbb{B} \Rightarrow y$` with `$\Pi x:\mathbb{B}, \text{if } x \text{ then } \mathbb{B} \text{ else } \mathbb{B} \rightarrow \mathbb{B}$` using the following derivation:

$$\begin{array}{c}
\text{T-LIT} \\
\hline
\text{T-CONV} \frac{x:\mathbb{B} \vdash \text{false}:\mathbb{B} \quad \mathbb{B} \equiv \text{if true then } \mathbb{B} \text{ else } \mathbb{B} \rightarrow \mathbb{B} \quad \dots}{x:\mathbb{B} \vdash \text{false}:\text{if true then } \mathbb{B} \text{ else } \mathbb{B} \rightarrow \mathbb{B}} \quad \dots \quad \text{K-BOOL} \\
\text{T-IF} \frac{\dots}{x:\mathbb{B} \vdash \text{if } x \text{ then false else } \lambda y:\mathbb{B} \Rightarrow y:\text{if } x \text{ then } \mathbb{B} \text{ else } \mathbb{B} \rightarrow \mathbb{B}} \quad \dots \quad \text{T-ABS} \\
\hline
\cdot \vdash \lambda x:\mathbb{B} \Rightarrow \text{if } x \text{ then false else } \lambda y:\mathbb{B} \Rightarrow y:\Pi x:\mathbb{B}, \text{if } x \text{ then } \mathbb{B} \text{ else } \mathbb{B} \rightarrow \mathbb{B}
\end{array}$$

Only the derivation of the `then` branch in the T-IF application is shown; the `else` branch is similar, and the condition is easy to type. We choose the motive `if z then \mathbb{B} else $\mathbb{B} \rightarrow \mathbb{B}$` when applying T-IF. After instantiating z with `true` in the type of the `then` branch, this type is converted to \mathbb{B} by applying T-CONV.

The expressiveness and strong guarantees provided by dependently typed languages come at a price, one of which being a much more intricate metatheory. In the rest of this section, we will prove the standard preservation property as an example of these complexities. As a

side note, all typing contexts (Γ) in this dissertation (as well as in the Coq mechanization) are modeled as maps with extensional equality. As a result, in contrast to modeling typing contexts as ordered lists, as is done in some works, our statements and proofs have a simpler treatment to typing contexts.

First, we must establish that parallel reduction is *confluent*, also known as the *Church-Rosser property*. To this end, we show substitution preserves parallel reduction.

Lemma 2.4.1. *If $s \Rightarrow s'$, then $[s/x]e \Rightarrow [s'/x]e$.*

Proof. By routine induction on the structure of e . □

Lemma 2.4.2 (Substitution preserves parallel reduction). *If $e \Rightarrow e'$ and $s \Rightarrow s'$, then $[s/x]e \Rightarrow [s'/x]e'$.*

Proof. By routine induction on the derivation of $e \Rightarrow e'$. □

[Lemma 2.4.2](#) can be extended to a lemma about type equivalence.

Lemma 2.4.3 (Substitution preserves type equivalence). *If $e \equiv e'$ and $s \equiv s'$, then $[s/x]e \equiv [s'/x]e'$.*

Proof. Easily obtained from [Lemma 2.4.2](#). □

We then prove an important diamond property [47] that says the reductions of a term can converge in one reduction step. The proof is fairly straightforward (albeit tedious).

Lemma 2.4.4 (Diamond). *If $e \Rightarrow e_1$ and $e \Rightarrow e_2$, then $e_1 \Rightarrow e'$ and $e_2 \Rightarrow e'$ for some e' .*

Proof. We proceed by induction on the first derivation and then inverting the second one.

Case R-APPCGR: Suppose $e_1 \ e_2 \Rightarrow e'_1 \ e'_2$ because $e_1 \Rightarrow e'_1$ and $e_2 \Rightarrow e'_2$. Inverting the second derivation gives us three possibilities.

First, $e_1 \ e_2 \Rightarrow e_1 \ e_2$ because of R-REFL. This case is trivial.

Second, $e_1 \ e_2 \Rightarrow u_1 \ u_2$ for some u_1 and u_2 because of R-APPCGR and $e_1 \Rightarrow u_1$ and $e_2 \Rightarrow u_2$. In this case, by the induction hypotheses, $u_1 \Rightarrow t_1$ and $e'_1 \Rightarrow t_1$ for some t_1 . Similarly, $u_2 \Rightarrow t_2$ and $e'_2 \Rightarrow t_2$ for some t_2 . Therefore, $e'_1 \ e'_2$ and $u_1 \ u_2$ reduce to the same expression $t_1 \ t_2$ by applying R-APPCGR.

Finally, $e_1 = \lambda x:\tau \Rightarrow u_1$ and $(\lambda x:\tau \Rightarrow u_1) e_2 \Rightarrow [u'_1/x]u'_2$ for some u_1, u'_1 and u'_2 because of R-APP and $u_1 \Rightarrow u'_1$ and $e_2 \Rightarrow u'_2$. By assumption, $\lambda x:\tau \Rightarrow u_1 \Rightarrow e'_1$, so $e'_1 = \lambda x:\tau_1 \Rightarrow s$ for some τ_1 and s such that $\tau \Rightarrow \tau_1$ and $u_1 \Rightarrow s$ because of R-ABSCGR or R-REFL. Since $\lambda x:\tau \Rightarrow u_1 \Rightarrow \lambda x:\tau \Rightarrow u'_1$ from R-ABSCGR, by the induction hypothesis, $\lambda x:\tau \Rightarrow u'_1 \Rightarrow \lambda x:\tau' \Rightarrow t_1$ and $\lambda x:\tau_1 \Rightarrow s \Rightarrow \lambda x:\tau' \Rightarrow t_1$ for some τ' and t_1 . We also have $u'_1 \Rightarrow t_1$ and $s \Rightarrow t_1$ by inverting these reductions. On the other hand, $e'_2 \Rightarrow t_2$ and $u'_2 \Rightarrow t_2$ for some t_2 by the induction hypothesis. It then follows that $e'_1 e'_2$, i.e., $(\lambda x:\tau_1 \Rightarrow s) e'_2$ reduces to $[t_1/x]t_2$ by R-APP, and $[u'_1/x]u'_2$ also reduces to $[t_1/x]t_2$ by [Lemma 2.4.2](#), completing the proof of R-APPCGR case.

The proofs of other cases are similar and left as an exercise for the reader. \square

From the diamond property, we can easily obtain confluence of parallel reduction.

Lemma 2.4.5 (Confluence). *If $e \Rightarrow^* e_1$ and $e \Rightarrow^* e_2$, then $e_1 \Rightarrow^* e'$ and $e_2 \Rightarrow^* e'$ for some e' .*

Proof. Proceed by induction on the reflexive and transitive closure of parallel reduction and [Lemma 2.4.4](#). \square

We can finally show that the type equivalence defined previously is indeed an equivalence relation.

Lemma 2.4.6 (Type equivalence). *The definition of type equivalence \equiv is an equivalence relation.*

Proof. The confluence property, i.e., [Lemma 2.4.5](#), is necessary for proving transitivity of \equiv . Reflexivity and symmetry are trivial. \square

Some standard lemmas are needed to prove preservation.

Lemma 2.4.7 (Weakening). *If $\Gamma \vdash e : \tau$ and $\Gamma \subseteq \Gamma'$, then $\Gamma' \vdash e : \tau$.*

*If $\Gamma \vdash \tau :: *$ and $\Gamma \subseteq \Gamma'$, then $\Gamma' \vdash \tau :: *$.*

Proof. We prove these two statements simultaneously by induction on the typing and kinding derivations. \square

Lemma 2.4.8 (Substitution). *If $\mathbf{x} : \tau', \Gamma \vdash \mathbf{e} : \tau$ and $\Gamma \vdash \mathbf{s} : \tau'$, then $\Gamma \vdash [\mathbf{s}/\mathbf{x}]\mathbf{e} : [\mathbf{s}/\mathbf{x}]\tau$.*

*If $\mathbf{x} : \tau', \Gamma \vdash \tau :: *$ and $\Gamma \vdash \mathbf{s} : \tau'$, then $\Gamma \vdash [\mathbf{s}/\mathbf{x}]\tau :: *$.*

Proof. By routine (mutual) induction on the typing and kinding derivations. The proof requires the weakening lemma, i.e., [Lemma 2.4.7](#). \square

Lemma 2.4.9 (Regularity). *If $\Gamma \vdash \mathbf{e} : \tau$, then $\Gamma \vdash \tau :: *$.*

Proof. By induction on the typing derivation. Most cases are trivial, but the case of T-APP depends on the kinding part of [Lemma 2.4.8](#). \square

Lemma 2.4.10 (Type conversion in context). *If $\mathbf{x} : \tau_1, \Gamma \vdash \mathbf{e} : \tau$ and $\tau_1 \equiv \tau_2$ with $\Gamma \vdash \tau_2 :: *$, then $\mathbf{x} : \tau_2, \Gamma \vdash \mathbf{e} : \tau$.*

*If $\mathbf{x} : \tau_1, \Gamma \vdash \tau :: *$ and $\tau_1 \equiv \tau_2$ with $\Gamma \vdash \tau_2 :: *$, then $\mathbf{x} : \tau_2, \Gamma \vdash \tau :: *$.*

Proof. By routine (mutual) induction on the typing and kinding derivations. [Lemma 2.4.7](#), [Lemma 2.4.8](#) and [Lemma 2.4.9](#) are needed. \square

Lemma 2.4.11 (Inversion). *If $\Gamma \vdash \lambda \mathbf{x} : \tau_1 \Rightarrow \mathbf{e} : \Pi \mathbf{x} : \tau'_1, \tau_2$, then $\tau_1 \equiv \tau'_1$ and $\mathbf{x} : \tau_1, \Gamma \vdash \mathbf{e} : \tau_2$ and $\Gamma \vdash \tau_1 :: *$.*

Proof. By assumption, $\Gamma \vdash \lambda \mathbf{x} : \tau_1 \Rightarrow \mathbf{e} : \tau$ for some τ such that $\tau \equiv \Pi \mathbf{x} : \tau'_1, \tau_2$. We proceed by induction on this new typing derivation. Most cases are vacuous except for T-ABS and T-CONV. The case of T-CONV is trivial by the induction hypothesis, and it relies on transitivity of type equivalence ([Lemma 2.4.6](#)). In the case of T-ABS, $\Gamma \vdash \lambda \mathbf{x} : \tau_1 \Rightarrow \mathbf{e} : \Pi \mathbf{x} : \tau_1, \tau'_2$, so $\tau_1 \equiv \tau'_1$ and $\tau_2 \equiv \tau'_2$. It suffices to prove $\mathbf{x} : \tau_1, \Gamma \vdash \mathbf{e} : \tau_2$, which follows from T-CONV and the induction hypothesis. \square

Finally, we prove that parallel reduction preserves types.

Theorem 2.4.12 (Preservation). *If $\Gamma \vdash \mathbf{e} : \tau$ and $\mathbf{e} \Rightarrow \mathbf{e}'$, then $\Gamma \vdash \mathbf{e}' : \tau$.*

*If $\Gamma \vdash \tau :: *$ and $\tau \Rightarrow \tau'$, then $\Gamma \vdash \tau' :: *$.*

Proof. We prove these two statements simultaneously by (mutual) induction on the typing and kinding derivations. The cases of T-LIT, T-CONV and K-BOOL are trivial. We omit

proofs of the side conditions about well-kindedness (e.g., when applying T-CONV), which are easily discharged by the kinding rules or by [Lemma 2.4.9](#).

Case T-ABS: By assumption, $\lambda x:\tau_1 \Rightarrow e \Rightarrow \lambda x:\tau'_1 \Rightarrow e'$ for some τ'_1 and e' due to R-REFL or R-ABSCGR, as well as $\tau_1 \Rightarrow \tau_2$ and $e \Rightarrow e'$. By the induction hypothesis, $x : \tau_1, \Gamma \vdash e' : \tau_2$, from which [Lemma 2.4.10](#) gives $x : \tau'_1, \Gamma \vdash e' : \tau_2$ because $\tau_1 \equiv \tau'_1$ and $\Gamma \vdash \tau'_1 :: *$ again by the induction hypothesis. We then have $\Gamma \vdash \lambda x:\tau'_1 \Rightarrow e' : \Pi x:\tau'_1, \tau_2$ by T-ABS. Finally, T-CONV gives us $\Gamma \vdash \lambda x:\tau'_1 \Rightarrow e' : \Pi x:\tau_1, \tau_2$ as required, because $\Pi x:\tau'_1, \tau_2 \equiv \Pi x:\tau_1, \tau_2$.

Case T-APP: We consider two possible derivations of parallel reduction.

First, $e_2 e_1 \Rightarrow e'_2 e'_1$ with $e_2 \Rightarrow e'_2$ and $e_1 \Rightarrow e'_1$ for some e'_2 and e'_1 , due to R-REFL or R-APPCGR. In this case, $\Gamma \vdash e'_2 e'_1 : [e'_1/x]\tau_2$ by T-APP and induction hypotheses. Since $[e'_1/x]\tau_2 \equiv [e_1/x]\tau_2$ by [Lemma 2.4.3](#), we have $\Gamma \vdash e'_2 e'_1 : [e_1/x]\tau_2$ from T-CONV as desired.

Second, $e_2 = \lambda x:\tau \Rightarrow s$, and $(\lambda x:\tau \Rightarrow s) e_1 \Rightarrow [e'_1/x]s'$ for some τ , e'_1 , s and s' , with $s \Rightarrow s'$ and $e_1 \Rightarrow e'_1$, due to R-APP. It follows that $\lambda x:\tau \Rightarrow s \Rightarrow \lambda x:\tau \Rightarrow s'$, from which $\Gamma \vdash \lambda x:\tau \Rightarrow s' : \Pi x:\tau_1, \tau_2$ by the induction hypothesis. We then have $x : \tau, \Gamma \vdash s' : \tau_2$ with $\tau \equiv \tau_1$ and $\Gamma \vdash \tau :: *$ by [Lemma 2.4.11](#) (inversion lemma). It then follows that $\Gamma \vdash [e'_1/x]s' : [e'_1/x]\tau_2$ by [Lemma 2.4.8](#), because $\Gamma \vdash e'_1 : \tau$ by T-CONV and the induction hypothesis. Finally, T-CONV gives us $\Gamma \vdash [e'_1/x]s' : [e_1/x]\tau_2$ since $[e'_1/x]\tau_2 \equiv [e_1/x]\tau_2$ by [Lemma 2.4.3](#).

Case T-IF: Again, we consider two possible derivations of parallel reduction.

First, $\text{if } e_0 \text{ then } e_1 \text{ else } e_2 \Rightarrow \text{if } e'_0 \text{ then } e'_1 \text{ else } e'_2$ for some e'_0 , e'_1 and e'_2 , with $e_0 \Rightarrow e'_0$, $e_1 \Rightarrow e'_1$ and $e_2 \Rightarrow e'_2$, because of R-REFL or R-IFCGR. By the induction hypotheses and T-IF, we obtain $\Gamma \vdash \text{if } e'_0 \text{ then } e'_1 \text{ else } e'_2 : [e'_0/z]\tau$. The kinding side condition of T-IF can be discharged by [Lemma 2.4.1](#) and the induction hypothesis. T-CONV then gives us $\Gamma \vdash \text{if } e'_0 \text{ then } e'_1 \text{ else } e'_2 : [e_0/z]\tau$ because $[e_1/z]\tau \equiv [e_0/z]\tau$ by [Lemma 2.4.3](#).

Second, $e_0 = \text{true}$ and $\text{if true then } e_0 \text{ else } e_1 \Rightarrow e'_1$ for some e'_1 with $e_1 \Rightarrow e'_1$, due to R-IFTRUE (the case of $e_0 = \text{false}$ due to R-IFFALSE is similar). In this case, we want to show $\Gamma \vdash e'_1 : [\text{true}/z]\tau$, but that is immediate from the induction hypothesis.

The remaining cases about kinding derivations are similar: the proof for $\mathsf{K}\text{-PI}$ is similar to the case of $\mathsf{T}\text{-ABS}$, and the proof for $\mathsf{K}\text{-IF}$ is similar to $\mathsf{T}\text{-IF}$ (but easier as substitutions do not occur in kinds). \square

Since parallel reduction is a more liberal relation than the call-by-value or call-by-name semantics, these standard semantics also inherit the preservation property.

3. OBLIVIOUS ALGEBRAIC DATA TYPES

As discussed in [Chapter 1](#), to support private data structures, we have to ensure no private information, including data’s structural information if policy makers choose to hide it, can be inferred by observing the representation or manipulation of private data. In this chapter, we propose a novel representation of structured data types, which we call *oblivious algebraic data types* (OADTs). Our solution combines dependent types with language constructs for oblivious computation, and a security-type system which ensures that adversaries learn nothing more than the output of the function and the input they provide.

In summary, this chapter presents the following contributions:

- We observe that public views of private ADTs can be naturally expressed using dependent types with large elimination, allowing for a clean specification of what information is released at runtime.
- Exploiting this observation, we develop λ_{OADT} , a core calculus for writing oblivious programs using OADTs, whose strong type system ensures computations are secure.

The core calculus λ_{OADT} and its metatheory have been mechanically formalized in the Coq proof assistant. An artifact containing both these developments is publicly available [48].

3.1 Overview

To illustrate our approach, consider the simple function in [Figure 3.1](#), which looks for an element in a search tree by recursing over the tree. Suppose that Alice, the owner of a search tree, and Bob, the owner of some integer, want to check whether Bob’s integer is a member of Alice’s tree, without revealing any information to each other beyond what each can learn from their private data and the output. We adopt a variation of the

```
data tree = Leaf | Node ℤ tree tree
fn lookup (x : ℤ) (t : tree) : ℬ =
  match t with
  | Leaf ⇒ false
  | Node y t1 tr ⇒
    if x ≤ y then if y ≤ x then true
                  else lookup x t1
    else lookup x tr
```

Figure 3.1. Lookup element in a search tree

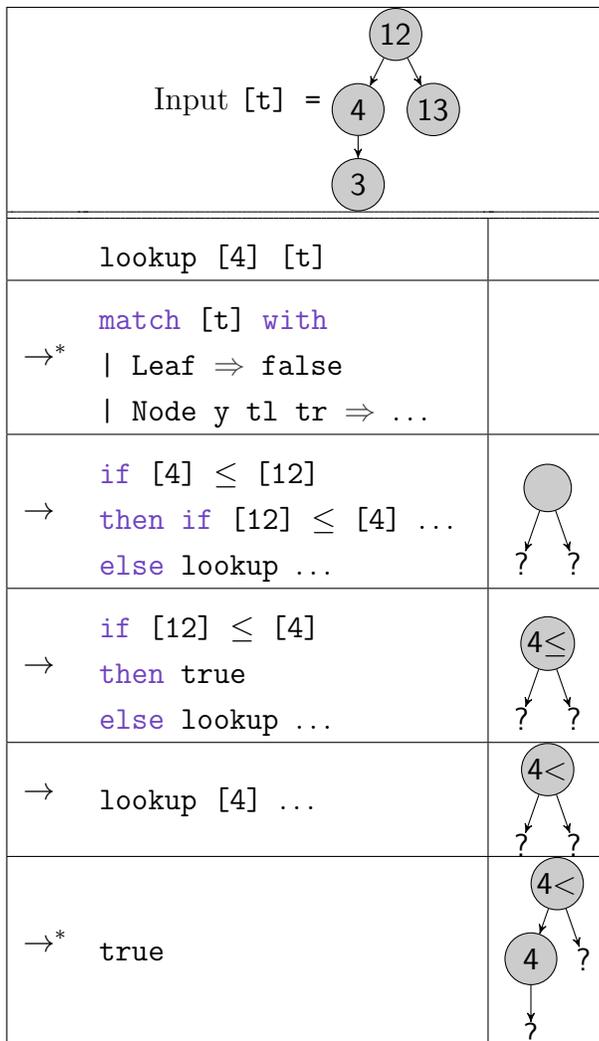


Figure 3.2. Execution trace of `lookup [4] [t]`. The columns show the current state of the program, and information learned by the owner of the lookup key, respectively.

standard semi-honest threat model from multiparty computation, where an untrusted party can observe every intermediate execution step of the program under a small-step operational semantics (Section 2.1). Protecting against such a powerful attacker inevitably impacts the performance of secure applications, a point we will discuss in more detail at the end of this overview. For now, let us consider the implications of this attack model on our current example.

Under this threat model, Bob can glean information about Alice’s tree just by examining how it affects the control flow of the program, *even if the tree is perfectly obfuscated*. To

see how, consider the execution trace of `lookup` [4] [t] shown in Figure 3.2, where [t] is the tree shown in the first row. The first column of each subsequent row shows the current execution step, while the second column shows what Bob can infer at that step. We use square brackets to denote that [t] is an *oblivious value*, i.e., it cannot be directly observed by a party. At each recursive call to `lookup`, there are two points that depend on the structure of the tree: the `match` statement that checks whether to recurse, and the `if` statement that decides which subtree to recurse on. As the fourth row illustrates, the branch `match` takes reveals some information about the structure of the current tree (it is non-empty) to Bob. The fifth and sixth row of the figure similarly show how the `if` statement reveals information about the relationship of the key to the value in the current node. By examining the program immediately following each such test, Bob adds to his knowledge of Alice’s tree. At the end, Bob learns not only output of the function, but also a partial view of the tree’s structure (including the exact node 4 is stored in); this view could be further refined by subsequent `lookup` operations.

Note that the participants of any terminating multiparty computation have to agree to share *some* public information: intuitively, simply knowing the number of intermediate steps in an execution of `lookup` leaks some upper bound on the number of nodes in the tree. Once that concession is made, the choice becomes *what* information to share: maybe the owner of the tree is okay with sharing its spine, but not the values stored in its internal nodes, or perhaps with revealing some upper bound on its depth¹. The goal then is to enable parties to compute functions over private data in a way that *only* depends on some mutually agreed upon public view of that data.

3.1.1 Encoding Private Data and Policies

The first component to our solution is our representation of both private data and the public information about the data that can be freely shared. We call this publicly shared information a *public view*, reflecting that it is some projection of the full data. Each public view corresponds to a *privacy policy* governing the data. Formally, policies are encoded

¹↑In the case the owner is okay with sharing the entire tree, the computation becomes quite efficient indeed!

<pre> obliv $\widehat{\text{tree}}$ (k : \mathbb{N}) = if k = 0 then 1 else 1 $\hat{+}$ $\widehat{\mathbb{Z}}$ \times $\widehat{\text{tree}}$ (k-1) \times $\widehat{\text{tree}}$ (k-1) </pre>	<pre> obliv $\widehat{\text{tree}}'$ (s : spine) = match s with SLeaf \Rightarrow 1 SNode s1 sr \Rightarrow 1 $\hat{+}$ $\widehat{\mathbb{Z}}$ \times $\widehat{\text{tree}}'$ s1 \times $\widehat{\text{tree}}'$ sr </pre>
(a) Maximum depth as public view	(b) Upper bound of spine as public view

Figure 3.3. Oblivious trees

as *oblivious algebraic data types* (OADTs), dependent types that take a public view as a parameter. The body of an OADT is the type of the private components of a data type, which are built using *oblivious* (i.e., secure) type formers, e.g., oblivious fixed-width integer ($\widehat{\mathbb{Z}}$) and oblivious sum ($\hat{+}$). By convention, we use $\hat{}$ to denote the oblivious version of something. Essentially, an OADT is a type-level function that maps the public view of a value to its *private representation*, i.e., the shape of its private component. [Section 3.2.4](#) formalizes oblivious data values, but the high-level intuition is that an observer of an execution trace cannot distinguish between the values of an oblivious type. When examining the trace in [Figure 3.2](#), the oblivious integer [4] is *indistinguishable* from [12], for example.

[Figure 3.3a](#) gives an example of an oblivious tree whose public view is its maximum depth. In general, a public view can be any public data type. We say tree is the *public type* or public counterpart of the OADT $\widehat{\text{tree}}$. The key idea behind oblivious ADTs is to construct a representation of private data from the public view. As a consequence, private values with the same public view are *indistinguishable* to an attacker, as their private representation is completely determined by the public view. For example, all oblivious trees with a maximum depth of two have the same private representation, regardless of the actual depth of the tree:

$$\widehat{\text{tree}}\ 2 \equiv 1 \hat{+} \widehat{\mathbb{Z}} \times (1 \hat{+} \widehat{\mathbb{Z}} \times 1 \times 1) \times (1 \hat{+} \widehat{\mathbb{Z}} \times 1 \times 1)$$

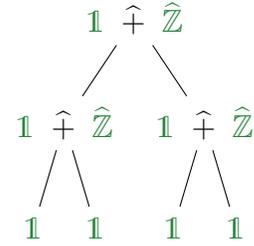


Figure 3.4. Oblivious tree with a maximum depth of two

Using data to compute a type is an example of *large elimination* from dependent type theory, where it is commonly used to recursively define propositions from terms. In this example, the type of an oblivious tree is computed from the public view 2, resulting in the *type value* on the right hand side, which stipulates the “shape” of the private data. This type roughly corresponds to the tree shown in [Figure 3.4](#). Every tree of this type is padded to depth 2, even a single “leaf”, to avoid leaking structural information. This padding is implied by the use of oblivious coproduct $\hat{+}$, as the left injection (e.g., a `Leaf`) and the right injection (e.g., a `Node`) of an oblivious sum will be indistinguishable. The adversaries can not tell them apart by inspecting the payload, even if the two components have different types. The “tag” of a sum value is of course obfuscated as well. Constructing oblivious data types in this way ensures that all private values corresponding to a particular view are indistinguishable to an attacker: an empty tree, singleton tree, a tree with two elements, or a complete tree of depth 2 all appear the same to an attacker.

[Figure 3.3b](#) shows the type of oblivious trees using an upper bound on its spine as the public view, where the spine is another user-defined ADT. This definition releases more public information than the one in [Figure 3.3a](#), but it also enjoys a more efficient representation, as it requires less padding than a complete tree.

[Figure 3.5](#) presents an oblivious type for a simplified version of the medical record with the “either-or” policy from the introduction ([Chapter 1](#)). A `patient` record consists of their ID, age, height and weight. The public view in this example consists of either a patient’s ID (`Known_id`), or their height and weight (`Known_data`); their age is always private. The corresponding oblivious type $\widehat{\text{patient}}$ is straightforward: it is the oblivious

```

// Id, age, height and weight
data patient = Patient ℤ ℤ ℤ ℤ
data patient_view = Known_id ℤ
                    | Known_data ℤ ℤ
obliv  $\widehat{\text{patient}}$  (v : patient_view) =
  match v with
  | Known_id _ =>  $\widehat{\mathbb{Z}} \times \widehat{\mathbb{Z}} \times \widehat{\mathbb{Z}}$ 
  | Known_data _ _ =>  $\widehat{\mathbb{Z}} \times \widehat{\mathbb{Z}}$ 

```

Figure 3.5. Either-or policy as an OADT

data that has been omitted from the public view. If the ID is disclosed, for example, then the oblivious type is essentially an encrypted version of the remaining 3 fields. Oblivious ADTs are expressive enough to directly support this kind of “either-or” policy.

Conceptually, OADTs generalize the notion of secure fixed-width integers to secure structured data, as illustrated in Figure 3.6. Every fixed-width integer (of type \mathbb{Z}) can be sent to its secure value in $\widehat{\mathbb{Z}}$ by “encryption”, and a secure integer can be converted back to \mathbb{Z} by “decryption”. We call these conversion functions *section* (e.g., $\widehat{\mathbb{Z}}\#s$) and *retraction* (e.g., $\widehat{\mathbb{Z}}\#r$). The names reflect their expected semantics: applying retraction to the section of a value should produce the same value.

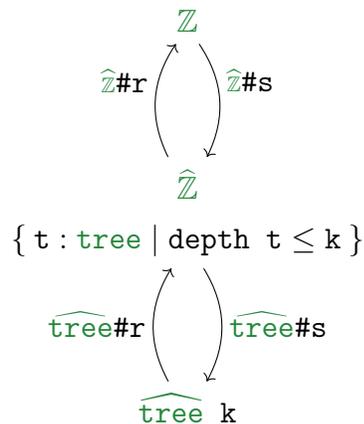


Figure 3.6. Public and oblivious types

Importantly, while the oblivious integer type $\widehat{\mathbb{Z}}$ does not appear to have much structure, it nonetheless has an implicit policy: the public view of an integer is its bit width. If we use 32-bit integers, for example, \mathbb{Z} is the set of all integers whose bit width is 32, and $\widehat{\mathbb{Z}}$ is the set of their “encrypted” values, related by a pair of conversion functions. Similarly, $\widehat{\text{tree}}\ k$ consists of the secure encodings of trees that have at most k layers. Like $\widehat{\mathbb{Z}}$, $\widehat{\text{tree}}\ k$ is equipped with a section function, $\widehat{\text{tree}}\#s$, and a retraction function, $\widehat{\text{tree}}\#r$, which convert public values of tree to their oblivious counterparts and back. Crucially, just as the oblivious integers in $\widehat{\mathbb{Z}}$ are indistinguishable, the elements of $\widehat{\text{tree}}\ k$ are also indistinguishable.

3.1.2 Enforcing Policies

Oblivious ADTs are only half the solution to secure computation; it still remains to ensure computations over private values are also oblivious. Even if an attacker cannot tell which values are being compared in `if [4] ≤ [12] then ... else ...`, they can still learn something about their relationship just by knowing the expression it steps to, as we saw in our previous example. To prevent these sorts of information leaks, we have designed λ_{OADT} , a pure functional language for writing secure computations over OADTs. λ_{OADT} is equipped with dependent types with large elimination to express OADTs, and type-based information flow control to guarantee oblivious computations.

Key to this calculus are its operations for securely constructing and destructing oblivious data values. As an example of these operations, consider the following λ_{OADT} expression, which compares two secure integers to determine what value to return:

`mux ([0] $\hat{\leq}$ [1]) ([2] $\hat{+}$ [3]) ([4] $\hat{+}$ [5])`

We use notations $\hat{+}$ and $\hat{\leq}$ for the oblivious versions of $+$ and \leq , such that $[4] \hat{+} [3] \longrightarrow [7]$ and $[4] \hat{\leq} [3] \longrightarrow [\text{false}]$.² Here, `mux` (short for *multiplexer*) is a special conditional which returns an oblivious value according to the value of an oblivious boolean. In order to avoid leaking information, `mux` generates the same evaluation trace regardless of the value of the private condition. To do so, it fully evaluates *both* branches before stepping to the final (oblivious) result:

$$\begin{aligned} & \text{mux } ([0] \hat{\leq} [1]) ([2] \hat{+} [3]) ([4] \hat{+} [5]) \\ \longrightarrow & \text{mux } [\text{true}] ([2] \hat{+} [3]) ([4] \hat{+} [5]) \\ \longrightarrow & \text{mux } [\text{true}] [5] ([4] \hat{+} [5]) \longrightarrow \text{mux } [\text{true}] [5] [9] \longrightarrow [5] \end{aligned}$$

Replacing `[0]` with `[6]` in the initial expression yields the same execution trace, modulo the private values at each step. Thus, nothing about the private information can be inferred by observing the execution:

$$\begin{aligned} & \text{mux } ([6] \hat{\leq} [1]) ([2] \hat{+} [3]) ([4] \hat{+} [5]) \\ \longrightarrow & \text{mux } [\text{false}] ([2] \hat{+} [3]) ([4] \hat{+} [5]) \\ \longrightarrow & \text{mux } [\text{false}] [5] ([4] \hat{+} [5]) \longrightarrow \text{mux } [\text{false}] [5] [9] \longrightarrow [9] \end{aligned}$$

The oblivious sum pattern matching statement (`match`) behaves similarly, with the additional wrinkle that the pattern variables of the “wrong” branch are bound to some arbitrary oblivious values, which [Section 3.2](#) explains in full detail.

λ_{OADT} is equipped with a security-type system [27], to ensure the correct use of its secure operations. The full details of this type system can be found in [Section 3.2.3](#), but at a high-level it enforces three key policies. First, oblivious types can only be built from oblivious types. For example, an oblivious coproduct cannot be built from public types, such as $\mathbb{B} \hat{+} \mathbb{Z}$.

²↑We abuse the notation $\hat{+}$ to mean both oblivious sum and oblivious integer addition.

```

fn lookup0 (x :  $\widehat{\mathbb{Z}}$ ) (k :  $\mathbb{N}$ ) :  $\widehat{\text{tree}}$  k  $\rightarrow$   $\widehat{\mathbb{B}}$  =
  if k = 0
  then  $\lambda\_ \Rightarrow \widehat{\mathbb{B}}\#s$  false
  else  $\lambda t \Rightarrow \widehat{\text{match}}$  t with
    |  $\widehat{\text{inl}}$  _  $\Rightarrow \widehat{\mathbb{B}}\#s$  false
    |  $\widehat{\text{inr}}$  (y, t1, tr)  $\Rightarrow$ 
      mux (x  $\widehat{\leq}$  y)
        (mux (y  $\widehat{\leq}$  x) ( $\widehat{\mathbb{B}}\#s$  true) (lookup0 x (k-1) t1))
        (lookup0 x (k-1) tr)

```

Figure 3.7. Oblivious lookup function in λ_{OADT}

If this were allowed, an adversary could infer whether a value of this type is a left or a right injection by observing the payload. Second, secure operations like `mux` can only be applied to oblivious terms. `mux [true] 1 2` is prohibited, for example, as knowing the public result of this `mux` reveals the oblivious discriminée. Third, types are treated as public information, otherwise the parties could not even agree on the data representation. Thus, oblivious types can only depend on public terms: `mux [true] $\widehat{\mathbb{B}}$ $\widehat{\mathbb{Z}}$` is not a valid type in λ_{OADT} .

Figure 3.7 presents an oblivious implementation of the `lookup` function for the oblivious tree from Figure 3.3a. While the high-level program logic is the same, extra care is needed to ensure correct use of the oblivious tree. First, the function takes an extra argument for the public view; the argument needs to be correctly passed to every recursive call. Second, the function eliminates the public view, following the definition of $\widehat{\text{tree}}$, before accessing any secure data. Third, public constants and operations are replaced by their secure counterparts: e.g., `if` is replaced by `mux`. Similarly, the constants `true` and `false` are wrapped by the $\widehat{\mathbb{B}}\#s$ operation (i.e., boolean section), which acts like a coercion from public booleans to oblivious booleans. This implementation is guaranteed to be secure, although it is not quite pleasing to write due to the intermixing of privacy policies and program logic. Chapter 4 will introduce a more ergonomic language that allows programmers to write functionality just like the “standard” implementation of `lookup` in Figure 3.1.

3.1.3 Performance Implication of the Threat Model

Before presenting a detailed accounting of our calculi for oblivious computation, we pause to discuss the consequences of our chosen threat model, where attackers can observe every program state in executions.

Protecting against such a strong attacker necessarily comes with a cost: many of the asymptotic efficiency benefits normally enjoyed by ADTs are lost in the MPC setting. While the `lookup` function from our running example provides a simple and familiar illustration of OADTs, it is also not as performant as its insecure counterpart. In order to avoid leaking private information via control flow channels, `lookup` *must* touch all the elements in the tree; there is no way to implement a logarithmic oblivious lookup function for this particular OADT in λ_{OADT} . For fold-like computations that touch the entire data structure (e.g., `map`), however, the right choice of a public view (e.g., a tree whose spine is its public view) allows OADTs to feature similar asymptotic behavior to standard ADTs.

While sacrificing some performance gains for security, OADTs provide other advantages over unstructured data, much like their non-oblivious counterparts. OADTs enable users to more easily write computations over data that is naturally represented using ADTs, such as file systems, organizational hierarchies, probability tree diagram, query languages, and decision trees. Complex policies, such as the either-or policy, can be encoded as OADTs as well. Using the language introduced in [Chapter 4](#), users can quickly prototype secure computation over structured data, and explore the impact of different public views on a computation.

3.2 λ_{OADT} , Formally

This section formalizes λ_{OADT} , a core calculus for programming with OADTs. The calculus described in this section has been mechanized in the Coq proof assistant and can be found in the publicly available artifact [\[48\]](#).

$e, \tau ::=$	$\mathbf{1} \mid \mathbb{B} \mid \widehat{\mathbb{B}} \mid \tau \times \tau \mid \tau + \tau \mid \tau \widehat{+} \tau$ $\Pi x : \tau, \tau$ x $() \mid \mathbf{true} \mid \mathbf{false}$ $\lambda x : \tau \Rightarrow e$ $\mathbf{let} \ x = e \ \mathbf{in} \ e$ $e \ e \mid \widehat{T} \ e$ $\mathbf{if} \ e \ \mathbf{then} \ e \ \mathbf{else} \ e$ $\mathbf{mux} \ e \ e \ e$ $(e, e) \mid \pi_b \ e$ $\iota_b \langle \tau \rangle \ e \mid \widehat{\iota}_b \langle \tau \rangle \ e$ $\mathbf{match} \ e \ \mathbf{with} \ x \Rightarrow e \mid x \Rightarrow e$ $\widehat{\mathbf{match}} \ e \ \mathbf{with} \ x \Rightarrow e \mid x \Rightarrow e$ $\mathbf{fold} \langle T \rangle \ e \mid \mathbf{unfold} \langle T \rangle \ e$ $\widehat{\mathbb{B}} \# s \ e$ $[b] \mid [\iota_b \langle \widehat{\omega} \rangle \ \widehat{v}]$	EXPRESSIONS: simple types dependent function type variable unit and boolean values function abstraction let binding expression and type application conditional atomic conditional pair and projection (oblivious) sum injection sum elimination oblivious sum elimination iso-recursive type intro. and elim. boolean section runtime boxed values
$D ::=$	$\mathbf{data} \ T = \tau$ $\mathbf{fn} \ x : \tau = e$ $\mathbf{obliv} \ \widehat{T} \ (x : \tau) = \tau$	GLOBAL DEFINITIONS: algebraic data type definition (recursive) function definition (recursive) oblivious type definition
$\widehat{\omega} ::=$	$\mathbf{1} \mid \widehat{\mathbb{B}} \mid \widehat{\omega} \times \widehat{\omega} \mid \widehat{\omega} \widehat{+} \widehat{\omega}$	OBLIVIOUS TYPE VALUES
$\widehat{v} ::=$	$() \mid [b] \mid (\widehat{v}, \widehat{v}) \mid [\iota_b \langle \widehat{\omega} \rangle \ \widehat{v}]$	OBLIVIOUS VALUES
$v ::=$	$\widehat{v} \mid b \mid (v, v) \mid \lambda x : \tau \Rightarrow e$ $\iota_b \langle \tau \rangle \ v \mid \mathbf{fold} \langle T \rangle \ v$	VALUES

Figure 3.8. λ_{OADT} syntax

3.2.1 Syntax

The core syntax of λ_{OADT} is shown in [Figure 3.8](#). For simplicity, the core calculus of λ_{OADT} does not include primitive fixed-width integers. We discuss how the language may be extended with primitive integers in [Section 4.3](#). As λ_{OADT} is dependently typed, types and terms belong to the same syntactic class, although by convention, we use the metavariable τ to refer to types, and e to terms. λ_{OADT} programs consist of an expression and a global context of public ADTs, oblivious ADTs, and functions. These are defined using **data**, **obliv**,

and `fn`, respectively. Using a global set of function definitions naturally supports general recursion and mutual recursion. When possible, we use `x` for function names, `T` for public ADT names, and \hat{T} for the names of oblivious ADTs.

Types in λ_{OADT} include dependent function types (Π), sums ($+$), products (\times), and booleans (\mathbb{B}); as well as oblivious sums ($\hat{+}$) and booleans ($\hat{\mathbb{B}}$). We do not include a type for oblivious products, as they can be encoded via normal products with oblivious components. In λ_{OADT} , the typing rules and semantics for oblivious types are quite different from their public counterparts, which is why we choose to assign them distinct syntax, as opposed to using security labels [29]. λ_{OADT} supports type-level computation via large elimination, allowing users to compute types from terms using application, `let`, `if`, and `match`. Sum and product types are also allowed to have both oblivious and public components, allowing types to contain a mixture of public and private data.

Terms in λ_{OADT} are largely standard. A subscript distinguishes between left or right injection (ι_b) and projection (π_b), where the metavariable `b` is either `true` or `false`. We also use the more conventional synonyms `inl` (`inr`) and π_1 (π_2) for ι_{true} (ι_{false}) and π_{true} (π_{false}). Injections are annotated with their full type, in order to completely determine its data representation. λ_{OADT} has a nominal type system, so `fold` and `unfold` take the name of a public ADT, instead of a recursive type definition (i.e., μ type). The atomic conditional `mux` is the core oblivious construct in λ_{OADT} ; as Section 3.1 discussed, `mux` fully evaluates both of its branches before taking a single atomic step to the correct branch. Other oblivious constructs include boolean section $\hat{\mathbb{B}}\#s$, which builds an oblivious boolean from its argument, and constructors ($\hat{\iota}_b$) and an eliminator ($\widehat{\text{match}}$) for oblivious sums.

In addition to the expected sorts of public values, λ_{OADT} also includes oblivious values for booleans (`[b]`) and sums (`$[\iota_b \langle \hat{\omega} \rangle \hat{v}]$`). In general, these oblivious values do not appear in the definitions in the global context: they are either provided by the data owner as the arguments to a global function at runtime, or created by evaluating $\hat{\mathbb{B}}\#s$ or $\hat{\iota}_b$. As Section 3.1 discussed, these “boxed” values represent secure data which cannot be observed by an adversary. Since λ_{OADT} has type-level computation, we also define a class of oblivious type values ($\widehat{\omega}$). Such values are built from a combination of oblivious base types and the other oblivious polynomial type formers.

3.2.2 Semantics

Figure 3.9 defines a relation for the small-step operational semantics of λ_{OADT} . The judgment of this relation has the form $\Sigma \vdash e \longrightarrow e'$, and is read as “ e steps to e' under the global context Σ ”. Since a λ_{OADT} program is evaluated under a fixed global context, we often abbreviate this judgment as $e \longrightarrow e'$, referring to Σ only when needed. The S-CTX rule uses the evaluation contexts (\mathcal{E}) defined at the bottom of Figure 3.9 to evaluate subexpressions. While these evaluation contexts are not inductively defined, it is possible to recursively apply S-CTX when evaluating subterms.

Most of the non-oblivious reduction rules are standard. For brevity, several of the rules use the `ite` meta-function, which returns e_1 when its first argument is `true`, and e_2 otherwise. S-IF is essentially the following two rules, for example:

$$\begin{array}{c}
 \text{S-IFTRUE} \\
 \hline
 \text{if true then } e_1 \text{ else } e_2 \longrightarrow e_1
 \end{array}
 \qquad
 \begin{array}{c}
 \text{S-IFFALSE} \\
 \hline
 \text{if false then } e_1 \text{ else } e_2 \longrightarrow e_2
 \end{array}$$

To ensure that oblivious rules avoid leaking information, they require that any subexpressions have been fully evaluated before an oblivious expression is reduced. As an example, S-CTX must be used to reduce the type and payload of an oblivious injection \hat{t} to values before the S-OINJ rule can be applied to obtain the oblivious value. The other oblivious rules (e.g., S-SEC and S-MUX) are similar.

The most interesting evaluation rule is S-OMATCH, which also ensures that an adversary can not infer anything about the oblivious value being eliminated. In contrast to other oblivious elimination rules like S-MUX, each branch binds the value stored in the sum to its pattern variables. This begs the question of how to instantiate this variable when evaluating the “wrong” branch. Since this branch is eventually discarded when the resulting `mux` is evaluated, we opt to simply instantiate this variable with an arbitrary payload of the right type. This value is synthesized using the auxiliary relation, $\hat{v} \leftarrow \hat{\omega}$, which is also shown in

$e \longrightarrow e'$		
$\frac{\text{S-CTX} \quad e \longrightarrow e'}{\mathcal{E}[e] \longrightarrow \mathcal{E}[e']}$	$\frac{\text{S-FUN} \quad \text{fn } x:\tau = e \in \Sigma}{x \longrightarrow e}$	$\frac{\text{S-OADT} \quad \text{obliv } \hat{T} \quad (x:\tau) = \tau' \in \Sigma}{\hat{T} \ v \longrightarrow [v/x]\tau'}$
$\frac{\text{S-APP}}{(\lambda x:\tau \Rightarrow e) \ v \longrightarrow [v/x]e}$	$\frac{\text{S-LET}}{\text{let } x = v \ \text{in } e \longrightarrow [v/x]e}$	
$\frac{\text{S-IF}}{\text{if } b \ \text{then } e_1 \ \text{else } e_2 \longrightarrow \text{ite}(b, e_1, e_2)}$		
$\frac{\text{S-MATCH}}{\text{match } \iota_b \langle \tau \rangle \ v \ \text{with } x \Rightarrow e_1 \ x \Rightarrow e_2 \longrightarrow \text{ite}(b, [v/x]e_1, [v/x]e_2)}$		
$\frac{\text{S-PROJ}}{\pi_b \ (v_1, v_2) \longrightarrow \text{ite}(b, v_1, v_2)}$	$\frac{\text{S-UNFOLD}}{\text{unfold}\langle T \rangle \ (fold\langle T' \rangle \ v) \longrightarrow v}$	$\frac{\text{S-SEC}}{\hat{\mathbb{B}}\#s \ b \longrightarrow [b]}$
$\frac{\text{S-OINJ}}{\hat{\iota}_b \langle \hat{\omega} \rangle \ \hat{v} \longrightarrow [\iota_b \langle \hat{\omega} \rangle \ \hat{v}]}$	$\frac{\text{S-MUX}}{\text{mux } [b] \ v_1 \ v_2 \longrightarrow \text{ite}(b, v_1, v_2)}$	
$\frac{\text{S-OMATCH} \quad \hat{v}_1 \Leftarrow \hat{\omega}_1 \quad \hat{v}_2 \Leftarrow \hat{\omega}_2}{\widehat{\text{match}} \ [\iota_b \langle \hat{\omega}_1 + \hat{\omega}_2 \rangle \ \hat{v}] \ \text{with } x \Rightarrow e_1 \ x \Rightarrow e_2 \longrightarrow \text{mux } [b] \ \text{ite}(b, [\hat{v}/x]e_1, [\hat{v}_1/x]e_1) \ \text{ite}(b, [\hat{v}_2/x]e_2, [\hat{v}/x]e_2)}$		
$\hat{v} \Leftarrow \hat{\omega}$		
EVALUATION CONTEXTS		
$\mathcal{E} ::=$	$\begin{array}{l} \ \square \times \tau \ \ \hat{\omega} \times \square \ \ \square \hat{+} \tau \ \ \hat{\omega} \hat{+} \square \\ \ \text{let } x = \square \ \text{in } e \ e \ \square \ \ \square \ v \ \ \hat{T} \ \square \\ \ (\square, e) \ \ (v, \square) \ \ \pi_b \ \square \\ \ \iota_b \langle \tau \rangle \ \square \ \ \hat{\iota}_b \langle \square \rangle \ e \ \ \hat{\iota}_b \langle \hat{\omega} \rangle \ \square \\ \ \text{fold}\langle T \rangle \ \square \ \ \text{unfold}\langle T \rangle \ \square \\ \ \text{if } \square \ \text{then } e \ \text{else } e \\ \ \text{match } \square \ \text{with } x \Rightarrow e \ x \Rightarrow e \\ \ \widehat{\text{match}} \ \square \ \text{with } x \Rightarrow e \ x \Rightarrow e \\ \ \text{mux } \square \ e \ e \ \ \text{mux } v \ \square \ e \ \ \text{mux } v \ v \ \square \\ \ \hat{\mathbb{B}}\#s \ \square \end{array}$	$\begin{array}{l} \text{OT-UNIT} \quad \text{OT-OBOOL} \\ \hline (\) \Leftarrow \mathbf{1} \quad [b] \Leftarrow \hat{\mathbb{B}} \\ \\ \text{OT-PROD} \\ \frac{\hat{v}_1 \Leftarrow \hat{\omega}_1 \quad \hat{v}_2 \Leftarrow \hat{\omega}_2}{(\hat{v}_1, \hat{v}_2) \Leftarrow \hat{\omega}_1 \times \hat{\omega}_2} \\ \\ \text{OT-OSUM} \\ \frac{\hat{v} \Leftarrow \text{ite}(b, \hat{\omega}_1, \hat{\omega}_2)}{[\iota_b \langle \hat{\omega}_1 + \hat{\omega}_2 \rangle \ \hat{v}] \Leftarrow \hat{\omega}_1 \hat{+} \hat{\omega}_2} \end{array}$

Figure 3.9. λ_{OADT} semantics

Figure 3.9. Equipped with this relation, S-OMATCH can be straightforwardly reduced to a `mux` expression. To see how, consider the rule corresponding to the case where `b` is `false`:

$$\frac{\hat{v}_1 \Leftarrow \hat{\omega}_1 \quad \hat{v}_2 \Leftarrow \hat{\omega}_2}{\widehat{\text{match}} \text{ [inr} \langle \hat{\omega}_1 \hat{+} \hat{\omega}_2 \rangle \hat{v}] \text{ with } x_1 \Rightarrow e_1 \mid x_2 \Rightarrow e_2 \longrightarrow \text{mux} \text{ [false] } [\hat{v}_1/x_1]e_1 \text{ } [\hat{v}/x_2]e_2}$$

In the `true` branch of the resulting `mux` expression, the pattern variable `x1` is instantiated with an arbitrary oblivious value, \hat{v}_1 , while the corresponding pattern variable in the `false` branch is instantiated with the actual payload \hat{v} . Using this rule, the expression on the top below can step to either of the (indistinguishable) expressions on the bottom:

$$\begin{array}{ccc} \widehat{\text{match}} \text{ [inr} \langle (\hat{\mathbb{B}} \times \hat{\mathbb{B}}) \hat{+} \hat{\mathbb{B}} \rangle \text{ [false]] with } x_1 \Rightarrow \pi_2 \ x_1 \mid x_2 \Rightarrow x_2 & & \\ \swarrow & & \searrow \\ \text{mux} \text{ [false] } & & \text{mux} \text{ [false] } \\ (\pi_2 \text{ ([false], [true])}) & & (\pi_2 \text{ ([true], [false])}) \\ \text{[false]} & & \text{[false]} \end{array}$$

The pattern variables in the first branch can be substituted by any pair of oblivious booleans, e.g., `([false], [true])` or `([true], [false])`.

3.2.3 Type System

The type system of λ_{OADT} ensures that well-typed programs are secure, in that adversaries cannot glean information about private data by observing public information. To guarantee this, kinds in λ_{OADT} are augmented with a *security label* which constrains how information flows through a program:

$$\begin{array}{l} \kappa ::= \\ | \ *^{\text{A}} \ \text{Any} \\ | \ *^{\text{P}} \ \text{Public} \\ | \ *^{\text{O}} \ \text{Oblivious} \\ | \ *^{\text{M}} \ \text{Mixed} \end{array}$$

Types that can be treated as either public or oblivious are labeled with \mathbf{A} . In practice, this is almost always the unit type, but it includes other singleton types, e.g., $\mathbf{1} \times \mathbf{1}$. Types which are entirely public or entirely private have the labels \mathbf{P} and \mathbf{O} , respectively. Finally, types with a mixture of public and private data, e.g., $\mathbb{B} \times \widehat{\mathbb{B}}$, are labeled with \mathbf{M} . This label is also used to classify function types, which we will discuss in more detail shortly. Kinds form a secure join semi-lattice, as shown in [Figure 3.10](#), with \mathbf{M} being the most restrictive label. Unlike most secure type systems where types with a public label can be promoted to their secure counterparts, in $\lambda_{\text{O}ADT}$ public and oblivious labels are not compatible. We elide the security label of a kind when it is not relevant, e.g., $\Gamma \vdash \tau :: *$.

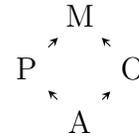


Figure 3.10. Semi-lattice on $\lambda_{\text{O}ADT}$ kinds

Programs in $\lambda_{\text{O}ADT}$ are typed using a pair of typing and kinding judgments; we denote these as $\Sigma; \Gamma \vdash e : \tau$ and $\Sigma; \Gamma \vdash \tau :: \kappa$, respectively. [Figure 3.11](#) and [Figure 3.12](#) give the kinding and typing rules for $\lambda_{\text{O}ADT}$. We elide Σ from these definitions, as they both assume a fixed global context. For brevity, we omit some side conditions about kinding from the typing rules; these can be found in the Coq development.

The kinding rules for $\lambda_{\text{O}ADT}$ are shown in [Figure 3.11](#). The rules for base types are straightforward. As previously mentioned, function types are assigned a mixed label. The reasons for this are two-fold: firstly, $\lambda_{\text{O}ADT}$ does not support oblivious function values. Secondly, this prevents function values from being used as the public view of oblivious types, making it easier for users to be sure oblivious types terminate. The subsumption rule $\mathbf{K-SUB}$ allows kinds to be converted to a more restricted label. This rule can be used with $\mathbf{K-PROD}$ to label a product type with the join of the labels of its components. $\mathbf{K-SUM}$ is similar, but it also includes the public label in the join, as the tag of a public sum is practically public. For example, $\mathbf{1} + \mathbf{1}$, which is equivalent to \mathbb{B} , should be kinded $*^{\mathbf{P}}$ instead of $*^{\mathbf{A}}$. Similarly, $\widehat{\mathbb{B}} + \widehat{\mathbb{B}}$ has to be kinded $*^{\mathbf{M}}$, the join of $*^{\mathbf{P}}$ and $*^{\mathbf{O}}$, as using it in an oblivious context risks leaking the tag. For similar reasons, $\mathbf{K-OSUM}$ requires the components of oblivious sums to also be oblivious. $\mathbf{K-OADT}$ requires the argument of an oblivious type to be well-typed according to its definition in the global context. It does not need to check the index is public, as it is

$$\boxed{\Gamma \vdash \tau :: \kappa}$$

$$\begin{array}{c}
\text{K-UNIT} \\
\hline
\Gamma \vdash \mathbb{1} :: *^{\mathbf{A}}
\end{array}
\quad
\begin{array}{c}
\text{K-BOOL} \\
\hline
\Gamma \vdash \mathbb{B} :: *^{\mathbf{P}}
\end{array}
\quad
\begin{array}{c}
\text{K-OB00L} \\
\hline
\Gamma \vdash \widehat{\mathbb{B}} :: *^{\mathbf{0}}
\end{array}
\quad
\begin{array}{c}
\text{K-ADT} \\
\text{data } \mathbb{T} = \tau \in \Sigma \\
\hline
\Gamma \vdash \mathbb{T} :: *^{\mathbf{P}}
\end{array}$$

$$\begin{array}{c}
\text{K-PI} \\
\Gamma \vdash \tau_1 :: * \quad x : \tau_1, \Gamma \vdash \tau_2 :: * \\
\hline
\Gamma \vdash \Pi_{x:\tau_1} \tau_2 :: *^{\mathbf{M}}
\end{array}
\quad
\begin{array}{c}
\text{K-PROD} \\
\Gamma \vdash \tau_1 :: \kappa \quad \Gamma \vdash \tau_2 :: \kappa \\
\hline
\Gamma \vdash \tau_1 \times \tau_2 :: \kappa
\end{array}
\quad
\begin{array}{c}
\text{K-SUM} \\
\Gamma \vdash \tau_1 :: \kappa \quad \Gamma \vdash \tau_2 :: \kappa \\
\hline
\Gamma \vdash \tau_1 + \tau_2 :: \kappa \sqcup *^{\mathbf{P}}
\end{array}$$

$$\begin{array}{c}
\text{K-OSUM} \\
\Gamma \vdash \tau_1 :: *^{\mathbf{0}} \quad \Gamma \vdash \tau_2 :: *^{\mathbf{0}} \\
\hline
\Gamma \vdash \tau_1 \widehat{+} \tau_2 :: *^{\mathbf{0}}
\end{array}
\quad
\begin{array}{c}
\text{K-OADT} \\
\text{obliv } \widehat{\mathbb{T}} (x:\tau) = \tau' \in \Sigma \quad \Gamma \vdash e : \tau \\
\hline
\Gamma \vdash \widehat{\mathbb{T}} e :: *^{\mathbf{0}}
\end{array}$$

$$\begin{array}{c}
\text{K-LET} \\
\Gamma \vdash e : \tau \quad x : \tau, \Gamma \vdash \tau' :: *^{\mathbf{0}} \\
\hline
\Gamma \vdash \text{let } x = e \text{ in } \tau' :: *^{\mathbf{0}}
\end{array}
\quad
\begin{array}{c}
\text{K-IF} \\
\Gamma \vdash e_0 : \mathbb{B} \quad \Gamma \vdash \tau_1 :: *^{\mathbf{0}} \quad \Gamma \vdash \tau_2 :: *^{\mathbf{0}} \\
\hline
\Gamma \vdash \text{if } e_0 \text{ then } \tau_1 \text{ else } \tau_2 :: *^{\mathbf{0}}
\end{array}$$

$$\begin{array}{c}
\text{K-MATCH} \\
\Gamma \vdash e_0 : \tau'_1 + \tau'_2 \quad x : \tau'_1, \Gamma \vdash \tau_1 :: *^{\mathbf{0}} \quad x : \tau'_2, \Gamma \vdash \tau_2 :: *^{\mathbf{0}} \\
\hline
\Gamma \vdash \text{match } e_0 \text{ with } x \Rightarrow \tau_1 \mid x \Rightarrow \tau_2 :: *^{\mathbf{0}}
\end{array}
\quad
\begin{array}{c}
\text{K-SUB} \\
\Gamma \vdash \tau :: \kappa \quad \kappa \sqsubseteq \kappa' \\
\hline
\Gamma \vdash \tau :: \kappa'
\end{array}$$

Figure 3.11. λ_{OADT} kinding rules

done when typing the global context. K-LET, K-IF and K-MATCH are the key components for large elimination. They both require the discriminée to be well-typed and the returned types to be obviously kinded. The K-MATCH rule is rather permissive in that it does not require the type of discriminée e_0 to be completely publicly typed. While it is unclear when a programmer would ever actually use a type-level discriminée with oblivious components, it does not leak any information either.

The typing rules for public constructs are largely standard. Since λ_{OADT} is dependently typed, T-IF and T-MATCH rely on an implicit motive that is specialized when typing branches³. This motive, (τ) , has a special free variable (z) which stands in for the term being eliminated. The type used for the **then** branch in T-IF $([\text{true}/z]\tau)$ concretizes the occurrences of this variable with **true**, for example. The typing rules for oblivious

³↑This strategy is in line with other dependently typed languages (e.g., Coq), which try to infer a motive when none is supplied by the programmer.

$\Gamma \vdash e : \tau$			
$\frac{\text{T-VAR}}{x : \tau \in \Gamma} \quad \Gamma \vdash x : \tau$	$\frac{\text{T-UNIT}}{\Gamma \vdash () : \mathbf{1}}$	$\frac{\text{T-LIT}}{\Gamma \vdash b : \mathbb{B}}$	$\frac{\text{T-FUN}}{\text{fn } x : \tau = e \in \Sigma} \quad \Gamma \vdash x : \tau$
$\frac{\text{T-ABS}}{x : \tau_1, \Gamma \vdash e : \tau_2} \quad \Gamma \vdash \tau_1 :: *$ $\Gamma \vdash \lambda x : \tau_1 \Rightarrow e : \Pi x : \tau_1, \tau_2$	$\frac{\text{T-APP}}{\Gamma \vdash e_2 : \Pi x : \tau_1, \tau_2} \quad \Gamma \vdash e_1 : \tau_1$ $\Gamma \vdash e_2 \ e_1 : [e_1/x] \tau_2$		
$\frac{\text{T-LET}}{\Gamma \vdash e_1 : \tau_1} \quad x : \tau_1, \Gamma \vdash e_2 : \tau_2$ $\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : [e_1/x] \tau_2$	$\frac{\text{T-PAIR}}{\Gamma \vdash e_1 : \tau_1} \quad \Gamma \vdash e_2 : \tau_2$ $\Gamma \vdash (e_1, e_2) : \tau_1 \times \tau_2$		
$\frac{\text{T-PROJ}}{\Gamma \vdash e : \tau_1 \times \tau_2}$ $\Gamma \vdash \pi_b \ e : \text{ite}(b, \tau_1, \tau_2)$	$\frac{\text{T-INJ}}{\Gamma \vdash e : \text{ite}(b, \tau_1, \tau_2)} \quad \Gamma \vdash \tau_1 + \tau_2 : *$ $\Gamma \vdash \iota_b \langle \tau_1 + \tau_2 \rangle \ e : \tau_1 + \tau_2$		
$\frac{\text{T-IF}}{\Gamma \vdash e_0 : \mathbb{B}}$ $\Gamma \vdash e_1 : [\text{true}/z] \tau \quad \Gamma \vdash e_2 : [\text{false}/z] \tau$ $\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : [e_0/z] \tau$			
$\frac{\text{T-MATCH}}{\Gamma \vdash e_0 : \tau_1 + \tau_2}$ $x : \tau_1, \Gamma \vdash e_1 : [\text{inl} \langle \tau_1 + \tau_2 \rangle \ x/z] \tau \quad x : \tau_2, \Gamma \vdash e_2 : [\text{inr} \langle \tau_1 + \tau_2 \rangle \ x/z] \tau$ $\Gamma \vdash \text{match } e_0 \text{ with } x \Rightarrow e_1 \mid x \Rightarrow e_2 : [e_0/z] \tau$			
$\frac{\text{T-FOLD}}{\text{data } T = \tau \in \Sigma} \quad \Gamma \vdash e : \tau$ $\Gamma \vdash \text{fold} \langle T \rangle \ e : T$	$\frac{\text{T-UNFOLD}}{\text{data } T = \tau \in \Sigma} \quad \Gamma \vdash e : T$ $\Gamma \vdash \text{unfold} \langle T \rangle \ e : \tau$	$\frac{\text{T-SEC}}{\Gamma \vdash e : \mathbb{B}}$ $\Gamma \vdash \widehat{\mathbb{B}} \# s \ e : \widehat{\mathbb{B}}$	$\frac{\text{T-MUX}}{\Gamma \vdash e_0 : \widehat{\mathbb{B}}} \quad \Gamma \vdash \tau :: *^0$ $\Gamma \vdash e_1 : \tau \quad \Gamma \vdash e_2 : \tau$ $\Gamma \vdash \text{mux } e_0 \ e_1 \ e_2 : \tau$
$\frac{\text{T-OINJ}}{\Gamma \vdash e : \text{ite}(b, \tau_1, \tau_2)} \quad \Gamma \vdash \tau_1 + \tau_2 :: *^0$ $\Gamma \vdash \widehat{\iota}_b \langle \tau_1 + \tau_2 \rangle \ e : \tau_1 + \tau_2$	$\frac{\text{T-OMATCH}}{\Gamma \vdash e_0 : \tau_1 + \tau_2} \quad \Gamma \vdash \tau :: *^0$ $x : \tau_1, \Gamma \vdash e_1 : \tau \quad x : \tau_2, \Gamma \vdash e_2 : \tau$ $\Gamma \vdash \widehat{\text{match}} \ e_0 \ \text{with } x \Rightarrow e_1 \mid x \Rightarrow e_2 : \tau$		$\frac{\text{T-BOXEDLIT}}{\Gamma \vdash [b] : \widehat{\mathbb{B}}}$
$\frac{\text{T-BOXEDINJ}}{[\iota_b \langle \widehat{\omega} \rangle \ \widehat{v}] \Leftarrow \widehat{\omega}}$ $\Gamma \vdash [\iota_b \langle \widehat{\omega} \rangle \ \widehat{v}] : \widehat{\omega}$		$\frac{\text{T-CONV}}{\Gamma \vdash e : \tau} \quad \tau \equiv \tau' \quad \Gamma \vdash \tau' :: *$ $\Gamma \vdash e : \tau'$	

Figure 3.12. λ_{OADT} typing rules

constructs are largely similar to their public counterparts, with the caveat that they place more constraints on their subterms: T-OINJ requires the type of its payload to have an oblivious kind, for example. In addition to requiring that their branches have oblivious kind, the typing rules for oblivious eliminators (T-MUX and T-OMATCH) are required to return types that do not depend on the discriminees, in order to avoid leaking information about the discriminees via their types. T-BOXEDLIT and T-BOXEDINJ type oblivious values, with the latter simply outsourcing it to the relation used in S-OMATCH.

The final typing rule, T-CONV, allows any well-typed term to be typed using an equivalent type, denoted $\Sigma \vdash \tau \equiv \tau'$. This equivalence is defined directly in terms of a *parallel reduction* relation, $\Sigma \vdash e \Rightarrow e'$, or simply $e \Rightarrow e'$. Parallel reduction is a more liberal version of our call-by-value semantics which allows, for example, reduction under binders and congruence rules. Two terms are then said to be equivalent when they can parallel reduce to the same term in zero or more steps:

$$\Sigma \vdash \tau_1 \equiv \tau_2 \triangleq \exists \tau. \Sigma \vdash \tau_1 \Rightarrow^* \tau \wedge \Sigma \vdash \tau_2 \Rightarrow^* \tau$$

Parallel reduction also plays an important role in the metatheory of λ_{OADT} , particularly in the proof of obliviousness ([Theorem 3.2.6](#)).

A subset of the parallel reduction rules are shown in [Figure 3.13](#); the remaining rules are similar to the rules in [Figure 2.6 \(Chapter 2\)](#) and can be found in our Coq development; despite their importance in the metatheory of λ_{OADT} , the parallel reduction rules are straightforward. As the figure shows, the rules are essentially more permissive versions of their counterparts in the step relation from [Figure 3.9](#). As an example, the parallel reduction rule for `mux`, R-MUX, does not reduce its branches to values, but immediately takes the corresponding branch, just like S-IF. While this rule would leak information if the condition of the `mux` could be reduced to an oblivious value, this does not occur in practice. The reason for this is that parallel reduction is only used for statically type checking programs, and this rule will therefore never be used at runtime, when private data is made available.

To type a λ_{OADT} program, we also check the definitions in the global context using the rules in [Figure 3.14](#). DT-FUN is straightforward: the type ascription needs to be well-kinded

e ⇒ e'

$$\begin{array}{c}
\text{R-REFL} \\
\frac{}{e \Rightarrow e} \\
\\
\text{R-APP} \\
\frac{e_1 \Rightarrow e'_1 \quad e_2 \Rightarrow e'_2}{(\lambda x:\tau \Rightarrow e_2) e_1 \Rightarrow [e'_1/x] e'_2} \\
\\
\text{R-FUN} \\
\frac{\text{fn } x:\tau = e \in \Sigma}{x \Rightarrow e} \\
\\
\text{R-OADT} \\
\frac{\text{obliv } \widehat{T} (x:\tau') = \tau \in \Sigma \quad e \Rightarrow e'}{\widehat{T} e \Rightarrow [e'/x] \tau} \\
\\
\text{R-MUX} \\
\frac{e_1 \Rightarrow e'_1 \quad e_2 \Rightarrow e'_2}{\text{mux } [b] e_1 e_2 \Rightarrow \text{ite}(b, e'_1, e'_2)} \\
\\
\text{R-SEC} \\
\frac{}{\widehat{\mathbb{B}}\#s b \Rightarrow [b]} \\
\\
\text{R-OINJ} \\
\frac{}{\widehat{v}_b \langle \widehat{\omega} \rangle \widehat{v} \Rightarrow [b] \langle \widehat{\omega} \rangle \widehat{v}} \\
\\
\text{R-OMATCH} \\
\frac{\widehat{v}_1 \leftarrow \widehat{\omega}_1 \quad \widehat{v}_2 \leftarrow \widehat{\omega}_2 \quad e_1 \Rightarrow e'_1 \quad e_2 \Rightarrow e'_2}{\widehat{\text{match}} [b] \langle \widehat{\omega}_1 + \widehat{\omega}_2 \rangle \widehat{v} \text{ with } x \Rightarrow e_1 \mid x \Rightarrow e_2 \Rightarrow \text{mux } [b] \text{ite}(b, [\widehat{v}/x] e'_1, [\widehat{v}_1/x] e'_1) \text{ite}(b, [\widehat{v}_2/x] e'_2, [\widehat{v}/x] e'_2)}
\end{array}$$

Figure 3.13. Subset of λ_{OADT} parallel reduction rules

$\Sigma \vdash D$

$$\begin{array}{c}
\text{DT-FUN} \\
\frac{\cdot \vdash \tau :: * \quad \cdot \vdash e : \tau}{\Sigma \vdash \text{fn } x:\tau = e} \\
\\
\text{DT-ADT} \\
\frac{\cdot \vdash \tau :: *^P}{\Sigma \vdash \text{data } T = \tau} \\
\\
\text{DT-OADT} \\
\frac{\cdot \vdash \tau :: *^P \quad x:\tau \vdash \tau' :: *^0}{\Sigma \vdash \text{obliv } \widehat{T} (x:\tau) = \tau'}
\end{array}$$

Figure 3.14. λ_{OADT} global definition typing rules

and the definition needs to be well-typed using an empty typing context. A definition may recursively refer to the name being defined, which is included in Σ . DT-ADT, the typing rule for public ADTs, simply requires the type to be completely public. The typing rule for oblivious ADTs, DT-OADT, requires that its index be completely public, as it is used as the public view. In contrast, the rest of the definition has to have an oblivious kind, under a context that includes the index x .

3.2.4 Type Safety and Obliviousness

This section presents sketches of the key metatheory proofs for λ_{OADT} 's type system. All the theorems in this section assume a well-typed global context. Firstly, λ_{OADT} enjoys the standard progress and preservation theorems:

Theorem 3.2.1 (Progress). *If $\cdot \vdash e : \tau$, then either $e \longrightarrow e'$ for some e' , or e is a value.*

*If $\cdot \vdash \tau : *^0$, then either $\tau \longrightarrow \tau'$ for some τ' , or τ is an oblivious type value.*

The proof of progress proceeds by mutual induction on typing and kinding derivation. The S-OMATCH case relies on the fact that every oblivious type value is inhabited, in order to find the oblivious value needed to reduce the “wrong” branch.

The preservation theorem also consists of two parts.

Theorem 3.2.2 (Preservation). *If $\Gamma \vdash e : \tau$, and $e \longrightarrow e'$, then $\Gamma \vdash e' : \tau$.*

If $\Gamma \vdash \tau :: \kappa$ and $\tau \longrightarrow \tau'$, then $\Gamma \vdash \tau' :: \kappa$.

The induction hypothesis for a direct proof of preservation is too weak to prove the T-IF and T-MATCH cases. Instead, we show that the step relation refines parallel reduction and then prove preservation for the more general relation.

Lemma 3.2.3 (Preservation for parallel reduction). *If $\Gamma \vdash e : \tau$, and $e \Rightarrow e'$, then $\Gamma \vdash e' : \tau$.*

If $\Gamma \vdash \tau :: \kappa$ and $\tau \Rightarrow \tau'$, then $\Gamma \vdash \tau' :: \kappa$.

The proof of [Lemma 3.2.3](#) depends on two additional lemmas. The first is a regularity lemma needed for the kinding constraints used by several typing rules.

Lemma 3.2.4 (Regularity). *If $\Gamma \vdash e : \tau$, then $\Gamma \vdash \tau :: \kappa$ for some κ .*

The second is that parallel reduction is confluent.

Lemma 3.2.5 (Confluence of parallel reduction). *If $e \Rightarrow^* e_1$ and $e \Rightarrow^* e_2$, then there exists e' such that $e_1 \Rightarrow^* e'$ and $e_2 \Rightarrow^* e'$.*

Interestingly, the regular call-by-value semantics of λ_{OADT} are not confluent, thanks to a combination of the (limited) nondeterminism in S-OMATCH and nontermination. Observe

that S-OMATCH can be applied with different choices of the arbitrary oblivious values. This is not a problem if the oblivious case expression terminates because the “wrong” branch will eventually be discarded. However, it is possible that the `mux` expression it steps to loops forever, such that the “wrong” branch is never discarded. Thankfully, the R-MUX rule is more liberal than S-MUX, ensuring that parallel reduction is confluent. Whenever $\widehat{\text{match}}$ parallel reduces to a `mux` expression, however, R-MUX will immediately discard the “wrong” branch, forcing both choices to converge within one step.

Obliviousness

Adversaries should not be able to infer any information about the private information (i.e., oblivious values) of well-typed λ_{OADT} programs by observing the whole execution of a λ_{OADT} program. To prove this, we first formalize a notion of *indistinguishability* for λ_{OADT} expressions:

Definition 3.2.1 (Indistinguishability). We say two expressions are *indistinguishable*, denoted by $e \approx e'$, if

1. they are both oblivious boolean values: $[b] \approx [b']$, or
2. they are both oblivious injections with the same type: $[\iota_b \langle \widehat{\omega} \rangle v] \approx [\iota_{b'} \langle \widehat{\omega} \rangle v']$, or
3. they are the same expression with indistinguishable sub-expressions.

Intuitively, two expressions are indistinguishable if they only differ in their oblivious values. Note that indistinguishability is a completely syntactic notion: two lambda abstractions are indistinguishable only if their bodies are indistinguishable. This is a direct consequence of our strong threat model: dishonest parties are capable of peeking “under the binders”, i.e., lambda abstractions are *not* black boxes to them. As an example, the functions $\lambda x \ y \Rightarrow x+y$ and $\lambda x \ y \Rightarrow y+x$ are not indistinguishable, even though their “big-step” behaviors are the same: if `mux [true] ($\lambda x \ y \Rightarrow x+y$) ($\lambda x \ y \Rightarrow y+x$)` were to step to $\lambda x \ y \Rightarrow x+y$, an attacker could learn about the private condition by inspecting the resulting function. More pleasantly, this syntactic definition enjoys a congruence property: plugging indistinguishable partial

programs into indistinguishable contexts is guaranteed to result in indistinguishable whole programs.

Equipped with this relation, we can now formally state the obliviousness theorem for λ_{OADT} :

Theorem 3.2.6 (Obliviousness). *If $e_1 \approx e_2$ and $\cdot \vdash e_1 : \tau_1$ and $\cdot \vdash e_2 : \tau_2$, then*

1. $e_1 \longrightarrow^n e'_1$ if and only if $e_2 \longrightarrow^n e'_2$ for some e'_2 .
2. if $e_1 \longrightarrow^n e'_1$ and $e_2 \longrightarrow^n e'_2$, then $e'_1 \approx e'_2$.

We write $e \longrightarrow^n e'$ to mean e reduces to e' in exactly n steps. The first piece of this theorem is a generalization of progress, and ensures that information is not leaked via a termination channel. The second piece says that for any two indistinguishable programs, an observer cannot learn anything about their oblivious values by examining the states they can step to. Taken together, these two properties ensure that an observer cannot learn anything about the private values in a well-typed λ_{OADT} program, even given the entire execution trace of that program. If we treat the observable parts of the intermediate execution states as a public channel, obliviousness provides a sort of noninterference property [21, 27], in that different private (i.e., high-security) inputs do not leak any information via this public channel.

The proof of [Theorem 3.2.6](#) is by induction on the derivation of $e_1 \longrightarrow^n e'_1$. The first part of the proof of obliviousness is a direct consequence of progress and the fact that well-typed values are only indistinguishable from other values. The second part is more involved, and requires the following two key lemmas to prove the S-MUX case:

Lemma 3.2.7. *If $\Gamma \vdash v : \tau$ and $\Gamma \vdash v' : \tau$, and $\Gamma \vdash \tau :: *^0$, then $v \approx v'$.*

Lemma 3.2.8. *If $v \approx v'$, $\Gamma \vdash v : \tau$, $\Gamma \vdash v' : \tau'$, and $\Gamma \vdash \tau :: *^0$, then $\tau \equiv \tau'$.*

[Lemma 3.2.7](#) states that all values of the same oblivious type are indistinguishable, and [Lemma 3.2.8](#) ensures that two indistinguishable, obviously-typed values have the same type up to type equivalence. The proofs of both lemmas proceed by induction on the typing derivation. Most of the proofs are straightforward, except for the case of T-CONV in both lemmas. Since applying the induction hypothesis requires that τ' also be oblivious, we need

to show that two equivalent, well-kinded types simultaneously have oblivious kinds, which follows from [Lemma 3.2.3](#):

Lemma 3.2.9. *If $\tau \equiv \tau'$, $\Gamma \vdash \tau :: *^0$, and $\Gamma \vdash \tau' :: *$, then $\Gamma \vdash \tau' :: *^0$.*

In practice, well-typed λ_{OADT} programs are functions that take arguments of oblivious types, such as `lookup0` from [Figure 3.7](#). The program built by supplying such a function with private inputs of the right types is indistinguishable from one built using different private inputs, thanks to the congruence property of indistinguishability and [Lemma 3.2.7](#). As a direct consequence of the obliviousness theorem, an attacker can not glean any information about the private inputs of such programs. This fact is captured in the following corollary about open λ_{OADT} programs:

Corollary 3.2.10. *If $x : \tau \vdash e : \tau$ with $\cdot \vdash \tau :: *^0$, then for any two values v_1 and v_2 of oblivious type τ' :*

1. $[v_1/x]e \longrightarrow^n e_1$ if and only if $[v_2/x]e \longrightarrow^n e_2$ for some e_2 .
2. $[v_1/x]e \longrightarrow^n e_1$ and $[v_2/x]e \longrightarrow^n e_2$ implies that e_1 and e_2 are indistinguishable, i.e., $e_1 \approx e_2$.

3.3 Conclusion

To our best knowledge, this work is the first programming language that supports hiding the structure of rich recursive data types in secure computations. We have presented λ_{OADT} , a core calculus for encoding oblivious programs over oblivious algebraic data types. λ_{OADT} combines dependent types with large elimination to represent oblivious algebraic data types, and provides a security-type system to ensure that computations reveal no private information over what is provided by the public view of the data. We have proved, mechanically, that our solution provides a strong and formal security guarantee: an adversary can not infer any private information, even given the entire execution trace of a program.

4. TAPE SEMANTICS: DYNAMIC ENFORCEMENT OF POLICIES

While λ_{OADT} provides a foundation for implementing secure computation that involves structured data and complex policies, writing secure programs directly in λ_{OADT} is challenging, especially for programmers who are not well-versed in dependent types or information flow control. This chapter introduces a semantics-based approach to enforce privacy policies dynamically. This novel execution model, called *tape semantics*, allows programs to include unsafe computations and repairs these unsafe computations at runtime. Tape semantics decouples privacy and programmatic concerns, enabling a sort of modular design that allows programmers to implement the functionality of their secure applications in a standard way. Using our language, clients can write a single function over private data, and then build an equivalent oblivious computation over some public view (i.e., policy) of that data. By switching views, users can explicitly trade off between how much information is leaked via public channels and the performance of the underlying computation.

In summary, this chapter presents the following contributions:

- To enable both a more pleasant programming experience and more modular programs, we develop an extension of λ_{OADT} , dubbed λ_{OADT^+} , which is equipped with a novel semantics which enables creating, from one single public program, oblivious programs with different public views.
- We present a reference semantics and connect it to tape semantics. The correspondence between tape semantics and a more standard semantics allows us to reason about programs in λ_{OADT^+} using traditional mindset and methods.
- To further reduce the user burden, we develop an algorithm to derive secure implementations in λ_{OADT^+} from public programs and their privacy policies as type signatures.

Similar to λ_{OADT} , λ_{OADT^+} and its metatheory have been mechanized in the Coq proof assistant [48].

4.1 Overview

Recall the secure implementation of `lookup0` from [Figure 3.7](#) in [Chapter 3](#). While this implementation is guaranteed to be secure, it is quite far from the “standard” implementation of `lookup` in [Figure 3.1](#), as its control flow has been restructured to only depend on public inputs and to meet the demands of a secure type system. As a consequence, a programmer must write distinct versions of `lookup` for each public view, despite the fact that the high-level program logic is

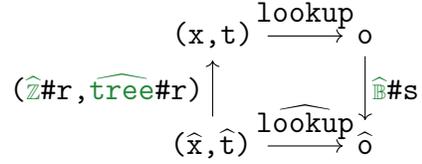


Figure 4.1. Sketch of a secure `lookup` function

exactly the same. Note that `lookup` is, in fact, a valid λ_{OADT} program, *as long as* it is applied to public data. This observation suggests the implementation of $\widehat{\text{lookup}}$ sketched in [Figure 4.1](#), which simply converts its private inputs to public versions, applies `lookup` to those arguments, and converts the result back to an oblivious value. Recall that these conversions are called section and retraction ([Section 3.1.1](#)). From a cryptographic perspective, we “decrypt” the secure inputs using retraction functions (e.g., $\widehat{z}\#r$ and $\widehat{\text{tree}}\#r$), and “encrypt” the computed result using section functions (e.g., $\widehat{\mathbb{B}}\#s$). There is a fundamental flaw with this approach, however: applying a retraction in this manner completely leaks the private inputs of $\widehat{\text{lookup}}$! Thus, this program must be rejected by λ_{OADT} ’s type system as insecure.

The ideal language for oblivious computation would permit implementations that combine the clarity of `lookup` with the security guarantees of `lookup0`. In pursuit of this goal, we have developed an extension of λ_{OADT} , called $\lambda_{\text{OADT}+}$, that allows implementations that follow the recipe sketched in [Figure 4.1](#) *without* compromising obliviousness. Our key idea is to have the semantics of $\lambda_{\text{OADT}+}$ repair or “tape up” potentially leaky expressions during execution. This allows users to write section and leaky retraction functions that convert between oblivious and public values, relying on the semantics to ensure oblivious execution of any program that uses those functions.

To understand how this works, consider the execution trace of the simple $\lambda_{\text{OADT}+}$ program shown in [Figure 4.2a](#). The new conditional `if` is similar to `mux` in λ_{OADT} , but it allows

$$\begin{array}{l}
\text{tape (if (\widehat{if} [\text{true}] \quad \text{tape (\widehat{if} [\text{true}] \\
\quad \text{then true} \quad \text{then if true then [5]} \\
\quad \text{else false})} \longrightarrow \quad \text{else [4]} \\
\text{then [5]} \quad \text{else if false then [5]} \\
\text{else [4])} \quad \text{else [4])} \\
\longrightarrow^* \text{tape (\widehat{if} [\text{true}] \text{ then [5]} \longrightarrow \text{mux [\text{true}] [5] [4]} \longrightarrow [5] \\
\quad \text{else [4])} \\
\text{(a) } \widehat{\text{if}} \text{ inside if} \\
\\
\text{tape ((\widehat{if} [\text{true}] \quad \text{tape (\widehat{if} [\text{true}] \\
\quad \text{then } (\lambda x \Rightarrow x \hat{+} [1]) \longrightarrow \quad \text{then } (\lambda x \Rightarrow x \hat{+} [1]) [4]} \\
\quad \text{else } (\lambda x \Rightarrow x)) [4])} \quad \text{else } (\lambda x \Rightarrow x) [4])} \\
\longrightarrow^* \text{tape (\widehat{if} [\text{true}] \text{ then [5]} \longrightarrow \text{mux [\text{true}] [5] [4]} \longrightarrow [5] \\
\quad \text{else [4])} \\
\text{(b) } \widehat{\text{if}} \text{ inside application} \\
\\
\text{tape (\widehat{Z}\#s (\widehat{Z}\#r [2] + \widehat{Z}\#r [3]))} \longrightarrow \text{tape (\widehat{Z}\#s (\widehat{Z}\#r ([2] \hat{+} [3]))} \\
\longrightarrow \text{tape (\widehat{Z}\#s (\widehat{Z}\#r [5]))} \longrightarrow \text{tape [5]} \longrightarrow [5] \\
\text{(c) retraction of integer}
\end{array}$$

Figure 4.2. Example λ_{OADT^+} execution traces

non-oblivious branches. Note that this $\widehat{\text{if}}$ would leak the value of its private condition if it was evaluated using the semantics of mux . Similar leaks occur for any $\widehat{\text{if}}$ expression whose branches can evaluate to a public value. The idea behind the semantics of λ_{OADT^+} is straightforward: since $\widehat{\text{if}}$ only leaks information when it is evaluated, we will simply not do that! Rather, the surrounding tape annotation ensures the expression will be *eventually oblivious*, and tells λ_{OADT^+} to defer reducing $\widehat{\text{if}}$ until it is safe to do so. This example makes progress by distributing the surrounding if statement into its branches and then evaluating both branches to oblivious values instead. Once both branches of an $\widehat{\text{if}}$ are evaluated to oblivious values, it can be securely reduced to a mux to produce the final result. Note that

```

fn  $\widehat{\text{tree}}\#s \{ \perp \} (k : \mathbb{N})_{\perp} (t : \text{tree})_{\top} : \widehat{\text{tree}} k =
  \text{if } k = 0
  \text{then } ()
  \text{else } \text{tape } (\text{match } t \text{ with}
    | \text{Leaf} \Rightarrow \widehat{\text{inl}} ()
    | \text{Node } x \text{ t1 } \text{tr} \Rightarrow
      \widehat{\text{inr}} (\text{tape } (\widehat{\mathbb{Z}}\#s \ x, \widehat{\text{tree}}\#s \ \text{t1} \ (k-1), \widehat{\text{tree}}\#s \ \text{tr} \ (k-1))))

fn  $\widehat{\text{tree}}\#r \{ \top \} (k : \mathbb{N})_{\perp} : (\widehat{\text{tree}} k)_{\perp} \rightarrow \text{tree} =
  \text{if } k = 0
  \text{then } \lambda\_ \Rightarrow \text{Leaf}
  \text{else } \lambda t \Rightarrow \widehat{\text{match}} t \text{ of}
    | \widehat{\text{inl}}\_ \Rightarrow \text{Leaf}
    | \widehat{\text{inr}} (x, \text{t1}, \text{tr}) \Rightarrow
      \text{Node } (\widehat{\mathbb{Z}}\#r \ x) (\widehat{\text{tree}}\#r \ (k-1) \ \text{t1}) (\widehat{\text{tree}}\#r \ (k-1) \ \text{tr})

fn  $\widehat{\text{lookup}} \{ \perp \} (k : \mathbb{N})_{\perp} (x : \widehat{\mathbb{Z}})_{\perp} (t : \widehat{\text{tree}} k)_{\perp} : \widehat{\mathbb{B}} =
  \text{tape } (\widehat{\mathbb{B}}\#s \ (\text{lookup } (\widehat{\mathbb{Z}}\#r \ x) (\widehat{\text{tree}}\#r \ k \ t)))$$$ 
```

Figure 4.3. Oblivious lookup function in $\lambda_{\text{OADT}\dagger}$

swapping [true] with [false] in this example produces the exact same trace, modulo oblivious values.

This example demonstrates the two key ideas behind the semantics of $\lambda_{\text{OADT}\dagger}$: avoid leaks by delaying evaluation of potentially insecure expressions, while still making progress by distributing the surrounding context into such expressions. This strategy works for contexts like function application as well, as the example in Figure 4.2b shows. Figure 4.2c includes an example of a potential leak of oblivious integers via the $\widehat{\mathbb{Z}}\#r$ operation, a leak that is ultimately patched using $\widehat{\mathbb{Z}}\#s$. The program first progresses by distributing the insecure addition operation into retraction, then obviously adding the result. After evaluating the oblivious addition, we have [5], and $\widehat{\mathbb{Z}}\#s$ and $\widehat{\mathbb{Z}}\#r$ can “cancel” each other, as the functions are effectively inverses. As [5] is already an oblivious value, `tape` becomes a no-op in this example.

Figure 4.3 shows the section and retraction functions for $\widehat{\text{tree}}$, along with a version of $\widehat{\text{lookup}}$ implemented using the recipe from Figure 4.1. While the section function is not used in $\widehat{\text{lookup}}$, it is needed for functions that return an oblivious tree. Function definitions in λ_{OADT^+} require an additional annotation which signals if the function body includes any potentially leaky operations (e.g., $\widehat{\text{if}}$) that needs to be patched by the context surrounding the function call. Section 4.2 discusses how the type system of λ_{OADT^+} uses these annotations in more detail. At a high level, its type system enforces two policies. First, as types are always public, they should not contain any potential leaks: any type which depends on $\widehat{\text{if}} [\text{true}] \mathbb{1} \widehat{\mathbb{B}}$ is disallowed, for example. Next, because only terms that evaluate to oblivious values can be patched up, our type system ensures that terms with potential leaks, e.g., a call to a retraction function or an $\widehat{\text{if}}$, are obliviously typed.

This strategy of decoupling program logic and privacy policies enjoys multiple benefits. First, the core program logic is easier to read, write and reason about, because it is simply a normal functional program, just like `lookup`. Second, these core functions are agnostic to a particular security policy. To share the spine of the tree, we only need to choose a different $\widehat{\text{tree}}\#r$ and $\widehat{\text{tree}}\#s$; `lookup` itself remains unchanged. This frees users from writing different versions of the same function for different security policies. Third, this approach allows users to experiment and trade off between performance and security guarantee. Sharing the exact spine of the tree will result in better performance, for example, if both parties agree to this policy.

4.2 λ_{OADT^+} , Formally

This section formalizes λ_{OADT^+} , an extension to λ_{OADT} that permits implementations in the vein of Figure 4.1.

4.2.1 Syntax

The extended syntax of λ_{OADT^+} is shown in Figure 4.4. These extensions permit λ_{OADT^+} expressions that *potentially* leak information *locally*, as long as they can eventually be repaired by the surrounding context. The new $\widehat{\text{if}}$ operation is similar to `mux`, but its branches are

$e, \tau ::=$	\dots $\widehat{\text{if}}\ e\ \text{then}\ e\ \text{else}\ e$ $\text{tape}\ e$ $\text{let}\ x : \iota\tau = e\ \text{in}\ e$ $\lambda x : \iota\tau \Rightarrow e \mid \Pi x : \iota\tau, \tau$	EXTENDED EXPRESSIONS: oblivious leaky conditional tape operation let binding with leakage label function and function types with leakage label
$D ::=$	\dots $\text{fn}\ x : \iota\tau = e$	EXTENDED GLOBAL DEFINITIONS: (recursive) function definition with leakage label
$l ::=$	$\top \mid \perp$	LEAKAGE LABEL

Figure 4.4. λ_{OADT^+} syntax

permitted to be non-oblivious, causing a potential leak if $\widehat{\text{if}}$ is evaluated naively. The new `tape` annotation acts as a boundary for potential leaks, and is used to ensure that they never occur during execution, as [Section 4.2.2](#) will discuss in more detail. Finally, λ_{OADT^+} updates the syntax for let bindings, anonymous functions, function types and function definitions with a *leakage label*. A leakage label is either \top or \perp , and signals either the presence or the absence of a potential leak, respectively.

4.2.2 Semantics

The semantics of λ_{OADT^+} are an extension of the semantics of λ_{OADT} . [Figure 4.5](#) shows the new and updated rules; the rest are identical to the rules in [Figure 3.9](#). This semantics introduces a new syntactic class of *weak values*, which are used to ensure that $\widehat{\text{if}}$ does not leak information when evaluated. Weak values simply extend the values in λ_{OADT} with $\widehat{\text{if}}$: a $\widehat{\text{if}}$ is a weak value if all its subexpressions are weak values. All references to v in the reduction rules (including those not shown in [Figure 4.5](#)) now refer to weak values unless explicitly identified as a value. The semantics also extend evaluation contexts to handle $\widehat{\text{if}}$ and `tape` expressions.

The S-OIF rule captures the key idea of distributing surrounding context into the branches of $\widehat{\text{if}}$. Like S-MUX, this rule requires its branches to first be evaluated to weak values using S-CTX. Note that not all contexts need to be distributed into these branches in order to

S-OMATCH

$$\frac{\hat{v}_1 \leftarrow \hat{\omega}_1 \quad \hat{v}_2 \leftarrow \hat{\omega}_2}{\widehat{\text{match}} \ [l_b \langle \hat{\omega}_1 + \hat{\omega}_2 \rangle \ \hat{v}] \ \text{with} \ x \Rightarrow e_1 \mid x \Rightarrow e_2 \longrightarrow \widehat{\text{if}} \ [b] \ \text{then} \ \text{ite}(b, [\hat{v}/x]e_1, [\hat{v}_1/x]e_1) \\ \text{else} \ \text{ite}(b, [\hat{v}_2/x]e_2, [\hat{v}/x]e_2)}$$

S-OIF

$$\frac{\widehat{\mathcal{E}}[\widehat{\text{if}} \ [b] \ \text{then} \ v_1 \ \text{else} \ v_2] \longrightarrow \widehat{\text{if}} \ [b] \ \text{then} \ \widehat{\mathcal{E}}[v_1] \ \text{else} \ \widehat{\mathcal{E}}[v_2]}$$

S-TAPEOIF

$$\frac{\text{tape} \ (\widehat{\text{if}} \ [b] \ \text{then} \ v_1 \ \text{else} \ v_2) \longrightarrow \text{mux} \ [b] \ (\text{tape} \ v_1) \ (\text{tape} \ v_2)}$$

S-TAPEOVAL

\hat{v} is oblivious value but not pair

$$\frac{}{\text{tape} \ \hat{v} \longrightarrow \hat{v}}$$

S-TAPEPAIR

$$\frac{}{\text{tape} \ (v_1, v_2) \longrightarrow (\text{tape} \ v_1, \text{tape} \ v_2)}$$

WEAK VALUES

$$v ::= \dots \\ \mid \widehat{\text{if}} \ [b] \ \text{then} \ v \ \text{else} \ v$$

EVALUATION CONTEXTS

$$\mathcal{E} ::= \dots \\ \mid \widehat{\text{if}} \ \square \ \text{then} \ e \ \text{else} \ e \\ \mid \widehat{\text{if}} \ v \ \text{then} \ \square \ \text{else} \ e \\ \mid \widehat{\text{if}} \ v \ \text{then} \ v \ \text{else} \ \square \\ \mid \text{tape} \ \square$$

LEAKY CONTEXTS

$$\widehat{\mathcal{E}} ::= \\ \mid \square \ v \\ \mid \pi_b \ \square \\ \mid \text{if} \ \square \ \text{then} \ e \ \text{else} \ e \\ \mid \text{match} \ \square \ \text{with} \ x \Rightarrow e \mid x \Rightarrow e \\ \mid \widehat{\mathbb{B}}\#s \ \square \\ \mid \text{unfold}\langle T \rangle \ \square$$

Figure 4.5. λ_{ADT^+} semantics

make progress; pushing `fold` into `if` in the expression `fold<tree> (if [true] ...)` does not gain us anything, for example, since the expression is already a weak value. Figure 4.5 defines the *leaky contexts* ($\widehat{\mathcal{E}}$) that can be distributed through `if`. For simplicity, we adopt a minimal set of leaky contexts, though allowing more contexts is a potential avenue for optimizing executions. This does not limit the expressivity of λ_{ADT^+} , for similar reasons to the fold example from above. The semantics of `match` are also updated to allow potential leaks, with S-OMATCH now evaluating to `if` instead of `mux`.

The last three rules in Figure 4.5 show how to evaluate `tape` annotations. The key idea is to use `tape` as a signal that the context surrounding an `if` expression has been sufficiently distributed to prevent leaks. Mechanically, whenever a `tape` annotation is applied to `if`

expression whose branches are weak values, it is safe to reduce expression to a secure `mux` using S-TAPEOIF. As an example, consider the following expression:

$$\begin{array}{l}
\text{tape } (\widehat{\mathbb{B}}\#s \ (\widehat{\text{if}} \ [\text{true}] \\
\qquad \qquad \text{then } \text{false} \ \longrightarrow \ \text{tape } (\widehat{\text{if}} \ [\text{true}] \\
\qquad \qquad \text{else } \text{true})) \qquad \qquad \text{then } (\widehat{\mathbb{B}}\#s \ \text{false}) \\
\qquad \qquad \qquad \qquad \qquad \qquad \text{else } (\widehat{\mathbb{B}}\#s \ \text{true})) \\
\longrightarrow^* \text{tape } (\widehat{\text{if}} \ [\text{true}] \qquad \text{mux } [\text{true}] \\
\qquad \qquad \text{then } [\text{false}] \ \longrightarrow \ (\text{tape } [\text{false}]) \\
\qquad \qquad \text{else } [\text{true}]) \qquad \qquad (\text{tape } [\text{true}]) \\
\longrightarrow^* \text{mux } [\text{true}] \ [\text{false}] \ [\text{true}] \ \longrightarrow \ [\text{false}]
\end{array}$$

After applying S-OIF to distribute the surrounding boolean section $\widehat{\mathbb{B}}\#s$, the $\widehat{\text{if}}$ expression is now annotated with `tape`, and S-TAPEOIF can be applied. The `tape` annotations are pushed inside the branches of `mux` to ensure any $\widehat{\text{if}}$ expressions they may contain are also repaired. The final two rules ensure `tape` annotations are eventually dropped from oblivious values. An oblivious (non-pair) value annotated with `tape` cannot leak any information, and S-TAPEOVAL can be applied to remove the extraneous `tape`. S-TAPEPAIR allows `tape` annotations to be distributed into the components of an oblivious pair, in order to eventually repair any $\widehat{\text{if}}$ expressions they may contain.

4.2.3 Type System

The typing judgment of λ_{OADT^+} now includes a leakage label for the typed expression: $\Gamma \vdash e ; \tau$, as do entries in typing contexts Γ . Figure 4.6 shows a subset of the typing rules of λ_{OADT^+} ; the omitted rules are copies of those from λ_{OADT} with straightforward leakage labels annotations. As mentioned in Section 4.2.1, leakage labels signal whether an expression might contain a potential leak. The reason for these labels is similar to the *security labels* found in other security-type systems [27, 29], where type-based information flow control is used to enforce noninterference between high- and low- security information. In λ_{OADT^+} , expressions with \top labels should not influence expressions with \perp labels. In order to minimize the extension to λ_{OADT} , we do not annotate every type with a leakage label, opting to only annotate top-level definitions and function parameters with leakage labels. While it is

$$\boxed{\Gamma \vdash e :_l \tau}$$

$$\begin{array}{c}
\text{T-VAR} \\
\frac{x :_l \tau \in \Gamma}{\Gamma \vdash x :_l \tau} \\
\\
\text{T-UNIT} \\
\frac{}{\Gamma \vdash () :_{\perp} \mathbb{1}} \\
\\
\text{T-PAIR} \\
\frac{\Gamma \vdash e_1 :_{l_1} \tau_1 \quad \Gamma \vdash e_2 :_{l_2} \tau_2 \quad l = l_1 \sqcup l_2}{\Gamma \vdash (e_1, e_2) :_l \tau_1 \times \tau_2} \\
\\
\text{T-PROJ} \\
\frac{\Gamma \vdash e :_l \tau_1 \times \tau_2}{\Gamma \vdash \pi_b e :_l \text{ite}(b, \tau_1, \tau_2)} \\
\\
\text{T-ABS} \\
\frac{x :_{l_1} \tau_1, \Gamma \vdash e :_{l_2} \tau_2 \quad \Gamma \vdash \tau_1 :: *}{\Gamma \vdash \lambda x :_{l_1} \tau_1 \Rightarrow e :_{l_2} \prod x :_{l_1} \tau_1, \tau_2} \\
\\
\text{T-APP} \\
\frac{\Gamma \vdash e_2 :_{l_2} \prod x :_{l_1} \tau_1, \tau_2 \quad \Gamma \vdash e_1 :_{l_1} \tau_1}{\Gamma \vdash e_2 e_1 :_{l_2} [e_1/x] \tau_2} \\
\\
\text{T-IF} \\
\frac{\Gamma \vdash e_0 :_{\perp} \mathbb{B} \quad l = l_1 \sqcup l_2 \quad \Gamma \vdash e_1 :_{l_1} [\text{true}/z] \tau \quad \Gamma \vdash e_2 :_{l_2} [\text{false}/z] \tau}{\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 :_l [e_0/z] \tau} \\
\\
\text{T-IFNODEP} \\
\frac{\Gamma \vdash e_0 :_{l_0} \mathbb{B} \quad l = l_0 \sqcup l_1 \sqcup l_2 \quad \Gamma \vdash e_1 :_{l_1} \tau \quad \Gamma \vdash e_2 :_{l_2} \tau}{\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 :_l \tau} \\
\\
\text{T-MATCH} \\
\frac{\Gamma \vdash e_0 :_{\perp} \tau_1 + \tau_2 \quad l = l_1 \sqcup l_2 \quad x :_{\perp} \tau_1, \Gamma \vdash e_1 :_{l_1} [\text{inl} \langle \tau_1 + \tau_2 \rangle x/z] \tau \quad x :_{\perp} \tau_2, \Gamma \vdash e_2 :_{l_2} [\text{inr} \langle \tau_1 + \tau_2 \rangle x/z] \tau}{\Gamma \vdash \text{match } e_0 \text{ with } x \Rightarrow e_1 \mid x \Rightarrow e_2 :_l [e_0/z] \tau} \\
\\
\text{T-MATCHNODEP} \\
\frac{\Gamma \vdash e_0 :_{l_0} \tau_1 + \tau_2 \quad l = l_0 \sqcup l_1 \sqcup l_2 \quad x :_{l_0} \tau_1, \Gamma \vdash e_1 :_{l_1} \tau \quad x :_{l_0} \tau_2, \Gamma \vdash e_2 :_{l_2} \tau}{\Gamma \vdash \text{match } e_0 \text{ with } x \Rightarrow e_1 \mid x \Rightarrow e_2 :_l \tau} \\
\\
\text{T-CONV} \\
\frac{\Gamma \vdash e :_l \tau \quad \tau \equiv \tau' \quad \Gamma \vdash \tau' :: * \quad l \sqsubseteq l'}{\Gamma \vdash e :_{l'} \tau'} \\
\\
\text{T-MUX} \\
\frac{\Gamma \vdash e_0 :_{\perp} \widehat{\mathbb{B}} \quad \Gamma \vdash \tau :: *^0 \quad \Gamma \vdash e_1 :_{\perp} \tau \quad \Gamma \vdash e_2 :_{\perp} \tau}{\Gamma \vdash \text{mux } e_0 e_1 e_2 :_{\perp} \tau} \\
\\
\text{T-OINJ} \\
\frac{\Gamma \vdash e :_{\perp} \text{ite}(b, \tau_1, \tau_2) \quad \Gamma \vdash \tau_1 \hat{+} \tau_2 :: *^0}{\Gamma \vdash \widehat{l}_b \langle \tau_1 \hat{+} \tau_2 \rangle e :_{\perp} \tau_1 \hat{+} \tau_2} \\
\\
\text{T-OIF} \\
\frac{\Gamma \vdash e_0 :_{\perp} \widehat{\mathbb{B}} \quad \Gamma \vdash e_1 :_{l_1} \tau \quad \Gamma \vdash e_2 :_{l_2} \tau}{\Gamma \vdash \widehat{\text{if}} e_0 \text{ then } e_1 \text{ else } e_2 :_{\top} \tau} \\
\\
\text{T-OMATCH} \\
\frac{\Gamma \vdash e_0 :_{\perp} \tau_1 \hat{+} \tau_2 \quad x :_{\perp} \tau_1, \Gamma \vdash e_1 :_{l_1} \tau \quad x :_{\perp} \tau_2, \Gamma \vdash e_2 :_{l_2} \tau}{\Gamma \vdash \widehat{\text{match}} e_0 \text{ with } x \Rightarrow e_1 \mid x \Rightarrow e_2 :_{\top} \tau} \\
\\
\text{T-TAPE} \\
\frac{\Gamma \vdash e :_l \tau \quad \Gamma \vdash \tau :: *^0}{\Gamma \vdash \text{tape } e :_{\perp} \tau}
\end{array}$$

Figure 4.6. Selected $\lambda_{\text{O}ADT+}$ typing rules

certainly possible to implement a more precise analysis, this coarse-grained analysis is strong enough for our purposes.

The leakage label of base types is always \perp , e.g., in T-UNIT. Leakage labels for local or global variables is taken directly from the context, e.g., T-VAR. For most public constructs, e.g., T-PAIR and T-PROJ, the label is the join (\sqcup) of the labels of all sub-expressions, where $\perp \sqcup \perp \equiv \perp$ and \top otherwise. T-PROJ shows why leakage is an overapproximation, as we cannot always tell which component of a pair labeled with \top is the source of the potential leak. In T-ABS, both the type and label of a parameter are added to the typing context when typing the function body. The label assigned to the function body is then propagated to the whole lambda abstraction. This strategy may seem a bit counterintuitive, as a lambda abstraction is irreducible, and thus cannot leak any information during further evaluation. Of course, while a lambda value will not leak any information on its own, it does have the *potential* to leak when applied to an argument. Because our leakage analysis is quite coarse, we simply consider an expression leaky if it may leak when it is “used”. T-APP requires a function to be applied to an argument whose label matches that of its parameter. Applying a function with a potentially leaky parameter to a non-leaky argument can be typed by first using the T-CONV rule, which allows the label of an expression to be downgraded. As an example, these rules ensure both $(\lambda x:_{\top} \mathbb{B} \Rightarrow \widehat{\mathbb{B}} \# s \ x) \ (\widehat{\text{if}} \ [\text{true}] \ \text{then true else false})$ and $(\lambda x:_{\top} \mathbb{B} \Rightarrow \widehat{\mathbb{B}} \# s \ x) \ \text{true}$ are well-typed expressions.

$\lambda_{\text{Oadt}+}$ has dependent and nondependent versions of the typing rule for `if`. In the dependent version, T-IF, the discriminée is not allowed to contain a potential leak, as it may appear in the type. In the nondependent version T-IFNODEP, there is no such restriction, but the type is not allowed to depend on the discriminée. The typing rules for `match`, T-MATCH and T-MATCHNODEP, are similar.

The remaining typing rules deal with expressions that either repair or introduce potential leaks. An expression annotated with `tape` is always assigned the \perp label, as long as that expression has an oblivious type. This is in line with the semantics of `tape`: when applied to an oblivious expression, it eventually evaluates to an oblivious value or a weak value ($\widehat{\text{if}}$). The former is already safe, and the latter can be repaired by S-TAPEOIF. The \perp label in the rule captures the idea that `tape` safely repairs a local leak, such that the surrounding

$$\boxed{\Gamma \vdash \tau :: \kappa}$$

$$\begin{array}{c}
\text{K-OADT} \\
\frac{\text{obliv } \widehat{T} (x:\tau) = \tau' \in \Sigma \quad \Gamma \vdash e :_{\perp} \tau}{\Gamma \vdash \widehat{T} e :: *^0}
\end{array}
\qquad
\begin{array}{c}
\text{K-IF} \\
\frac{\Gamma \vdash e_0 :_{\perp} \mathbb{B} \quad \Gamma \vdash \tau_1 :: *^0 \quad \Gamma \vdash \tau_2 :: *^0}{\Gamma \vdash \text{if } e_0 \text{ then } \tau_1 \text{ else } \tau_2 :: *^0}
\end{array}$$

$$\begin{array}{c}
\text{K-MATCH} \\
\frac{\Gamma \vdash e_0 :_{\perp} \tau'_1 + \tau'_2 \quad x :_{\perp} \tau'_1, \Gamma \vdash \tau_1 :: *^0 \quad x :_{\perp} \tau'_2, \Gamma \vdash \tau_2 :: *^0}{\Gamma \vdash \text{match } e_0 \text{ with } x \Rightarrow \tau_1 \mid x \Rightarrow \tau_2 :: *^0}
\end{array}$$

Figure 4.7. Selected λ_{OADT^+} kinding rules

computation can treat it as non-leaky. The rules for $\widehat{\text{if}}$ and $\widehat{\text{match}}$ reflect the fact that they are sources of potential leaks, as both expressions are labeled with \top . Both rules require their discriminates to be free of potential leaks, but this does not affect expressiveness, since their discriminates can always be wrapped with `tape`. T-OINJ and T-MUX feature similar requirements.

Figure 4.7 shows the updated kinding rules for λ_{OADT^+} ; the other kinding rules are identical to those in Figure 3.11. The updated rules require types to only depend on terms that do not contain potential leaks, i.e., those assigned the \perp label. To see why, consider the following ill-kinded type:

`if ($\widehat{\text{if}}$ [true] then true else false) then 1 else $\widehat{\mathbb{B}}$`

After distributing the surrounding `if` into $\widehat{\text{if}}$, this reduces to $\widehat{\text{if}}$ [true] then 1 else $\widehat{\mathbb{B}}$. Similar expressions at the term level can be repaired by, e.g., distributing $\widehat{\mathbb{B}}\#s$ through the branches to secure the result of the $\widehat{\text{if}}$. At the type level we have no such recourse, however: since types are always public, there is no corresponding way to repair this type by securing its branches.

Figure 4.8 shows a subset of the updated and new parallel reduction rules. Again, the rules for oblivious constructs are similar to the corresponding step rules. In R-OIFCTX, we write $\widehat{\mathcal{E}} \Rightarrow \widehat{\mathcal{E}}'$ to mean all the subexpressions in the leaky context take a parallel reduction step. R-OIF is required for confluence, similar to R-MUX. We say a λ_{OADT^+} program is well-typed if the global context is well-typed (the updated typing rules for the global context

$e \Rightarrow e'$

$$\begin{array}{c}
\text{R-OIFCTX} \\
\frac{e_1 \Rightarrow e'_1 \quad e_2 \Rightarrow e'_2 \quad \widehat{\mathcal{E}} \Rightarrow \widehat{\mathcal{E}}'}{\widehat{\mathcal{E}}[\widehat{\text{if}} [b] \text{ then } e_1 \text{ else } e_2] \Rightarrow \widehat{\text{if}} [b] \text{ then } \widehat{\mathcal{E}}'[e'_1] \text{ else } \widehat{\mathcal{E}}'[e'_2]} \\
\\
\text{R-TAPEOIF} \\
\frac{e_1 \Rightarrow e'_1 \quad e_2 \Rightarrow e'_2}{\text{tape} (\widehat{\text{if}} [b] \text{ then } e_1 \text{ else } e_2) \Rightarrow \text{mux} [b] (\text{tape } e'_1) (\text{tape } e'_2)} \\
\\
\begin{array}{cc}
\text{R-TAPEPAIR} & \text{R-TAPEOVAL} \\
\frac{e_1 \Rightarrow e'_1 \quad e_2 \Rightarrow e'_2}{\text{tape} (e_1, e_2) \Rightarrow (\text{tape } e'_1, \text{tape } e'_2)} & \frac{\widehat{v} \text{ is oblivious value but not pair}}{\text{tape } \widehat{v} \Rightarrow \widehat{v}}
\end{array} \\
\\
\text{R-OIF} \\
\frac{e_1 \Rightarrow e'_1 \quad e_2 \Rightarrow e'_2}{\widehat{\text{if}} [b] \text{ then } e_1 \text{ else } e_2 \Rightarrow \text{ite}(b, e'_1, e'_2)}
\end{array}$$

Figure 4.8. Selected λ_{OADT^+} parallel reduction rules

are trivial) and the expression is well-typed with \perp label. The latter restriction ensures that all potential leaks in a λ_{OADT^+} program are eventually repaired.

4.2.4 Type Safety and Obliviousness

The guarantees of the type system of λ_{OADT^+} are quite similar to those of λ_{OADT} , although they have been adapted slightly to account for leakage labels. The statement of progress for λ_{OADT^+} , for example, is limited to expressions without potential leaks:

Theorem 4.2.1 (Progress). *If $\cdot \vdash e :_{\perp} \tau$, then either $e \longrightarrow e'$ for some e' , or e is a value. If $\cdot \vdash \tau : *^0$, then either $\tau \longrightarrow \tau'$ for some τ' , or τ is an oblivious type value.*

This updated statement reflects the fact that leaky expressions only reduce to *weak values*. The proof of this theorem is a consequence of a stronger lemma which also accounts for potentially leaky expressions:

Lemma 4.2.2. *If $\cdot \vdash e :_l \tau$, then either $e \longrightarrow e'$ for some e' , or e is a weak value.*

The proof of this stronger lemma proceeds similarly to the proof of progress for λ_{OADT} , with the canonical form lemmas extended to weak values. One technicality needed by this proof is a notion of *weak oblivious value*, which extends oblivious values to include $\widehat{\text{if}}$ expressions. For the T-TAPE case, we have to show that a weak value with an oblivious type is also a weak oblivious value, as `tape` can only be reduced when it is applied to weak oblivious values. This extra lemma requires an updated version of [Lemma 3.2.9](#), so the proof of progress for λ_{OADT^+} now depends on type preservation for parallel reduction. With this lemma in hand, the progress theorem immediately follows from the fact that a weak value is a value if it is labeled with \perp .

The statements of preservation and obliviousness must also be updated to deal with leakage labels, but are otherwise identical:

Theorem 4.2.3 (Preservation). *If $\Gamma \vdash e :_l \tau$, and $e \longrightarrow e'$, then $\Gamma \vdash e' :_l \tau$.*

If $\Gamma \vdash \tau :: \kappa$ and $\tau \longrightarrow \tau'$, then $\Gamma \vdash \tau' :: \kappa$.

Theorem 4.2.4 (Obliviousness). *If $e_1 \approx e_2$ and $\cdot \vdash e_1 :_{l_1} \tau_1$ and $\cdot \vdash e_2 :_{l_2} \tau_2$, then*

1. $e_1 \longrightarrow^n e'_1$ if and only if $e_2 \longrightarrow^n e'_2$ for some e'_2 .
2. if $e_1 \longrightarrow^n e'_1$ and $e_2 \longrightarrow^n e'_2$, then $e'_1 \approx e'_2$.

Proofs of both theorems follow the same structure as their counterparts in λ_{OADT} , although many of the lemmas used in the proof of obliviousness now use weak values instead of values.

4.3 Extending λ_{OADT^+}

This section considers how additional base types might be added to the core calculus of λ_{OADT^+} , using fixed-width integers as an example. [Figure 4.9](#) shows a subset of the syntax, semantics and typing rules needed for this new primitive type. The extended language includes public and oblivious versions of integer types, literals, and operators. For simplicity, we only consider a comparison operation, but additional operators could be added in a similar manner. In order to move between the public and oblivious types, section ($\widehat{\text{Z}}\#s$) and retraction ($\widehat{\text{Z}}\#r$) operations for integers are also added; both have similar semantics to their

		$\Gamma \vdash e :_l \tau$
$e, \tau ::=$	<p style="text-align: center;">EXTENDED EXPRESSIONS:</p> <ul style="list-style-type: none"> \dots $\widehat{\mathbb{Z}} \mid \widehat{\mathbb{Z}}$ primitive integer types $i \mid [i]$ (runtime) integer literals $e \leq e \mid e \widehat{\leq} e$ integer operators $\widehat{\mathbb{Z}}\#s \ e \mid \widehat{\mathbb{Z}}\#r \ e$ integer section and retraction 	<p style="text-align: center;">(a) Extended syntax</p>
		<p style="text-align: center;">(b) Extended typing rules</p>

		$e \longrightarrow e'$
	<p>SI-SEC</p> $\frac{}{\widehat{\mathbb{Z}}\#s \ i \longrightarrow [i]}$	<p>SI-SECRET</p> $\frac{}{\widehat{\mathbb{Z}}\#s \ (\widehat{\mathbb{Z}}\#r \ [i]) \longrightarrow [i]}$
SI-RETLE ₁	$\frac{}{\widehat{\mathbb{Z}}\#r \ [i_1] \leq \widehat{\mathbb{Z}}\#r \ [i_2] \longrightarrow \widehat{\mathbb{B}}\#r \ ([i_1] \widehat{\leq} [i_2])}$	SI-RETLE ₂
	<p>SI-RETLE₃</p> $\frac{}{i_1 \leq \widehat{\mathbb{Z}}\#r \ [i_2] \longrightarrow \widehat{\mathbb{B}}\#r \ (\widehat{\mathbb{Z}}\#s \ i_1 \widehat{\leq} [i_2])}$	
	(c) Extended semantics	

Figure 4.9. A subset of extended language for fixed-width integers

boolean counterparts⁴. $\widehat{\mathbb{Z}}\#r$ always introduces a potential leak, and $\widehat{\mathbb{Z}}\#r \ \hat{v}$ is considered weak value.

When defining the semantics of potentially leaky expressions like \leq , it is important that the semantics does not leak information via the execution trace. When comparing oblivious values with \leq , for example, SI-RETLE₁ combines $\widehat{\leq}$ and $\widehat{\mathbb{B}}\#r$ to first securely compare the operands before retracting the resulting oblivious boolean. SI-RETLE₂ and SI-RETLE₃ are similar, but they apply to cases when one of the operands is not a retraction of an oblivious value by lifting it to oblivious values first. The semantics of other operators can be defined

⁴↑Although λ_{OADT^+} does not include boolean retraction $\widehat{\mathbb{B}}\#r$ as a primitive, it is easily defined in terms of $\widehat{\text{if}}$: $\widehat{\mathbb{B}}\#r \ e \triangleq \widehat{\text{if}} \ e \ \text{then true else false}$.

through similar uses of section and retraction functions. As an example, integer addition returns an integer instead of boolean, so we apply $\widehat{\mathbb{Z}}\#r$ to the result of oblivious addition. If a leaky integer expression is used in a well-typed context, then $\widehat{\mathbb{Z}}\#r$ will eventually meet $\widehat{\mathbb{Z}}\#s$ and they can be canceled out via SI-SECRET. Updated versions of evaluation contexts, leaking contexts and the other reduction rules are omitted, as they are straightforward extensions of their counterparts in λ_{OADT^+} . The extended typing and kinding rules are also straightforward, and are similar to those for the primitive types in λ_{OADT^+} . Figure 4.9 gives the rules for integer retraction (TI-RET) and oblivious less-than (TI-OLE) as examples.

4.4 λ_{OADT^+} in Action

To demonstrate the expressiveness of λ_{OADT^+} , we have written some example oblivious functions and oblivious types with different public views. We have directly encoded these in our Coq development, as well as some accompanying typing and evaluation derivations. All of the examples described in this section are included in our public artifact [48].

We have encoded the following OADTs for lists and trees. Each oblivious type consists of its type definition, a section function and a retraction function.

- List with the upper bound of its length.
- Tree with the upper bound of its depth.
- Tree with the upper bound of its spine.
- Tree with the upper bound of the number of its vertices (including leaves and nodes).

The second and third of these examples were presented in Section 3.1. The oblivious tree with the upper bound of its total vertices is the most complicated: while its type definition is effectively an oblivious list, its section and retraction functions correspond to flattening a tree and rebuilding a tree from a list.

In addition to the lookup function from Section 3.1, we have also written a tree insertion function as a demonstration of how oblivious ADTs are constructed. A more interesting example is a standard map function over oblivious trees, which shows that higher-order

functions can be naturally written in λ_{OADT^+} . The following code snippet for an oblivious map function follows the recipe in Figure 4.1. Label annotations are omitted for brevity, and we use a boolean payload for simplicity.

```
fn  $\widehat{\text{map}}$  (f :  $\mathbb{B}$   $\rightarrow$   $\mathbb{B}$ ) (k :  $\mathbb{N}$ ) (t :  $\widehat{\text{tree}}$  k) :  $\widehat{\text{tree}}$  k =
   $\widehat{\text{tree\#s}}$  (map f ( $\widehat{\text{tree\#r}}$  k t)) k
```

The function argument of $\widehat{\text{map}}$ takes a public boolean to public boolean, but $\widehat{\text{map}}$ could be adapted to accept a function from oblivious boolean to oblivious boolean by composing boolean section and retraction to appropriately “transport” the function argument. The $\widehat{\text{map}}$ function could also be adapted any oblivious tree definition by simply replacing $\widehat{\text{tree}}$, $\widehat{\text{tree\#s}}$, and $\widehat{\text{tree\#r}}$.

4.5 An Unsafe Reference Semantics

Reasoning about program behaviors directly under tape semantics is not ideal, as tape semantics is a nonstandard semantics: we cannot directly apply our experience, theoretical frameworks or practical tools developed for standard functional languages. This section shows that most of the reasoning principles can be recovered by connecting tape semantics to a more standard, reference semantics, called *reveal semantics*.

Reveal semantics is an *unsafe*, big-step operational semantics for λ_{OADT^+} , whose rules are presented in Figure 4.10. Its judgment $\Sigma \vdash e \Downarrow v$ evaluates an expression e to a value v , under a global context Σ which is elided for brevity in these rules. As λ_{OADT^+} is dependently typed, this relation also evaluates a type to an oblivious type value.

Most rules are straightforward, similar to the big-step semantics of other systems. However, unlike other languages, in E-VAL a value v in reveal semantics is evaluated to the erasure of v , $[v]$, instead of itself. This erasure operation reduces all “dummy” leaky conditionals, i.e., those that are already weak values, even under the binder of a lambda abstraction. Its definition is simple: $[\widehat{\text{if}} e_0 \text{ then } e_1 \text{ else } e_2] = \text{ite}(b, [e_1], [e_2])$ if $[e_0] = [b]$ and both $[e_1]$ and $[e_2]$ are weak values, or otherwise the erasure is $\widehat{\text{if}} [e_0] \text{ then } [e_1] \text{ else } [e_2]$. Other expressions simply erase their subterms recursively. This erasure operation is mainly used for establishing the equivalence to tape semantics. To see why this is necessary,

$e \Downarrow v$

$\frac{}{v \Downarrow [v]}$	$\frac{}{\widehat{\omega} \Downarrow \widehat{\omega}}$	$\frac{\tau_1 \Downarrow \widehat{\omega}_1 \quad \tau_2 \Downarrow \widehat{\omega}_2}{\tau_1 \times \tau_2 \Downarrow \widehat{\omega}_1 \times \widehat{\omega}_2}$	$\frac{\tau_1 \Downarrow \widehat{\omega}_1 \quad \tau_2 \Downarrow \widehat{\omega}_2}{\tau_1 \hat{+} \tau_2 \Downarrow \widehat{\omega}_1 \hat{+} \widehat{\omega}_2}$
$\frac{\text{obliv } \widehat{T} \quad (x:\tau) = \tau' \in \Sigma \quad e \Downarrow v \quad [v/x]\tau' \Downarrow \widehat{\omega}}{\widehat{T} \quad e \Downarrow \widehat{\omega}}$	$\frac{\text{fn } x:\tau = e \in \Sigma \quad e \Downarrow v}{x \Downarrow v}$	$\frac{e_2 \Downarrow \lambda x:\tau \Rightarrow e \quad e_1 \Downarrow v_1}{[v_1/x]e \Downarrow v} \quad e_2 \quad e_1 \Downarrow v$	
$\frac{e_1 \Downarrow v_1 \quad [v_1/x]e_2 \Downarrow v}{\text{let } x = e_1 \text{ in } e_2 \Downarrow v}$	$\frac{e_0 \Downarrow b \quad \text{ite}(b, e_1, e_2) \Downarrow v}{\text{if } e_0 \text{ then } e_1 \text{ else } e_2 \Downarrow v}$		
$\frac{e_0 \Downarrow [b] \quad \text{ite}(b, e_1, e_2) \Downarrow v}{\text{mux } e_0 \quad e_1 \quad e_2 \Downarrow v}$	$\frac{e_0 \Downarrow [b] \quad \text{ite}(b, e_1, e_2) \Downarrow v}{\widehat{\text{if}} \quad e_0 \text{ then } e_1 \text{ else } e_2 \Downarrow v}$		
$\frac{e \Downarrow v}{\iota_b \langle \tau \rangle \quad e \Downarrow \iota_b \langle [\tau] \rangle \quad v}$	$\frac{e_0 \Downarrow \iota_b \langle \tau \rangle \quad v \quad \text{ite}(b, [v/x]e_1, [v/x]e_2) \Downarrow v}{\text{match } e_0 \text{ with } x \Rightarrow e_1 \mid x \Rightarrow e_2 \Downarrow v}$		
$\frac{\tau \Downarrow \widehat{\omega} \quad e \Downarrow \widehat{v}}{\widehat{\iota}_b \langle \tau \rangle \quad e \Downarrow [\iota_b \langle \widehat{\omega} \rangle \widehat{v}]}$	$\frac{e_0 \Downarrow [\iota_b \langle \widehat{\omega} \rangle \widehat{v}] \quad \text{ite}(b, [v/x]e_1, [v/x]e_2) \Downarrow v}{\widehat{\text{match}} \quad e_0 \text{ with } x \Rightarrow e_1 \mid x \Rightarrow e_2 \Downarrow v}$		
$\frac{e_1 \Downarrow v_1 \quad e_2 \Downarrow v_2}{(e_1, e_2) \Downarrow (v_1, v_2)}$	$\frac{e \Downarrow (v_1, v_2)}{\pi_b \quad e \Downarrow \text{ite}(b, v_1, v_2)}$	$\frac{e \Downarrow v}{\text{fold} \langle T \rangle \quad e \Downarrow \text{fold} \langle T \rangle \quad v}$	
$\frac{e \Downarrow \text{fold} \langle T' \rangle \quad v}{\text{unfold} \langle T \rangle \quad e \Downarrow v}$	$\frac{e \Downarrow b}{\widehat{B} \# s \quad e \Downarrow [b]}$	$\frac{e \Downarrow \widehat{v}}{\text{tape} \quad e \Downarrow \widehat{v}}$	

Figure 4.10. λ_{OAdT^+} reveal semantics

consider the expression $\widehat{\text{if}} \text{ [true] then } 1 \text{ else } 2$. This expression evaluates to 1 under reveal semantics, but it is a stuck term in tape semantics to avoid revealing its private condition [true] . To connect reveal semantics to tape semantics, we have to consider a weaker “sameness” up to erasure. We further bake erasure into reveal semantics, so that an expression always evaluates to a sort of “canonical” value free of dummy leaky conditionals. E-INJ also uses this operation to erase a sum injection’s type annotation, similar to E-VAL.

E-OIF executes an $\widehat{\text{if}}$ similarly to the standard semantics of if , which allows us to reason about $\widehat{\text{if}}$ easily by treating it as a normal conditional. This behavior of course reveals the private conditions (in addition to the erasure), so reveal semantics is an unsafe semantics which should not be used for secure computation. However, it is suitable for executing retraction functions in order to reveal the computed secure output to the privileged parties. Note that reveal semantics does not guarantee to have the same termination behavior as tape semantics; it is entirely possible that a program terminates in reveal semantics but does not terminate in tape semantics, as tape semantics executes both branches of an $\widehat{\text{if}}$ regardless of its private condition. Readers may wonder if equi-termination can be achieved by forcing $\widehat{\text{if}}$ to similarly evaluate both branches in E-IF. Unfortunately, $\widehat{\text{if}}$ can still indirectly do more computations under tape semantics. To see how, consider a program that has anonymous functions as branches: $(\widehat{\text{if}} \text{ [true] then } (\lambda_ \Rightarrow 1) \text{ else } (\lambda_ \Rightarrow \text{loop})) ()$. Under reveal semantics, the $\widehat{\text{if}}$ expression evaluates to $\lambda_ \Rightarrow 1$, and then the whole program terminates at 1, even though both branches are evaluated. On the other hand, the function application is distributed into both branches under tape semantics, and the program reduces to $\widehat{\text{if}} \text{ [true] then } 1 \text{ else } \text{loop}$ which diverges due to the second branch. We made a similar design decision for mux in E-MUX, although it is also reasonable to make this rule closer to tape semantics.

4.5.1 Metatheory of Reveal Semantics

Reveal semantics enjoys several metatheoretic properties. First, it is deterministic.

Theorem 4.5.1 (Determinism of reveal semantics). *If $e \Downarrow v_1$ and $e \Downarrow v_2$, then $v_1 = v_2$.*

Reveal semantics also guarantees type preservation with respect to $\lambda_{\text{O}ADT^+}$ ’s type system.

Theorem 4.5.2 (Preservation of reveal semantics). *If $\Gamma \vdash e :_l \tau$ and $e \Downarrow v$, then $\Gamma \vdash v :_l \tau$.*

Furthermore, the evaluated result v must be a value if the typing context Γ is empty.

A direct proof of [Theorem 4.5.2](#) is tricky. We instead reduce the proof of preservation of reveal semantics to that of parallel reduction, by first showing parallel reduction refines reveal semantics.

Lemma 4.5.3 (Parallel reduction refines reveal semantics). *If $e \Downarrow v$, then $e \Rightarrow^* v$.*

Most importantly, we are able to relate reveal semantics to tape semantics using the following simulation theorem.

Theorem 4.5.4 (Simulation of reveal semantics). *If $e \longrightarrow^* v$, then $e \Downarrow [v]$.*

If $\tau \longrightarrow^ \widehat{\omega}$, then $\tau \Downarrow \widehat{\omega}$.*

[Theorem 4.5.4](#) only requires v to be a weak value. This theorem relies on a crucial lemma that says the equivalence up to erasure preserves the reveal semantics relation, as follows.

Lemma 4.5.5. *If $e_1 \Downarrow v$ and $[e_1] = [e_2]$, then $e_2 \Downarrow v$.*

Note that [Theorem 4.5.4](#) only considers one direction: if e evaluates to a value under reveal semantics, e does not necessarily reduce to any equivalent values under tape semantics, due to different termination behaviors exhibited by these two semantics. Nonetheless, if an expression terminates under both semantics, the results are equivalent.

Corollary 4.5.6. *If $e \longrightarrow^* v$ and $e \Downarrow v'$, then $[v] = v'$.*

If $\tau \longrightarrow^ \widehat{\omega}$ and $\tau \Downarrow \widehat{\omega}'$, then $\widehat{\omega} = \widehat{\omega}'$.*

As a bonus, the determinism of reveal semantics allows us to prove that tape semantics is also deterministic (up to erasure) if the computation terminates (at some weak values).

Corollary 4.5.7 (Weak determinism of tape semantics). *If $e \longrightarrow^* v_1$ and $e \longrightarrow^* v_2$, then $[v_1] = [v_2]$.*

The formalization of reveal semantics and all the proofs in this section are mechanized in the Coq proof assistant.

4.6 Deriving Secure Implementations

Following the recipe in Figure 4.1, one can easily implement a secure version of a public program written in a standard way. This strategy of deriving secure implementations can be automated, further levitating the user burden. This section develops an algorithm, `chase`, that converts a given public program to an equivalent secure version, by chasing the commuting diagrams as in Figure 4.1.

The algorithm takes as input a public program e (source program), usually a function, its type τ (source type), and a target type $\dot{\tau}$ which also serves as the security specification. It then outputs a secure program \dot{e} of type $\dot{\tau}$. As an example, given the source program `lookup` of source type $\mathbb{Z} \rightarrow \text{tree} \rightarrow \mathbb{B}$ (with \top labels in all arguments elided) from Figure 3.1 and the target type $\Pi x:\mathbb{N}, \widehat{\mathbb{Z}} \rightarrow \widehat{\text{tree}} \text{ k} \rightarrow \widehat{\mathbb{B}}$ (with \perp labels in all arguments elided), `chase`($e;\tau;\dot{\tau}$) generates the oblivious function `lookup` in Figure 4.3. For simplicity, we require that the target type be in *prenex normal form*: $\dot{\tau}$ has a prefix of public views (e.g., $\Pi x:\mathbb{N}, \dots$), followed by a non-dependent (i.e., quantifier-free) component (e.g., $\widehat{\mathbb{Z}} \rightarrow \widehat{\text{tree}} \text{ k} \rightarrow \widehat{\mathbb{B}}$). We also assume the source type τ is non-dependent, as the source program is supposed to be implemented in the standard fragment of λ_{OADT^+} . The labels are elided in the rest of this section for brevity, unless they are unclear from the context: a source type always has \top labels for all of its arguments, while a target type has \perp labels.

Before we present the algorithm formally, we explore a few examples to demonstrate some design decisions of the algorithm. The first example converts a tree insertion function.

Example 4.6.1 (`insert`).

$$\text{chase}(\text{insert}; \mathbb{Z} \rightarrow \text{tree} \rightarrow \text{tree}; \Pi x:\mathbb{N}, \widehat{\mathbb{Z}} \rightarrow \widehat{\text{tree}} \text{ k} \rightarrow \widehat{\text{tree}} (\text{k}+1)) \equiv \\ \lambda \text{k x t} \Rightarrow \widehat{\text{tree}}\#\text{s} (\text{k}+1) (\text{insert} (\widehat{\mathbb{Z}}\#\text{r x}) (\widehat{\text{tree}}\#\text{r k t}))$$

As the result is an OADT, `chase` uses `tree#s` to convert the output of `insert` back to an oblivious tree. The `tape` keyword is not necessary in this example, because `tree#s` should have label \perp , implicitly having `taped` the result in its definition. It is crucial that the programmers provide the correct public views in the target type signatures: if the return type in this example was `tree k` instead, the resulting oblivious tree could get truncated

due to possible increase in the tree depth. Our algorithm assumes the provided public views are always “big” enough to hold the data.

The next example illustrates that our algorithm is able to handle polynomial types.

Example 4.6.2 (add1).

```
chase(add1;  $\mathbb{Z} + \mathbb{1} \rightarrow \mathbb{Z} + \mathbb{1}$ ;  $\widehat{\mathbb{Z}} \hat{+} \mathbb{1} \rightarrow \widehat{\mathbb{Z}} \hat{+} \mathbb{1}$ )  $\equiv$ 
   $\lambda n \Rightarrow$  let m =  $\widehat{\text{match}}$  n with
    |  $\widehat{\text{inl}}$  x  $\Rightarrow$   $\text{inl}$  ( $\widehat{\mathbb{Z}}\#r$  x)
    |  $\widehat{\text{inr}}$  x  $\Rightarrow$   $\text{inr}$  x
  in
  tape
    ( $\text{match}$  add1 m with
      |  $\text{inl}$  x  $\Rightarrow$   $\widehat{\text{inl}}$  ( $\widehat{\mathbb{Z}}\#s$  x)
      |  $\text{inr}$  x  $\Rightarrow$   $\widehat{\text{inr}}$  x)
```

The function `add1` adds 1 to the left injection of an option type $\mathbb{Z} + \mathbb{1}$ (i.e., an `fmap`). The generated program converts between sum type ($\mathbb{Z} + \mathbb{1}$) and oblivious sum type ($\widehat{\mathbb{Z}} \hat{+} \mathbb{1}$), similar to the previous examples. However, the conversions are derived from the polynomial types (e.g., sum types), instead of provided by the users as section and retraction functions.

The last example is a standard higher-order function `map`.

Example 4.6.3 (map).

```
chase(map; ( $\mathbb{Z} \rightarrow \mathbb{Z}$ )  $\rightarrow$   $\text{tree} \rightarrow \text{tree}$ ;  $\prod k:\mathbb{N}$ , ( $\mathbb{Z} \rightarrow \mathbb{Z}$ )  $\rightarrow$   $\widehat{\text{tree}} k \rightarrow \widehat{\text{tree}} k$ )  $\equiv$ 
   $\lambda k f t \Rightarrow$   $\widehat{\text{tree}}\#s$  k (map f ( $\widehat{\text{tree}}\#r$  k t))
```

While this secure implementation is fairly easy to derive by chasing the commuting diagram, the higher-order argument ($\mathbb{Z} \rightarrow \mathbb{Z}$) in the target type is required to remain public. Converting higher-order arguments is possible for the primitive types in this example by “transporting” these arguments properly. However, it becomes challenging when the higher-order arguments use OADTs, which may force us to infer the public views of these OADTs. Thus, our algorithm opts to not convert higher-order arguments. This limitation is nonetheless insignificant in practice, because it is easier for a client of the generated secure `map` function to supply a public higher-order argument anyway, which also avoids unnecessary conversions.

$\text{gsec}(e; \tau; \dot{\tau}) = \dot{e}$	
$\frac{\text{GS-UNIT}}{\text{gsec}(e; \mathbf{1}; \mathbf{1}) = \text{tape } e}$	$\frac{\text{GS-BOOL}}{\text{gsec}(e; \mathbb{B}; \widehat{\mathbb{B}}) = \text{tape } (\widehat{\mathbb{B}}\#s \ e)}$
$\frac{\text{GS-OADT}}{\text{gsec}(e; \widehat{T}; \widehat{T} \ k) = \widehat{T}\#s \ k \ e}$	$\frac{\text{GS-PROD} \quad \text{gsec}(\pi_1 \ e; \tau_1; \dot{\tau}_1) = \dot{e}_1 \quad \text{gsec}(\pi_2 \ e; \tau_2; \dot{\tau}_2) = \dot{e}_2}{\text{gsec}(e; \tau_1 \times \tau_2; \dot{\tau}_1 \times \dot{\tau}_2) = (\dot{e}_1, \dot{e}_2)}$
$\frac{\text{GS-SUM} \quad \text{gsec}(x; \tau_1; \dot{\tau}_1) = \dot{e}_1 \quad \text{gsec}(x; \tau_2; \dot{\tau}_2) = \dot{e}_2}{\text{gsec}(e; \tau_1 + \tau_2; \dot{\tau}_1 + \dot{\tau}_2) = \text{tape } (\widehat{\text{match}} \ e \ \text{with } x \Rightarrow \widehat{\text{inl}} \ \dot{e}_1 \mid x \Rightarrow \widehat{\text{inr}} \ \dot{e}_2)}$	
$\text{gret}(e; \tau; \dot{\tau}) = \dot{e}$	
$\frac{\text{GR-ID}}{\text{gret}(e; \tau; \tau) = e}$	$\frac{\text{GR-BOOL}}{\text{gret}(e; \widehat{\mathbb{B}}; \mathbb{B}) = \widehat{\text{if}} \ e \ \text{then } \text{true} \ \text{else } \text{false}}$
$\frac{\text{GR-OADT}}{\text{gret}(e; \widehat{T} \ k; T) = \widehat{T}\#r \ k \ e}$	$\frac{\text{GR-PROD} \quad \text{gret}(\pi_1 \ e; \tau_1; \dot{\tau}_1) = \dot{e}_1 \quad \text{gret}(\pi_2 \ e; \tau_2; \dot{\tau}_2) = \dot{e}_2}{\text{gret}(e; \tau_1 \times \tau_2; \dot{\tau}_1 \times \dot{\tau}_2) = (\dot{e}_1, \dot{e}_2)}$
$\frac{\text{GR-SUM} \quad \text{gret}(x; \tau_1; \dot{\tau}_1) = \dot{e}_1 \quad \text{gret}(x; \tau_2; \dot{\tau}_2) = \dot{e}_2}{\text{gret}(e; \tau_1 + \tau_2; \dot{\tau}_1 + \dot{\tau}_2) = \text{match } e \ \text{with } x \Rightarrow \text{inl} \ \dot{e}_1 \mid x \Rightarrow \text{inr} \ \dot{e}_2}$	
$\frac{\text{GR-OSUM} \quad \text{gret}(x; \tau_1; \dot{\tau}_1) = \dot{e}_1 \quad \text{gret}(x; \tau_2; \dot{\tau}_2) = \dot{e}_2}{\text{gret}(e; \tau_1 + \tau_2; \dot{\tau}_1 + \dot{\tau}_2) = \widehat{\text{match}} \ e \ \text{with } x \Rightarrow \text{inl} \ \dot{e}_1 \mid x \Rightarrow \text{inr} \ \dot{e}_2}$	

Figure 4.11. Generalized section and retraction

Chapter 6 will describe a more general static approach to convert a public program to an equivalent secure program, which has better support for higher-order functions.

4.6.1 Derivation Algorithm

We first define the procedures that convert between public and oblivious polynomial types. The *generalized section* and *generalized retraction* procedures, $\text{gsec}(e; \tau; \dot{\tau})$ and $\text{gret}(e; \tau; \dot{\tau})$ respectively, generate a section computation or retraction computation that converts e from

$$\boxed{\text{chase}(e; \tau; \dot{\tau}) = \dot{e}}$$

$$\frac{\text{C-PREFIX} \quad \frac{x:\tau' \text{ is the public view prefix} \quad \text{chase}(e; \tau; \dot{\tau}) = \dot{e}}{\text{chase}(e; \tau; \Pi x:\tau', \dot{\tau}) = \lambda x:\tau' \Rightarrow \dot{e}}}{\text{C-RET} \quad \frac{\text{gret}(x; \dot{\tau}_1; \tau_1) = r \quad \text{chase}(e \ r; \tau_2; \dot{\tau}_2) = \dot{e}}{\text{chase}(e; \tau_1 \rightarrow \tau_2; \dot{\tau}_1 \rightarrow \dot{\tau}_2) = \lambda x:\dot{\tau}_1 \Rightarrow \dot{e}}} \quad \text{C-SEC} \quad \frac{\text{gsec}(e; \tau; \dot{\tau}) = \dot{e}}{\text{chase}(e; \tau; \dot{\tau}) = \dot{e}}$$

Figure 4.12. Commuting diagram chasing

type τ to type $\dot{\tau}$. The result of `gsec` is always safe (i.e., with \perp label), while that of `gret` may be leaky (i.e., with \top label). On the other hand, the input e to `gret` is assumed safe. [Figure 4.11](#) shows the definitions of these two procedures as inference rules. Note that these generators are partially defined; trying to convert a $\mathbb{1}$ to \mathbb{B} will fail, for example. Most of these rules are straightforward. `GS-OADT` and `GR-OADT` outsource the conversions between ADTs and OADTs to the user-defined section and retraction functions, assuming \hat{T} is an OADT of T . Generalized retraction also permits an identity conversion per `GR-ID`, if the source and target types are the same.

The main `chase` algorithm, presented in [Figure 4.12](#) as inference rules, captures the key idea of the recipe from [Figure 4.1](#). After introducing the public view arguments using `C-PREFIX`, `chase` “decrypts” every function argument using the retraction computation generated by `gret` in `C-RET`, and finally “encrypt” the result using the section computation generated by `gsec` in `C-SEC`. Similar to `gsec` and `gret`, `chase` is also partial: the algorithm fails if the source and target types do not match.

4.6.2 Metatheory of the Derivation Algorithm

The `chase` algorithm always generates well-typed and equivalent programs. The formalization of the algorithm and the proofs in this section are mechanized in the Coq proof assistant.

The well-typedness theorem, stated as follows, also guarantees that the generated programs are secure, thanks to [Theorem 4.2.4](#).

Theorem 4.6.1 (Well-typedness of the derivation algorithm). *Given a well-kinded target type $\dot{\tau}$, if $\text{chase}(e, \tau, \dot{\tau}) = \dot{e}$ and $\cdot \vdash e :_{\perp} \tau$, then $\cdot \vdash \dot{e} :_{\perp} \dot{\tau}$.*

To establish the correctness of this algorithm, we consider two programs equivalent *up to revelation*, defined as follows in a logical relation style [30, 31].

Definition 4.6.1 (Equivalence of source and target programs). The equivalence between a source program and a target program is defined as set-valued denotations indexed by their types: a value interpretation $\mathcal{V}[\tau; \dot{\tau}]$ and an expression interpretation $\mathcal{E}[\tau; \dot{\tau}]$. We say two values v of type τ and \dot{v} of type $\dot{\tau}$ are equivalent up to revelation if $(v, \dot{v}) \in \mathcal{V}[\tau; \dot{\tau}]$. The equivalence of two expressions are similarly defined using $\mathcal{E}[\tau; \dot{\tau}]$.

$$\mathcal{V}[\tau; \tau] = \{ (v, v) \} \quad \mathcal{V}[\mathbb{B}; \widehat{\mathbb{B}}] = \{ (b, [b]) \} \quad \mathcal{V}[\mathbb{T}; \widehat{\mathbb{T}} \ k] = \{ (v, \dot{v}) \mid \widehat{\mathbb{T}}\#_{\mathbf{r}} \ k \ \dot{v} \Downarrow v \}$$

$$\mathcal{V}[\tau_1 \times \tau_2; \dot{\tau}_1 \times \dot{\tau}_2] = \{ ((v_1, v_2), (\dot{v}_1, \dot{v}_2)) \mid (v_1, \dot{v}_1) \in \mathcal{V}[\tau_1; \dot{\tau}_1] \wedge (v_2, \dot{v}_2) \in \mathcal{V}[\tau_2; \dot{\tau}_2] \}$$

$$\mathcal{V}[\tau_1 + \tau_2; \dot{\tau}_1 + \dot{\tau}_2] = \cup \left\{ \begin{array}{l} \{ (\text{inl } v, \text{inl } \dot{v}) \mid (v, \dot{v}) \in \mathcal{V}[\tau_1; \dot{\tau}_1] \} \\ \{ (\text{inr } v, \text{inr } \dot{v}) \mid (v, \dot{v}) \in \mathcal{V}[\tau_2; \dot{\tau}_2] \} \end{array} \right.$$

$$\mathcal{V}[\tau_1 + \tau_2; \dot{\tau}_1 \widehat{+} \dot{\tau}_2] = \cup \left\{ \begin{array}{l} \{ (\text{inl } v, [\text{inl } \dot{v}]) \mid (v, \dot{v}) \in \mathcal{V}[\tau_1; \dot{\tau}_1] \} \\ \{ (\text{inr } v, [\text{inr } \dot{v}]) \mid (v, \dot{v}) \in \mathcal{V}[\tau_2; \dot{\tau}_2] \} \end{array} \right.$$

$$\mathcal{E}[\tau; \overline{\Pi_{\mathbf{x}:\tau'}}, \dot{\tau}] = \left\{ (e, \dot{e}) \mid \begin{array}{l} \forall \bar{v}. (\forall i. \cdot \vdash v_i :_{\perp} \tau'_i) \implies \\ (e, \dot{e} \ \bar{v}) \in \mathcal{E}[\tau; \dot{\tau}] \end{array} \right\} \quad \text{if } \overline{\Pi_{\mathbf{x}:\tau'}} \text{ is the public view prefix}$$

$$\mathcal{E}[\tau_1 \rightarrow \tau_2; \dot{\tau}_1 \rightarrow \dot{\tau}_2] = \left\{ (e, \dot{e}) \mid \begin{array}{l} \forall v \dot{v}. \cdot \vdash v :_{\perp} \tau_1 \wedge \cdot \vdash \dot{v} :_{\perp} \dot{\tau}_1 \wedge (v, \dot{v}) \in \mathcal{V}[\tau_1; \dot{\tau}_1] \implies \\ (e \ v, \dot{e} \ \dot{v}) \in \mathcal{E}[\tau_2; \dot{\tau}_2] \end{array} \right\}$$

$$\mathcal{E}[\tau; \dot{\tau}] = \{ (e, \dot{e}) \mid \forall v \dot{v}. e \longrightarrow^* v \wedge \dot{e} \longrightarrow^* \dot{v} \implies ([v], [\dot{v}]) \in \mathcal{V}[\tau; \dot{\tau}] \}$$

The value denotation of OADT shows why this equivalence is up to revelation: an OADT value \dot{v} is equivalent to a public value v if \dot{v} can be “decrypted” to v . Note that we use reveal semantics for this “decryption”. Unlike most logical relations in the literatures, two function

values are equivalent only if they are equal, reflecting the design decision of not converting higher-order arguments. The first two cases of expression denotation consume function arguments (including the public view arguments), and the last case says two (non-functional) expressions are equivalent if they reduce to equivalent values after erasure. The erasure is necessary here because expressions may reduce to weak values under tape semantics, and we consider two weak values, e.g., leaky conditionals, equivalent if their erasures (i.e., revelation) are equivalent. The last case of expression denotation also has an elided side condition that assumes the correctness of the public views in the target type; we refer interested readers to the Coq development for the full details.

We are now ready to state the correctness theorem for our diagram chasing algorithm.

Theorem 4.6.2 (Correctness of the derivation algorithm). *If $\text{chase}(e; \tau; \dot{\tau}) = \dot{e}$, then $(e, \dot{e}) \in \mathcal{E}[\tau; \dot{\tau}]$.*

The direct proof of this theorem involves a lot of tedious and complex reasoning about tape semantics. We instead reduce the expression denotation (specifically the last case in [Definition 4.6.1](#)) to the one that uses reveal semantics from the previous section, which greatly simplifies the proofs.

4.7 Conclusion

While λ_{OADT} enables secure applications that use private data structures and complex privacy policies, the lack of modularity forces users to implement ad-hoc secure versions of their programs for each desired policy. To allow programmers to write a single function and easily build secure programs with different public views, we have developed $\lambda_{\text{OADT}^\dagger}$. This language is equipped with a novel semantics that repairs potential leaks without compromising the security guarantees of λ_{OADT} . An obliviousness theorem analogous to the one for λ_{OADT} is mechanically proved. Despite being non-standard, tape semantics is equivalent to a standard semantics in a way that enables familiar reasoning. We have also developed an algorithm that derives an equivalent secure program automatically from a given policy specification and functionality written in a conventional way.

5. TAYPE: A POLICY-AGNOSTIC OBLIVIOUS LANGUAGE

The calculus introduced in [Chapter 4](#), λ_{OADT^+} , permits security concerns to be decoupled from the program logic of an oblivious computation. The first key component is a dependent security type system in which *oblivious algebraic data types* could be encoded. These types equip private data with a *public view*, and guarantee that every private value with the same view is indistinguishable, i.e., an attacker can learn nothing about private data other than what its public view entails. The second key component is a novel *tape semantics* that uses security information provided by the type system to dynamically *repair* any potential information leaks at runtime. These components allow the programmer to write a program as normal, and then combine it with the desired public view, relying on the tape semantics to patch any information leaks. Unfortunately, λ_{OADT^+} lacked an accompanying implementation.

This chapter presents a programming language for writing oblivious computations, TAYPE, that implements both the oblivious algebraic data types and tape semantics proposed in λ_{OADT^+} . TAYPE is equipped with a bidirectional type checker that enforces correct use of secure operations, and automatically infers annotations that enable potential leaks to be repaired. Our implementation is realized in a compilation pipeline, shown in [Figure 5.1](#), that translates a TAYPE program and privacy policy (in the form of a public view) to an OCaml implementation which, when linked with a cryptographic backend, can be used by client programs to securely compute functions over private data. The main challenge that this toolchain must overcome is how to securely implement these two language features in a standard functional programming language. To implement oblivious types, our key idea is to represent dependently typed oblivious data using an *oblivious array*, and the types themselves as sizes indexing into an array. To implement tape semantics, we equip each type, including function types, with a *leaky structure* that reifies potentially leaky operations into a distinguished data type and inserts repairs when values of this type are used.

To summarize, the contributions of this chapter are as follows:

- We implement a bidirectional type checker for an extension of λ_{OADT^+} . Given a source program, this checker outputs a fully-annotated version in a typed core language called `core TAYPE`.

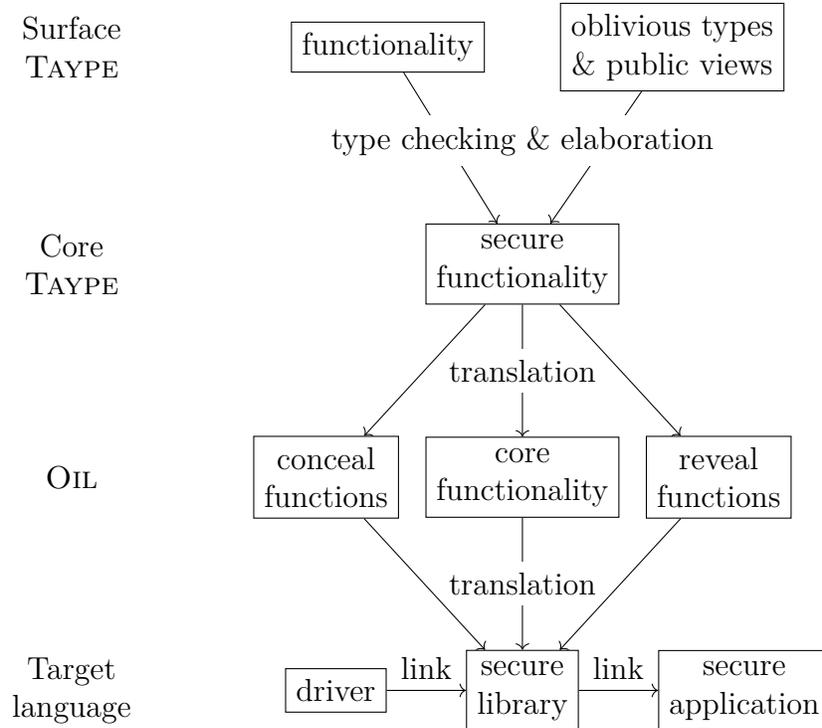


Figure 5.1. Compilation pipeline

- We present a translation from core TAYPE to OIL, an ML-style functional language with rank-1 polymorphism, built-in oblivious arrays, and secure array operations. In addition to translating the core functionality of the application, our translation also produces routines for concealing and revealing private data, which clients need to build a complete MPC application.
- We evaluate our implementation against several case studies and microbenchmarks. Our experiments feature a diverse set of computations and a range of security policies, including the medical record example in [Chapter 1](#), and also demonstrate that tradeoffs between privacy and performance can be made easily with our approach.

An artifact containing the Coq mechanization of core TAYPE, the implementation of TAYPE, its source code, and the source for all the benchmarks in our experiments with instructions is publicly available [\[49\]](#).

5.1 Overview

To demonstrate our approach, consider a simple list membership predicate written in TAYPE, `elem`, shown in Figure 5.2. Suppose Alice, the owner of a list, and Bob, the owner of an integer, want to check if Bob’s integer occurs in Alice’s list, without revealing any information beyond their own input and the result. Similar to λ_{OADT^+} , TAYPE allows participants to choose what public information to share as part of the security policy, and its type system ensures that all public data and computation only depend on this public view.

```

data list = Nil | Cons ℤ list
fn elem : ℤ → list → ℬ =
  λy xs ⇒
    match xs with
    | Nil ⇒ false
    | Cons x xs' ⇒
      if x = y then true
      else elem y xs'

```

Figure 5.2. List membership predicate

For example, Alice may be okay with sharing the size of the list, or with releasing some upper bound on its length.

Using oblivious algebraic data types (Chapter 3) and tape semantics (Chapter 4), we can implement a secure version of this membership predicate by composing the standard implementation in Figure 5.2 and the desired privacy policy (i.e., OADTs and their section and retraction functions). Figure 5.3 shows the full implementation of an oblivious list $\widehat{\text{list}}$ and its the section and retraction functions, using the maximum length of the list as its public view.

While Chapter 4 formalized a core calculus of oblivious algebraic data types and tape semantics, λ_{OADT^+} , it lacked both an algorithmic type checker and implementation; this chapter presents the design and implementation of a language for oblivious computation with both.

5.1.1 Type Checking and Core Taype

The input to our compiler is a program written in *surface* TAYPE, such as `elem`, $\widehat{\text{list}}\#r$, and $\widehat{\text{list}}\#s$ from Figure 5.2 and Figure 5.3. This language is equipped with a *bidirectional type checker* [50] that enforces correct use of secure and leaky operations, which ensures that

```

obliv  $\widehat{\text{list}}$  (k :  $\mathbb{Z}$ ) =
  if k = 0 then 1
  else 1  $\hat{+}$   $\widehat{\mathbb{Z}}$   $\hat{\times}$   $\widehat{\text{list}}$  (k-1)

#[section]
fn  $\widehat{\text{list}}\#s$  : (k :  $\mathbb{Z}$ )  $\rightarrow$  list  $\rightarrow$   $\widehat{\text{list}}$  k =  $\lambda k$  xs  $\Rightarrow$ 
  if k = 0 then ()
  else tape (match xs with
    | Nil  $\Rightarrow$   $\widehat{\text{inl}}$  ()
    | Cons x xs'  $\Rightarrow$ 
       $\widehat{\text{inr}}$  [tape ( $\widehat{\mathbb{Z}}\#s$  x),  $\widehat{\text{list}}\#s$  (k-1) xs'])

#[retraction]
fn  $\widehat{\text{list}}\#r$  : (k :  $\mathbb{Z}$ )  $\rightarrow$   $\widehat{\text{list}}$  k  $\rightarrow$  list =  $\lambda k$   $\Rightarrow$ 
  if k = 0 then  $\lambda \_ \Rightarrow$  Nil
  else  $\lambda xs \Rightarrow$   $\widehat{\text{match}}$  xs with
    |  $\widehat{\text{inl}}$  _  $\Rightarrow$  Nil
    |  $\widehat{\text{inr}}$  [x, xs']  $\Rightarrow$ 
      Cons ( $\widehat{\mathbb{Z}}\#r$  x) ( $\widehat{\text{list}}\#r$  (k-1) xs')

#[safe]
fn  $\widehat{\text{elem}}$  : (k :  $\mathbb{Z}$ )  $\rightarrow$   $\widehat{\mathbb{Z}}$   $\rightarrow$   $\widehat{\text{list}}$  k  $\rightarrow$   $\widehat{\mathbb{B}}$  =  $\lambda k$  x xs  $\Rightarrow$ 
  tape ( $\widehat{\mathbb{B}}\#s$  (elem ( $\widehat{\mathbb{Z}}\#r$  x) ( $\widehat{\text{list}}\#r$  k xs)))

```

Figure 5.3. An oblivious implementation of $\widehat{\text{elem}}$

all well-typed programs are oblivious. After type checking, programs in this language are elaborated into an intermediate language called *core TAYPE* (Section 5.2).

Core TAYPE programs are fully annotated with types and, crucially, *leakage labels* (Chapter 4). Leakage labels track whether an expression contains potential leaks, i.e., whether it contains any leaky operations: we say an expression is *leaky* (labelled \top) if so, and *safe* (labelled \perp) otherwise. For example, $\widehat{\text{if}}$ x then 1 else 2 is obviously leaky, as executing it naively leaks the private condition, while mux x [1] [2] and false are safe. In contrast to λ_{OADT^+} , an addition in core TAYPE is its *promotion* operation \uparrow , which explicitly casts a

<pre> fn elem_c :_T Z_T → list_T → B = λ(y :_T Z) (xs :_T list) ⇒ match xs with Nil ⇒ ↑false Cons x xs' ⇒ if x = y then ↑true else elem_c y xs' </pre>	<pre> fn elem_o : Z̃ → list̃ → B̃ = λy xs ⇒ match_{list} if_B xs (prom_B false) (λx xs' ⇒ if̃ if_B (x ≅ y) (prom_B true) (elem_o y xs')) </pre>
--	---

Figure 5.4. A fully annotated implementation of `elem` in core TAYPE (`elemc`) and its translation in OIL (`elemo`)

safe expression to a leaky one, to help with the translation: the `↑false` on the fourth line of `elemc` in Figure 5.4 is treated as a leaky expression, for example. Note that none of the label annotations or promotion are required in the surface language: they are either inferred or automatically inserted during elaboration.

5.1.2 Translating to Oil

The next compilation phase translates programs in core TAYPE into OIL, the *OADT intermediate language* (Section 5.3). OIL is an ML-style functional language with rank-1 polymorphism, extended with an *oblivious array* and its operations. These oblivious “primitives” will eventually be implemented by a cryptographic backend in the target language (Section 5.4), and OIL is agnostic to the particular implementation. OIL is designed to be a common subset of most standard functional languages, so that translating OIL to a particular language, e.g., OCaml, is straightforward. The main challenge in this phase is expressing the unique features of TAYPE that do not appear in conventional languages, particularly oblivious types (i.e., dependent types) and its tape semantics. Many of the core ingredients of this translation can be seen in Figure 5.4, which gives the fully elaborated implementation of `elem` in core TAYPE and its corresponding OIL version.

```

data  $\tilde{\mathcal{A}}$  = prom $_{\mathcal{A}}$   $\mathcal{A}$  | if $_{\mathcal{A}}$   $\mathcal{A}$   $\tilde{\mathcal{A}}$   $\tilde{\mathcal{A}}$ 
data  $\tilde{\mathbb{Z}}$  = r $_{\mathbb{Z}}$   $\mathcal{A}$  | prom $_{\mathbb{Z}}$   $\mathbb{Z}$  | if $_{\mathbb{Z}}$   $\mathcal{A}$   $\tilde{\mathbb{Z}}$   $\tilde{\mathbb{Z}}$ 

fn  $\tilde{\mathbb{Z}}\#s$  :  $\tilde{\mathbb{Z}}$   $\rightarrow$   $\tilde{\mathcal{A}}$  =  $\lambda\tilde{n} \Rightarrow$ 
  match  $\tilde{n}$  with
  | r $_{\mathbb{Z}}$   $\hat{n} \Rightarrow$  prom $_{\mathcal{A}}$   $\hat{n}$ 
  | prom $_{\mathbb{Z}}$   $n \Rightarrow$  prom $_{\mathcal{A}}$  ( $\tilde{\mathbb{Z}}\#s$   $n$ )
  | if $_{\mathbb{Z}}$   $\hat{b}$   $\tilde{n}_1$   $\tilde{n}_2 \Rightarrow$ 
    if $_{\mathcal{A}}$   $\hat{b}$  ( $\tilde{\mathbb{Z}}\#s$   $\tilde{n}_1$ ) ( $\tilde{\mathbb{Z}}\#s$   $\tilde{n}_2$ )

fn  $\widehat{\text{list}}$  :  $\mathbb{Z} \rightarrow \mathbb{N} = \lambda k \Rightarrow$ 
  if  $k = 0$  then 0
  else 1 + max 0 (1 +  $\widehat{\text{list}}$  ( $k-1$ ))

fn  $\widehat{\text{tape}}$  :  $\tilde{\mathcal{A}} \rightarrow \mathcal{A} = \lambda\tilde{a} \Rightarrow$ 
  match  $\tilde{a}$  with
  | prom $_{\mathcal{A}}$   $\hat{a} \Rightarrow$   $\hat{a}$ 
  | if $_{\mathcal{A}}$   $\hat{b}$   $\tilde{a}_1$   $\tilde{a}_2 \Rightarrow$ 
    mux  $\hat{b}$  ( $\widehat{\text{tape}}$   $\tilde{a}_1$ ) ( $\widehat{\text{tape}}$   $\tilde{a}_2$ )

```

Figure 5.5. Selected leaky types and functions in OIL

Translating Oblivious Type Definitions

As OIL does not have type-level computation, the definition of the oblivious type $\widehat{\text{list}}$ is translated into a function from the public view to its *size* \mathbb{N} , shown in Figure 5.5. As we will see shortly, the size of an oblivious type can be used to access secure data residing in an oblivious array.

Translating Oblivious Types and Operations

We represent every oblivious type as a single uniform type, the oblivious array \mathcal{A} . This array is essentially a secure “buffer” holding the private data. For example, $\widehat{\text{list}}\ k$ in the type signature of $\widehat{\text{list}}\#s$, in Figure 5.3, is translated to this array type \mathcal{A} , regardless of the public view k . Even though oblivious types are all flat arrays in OIL, the rich typing information is not lost: we can still extract the needed private information by (securely) accessing the array using the sizes of oblivious types, such as the aforementioned $\widehat{\text{list}}$. Oblivious operations are translated into corresponding oblivious array operations. For example, an oblivious pair of private data is simply the concatenation of the two corresponding arrays, and destructing an oblivious pair amounts to taking a slice of the array using the two components’ sizes. Section 5.3.2 describes the translation of other oblivious constructs, including injections into oblivious sums.

Translating Tape Semantics

Implementing the tape semantics is the main challenge in translating from TAYPE to OIL. Recall the three key ideas of the tape semantics. First, leaky operations, such as $\widehat{\text{if}}$, are themselves irreducible. Second, the surrounding context of $\widehat{\text{if}}$ is distributed into both branches. Third, the `tape` operation repairs potential leaks, by turning $\widehat{\text{if}}$ into `mux`.

To implement the first idea, we translate leaky types, e.g., \mathbb{Z}_\top , into a *leaky representation*, e.g., $\widetilde{\mathbb{Z}}$, a data type that explicitly represents expressions that may contain potential leaks. By convention, we use $\widetilde{\cdot}$ as a visual cue for a leaky representation, its associated functions and variables. The leaky representations of oblivious arrays ($\widetilde{\mathbb{A}}$) and integers ($\widetilde{\mathbb{Z}}$) are shown in Figure 5.5. The only way to build a leaky oblivious data type is to promote a safe one or using a leaky conditional, so we simply encode these leaky operations as the constructors of its leaky representation $\widetilde{\mathbb{A}}$. $\widetilde{\mathbb{Z}}$ also includes both of these constructors, as well as its own retraction operation. This encoding trivially makes the leaky operations irreducible. Leaky representations of ADTs are built using a similar strategy (Section 5.3.2). Every leaky representation needs to have reified versions of `prom` and $\widehat{\text{if}}$, because \uparrow and $\widehat{\text{if}}$ can be applied to any TAYPE type. During translation, they are instantiated using a process similar to typeclass resolution: the promotion of `false` in `elemc`, for example, is resolved to `promB`.

To distribute surrounding contexts into leaky constructs, we instrument the possible surrounding contexts to handle the leaky operations, by translating them into recursive functions following the tape semantics. For example, $\widehat{\mathbb{Z}}\#s$ is translated to $\widetilde{\mathbb{Z}}\#s$, also shown in Figure 5.5. Observe the last case of $\widetilde{\mathbb{Z}}\#s$, which has recursive calls to itself in both $\widehat{\text{if}}_{\mathbb{Z}}$ branches: this aligns with our intuition from the execution trace of the following example:

```

    tape ( $\widehat{\mathbb{Z}}\#s$  ( $\widehat{\text{if}}$  [true] then 3 else 4))
  → tape ( $\widehat{\text{if}}$  [true] then  $\widehat{\mathbb{Z}}\#s$  3 else  $\widehat{\mathbb{Z}}\#s$  4)
  →* tape ( $\widehat{\text{if}}$  [true] then [3] else [4]) → mux [true] [3] [4] → [3]

```

On the other hand, $\widetilde{\mathbb{Z}\#\mathbf{s}}$'s handling of \mathbf{r}_Z matches the following execution trace, for example:

$$\begin{aligned} & \text{tape } (\widehat{\mathbb{Z}\#\mathbf{s}} (\widehat{\mathbb{Z}\#\mathbf{r}} [3] + \widehat{\mathbb{Z}\#\mathbf{r}} [2])) \\ \longrightarrow & \text{tape } (\widehat{\mathbb{Z}\#\mathbf{s}} (\widehat{\mathbb{Z}\#\mathbf{r}} ([3] \hat{+} [2]))) \\ \longrightarrow & \text{tape } (\widehat{\mathbb{Z}\#\mathbf{s}} (\widehat{\mathbb{Z}\#\mathbf{r}} [5])) \longrightarrow \text{tape } [5] \longrightarrow [5] \end{aligned}$$

Our solution to patching leaky computation without `tape` is encapsulated in the definition of $\widetilde{\text{tape}}$, shown in [Figure 5.5](#). The function simply converts all reified `ifs` into `muxs`, and is essentially a transcription of `tape`'s evaluation rule.

To see how all these fit together, the initial expression from the previous example, `tape` $(\widehat{\mathbb{Z}\#\mathbf{s}} (\widehat{\text{if}} [\text{true}] \text{ then } 3 \text{ else } 4))$, is translated into the following OIL program:

$\widetilde{\text{tape}} (\widetilde{\mathbb{Z}\#\mathbf{s}} (\widehat{\text{if}}_Z [\text{true}] (\text{prom}_Z 3) (\text{prom}_Z 4)))$

Readers can verify that evaluating this program using a standard semantics produces the same behavior seen in the previous execution trace generated by `tape` semantics.

5.2 Taype, Formally

This section describes the fully annotated core TAYPE language. This language is inspired by the core calculus $\lambda_{\text{OADT}\dagger}$, but adds several features to aid its translation to OIL, including oblivious products, label promotion, ML-style ADT definitions, and explicit and uniform label checking. The user-facing version of TAYPE allows for many annotations to be omitted; these annotations are automatically inferred by our bidirectional type checker ([Section 5.2.5](#)) before translation to OIL ([Section 5.3](#)).

5.2.1 Syntax

[Figure 5.6](#) shows the syntax of core TAYPE. Types and terms are in the same syntactic class, similar to $\lambda_{\text{OADT}\dagger}$. By convention, we use \mathbf{e} for terms and τ for types whenever possible. A core TAYPE program consists of a global context of ADTs, functions and oblivious types, defined using `data`, `fn` and `obliv` respectively. We use lower case \mathbf{x} for function and variable names, \mathbf{C} for constructors, \mathbf{T} for ADT names and $\widehat{\mathbf{T}}$ for OADT names. Each constructor of an

$e, \tau ::=$	$\mathbb{B} \mid \mathbb{Z} \mid \tau \times \tau$ $\mathbf{1} \mid \widehat{\mathbb{B}} \mid \widehat{\mathbb{Z}} \mid \tau \widehat{\times} \tau \mid \tau \widehat{+} \tau$ $() \mid \mathbf{b} \mid \mathbf{n} \mid \mathbf{x} \mid \mathbb{T}$ $\Pi x: {}_l \tau, \tau \mid \lambda x: {}_l \tau \Rightarrow e$ $e \oplus e \mid e \widehat{\oplus} e$ $e \ e \mid \mathbb{C} \ e \mid \widehat{\mathbb{T}} \ e$ $\text{let } x: {}_l \tau = e \text{ in } e$ $\text{if}_\tau e \text{ then } e \text{ else } e$ $\text{mux } e \ e \ e$ $(e, e) \mid [e, e]$ $\widehat{l}_b \langle \tau \rangle \ e$ $\widehat{\text{match}}_\tau e \text{ with } (x, x) \Rightarrow e$ $\widehat{\text{match}} e: \tau \widehat{\times} \tau \text{ with } [x, x] \Rightarrow e$ $\widehat{\text{match}}_\tau e \text{ with } \overline{\mathbb{C}} \ x \Rightarrow e$ $\widehat{\text{match}}_\tau e: \tau \widehat{+} \tau \text{ with } x \Rightarrow e \mid x \Rightarrow e$ $\widehat{\mathbb{B}} \# s \ e \mid \widehat{\mathbb{Z}} \# s \ e$ $\widehat{\mathbb{Z}} \# r \ e$ $\widehat{\text{if}} e \text{ then } e \text{ else } e$ $\uparrow e$ $\text{tape } e$ $[b] \mid [n] \mid [l_b \langle \widehat{\omega} \rangle \ \widehat{v}]$	<p>EXPRESSIONS:</p> <p>standard types</p> <p>oblivious types</p> <p>literals and variables</p> <p>dependent function</p> <p>(oblivious) integer operations</p> <p>applications</p> <p>let binding</p> <p>conditional</p> <p>atomic conditional</p> <p>(oblivious) pair</p> <p>oblivious sum injection</p> <p>product elimination</p> <p>oblivious product elimination</p> <p>ADT elimination</p> <p>oblivious sum elimination</p> <p>primitive sections</p> <p>primitive integer retraction</p> <p>leaky conditional</p> <p>promotion</p> <p>tape operation</p> <p>runtime boxed values</p>
$D ::=$	$\text{data } \mathbb{T} = \overline{\mathbb{C}} \ \tau$ $\text{fn } x: {}_l \tau = e$ $\text{obliv } \widehat{\mathbb{T}} (x: \tau) = \tau$	<p>GLOBAL DEFINITIONS:</p> <p>algebraic data type definition</p> <p>(recursive) function definition</p> <p>(recursive) oblivious type definition</p>
$l ::=$	$\top \mid \perp$	<p>LEAKAGE LABEL</p>
$\widehat{\omega} ::=$	$\mathbf{1} \mid \mathbb{B} \mid \widehat{\mathbb{Z}} \mid \widehat{\omega} \widehat{\times} \widehat{\omega} \mid \widehat{\omega} \widehat{+} \widehat{\omega}$	<p>OBLIVIOUS TYPE VALUES</p>
$\widehat{v} ::=$	$() \mid [b] \mid [n] \mid [\widehat{v}, \widehat{v}] \mid [l_b \langle \widehat{\omega} \rangle \ \widehat{v}]$	<p>OBLIVIOUS VALUES</p>
$v ::=$	$\widehat{\text{if}} [b] \text{ then } v \text{ else } v \mid \uparrow v \mid \widehat{\mathbb{Z}} \# r \ v$ $\widehat{v} \mid \mathbf{b} \mid \mathbf{n} \mid (v, v) \mid \lambda x: {}_l \tau \Rightarrow e \mid \mathbb{C} \ v$	<p>WEAK VALUES</p> <p>VALUES</p>

Figure 5.6. Core TAYPE syntax: the annotations marked in gray are either omitted (e.g., promotion, labels) or optional (e.g., argument types to dependent functions) in the user-facing surface language; the expressions marked in brown are restricted to be variables in administrative normal form.

ADT definition takes exactly one argument for simplicity, but this argument can be $\mathbb{1}$ for constructors that takes no argument, or a tuple of types for constructors that have multiple arguments..

TAYPE features a number of oblivious types and constructs, including oblivious integers, booleans, sums, and conditionals. The primitive section functions $\widehat{\mathbb{B}}\#s$ and $\widehat{\mathbb{Z}}\#s$ “encrypt” boolean and integer values respectively. Unlike $\lambda_{\text{OADT}\dagger}$, TAYPE also includes oblivious product types ($\widehat{\times}$), which are built using $[\cdot, \cdot]$ and require both of their components to be oblivious and non-leaky. The atomic conditional `mux`, discussed in [Chapter 3](#), fully evaluates both of its branches before taking an atomic step to its final result.

In core TAYPE, the arguments of dependent function types and lambda abstractions are annotated with a *leakage label* that indicates if they accept leaky inputs. We say that an TAYPE expression is *leaky* (i.e., has the label \top) if it contain potential leaks, e.g., uses some leaky operations, and say that it is *safe* otherwise. Standard conditional, product and ADT pattern matching expressions are annotated with the result type, while the elimination forms for oblivious products and sums are also annotated with the type of the discriminée, to help with their translation. All of these annotations are inferred by our type checker. For brevity, we omit them from now if they can be inferred from the context. Note that product elimination in TAYPE is defined (positively) using a pattern matching expression, instead of (negatively) using projection as in $\lambda_{\text{OADT}\dagger}$. Similarly, ADT introduction and elimination forms are defined in ML-style, using constructors and pattern matching, instead of `fold` and `unfold`.

Leaky conditionals, $\widehat{\text{match}}$ expressions, and `tape` operations play a key role in the semantics of TAYPE. The leaky conditional $\widehat{\text{if}}$ is similar to `mux`, but it allows its branches to be non-oblivious; $\widehat{\text{match}}$ analysis for oblivious sums is similar. The promotion operation \uparrow explicitly converts a safe expression to a leaky one. Integer retraction $\widehat{\mathbb{Z}}\#r$ would reveal its oblivious argument if implemented naively (as would the other leaky operations), but this is disallowed by the semantics of TAYPE.

Oblivious type values ($\widehat{\omega}$), oblivious values (\widehat{v}) and runtime boxed values like `[b]` are all identical to the definitions in $\lambda_{\text{OADT}\dagger}$. Weak values (\widehat{v}) extends the one in $\lambda_{\text{OADT}\dagger}$ with promotion, primitive integer retraction and ML-style ADT values.

5.2.2 Semantics

Figure 5.7 shows a selection of the small-step operational semantics rules of core TAYPE (the full rules are in Appendix A.1). The judgment $e \longrightarrow e'$ means e steps to e' under a fixed global context of definitions, which we elide. S-CTX takes a step in a subexpression according to an evaluation context \mathcal{E} , also given in Figure 5.7. Oblivious types are subject to reduction, as seen in the evaluation contexts involving $\widehat{\times}$, $\widehat{+}$ and $\widehat{!}$. To prevent information leaks, all subexpressions of a `mux` are fully evaluated by first applying the S-CTX rule with the corresponding evaluation contexts, before `mux` itself can be reduced by the S-MUX rule. The semantics of `if` is similar. Note that an `if` expression is in normal form once all its components are normalized, in order to avoid revealing its private condition.

The evaluation rules involving `tape` are one of the distinguishing features of TAYPE. S-OIF captures the idea that the leaky conditional `if`, while in normal form, can still make progress by distributing its context into both branches. *Leaky contexts* $\widehat{\mathcal{E}}$ define what contexts can be distributed in this manner: other contexts are either ruled out by the type system or not useful. S-TAPEOIF and S-TAPEPROM show how the `tape` operation repairs an expression with potential leaks. In addition to turning `if` into `mux`, the enclosing `tape` is pushed inside the branches of `mux` in order to ensure any leaks they contain are also patched. On the other hand, S-TAPEPROM simply extracts the safe oblivious value from a promotion. In contrast to λ_{OADT^+} , TAYPE also includes new rules for promoted expressions. S-SECRETINT repairs the leaky operation $\widehat{\mathbb{Z}}\#r$ by canceling it with $\widehat{\mathbb{Z}}\#s$, for example, but it also promotes the resulting oblivious integer in order to preserve the leakage label. S-SECINTPROM shows how promotion interacts with $\widehat{\mathbb{Z}}\#s$.

$e \longrightarrow e'$

$$\begin{array}{c}
\text{S-CTX} \\
\frac{e \longrightarrow e'}{\mathcal{E}[e] \longrightarrow \mathcal{E}[e']} \\
\\
\text{S-MUX} \\
\frac{}{\text{mux } [b] \ v_1 \ v_2 \longrightarrow \text{ite}(b, v_1, v_2)} \\
\\
\text{S-SECRETINT} \\
\frac{}{\widehat{\mathbb{Z}}\#s \ (\widehat{\mathbb{Z}}\#r \ [n]) \longrightarrow \uparrow[n]} \\
\\
\text{S-SECINTPROM} \\
\frac{}{\widehat{\mathbb{Z}}\#s \ (\uparrow n) \longrightarrow \uparrow(\widehat{\mathbb{Z}}\#s \ n)} \\
\\
\text{S-OMATCH} \\
\frac{\widehat{v}_1 \Leftarrow \widehat{\omega}_1 \quad \widehat{v}_2 \Leftarrow \widehat{\omega}_2}{\widehat{\text{match}} \ [l_b \langle \widehat{\omega}_1 \widehat{\omega}_2 \rangle \ \widehat{v}] \ \text{with } x \Rightarrow e_1 \mid x \Rightarrow e_2 \longrightarrow \widehat{\text{if}} \ [b] \ \text{then } \text{ite}(b, [\widehat{v}/x]e_1, [\widehat{v}_1/x]e_1) \\ \text{else } \text{ite}(b, [\widehat{v}_2/x]e_2, [\widehat{v}/x]e_2)} \\
\\
\text{S-OIF} \\
\frac{}{\widehat{\mathcal{E}}[\widehat{\text{if}} \ [b] \ \text{then } v_1 \ \text{else } v_2] \longrightarrow \widehat{\text{if}} \ [b] \ \text{then } \widehat{\mathcal{E}}[v_1] \ \text{else } \widehat{\mathcal{E}}[v_2]} \\
\\
\text{S-TAPEOIF} \\
\frac{}{\text{tape} \ (\widehat{\text{if}} \ [b] \ \text{then } v_1 \ \text{else } v_2) \longrightarrow \text{mux} \ [b] \ (\text{tape} \ v_1) \ (\text{tape} \ v_2)} \\
\\
\text{S-TAPEPROM} \\
\frac{}{\text{tape} \ (\uparrow v) \longrightarrow v} \\
\\
\text{EVALUATION CONTEXTS} \\
\mathcal{E} ::= \square \widehat{\times} \tau \mid \widehat{\omega} \widehat{\times} \square \mid \square \widehat{+} \tau \mid \widehat{\omega} \widehat{+} \square \\
\mid \text{mux } \square \ e \ e \mid \text{mux } v \ \square \ e \\
\mid \text{mux } v \ v \ \square \\
\mid \widehat{l}_b \langle \square \rangle \ e \mid \widehat{l}_b \langle \widehat{\omega} \rangle \ \square \\
\mid \widehat{\text{if}} \ \square \ \text{then } e \ \text{else } e \\
\mid \widehat{\text{if}} \ v \ \text{then } \square \ \text{else } e \\
\mid \widehat{\text{if}} \ v \ \text{then } v \ \text{else } \square \\
\mid \dots \\
\\
\text{LEAKY CONTEXTS} \\
\widehat{\mathcal{E}} ::= \square \ v \\
\mid \text{if } \square \ \text{then } e \ \text{else } e \\
\mid \text{match } \square \ \text{with } \overline{C} \ x \Rightarrow e \\
\mid \text{match } \square \ \text{with } (x_1, x_2) \Rightarrow e \\
\mid \widehat{\mathbb{B}}\#s \ \square \mid \widehat{\mathbb{Z}}\#s \ \square \\
\mid \square \oplus v \mid v \oplus \square
\end{array}$$

Figure 5.7. Selected small-step semantics rules of core TAYPE

To see how these rules work, consider a core TAYPE version of the example from [Section 5.1](#), which produces the following execution trace:

```

      tape ( $\widehat{\mathbb{Z}}\#s$  ( $\widehat{\text{if}}$  [true] then  $\uparrow 3$  else  $\uparrow 4$ ))
→ tape ( $\widehat{\text{if}}$  [true] then  $\widehat{\mathbb{Z}}\#s \uparrow 3$  else  $\widehat{\mathbb{Z}}\#s \uparrow 4$ )
→ tape ( $\widehat{\text{if}}$  [true] then  $\uparrow(\widehat{\mathbb{Z}}\#s \ 3)$  else  $\widehat{\mathbb{Z}}\#s \uparrow 4$ )
→ tape ( $\widehat{\text{if}}$  [true] then  $\uparrow[3]$  else  $\widehat{\mathbb{Z}}\#s \uparrow 4$ )
→* tape ( $\widehat{\text{if}}$  [True] then  $\uparrow[3]$  else  $\uparrow[4]$ )
→ mux [true] (tape  $\uparrow[3]$ ) (tape  $\uparrow[4]$ )
→* mux [true] [3] [4] → [3]

```

The S-OMATCH rule reduces a leaky case analysis of an oblivious sum to an $\widehat{\text{if}}$ using the discriminatee’s private tag. Similar to λ_{OADT^+} , the pattern variable x in the “correct” branch is of course instantiated with the injection payload, while the one in the “wrong” branch is instantiated with an arbitrary oblivious value of the right type. For example, if the discriminatee of a leaky case is $[\text{inl}\langle\widehat{\mathbb{Z}}\widehat{+}\widehat{\mathbb{Z}}\widehat{\times}\widehat{\mathbb{Z}}\rangle 1]$, the pattern variable in the second branch can be substituted by $[[0], [0]]$, $[[0], [1]]$, or any other oblivious pair of oblivious integers.

5.2.3 Type System

[Figure 5.8](#) shows an illustrative subset of the typing rules of core TAYPE (the full set of typing and kinding rules are in [Appendix A.2](#)). The judgment $\Gamma \vdash e :_l \tau$ types the expression e with type τ and leakage label l , under the typing context Γ (and an elided global typing context). Some typing rules refer to the kinding judgment $\Gamma \vdash \tau :: \kappa$, which also classifies the security of a type; oblivious types have the kind $*^0$, for example.

TAYPE features a security-type system [27] that ensures well-typed programs protect their private data. To do so, this type system enforces a few key policies. First, oblivious types can only be built from oblivious types, which is enforced by the kinding rules. Otherwise, an attacker can infer the private tag of an oblivious sum, such as $\mathbb{B}\widehat{+}\mathbb{Z}$, by observing the payload. Oblivious products have the same requirement, although this is mainly to aid in

$$\boxed{\Gamma \vdash e :_l \tau}$$

$\text{T-CONV} \quad \frac{\Gamma \vdash e :_l \tau \quad \tau \equiv \tau' \quad \Gamma \vdash \tau' :: *}{\Gamma \vdash e :_l \tau'}$	$\text{T-ABS} \quad \frac{x :_{l_1} \tau_1, \Gamma \vdash e :_{l_2} \tau_2 \quad \Gamma \vdash \tau_1 :: *}{\Gamma \vdash \lambda x :_{l_1} \tau_1 \Rightarrow e :_{l_2} \Pi x :_{l_1} \tau_1, \tau_2}$	$\text{T-APP} \quad \frac{\Gamma \vdash e_2 :_{l_2} \Pi x :_{l_1} \tau_1, \tau_2 \quad \Gamma \vdash e_1 :_{l_1} \tau_1}{\Gamma \vdash e_2 \ e_1 :_{l_2} [e_1/x] \tau_2}$
$\text{T-PAIR} \quad \frac{\Gamma \vdash e_1 :_l \tau_1 \quad \Gamma \vdash e_2 :_l \tau_2}{\Gamma \vdash (e_1, e_2) :_l \tau_1 \times \tau_2}$	$\text{T-PMATCHNODEP} \quad \frac{\Gamma \vdash e_0 :_{l_0} \tau_1 \times \tau_2 \quad l_0 \sqsubseteq l \quad x_1 :_{l_0} \tau_1, x_2 :_{l_0} \tau_2, \Gamma \vdash e :_l \tau}{\Gamma \vdash \text{match}_\tau e_0 \ \text{with} \ (x_1, x_2) \Rightarrow e :_l \tau}$	
$\text{T-IFNODEP} \quad \frac{\Gamma \vdash e_0 :_{l_0} \mathbb{B} \quad l_0 \sqsubseteq l \quad \Gamma \vdash e_1 :_l \tau \quad \Gamma \vdash e_2 :_l \tau}{\Gamma \vdash \text{if}_\tau e_0 \ \text{then} \ e_1 \ \text{else} \ e_2 :_l \tau}$	$\text{T-MUX} \quad \frac{\Gamma \vdash e_0 :_\perp \widehat{\mathbb{B}} \quad \Gamma \vdash \tau :: *^0 \quad \Gamma \vdash e_1 :_\perp \tau \quad \Gamma \vdash e_2 :_\perp \tau}{\Gamma \vdash \text{mux} e_0 \ e_1 \ e_2 :_\perp \tau}$	
$\text{T-OPAIR} \quad \frac{\Gamma \vdash e_1 :_\perp \tau_1 \quad \Gamma \vdash e_2 :_\perp \tau_2 \quad \Gamma \vdash \tau_1 :: *^0 \quad \Gamma \vdash \tau_2 :: *^0}{\Gamma \vdash [e_1, e_2] :_\perp \tau_1 \widehat{\times} \tau_2}$	$\text{T-SECINT} \quad \frac{\Gamma \vdash e :_l \mathbb{Z}}{\Gamma \vdash \widehat{\mathbb{Z}}\#s \ e :_l \widehat{\mathbb{Z}}}$	$\text{T-RETINT} \quad \frac{\Gamma \vdash e :_\perp \widehat{\mathbb{Z}}}{\Gamma \vdash \widehat{\mathbb{Z}}\#r \ e :_\top \widehat{\mathbb{Z}}}$
$\text{T-OIF} \quad \frac{\Gamma \vdash e_0 :_\perp \widehat{\mathbb{B}} \quad \Gamma \vdash e_1 :_\top \tau \quad \Gamma \vdash e_2 :_\top \tau}{\Gamma \vdash \widehat{\text{if}} e_0 \ \text{then} \ e_1 \ \text{else} \ e_2 :_\top \tau}$	$\text{T-PROMOTE} \quad \frac{\Gamma \vdash e :_\perp \tau}{\Gamma \vdash \uparrow e :_\top \tau}$	$\text{T-TAPE} \quad \frac{\Gamma \vdash e :_\top \tau \quad \Gamma \vdash \tau :: *^0}{\Gamma \vdash \text{tape} \ e :_\perp \tau}$

Figure 5.8. Selected core TAYPE typing rules

translation. Second, oblivious control flow constructs like `mux` can only be applied to oblivious terms, otherwise their public result could reveal information about their condition. As an example, `mux [b] 1 2` is ill-typed, because an attacker can learn the value of `b` by observing its result. This policy is enforced by the kinding assumptions of the form $\Gamma \vdash \tau :: *^0$ in T-MUX and T-OPAIR. Third, types are not allowed to depend on leaky terms. The type `if $\widehat{\mathbb{Z}}\#r [0] = 0$ then 1 else $\widehat{\mathbb{B}}$` is not valid, for example, since the leaks in the condition can not be repaired. Thus, we require that any terms appearing in types to be labeled as non-leaky (\perp). Fourth, the argument to `tape` must be oblivious (T-TAPE). This ensures that leaky terms will *eventually* reduce to an oblivious value or a `if` tree of oblivious values that can then be repaired by, e.g., S-TAPEPROM or S-TAPEOIF. Intuitively, the \perp label in

the conclusion of T-TAPE signifies that the taped expression can be treated as non-leaky by its surrounding computation, as all leaks have been “patched up”. Finally, all oblivious components in the typing rules have the \perp label. All the labels in T-MUX and T-OPAIR are \perp , and the oblivious condition of $\widehat{\text{if}}$ (T-OIF) is also safe, for example. While this requirement is not crucial for security, it simplifies the type system and aids in our translation to OIL. Note that we can always apply `tape` to leaky oblivious expressions to make them safe, so this design does not harm the expressivity of well-typed TAYPE programs.

In addition to the above policies inherited from $\lambda_{\text{OADT}\blacktriangleleft}$, TAYPE’s type system imposes three more requirements that help our translation. First, safe terms must be explicitly converted to leaky ones using \uparrow . Thus, T-CONV requires convertible expressions to have the same label. Second, we usually require subexpressions to have the same label: the two components in T-PAIR have the same label l in TAYPE, for example. T-IFNODEP similarly requires both branches to have the same label. Its condition, however, is permitted to have a lower label. A similar requirement is particularly important for the case analysis of products and ADTs: each branch needs to use its pattern variables in a manner that is at least as safe as the discriminatee. Third, we require all *possibly* leaky subexpressions to be labelled as leaky. The branches in T-OIF and the argument to T-TAPE have label \top , even though they can technically also be typed at \perp : applying these rules to an expression with a safe subterm requires explicit promotion. Note that programmers do not need to do these explicit label conversion in the surface language, as \uparrow is automatically inserted by the typing algorithm presented in [Section 5.2.5](#).

5.2.4 Type Safety and Obliviousness

Given a well-typed global context, core TAYPE enjoys standard progress and preservation properties. Its type system also provides a strong security guarantee similar to $\lambda_{\text{OADT}\blacktriangleleft}$: an adversary cannot infer any private information from a well-typed core TAYPE program, even when they can observe each of its execution steps.

Theorem 5.2.1 (Obliviousness). *If $e_1 \approx e_2$ and $\cdot \vdash e_1 :_{l_1} \tau_1$ and $\cdot \vdash e_2 :_{l_2} \tau_2$, then*

1. $e_1 \longrightarrow^n e'_1$ if and only if $e_2 \longrightarrow^n e'_2$ for some e'_2 .

2. if $e_1 \longrightarrow^n e'_1$ and $e_2 \longrightarrow^n e'_2$, then $e'_1 \approx e'_2$.

Here, $e \approx e'$ means the two expressions are *indistinguishable*, i.e., they only differ in their oblivious values, and $e \longrightarrow^n e'$ means e reduces to e' in exactly n steps. Intuitively, the obliviousness theorem says that a pair of well-typed core TAYPE programs that are indistinguishable produce traces that are pairwise indistinguishable.

We have formalized a version of core TAYPE in Coq, including proofs of soundness and obliviousness for the calculus, based on the mechanization from [Chapter 4](#). In contrast to that development, this calculus includes the new features of TAYPE: oblivious products, label promotion and explicit and uniform label checking.

5.2.5 Surface Language and Bidirectional Type Checker

The source language of our compiler is a more user friendly version of core TAYPE. This language allows type annotations to be omitted, and does not require label annotations or explicit promotion operations. Its syntax is effectively that of [Figure 5.6](#) with the gray annotations removed and with an additional type ascription ($e:\tau$). Our type checker elaborates programs in this surface language into fully annotated core TAYPE programs in ANF [51]. Our inference algorithm is not sophisticated: unlike other dependent type systems, it does not support unification, for example. Nevertheless, it is capable of checking all the case studies and benchmarks in our experiments ([Section 5.5](#)) without any type or label annotations, except for top-level definitions.

At a high level, like standard bidirectional type checkers, our type checker operates in an *inference mode* and a *checking mode*. In inference mode, the algorithm infers the type of an expression (bottom-up), while in checking mode, the algorithm checks the expression against an expected type by propagating information to subexpressions as deeply as possible (top-down). Our type checker always starts with checking mode, as all top-level definitions are annotated with their type. [Figure 5.9](#) shows a representative selection of our bidirectional type checking rules, using the inference judgment $\Gamma \vdash e \Longrightarrow_l \tau \triangleright \dot{e}$ and the checking judgment $\Gamma \vdash e \Longleftarrow_l \tau \triangleright \dot{e}$, both of which output a fully-elaborated expression \dot{e} in core TAYPE and in ANF. As explained previously, the inference judgment generates the type τ as an output, while

$$\begin{array}{c}
\boxed{\Gamma \vdash e \Longrightarrow_l \tau \triangleright \dot{e}} \quad \boxed{\Gamma \vdash e \Leftarrow_l \tau \triangleright \dot{e}} \\
\\
\text{TI-UNIT} \quad \frac{}{\Gamma \vdash () \Longrightarrow_{\perp} \mathbf{1} \triangleright ()} \quad \text{TI-ASC} \quad \frac{\Gamma \vdash e \Leftarrow_l \tau \triangleright \dot{e} \quad \Gamma \vdash \tau :: *}{\Gamma \vdash (e:\tau) \Longrightarrow_l \tau \triangleright \dot{e}} \quad \text{TC-INFER} \quad \frac{\Gamma \vdash e \Longrightarrow_l \tau' \triangleright \dot{e} \quad \tau \equiv \tau'}{\Gamma \vdash e \Leftarrow_l \tau \triangleright \dot{e}} \\
\\
\text{TC-ABS} \quad \frac{x :_{l_1} \tau_1, \Gamma \vdash e \Leftarrow_l \tau_2 \triangleright \dot{e} \quad \Gamma \vdash \tau_1 :: *}{\Gamma \vdash \lambda x \Rightarrow e \Leftarrow_l \prod x :_{l_1} \tau_1, \tau_2 \triangleright \lambda x :_{l_1} \tau_1 \Rightarrow \dot{e}} \\
\\
\text{TI-APP} \quad \frac{\Gamma \vdash e_2 \Longrightarrow_l \prod x :_{l_1} \tau_1, \tau_2 \triangleright \dot{e}_2 \quad \Gamma \vdash e_1 \Leftarrow_{l'_1} \tau_1 \triangleright \dot{e}_1 \quad \dot{e}_1 l'_1 \triangleright_{l_1} \dot{e}'_1}{\Gamma \vdash e_2 \ e_1 \Longrightarrow_l [e_1/x] \tau_2 \triangleright \text{let } x_2 :_{l_1} \tau_1, \tau_2 = \dot{e}_2 \text{ in } \text{let } x_1 :_{l_1} \tau_1 = \dot{e}'_1 \text{ in } x_2 \ x_1} \\
\\
\text{TI-PAIR} \quad \frac{\Gamma \vdash e_1 \Longrightarrow_{l_1} \tau_1 \triangleright \dot{e}_1 \quad \Gamma \vdash e_2 \Longrightarrow_{l_2} \tau_2 \triangleright \dot{e}_2 \quad l = l_1 \sqcup l_2 \quad \dot{e}_1 l_1 \triangleright_l \dot{e}'_1 \quad \dot{e}_2 l_2 \triangleright_l \dot{e}'_2}{\Gamma \vdash (e_1, e_2) \Longrightarrow_l \tau_1 \times \tau_2 \triangleright \text{let } x_1 :_{l_1} \tau_1 = \dot{e}'_1 \text{ in } \text{let } x_2 :_{l_2} \tau_2 = \dot{e}'_2 \text{ in } (x_1, x_2)} \\
\\
\text{TI-IF} \quad \frac{\Gamma \vdash e_0 \Leftarrow_{\perp} \mathbb{B} \triangleright \dot{e}_0 \quad \Gamma \vdash e_1 \Longrightarrow_{l_1} \tau_1 \triangleright \dot{e}_1 \quad \Gamma \vdash e_2 \Longrightarrow_{l_2} \tau_2 \triangleright \dot{e}_2 \quad l = l_1 \sqcup l_2 \quad \Gamma \vdash \text{if } e_0 \text{ then } [\cdot; \cdot \vdash \tau_1] \text{ else } [\cdot; \cdot \vdash \tau_2] \triangleright \dot{\tau}}{\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \Longrightarrow_l \dot{\tau} \triangleright \dots} \\
\\
\text{TI-MATCH} \quad \frac{\text{data } \Gamma = \overline{\mathbb{C}} \ \tau \in \Sigma \quad \Gamma \vdash e_0 \Longrightarrow_{\perp} \Gamma \triangleright \dot{e}_0 \quad \forall i. x :_{\perp} \tau_i, \Gamma \vdash e_i \Longrightarrow_{l_i} \tau'_i \triangleright \dot{e}_i \quad l = \sqcup l_i \quad \Gamma \vdash \text{match } e_0 \text{ with } \overline{\mathbb{C}} \ x \Rightarrow [\cdot; x :_{\perp} \tau \vdash \tau'] \triangleright \dot{\tau}}{\Gamma \vdash \text{match } e_0 \text{ with } \overline{\mathbb{C}} \ x \Rightarrow e \Longrightarrow_l \dot{\tau} \triangleright \dots}
\end{array}$$

Figure 5.9. Selected surface TAYPE bidirectional typing rules

$$\boxed{\Gamma \vdash \mathcal{E}[\overline{\Delta}; \mathcal{P} \vdash \tau] \triangleright \dot{\tau}}$$

$$\begin{array}{c}
\text{DI-NODEP} \\
\frac{\forall i. \tau_i \equiv \dot{\tau} \quad \Delta, \Gamma \vdash \dot{\tau} :: *}{\Gamma \vdash \mathcal{E}[\overline{\Delta}; \mathcal{P} \vdash \tau] \triangleright \dot{\tau}}
\end{array}
\qquad
\begin{array}{c}
\text{DI-DEP} \\
\frac{\forall i. \Delta, \mathcal{P}, \Gamma \vdash \tau_i :: *^0}{\Gamma \vdash \mathcal{E}[\overline{\Delta}; \mathcal{P} \vdash \tau] \triangleright \mathcal{E}[\overline{\tau}]}
\end{array}$$

$$\begin{array}{c}
\text{DI-PROD} \\
\frac{\Gamma \vdash \mathcal{E}[\overline{\Delta}; \mathcal{P} \vdash \tau_1] \triangleright \dot{\tau}_1 \quad \Gamma \vdash \mathcal{E}[\overline{\Delta}; \mathcal{P} \vdash \tau_2] \triangleright \dot{\tau}_2}{\Gamma \vdash \mathcal{E}[\overline{\Delta}; \mathcal{P} \vdash \tau_1 \times \tau_2] \triangleright \dot{\tau}_1 \times \dot{\tau}_2}
\end{array}$$

$$\begin{array}{c}
\text{DI-PI} \\
\frac{\Gamma \vdash \mathcal{E}[\overline{\Delta}; \mathcal{P} \vdash \tau_1] \triangleright \dot{\tau}_1 \quad \Gamma \vdash \mathcal{E}[\overline{x : \iota \tau_1, \Delta}; \mathcal{P} \vdash \tau_2] \triangleright \dot{\tau}_2}{\Gamma \vdash \mathcal{E}[\overline{\Delta}; \mathcal{P} \vdash \Pi x : \iota \tau_1, \tau_2] \triangleright \Pi x : \iota \dot{\tau}_1, \dot{\tau}_2}
\end{array}$$

Figure 5.10. Selected inference rules for dependent contexts

the checking judgment takes τ as an input. For simplicity, the rules presented in [Figure 5.9](#) always infer the labels, but our implementation also checks labels bidirectionally; we will discuss how labels are handled shortly. Our implementation also includes a bidirectional kind checker that generates fully-elaborated types in ANF.

As it is standard, we switch between inference and checking modes using TI-ASC and TC-INFER. To automatically promote expressions, several rules (e.g., TI-PAIR and TI-APP) use an auxiliary relation $e_{l_1} \triangleright_{l_2} \dot{e}$, which potentially inserts a promotion operation to e according to its label l_1 and the target label l_2 . The definition of this relation is straightforward: $e_{l_1} \triangleright_{l_2} \mathbf{let} \ x = e \ \mathbf{in} \ \uparrow x$ if $l_1 \sqsubset l_2$, $e_{l_1} \triangleright_{l_2} e$ if $l_1 = l_2$, and fail otherwise.

The main challenge to algorithmic type checking are dependent conditionals and ADT case analysis, specifically inferring their *implicit motives*. Since dependent types in TAYPE are oblivious types, they are more restricted than the ones in most dependent type systems. To see how, consider the expression: `if x then ($\lambda b : \widehat{\mathbb{B}} \Rightarrow \widehat{\mathbb{I}}$ if b then 1 else 0) else ($\lambda n : \widehat{\mathbb{Z}} \Rightarrow \widehat{\mathbb{Z}}$ #r n)`. Ignoring labels, the left branch of this conditional has type $\widehat{\mathbb{B}} \rightarrow \mathbb{Z}$, while the right one has type $\widehat{\mathbb{Z}} \rightarrow \mathbb{Z}$. In many dependent type systems, this expression can simply be typed with `if x then $\widehat{\mathbb{B}} \rightarrow \mathbb{Z}$ else $\widehat{\mathbb{Z}} \rightarrow \mathbb{Z}$` . However, this type is not well-kinded in TAYPE! The type-level computation, `if` in this case, is only defined over oblivious types, which the types in the branches are clearly not. The correct type of this expression has to be

$(\text{if } x \text{ then } \hat{\mathbb{B}} \text{ else } \hat{\mathbb{Z}}) \rightarrow \mathbb{Z}$. The same problem also occurs in dependent ADT pattern matching, and when the branches have product types. Our type checker is equipped with special inference and checking rules for handling these cases. Figure 5.10 shows the auxiliary relation for inferring the well-kinded types that may depend on discriminates, which is invoked in the bidirectional typing rules for dependent conditionals and pattern matching, e.g., TI-IF and TI-MATCH. The rules presented in Figure 5.10 are used in the inference mode, but our implementation also has similar rules for the checking mode. The judgment $\Gamma \vdash \mathcal{E}[\overline{\Delta}; \mathcal{P} \vdash \tau] \triangleright \hat{\tau}$ infers a type $\hat{\tau}$ for the dependent context \mathcal{E} , under the typing context Γ . The context \mathcal{E} has multiple holes corresponding to the branches of this dependent pattern matching. Each of these holes also has a description consisting of its branch's type τ , the pattern variables with their types \mathcal{P} , and a local typing context Δ for collecting arguments in a Π -type. For example, TI-IF calls this judgment using the context $\text{if } e_0 \text{ then } \square \text{ else } \square$ with two branches of types τ_1 and τ_2 , and empty pattern variable contexts. In the previous example, these two branch types are $\hat{\mathbb{B}} \rightarrow \mathbb{Z}$ and $\hat{\mathbb{Z}} \rightarrow \mathbb{Z}$ respectively. Starting with an empty local context Δ , DI-PI decomposes these branch types and infers each component recursively. The argument types $\hat{\mathbb{B}}$ and $\hat{\mathbb{Z}}$ are obviously kinded, so DI-DEP generates a dependent type that combines them using the same dependent context: $\text{if } e_0 \text{ then } \hat{\mathbb{B}} \text{ else } \hat{\mathbb{Z}}$. On the other hand, the return type is simply \mathbb{Z} by DI-NODEP, as both branches have the same return type (according to type equivalence \equiv). The side condition of DI-NODEP also ensures the inferred type does not refer to pattern variables, which are not in the scope of this type. The return type may refer to the arguments of a Π -type though, which are collected in the context Δ in the rule DI-PI. DI-PROD is similar to DI-PI, although it does not need to modify any contexts.

Our bidirectional type checker also infers leakage labels. In contrast to type annotations, label annotations are not required even for top-level function signatures:⁵ they are instead derived from a function *attribute*, which indicates the purpose of a function. A function can be marked as either *section*, *retraction*, or *safe* using the $\#[\text{attribute}]$ syntax, as shown in Figure 5.3. A function without an attribute, such as `elem`, implements the program logic in the conventional fragment of TAYPE. Such functions and their arguments are always labelled

⁵↑In fact, our surface syntax does not allow users to provide label annotations.

as leaky, since they have to accept retracted values to work with the recipe from [Figure 4.1](#). Any intermediate labels in the bodies of such functions can also be reliably inferred to be leaky, as these functions do not mention oblivious types or public views directly. As a result, a programmer can write the functionality as in a conventional functional language. On the other hand, a function annotated with `#[section]`, e.g., `list#s`, defines a section function. Its public view argument obviously has a safe label, while its data argument (e.g., `list`) has a leaky label. A section function itself is safe, signaling that all potential leaks have been patched. Conversely, a function, annotated with `#[retraction]`, labels its arguments as safe, but itself has leaky label. Lastly, functions annotated with `#[safe]`, e.g., `elem`, are secure functions. These constitute the API of a secure library, so their arguments and the functions themselves are assigned safe labels. Note that while the labels in a function’s signature are determined by its attribute, any intermediate labels in its body need to be inferred. Similar to types, labels are inferred bidirectionally. When an inferred label is checked against a label, the checker will insert a promotion if the expected label is more restrictive than the inferred label, and reject the program when the expected label is less restrictive than the inferred label, using the auxiliary relation $e_{l_1} \triangleright_{l_2} \dot{e}$.

5.3 Oil and Translation

This section describes the OADT intermediate language, OIL, and its translation from TAYPE. The main challenge is how to encode the features of TAYPE that OIL lacks, including dependent types, leaky operations (`if` and `tape`), and, most importantly, its tape semantics.

5.3.1 Syntax, Semantics and Type System

[Figure 5.11](#) shows the syntax of OIL. It is mostly a standard ML-style language with rank-1 polymorphism, extended with an *oblivious array* type and its operations. An oblivious array \mathcal{A} is essentially a “buffer” holding all the private data in a joint computation. The elements of this array are the oblivious representation (usually encrypted values) of members of some fixed finite field. To remain agnostic to the underlying cryptographic protocol, OIL does not place any restrictions on the oblivious representation or the finite field, so the array

$e ::=$	$() \mid b \mid n \mid x$ $e \oplus e \mid e \hat{\oplus} e$ $\lambda x \Rightarrow e$ $\text{let } x = e \text{ in } e$ $e e \mid C e$ $\text{if } e \text{ then } e \text{ else } e$ $\text{mux } e e e$ (e, e) $\text{match } e \text{ with } (x, x) \Rightarrow e$ $\text{match } e \text{ with } \overline{C} x \Rightarrow e$ $\hat{\mathbb{B}}\#s e \mid \hat{\mathbb{Z}}\#s e$ $\mathcal{A}(e) \mid e ++ e \mid e(e, e)$ \dots	EXPRESSIONS: literals and variables (oblivious) integer operations function abstraction let binding function and constructor applications conditional atomic conditional pair product elimination ADT elimination primitive sections oblivious array operations size (\mathbb{N}) operations omitted
$\tau ::=$	$\mathbb{1} \mid \mathbb{B} \mid \mathbb{Z}$ \mathcal{A} \mathbb{N} α T $\tau \times \tau$ $\tau \rightarrow \tau$	TYPES base types oblivious array size type type variable ADT variable product type function type
$D ::=$	$\text{data } T[\overline{\alpha}] = \overline{C} \tau$ $\text{fn } x[\overline{\alpha}] : \tau = e$	GLOBAL DEFINITIONS: algebraic data type definition (recursive) function definition

Figure 5.11. OIL source syntax

can hold the encryption of bits, or shared secrets of 64-bit integers, for example. Conceptually, each array element is simply an oblivious integer that encodes a piece of the private data, such as an oblivious integer, the tag of an oblivious injection, or an oblivious boolean. Programs create an array of size n using $\mathcal{A}(n)$, concatenate two arrays using $++$, and take a slice of n elements starting at offset m in array a via $a(m, n)$.

Like TAYPE, OIL includes oblivious operations, but these operations are restricted to take and produce oblivious arrays, as this is the *only* oblivious type in OIL. The section operations for base types, $\hat{\mathbb{B}}\#s$ and $\hat{\mathbb{Z}}\#s$, for example, return a singleton array containing the “encrypted” result.

Types and global definitions are also standard, but OIL also includes the *size type*, \mathbb{N} , for array offsets and lengths. OIL has a standard CBV semantics and type system. The semantics of array operations that use out-of-bound indices (e.g., slicing) is undefined: this should never happen if translated from a well-typed TAYPE program.

5.3.2 Translating from Taype to Oil

Our translation from TAYPE to OIL is syntax- and type-directed, and uses the leakage label to identify and repair potential leaks in the program. The translation assumes the source program is in *administrative normal form* (ANF), restricting the brown-colored expressions e in [Figure 5.6](#) to be variables. The algorithm roughly consists of three components: translating TAYPE types to OIL types ([Figure 5.12](#)), translating TAYPE expressions to OIL expressions ([Figure 5.16](#) and [Figure 5.17](#)), and translating TAYPE oblivious types to OIL expressions of the size type ([Figure 5.18](#)). The full set of these rules are included in [Appendix A.3](#).

Translating Types

[Figure 5.12](#) shows the translation of a TAYPE type τ to an OIL type, guided by a leakage label l . With the \perp label, public types are translated as they are or congruently, as expected. Oblivious types, in contrast, are always converted to an oblivious array in OIL. The rich typing information of an oblivious type is not thrown away however: as we shall see, this information is used to implement oblivious array operations. Dependent function types are translated to their nondependent counterpart, with the label on the parameter type dictating its translation.

The translation of types under the \top label is more involved. To understand why, recall that an expression with this label may contain a potentially leaky subexpression which should be repaired via `tape`. Thus, its OIL counterpart must be equipped with a similar mechanism capable of patching leaks. Our solution is to explicitly capture the insecure operations associated with a particular leaky type in its OIL representation, and to insert repairs for each kind of leak when translating a leaky expression. We call this first component a *leaky representation*. As an example, an integer expression can have three kinds of leaks:

$[[\tau]]_{\perp}$

$$\begin{aligned}
[[\mathbb{B}]]_{\perp} &= \mathbb{B} & [[\mathbb{Z}]]_{\perp} &= \mathbb{Z} & [[\mathbb{T}]]_{\perp} &= \mathbb{T} & [[\tau_1 \times \tau_2]]_{\perp} &= [[\tau_1]]_{\perp} \times [[\tau_2]]_{\perp} \\
[[\Pi_{\mathbf{x}:l\tau_1, \tau_2}]]_{\perp} &= [[\tau_1]]_{l \rightarrow} [[\tau_2]]_{\perp} & [[\mathbb{1}]]_{\perp} &= [[\hat{\mathbb{B}}]]_{\perp} = [[\hat{\mathbb{Z}}]]_{\perp} = [[\tau_1 \hat{\times} \tau_2]]_{\perp} = [[\tau_1 \hat{+} \tau_2]]_{\perp} = \mathcal{A} \\
[[\hat{\mathbb{T}} \ e]]_{\perp} &= [[\text{if } \dots]]_{\perp} = [[\text{let } \dots]]_{\perp} = [[\text{match } \dots]]_{\perp} = \mathcal{A}
\end{aligned}$$

$[[\tau]]_{\top}$

$$\begin{aligned}
[[\mathbb{B}]]_{\top} &= \tilde{\mathbb{B}} & [[\mathbb{Z}]]_{\top} &= \tilde{\mathbb{Z}} & [[\mathbb{T}]]_{\top} &= \tilde{\mathbb{T}} & [[\tau_1 \times \tau_2]]_{\top} &= [[\tau_1]]_{\top} \tilde{\times} [[\tau_2]]_{\top} \\
[[\Pi_{\mathbf{x}:l\tau_1, \tau_2}]]_{\top} &= [[\tau_1]]_{l \rightarrow} [[\tau_2]]_{\top} & [[\mathbb{1}]]_{\top} &= [[\hat{\mathbb{B}}]]_{\top} = [[\hat{\mathbb{Z}}]]_{\top} = [[\tau_1 \hat{\times} \tau_2]]_{\top} = [[\tau_1 \hat{+} \tau_2]]_{\top} = \tilde{\mathcal{A}} \\
[[\hat{\mathbb{T}} \ e]]_{\top} &= [[\text{if } \dots]]_{\top} = [[\text{let } \dots]]_{\top} = [[\text{match } \dots]]_{\top} = \tilde{\mathcal{A}}
\end{aligned}$$

Figure 5.12. Rules for translating core TAYPE types to OIL types

it could be a retraction of a secure integer $\hat{\mathbb{Z}}\#r$, it could be a leaky conditional $\hat{\text{if}}$, or it could be the promotion of a plaintext integer \uparrow . The corresponding leaky representation, $\tilde{\mathbb{Z}}$, is shown in [Figure 5.5](#), and contains a constructor for each of these cases. As every leaky type can leak information via \uparrow and $\hat{\text{if}}$, all leaky representations should be equipped with a reified form of these leaky expressions. Thus, every leaky representation (with its safe counterpart) forms a *leaky structure*, with operations `prom` and $\hat{\text{if}}$ for \uparrow and $\hat{\text{if}}$ respectively. From an implementation perspective, the leaky structure operations define a *typeclass*, so we call a particular `prom` and $\hat{\text{if}}$ *instances* of this typeclass. As one example, the constructors of $\tilde{\mathbb{Z}}$ trivially provide the necessary instances. As another example, [Figure 5.5](#) also shows $\tilde{\mathcal{A}}$, the leaky representation of an oblivious array; its leaky instance is similarly defined by the two constructors of this type. The leaky representation of function types, presented in [Figure 5.13](#), is slightly more complicated. Its two operations are essentially outsourced to the leaky instances of the codomain type β (`prom $_{\beta}$` and $\hat{\text{if}}_{\beta}$), taken as an extra argument that will be resolved at callsites.

In general, the $\hat{\text{if}}$ instances are usually constructors, as a leaky conditional needs to be irreducible to avoid leaking its private condition. The promotion instances are also constructors

```

fn prom→ [α β β̃] : (β → β̃) → (α → β) → (α → β̃) =
  λpromβ f x ⇒ promβ (f x)

fn if→ [α β̃] : (A → β̃ → β̃ → β̃) →
  A → (α → β̃) → (α → β̃) → (α → β̃) =
  λifβ b̂ f1 f2 x ⇒ ifβ b̂ (f1 x) (f2 x)

```

Figure 5.13. Leaky structures for function types

```

data list = Nil | Cons Z list
data list̃ = Nil̃ | Cons̃ Z̃ list̃ | promlist list | iflist A list̃ list̃

fn matchlist [γ̃] : (A → γ̃ → γ̃ → γ̃) → list̃ →
  γ̃ → (Z̃ → list̃ → γ̃) → γ̃ =
  λifγ x̃s f1 f2 ⇒
  match x̃s with
  | Nil ⇒ f1
  | Cons̃ x̃ x̃s' ⇒ f2 x̃ x̃s'
  | promlist xs ⇒
  match xs with
  | Nil ⇒ f1
  | Cons x xs' ⇒ f2 (promZ x) (promlist xs')
  | iflist b̂ x̃s1 x̃s2 ⇒
  ifγ b̂ (matchlist ifγ x̃s1 f1 f2) (matchlist ifγ x̃s2 f1 f2)

```

Figure 5.14. Leaky structures for `list`s

in our translation, although in general they need not be.⁶ Of course, our translation must also explain how to use leaky values, i.e., how to interpret the corresponding *elimination forms* of τ . To illustrate this, consider the leaky structure for `list` shown in [Figure 5.14](#).

The leaky representation of lists includes constructors for `Cons` and `Nil`, i.e., the introduction forms of `list`. Its leaky elimination form, $\widetilde{\text{match}}_{\text{list}}$, is straightforward: the `promlist` branch promotes the arguments of each constructor before applying the “alternative functions”, and

⁶↑Intuitively, these two instances are generated “for free”, though not in the algebraic sense. The only free leaky structure is the one for oblivious arrays $\widetilde{\mathcal{A}}$.

```

data T =  $\overline{C \llbracket \tau \rrbracket_{\perp}}$ 
data  $\tilde{T} = \overline{C \llbracket \tau \rrbracket_{\top}} \mid \text{prom}_{\top} T \mid \widehat{\text{if}}_{\top} \mathcal{A} \tilde{T} \tilde{T}$ 

fn  $\widetilde{\text{match}}_{\top} [\tilde{\gamma}] : (\mathcal{A} \rightarrow \tilde{\gamma} \rightarrow \tilde{\gamma} \rightarrow \tilde{\gamma}) \rightarrow \tilde{T} \rightarrow (\overline{\llbracket \tau \rrbracket_{\top}} \rightarrow \tilde{\gamma}) \rightarrow \tilde{\gamma} =$ 
   $\lambda \widehat{\text{if}}_{\gamma} \tilde{x} \tilde{f} \Rightarrow$ 
   $\text{match } \tilde{x} \text{ with}$ 
   $\mid \overline{C} x \Rightarrow f x$ 
   $\mid \text{prom}_{\top} x \Rightarrow \text{match } x \text{ with } \overline{C} x \Rightarrow f (\text{prom}(\tau) x)$ 
   $\mid \widehat{\text{if}}_{\top} \tilde{b} \tilde{x}_1 \tilde{x}_2 \Rightarrow \widehat{\text{if}}_{\gamma} \tilde{b} (\widetilde{\text{match}}_{\top} \widehat{\text{if}}_{\gamma} \tilde{x}_1 \tilde{f}) (\widetilde{\text{match}}_{\top} \widehat{\text{if}}_{\gamma} \tilde{x}_2 \tilde{f})$ 

```

Figure 5.15. Generating leaky ADT definitions

the $\widehat{\text{if}}_{\text{list}}$ branch essentially encodes the tape semantics rule S-OIF, specialized to the leaky context of `match` expressions, `match \square with $\overline{C} x \Rightarrow e$` .

A similar recipe is used to derive the leaky representation and its associated functions for other types: the introduction forms are encoded as constructors with the $\widehat{\text{if}}$ and `prom` instances, and the elimination forms capture the idea of distributing the corresponding leaky context into the $\widehat{\text{if}}$ branches and how \uparrow interacts with this context. While the leaky structures of builtin and arrow types are defined in the OIL prelude, the ones for user-defined ADTs are generated using the algorithm in Figure 5.15. This is how the leaky definition of `list` in Figure 5.14 was generated, for example. An ADT’s introduction forms are its constructors, so the leaky representation just renames them, with the constructor argument types translated with label \top . The $\widetilde{\text{match}}_{\top}$ function encodes the elimination form of ADTs, using a list of functions corresponding to branches of a case expression. The `prom $_{\top}$` branch relies on the instance resolution procedure `prom(\cdot)` to promote constructor arguments.

Translating Expressions

We now present our translation from TAYPE to OIL expressions. As with our translation of types, the translation of expressions is given as a judgment $\Gamma \vdash e \rightsquigarrow_l \dot{e}$, that is indexed by a leakage label l which guides the translation. Figure 5.16 illustrates how l drives the translation of standard constructs: if l identifies an expression as safe, it is simply translated congruently. On the other hand, if an expression is marked as leaky, the translation relies on the leaky

$$\boxed{\Gamma \vdash e \rightsquigarrow_l \dot{e}}$$

$$\begin{array}{c}
\text{TR-SECINT} \\
\hline
\Gamma \vdash \widehat{\mathbb{Z}\#\mathbf{s}} \ x \rightsquigarrow_l \begin{cases} \widehat{\mathbb{Z}\#\mathbf{s}} \ x & \text{if } l = \perp \\ \widetilde{\mathbb{Z}\#\mathbf{s}} \ x & \text{if } l = \top \end{cases} \\
\\
\text{TR-APP} \\
\hline
\Gamma \vdash \mathbf{x}_2 \ \mathbf{x}_1 \rightsquigarrow_l \mathbf{x}_2 \ \mathbf{x}_1 \\
\\
\text{TR-PAIR} \\
\hline
\Gamma \vdash (\mathbf{x}_1, \mathbf{x}_2) \rightsquigarrow_l \begin{cases} (\mathbf{x}_1, \mathbf{x}_2) & \text{if } l = \perp \\ \widetilde{\text{pair}} \ \mathbf{x}_1 \ \mathbf{x}_2 & \text{if } l = \top \end{cases} \\
\\
\text{TR-ABS} \\
\hline
\mathbf{x} :_{l_1} \tau_1, \Gamma \vdash e \rightsquigarrow_l \dot{e} \\
\hline
\Gamma \vdash \lambda \mathbf{x} :_{l_1} \tau_1 \Rightarrow e \rightsquigarrow_l \lambda \mathbf{x} \Rightarrow \dot{e} \\
\\
\text{TR-IF} \\
\hline
\mathbf{x}_0 :_{l_0} \mathbb{B} \in \Gamma \quad \Gamma \vdash e_1 \rightsquigarrow_l \dot{e}_1 \quad \Gamma \vdash e_2 \rightsquigarrow_l \dot{e}_2 \\
\hline
\Gamma \vdash \text{if}_\tau \ \mathbf{x}_0 \ \text{then } e_1 \ \text{else } e_2 \rightsquigarrow_l \begin{cases} \text{if } \mathbf{x}_0 \ \text{then } \dot{e}_1 \ \text{else } \dot{e}_2 & \text{if } l_0 = \perp \\ \widetilde{\text{if}} \ \widehat{\text{if}}(\tau) \ \mathbf{x}_0 \ \dot{e}_1 \ \dot{e}_2 & \text{if } l_0 = \top \end{cases}
\end{array}$$

Figure 5.16. Selected rules for translating core TAYPE standard expressions to OIL expressions

context of the expression to patch any leaks. This strategy can be seen in the TR-SECINT rule: using this rule to translate a leaky $\widehat{\mathbb{Z}\#\mathbf{s}} \ e$ expression delegates any repairs to $\widetilde{\mathbb{Z}\#\mathbf{s}}$. Translating lambda abstractions (TR-ABS) and applications (TR-APP) is straightforward. TR-PAIR shows why we require uniform labels in subexpressions: the components of a pair marked as leaky must also be leaky, as $\widetilde{\text{pair}}$ takes the leaky representations as arguments, similar to $\widetilde{\text{Cons}}$ from Figure 5.14. The translation of **if** (TR-IF) differs from the other rules in that the label of its discriminée dictates when its leaky counterpart is used, rather than the label of the whole expression. To see why, recall the typing rule T-IFNODEP from Figure 5.8: if the label of the discriminée is \top , the label of the whole expression must also be \top . On the other hand, we do allow the discriminée to be non-leaky, even if the whole expression is leaky. In this case, we simply use the standard **if** statement, as leaks can only occur in a subexpression. A similar strategy applies when translating **match**. The TR-IF rule illustrates why we annotate conditionals and case statements with their result type τ : this type is used to resolve the leaky if instances associated with τ via a call to the metafunction $\widehat{\text{if}}$.

Figure 5.17 presents some of the translation rules for leaky and oblivious constructs. This translation is more involved, as it needs to account for the switch to OIL's oblivious

tp

$$\boxed{\Gamma \vdash e \rightsquigarrow_l \dot{e}}$$

$$\begin{array}{c}
\text{TR-UNIT} \\
\hline
\Gamma \vdash () \rightsquigarrow_{\perp} \mathcal{A}(0)
\end{array}
\qquad
\begin{array}{c}
\text{TR-OPAIR} \\
\hline
\Gamma \vdash [x_1, x_2] \rightsquigarrow_{\perp} x_1 \# x_2
\end{array}$$

$$\begin{array}{c}
\text{TR-OINJ} \\
\hline
\Gamma \vdash \tau_1 \rightsquigarrow s_1 \quad \Gamma \vdash \tau_2 \rightsquigarrow s_2 \\
\Gamma \vdash \widehat{\iota}_b \langle \tau_1 \# \tau_2 \rangle x \rightsquigarrow_{\perp} \text{ite}(b, \widehat{\text{inl}}, \widehat{\text{inr}}) s_1 s_2 x
\end{array}
\qquad
\begin{array}{c}
\text{TR-RETINT} \\
\hline
\Gamma \vdash \widehat{\mathbb{Z}}\#r x \rightsquigarrow_{\top} r_{\mathbb{Z}} x
\end{array}$$

$$\begin{array}{c}
\text{TR-TAPE} \\
\hline
\Gamma \vdash \text{tape } x \rightsquigarrow_{\perp} \widetilde{\text{tape } x}
\end{array}
\qquad
\begin{array}{c}
\text{TR-PROMOTE} \\
\hline
x :_{\perp} \tau \in \Gamma \\
\Gamma \vdash \uparrow x \rightsquigarrow_{\top} \text{prom}(\tau) x
\end{array}$$

$$\begin{array}{c}
\text{TR-OIF} \\
\hline
x_1 :_{\top} \tau \in \Gamma \\
\Gamma \vdash \widehat{\text{if}} x_0 \text{ then } x_1 \text{ else } x_2 \rightsquigarrow_{\top} \widehat{\text{if}}(\tau) x_0 x_1 x_2
\end{array}$$

$$\begin{array}{c}
\text{TR-OPMATCH} \\
\hline
x_1 :_{\perp} \tau_1, x_2 :_{\perp} \tau_2, \Gamma \vdash e \rightsquigarrow_l \dot{e} \quad \Gamma \vdash \tau_1 \rightsquigarrow s_1 \quad \Gamma \vdash \tau_2 \rightsquigarrow s_2 \\
\Gamma \vdash \widehat{\text{match}} x_0 : \tau_1 \widehat{\times} \tau_2 \text{ with } [x_1, x_2] \Rightarrow e \rightsquigarrow_l \text{let } x_1 = x_0(0, s_1) \text{ in} \\
\text{let } x_2 = x_0(s_1, s_2) \text{ in} \\
\dot{e}
\end{array}$$

$$\begin{array}{c}
\text{TR-OMATCH} \\
\hline
x :_{\perp} \tau_1, \Gamma \vdash e_1 \rightsquigarrow_{\top} \dot{e}_1 \quad x :_{\perp} \tau_2, \Gamma \vdash e_2 \rightsquigarrow_{\top} \dot{e}_2 \quad \Gamma \vdash \tau_1 \rightsquigarrow s_1 \quad \Gamma \vdash \tau_2 \rightsquigarrow s_2 \\
\Gamma \vdash \widehat{\text{match}}_{\tau} x_0 : \tau_1 \widehat{\#} \tau_2 \text{ with } x \Rightarrow e_1 \mid x \Rightarrow e_2 \rightsquigarrow_{\top} \text{let tag} = x_0(0, 1) \text{ in} \\
\widehat{\text{if}}(\tau) \text{ tag } (\text{let } x = x_0(1, s_1) \text{ in } \dot{e}_1) \\
(\text{let } x = x_0(1, s_2) \text{ in } \dot{e}_2)
\end{array}$$

Figure 5.17. Selected rules for translating core TAYPE leaky and oblivious expressions to OIL expressions

$$\boxed{\Gamma \vdash \tau \rightsquigarrow \mathbf{s}}$$

$$\begin{array}{c}
\text{TR-UNITT} \\
\hline
\Gamma \vdash \mathbb{1} \rightsquigarrow 0
\end{array}
\qquad
\begin{array}{c}
\text{TR-OINT} \\
\hline
\Gamma \vdash \widehat{\mathbb{Z}} \rightsquigarrow 1
\end{array}
\qquad
\begin{array}{c}
\text{TR-OPROD} \\
\hline
\Gamma \vdash \tau_1 \rightsquigarrow \mathbf{s}_1 \quad \Gamma \vdash \tau_2 \rightsquigarrow \mathbf{s}_2 \\
\Gamma \vdash \tau_1 \widehat{\times} \tau_2 \rightsquigarrow \mathbf{s}_1 + \mathbf{s}_2
\end{array}$$

$$\begin{array}{c}
\text{TR-OSUM} \\
\hline
\Gamma \vdash \tau_1 \rightsquigarrow \mathbf{s}_1 \quad \Gamma \vdash \tau_2 \rightsquigarrow \mathbf{s}_2 \\
\Gamma \vdash \tau_1 \widehat{+} \tau_2 \rightsquigarrow 1 + \max \mathbf{s}_1 \ \mathbf{s}_2
\end{array}
\qquad
\begin{array}{c}
\text{TR-TAPP} \\
\hline
\Gamma \vdash \widehat{\mathbb{T}} \ x \rightsquigarrow \widehat{\mathbb{T}} \ x
\end{array}$$

$$\begin{array}{c}
\text{TR-TLET} \\
\hline
\Gamma \vdash e \rightsquigarrow_{\perp} \dot{e} \quad x :_{\perp} \tau_1, \Gamma \vdash \tau \rightsquigarrow \mathbf{s} \\
\Gamma \vdash \text{let } x :_{\perp} \tau_1 = e \text{ in } \tau \rightsquigarrow \text{let } x = \dot{e} \text{ in } \mathbf{s}
\end{array}$$

Figure 5.18. Selected rules for translating core TAYPE oblivious types to sizes in OIL

arrays. This is straightforward for simple data types: unit values are simply encoded as an empty array (TR-UNIT), while the translation of an oblivious pair simply concatenates the arrays produced by the translation of its two components (TR-OPAIR). Translating the destructor for oblivious pairs is more interesting (TR-OPMATCH), as it needs to extract each component from a flat array. To see how this is possible, observe that the “size” of an oblivious value is determined by its type, otherwise we risk leaking private information through this side-channel: thus, we can determine the location of each component of a pair based solely on their types. We do so via an auxiliary relation, $\Gamma \vdash \tau \rightsquigarrow \mathbf{s}$, given in Figure 5.18, which translate TAYPE types to OIL size expressions.

The translation of oblivious injections provide another example of how this relation is used. The TR-OINJ rule relies on the auxiliary function (and a similar right injection function) shown in Figure 5.19. This function takes as input the sizes of the left and right components and the injection payload, and produces an oblivious array containing the tag and payload, padding it out to the size of the larger component to avoid leaking information through its representation.

```

fn inl : N → N → A → A =
  λm n â ⇒
    let tag = B#s true in
    let payload =
      if n ≤ m then â
      else â ++ A(n-m)
    in tag ++ payload

```

Figure 5.19. Oblivious injection

$$\begin{array}{c}
\boxed{D \rightsquigarrow \dot{D}} \\
\text{TR-OADT} \quad \frac{x :_{\perp} \tau', \Gamma \vdash \tau \rightsquigarrow s}{\text{obliv } \widehat{T} (x:\tau') = \tau \rightsquigarrow \text{fn } \widehat{T} : [\tau']_{\perp} \rightarrow \mathbb{N} = \lambda x \Rightarrow s} \\
\text{TR-FUN} \quad \frac{\cdot \vdash e \rightsquigarrow_l \dot{e}}{\text{fn } x :_l \tau = e \rightsquigarrow \text{fn } x : [\tau]_l = \dot{e}}
\end{array}$$

Figure 5.20. Selected rules for translating core TAYPE definitions to OIL definitions

$$\begin{array}{c}
\boxed{\widehat{\text{if}}(\tau)} \\
\widehat{\text{if}}(\mathcal{A}) = \widehat{\text{if}}_{\mathcal{A}} \quad \widehat{\text{if}}(\mathbb{B}) = \widehat{\text{if}}_{\mathbb{B}} \quad \widehat{\text{if}}(\mathbb{Z}) = \widehat{\text{if}}_{\mathbb{Z}} \quad \widehat{\text{if}}(\mathbb{T}) = \widehat{\text{if}}_{\mathbb{T}} \quad \widehat{\text{if}}(\alpha \times \beta) = \widehat{\text{if}}_{\times} \\
\widehat{\text{if}}(\alpha \rightarrow \beta) = \widehat{\text{if}}_{\rightarrow} \widehat{\text{if}}(\beta) \\
\boxed{\text{prom}(\tau)} \\
\text{prom}(\mathcal{A}) = \text{prom}_{\mathcal{A}} \quad \text{prom}(\mathbb{B}) = \text{prom}_{\mathbb{B}} \quad \text{prom}(\mathbb{Z}) = \text{prom}_{\mathbb{Z}} \quad \text{prom}(\mathbb{T}) = \text{prom}_{\mathbb{T}} \\
\text{prom}(\alpha \times \beta) = \text{prom}_{\times} \quad \text{prom}(\alpha \rightarrow \beta) = \text{prom}_{\rightarrow} \text{prom}(\beta)
\end{array}$$

Figure 5.21. Resolving leaky instances

For example, the translation of $\widehat{\text{inl}} \langle \widehat{\mathbb{Z}} + \widehat{\mathbb{Z}} \times \widehat{\mathbb{Z}} \rangle$ [2] computes to oblivious array [1,2,0], while $\widehat{\text{inr}} \langle \widehat{\mathbb{Z}} + \widehat{\mathbb{Z}} \times \widehat{\mathbb{Z}} \rangle$ [[3],[4]] computes to [0,3,4].

The remaining rules in the figure adopt similar strategies; relying on a combination of leaky structures to patch up leaky constructs and the size relation to bridge the gap between oblivious types in TAYPE and oblivious arrays in OIL: the rule for **tape** (TR-TAPE), for example, simply delegates the repair to $\widetilde{\text{tape}}$, from Figure 5.5, which encodes the tape rules S-TAPEOIF and S-TAPEPROM. Similarly, the TR-OMATCH uses the size to extract the tag of a sum type, before processing both branches with a leaky $\widehat{\text{if}}$ expression which eventually discards the unused branch. Note that the payload x extracted from the injection in the “wrong” branch always uses the right size for that type: when matching on the previous example, x will be [2,0] in the second branch.

The translation of top-level definitions is straightforward; Figure 5.20 provides the rules for oblivious ADTs and functions of this translation. The elided translation of ADTs simply relies on the generation algorithm from Figure 5.15. The resolution procedures, $\widehat{\text{if}}$ and **prom**, are

also straightforward, shown in [Figure 5.21](#). Most instances are resolved to the corresponding definitions in the OIL prelude, or to the generated constructors for ADTs. Resolving function type instances also requires resolving the instances of a function’s codomain, as suggested by [Figure 5.13](#). Note that `prom` and `if` in [Figure 5.21](#) take an OIL type, while TAYPE types are given to the calls to these meta-functions in the translation rules; we implicitly apply $[\cdot]_{\perp}$ ([Figure 5.12](#)) to convert TAYPE types to OIL types for the leaky instance resolution.

Our translation algorithm is guaranteed to terminate, even when the source program does not. The reason can be seen in our translation rules, as every (mutually) recursive call to the translation judgment is applied to a structurally smaller core TAYPE sub-expression. The algorithm enjoys a stronger totality property: translation of a well-typed core TAYPE program never fails, i.e., a well-typed program satisfies all the side-conditions of the translation rules:

Theorem 5.3.1 (Totality of Translation). *If $\Gamma \vdash e :_l \tau$ and e is in ANF, then there exists an OIL expression \hat{e} such that $\Gamma \vdash e \rightsquigarrow_l \hat{e}$.*

The proof (and an analogous theorem for the type-to-size translation) is given in [Appendix A.4](#).

5.3.3 Translation for Conceal and Reveal Phases

Secure multiparty computations typically consist of three phases: a *conceal phase*, an *oblivious computation phase*, and a *reveal phase*. In the conceal phase, private data owners “encrypt” and share their data before the core computation takes place, while the oblivious output is revealed to all (or the privileged) parties in the reveal phase. In order to provide a complete solution, we also produce secure implementations of these two phases. Thankfully, section and retraction functions provide templates for concealing and revealing private data. Our toolchain thus translates section and retraction functions to special versions that implement each phase.

Translating the retraction functions needed for the reveal phase is simple: we simply make all the leaky operations “leak” by renaming all leaky operations in a retraction function, and link them to the revealing versions. For example, $\hat{Z}\#r$ is renamed to `Reveal. $\hat{Z}\#r$` (in an

OCaml module named `Reveal`). The retraction functions themselves are also renamed so client programs can use them to reveal the results of the computation.

Translating the section functions needed for the conceal phase is more involved. The main problem is that, unlike the core computation, only the private data owner can run the conceal function, as other parties do not have the data. But many MPC protocols, e.g., ones based on secret-sharing, require all parties to help create the encryption of the private information, so this has to be done synchronously. In our setting, since private data is encoded as oblivious arrays, all parties have to encrypt the elements of the same index at the same time. For example, during the conceal phase, if Alice is encrypting the third element of the array, Bob needs to do the same. However, this is not trivial to enforce: how does Bob know which element Alice is currently dealing with, when he does not have the data? Naturally, Bob may only construct the oblivious array from left to right, which means that Alice needs to do the same. Our implementation conceals private data in two steps. First, the private data owners (e.g., Alice) run the section functions *locally* using a plaintext backend for cryptographic operations (Section 5.4), resulting in an “oblivious” array whose elements are not encrypted, i.e., a plaintext array. Then, all parties (jointly) encrypt the input array from left to right, using the underlying cryptographic protocol, and obtain the actual oblivious array needed for the oblivious computation phase.

5.4 Implementation

We have implemented the above approach as a compiler that takes as input a `TAYPE` program containing the functions to be computed (as well as any auxiliary functions), the public views, and section and retraction functions. After type checking these pieces, our toolchain produces OCaml implementations of the conceal and reveal phases, as well as an OCaml implementation of the multiparty computation, all of which are specialized to the desired public view. The output programs are clients of a module that provides an implementation of OIL’s oblivious arrays and oblivious operations. Linking the generated programs with an implementation of this interface, or *driver*, produces a library that a programmer can use to build a secure application: they simply gather the private data,

“encrypt” the data using the generated conceal functions, call the multiparty functionality from the library, and finally reveal the result using the generated reveal functions. As the calculator case study in the next section demonstrates, programmers can also implement multi-round computation by chaining together calls to this library.

Our current implementation features two drivers: a plaintext driver and a cryptography-backed driver. The plaintext driver computes its results in the clear, and is intended for testing purposes and for establishing a performance baseline without any cryptographic overhead. This driver is also necessary for generating section functions, as explained in the previous section. The cryptographic driver uses the popular open-source EMP toolkit [52] to implement secure computations. This library is based on Yao’s Garbled Circuit [1] for semi-honest 2-party MPC. Integrating a new backend into our framework is conceptually simple: the driver just needs to implement an interface consisting of oblivious integer encryption, decryption and its arithmetic. Our EMP toolkit backend consists of boilerplate code for FFI (foreign function interface), for example. Other aspects of the driver, such as array operations, are independent of the cryptographic backends, and can thus be shared among all drivers.

5.4.1 Optimizations

We have implemented two optimizations to improve the performance of generated code.

First, the unoptimized programs produced naively by the strategy in Figure 5.16 and Figure 5.17 can suffer from exponential blowup when leaks are “taped” too late. To see the issue, consider a recursive function `f` from `list` to \mathbb{Z} , whose recursive calls on the tail of the input list appear in both branches of a conditional, either directly or indirectly via variables: e.g., `if x then f xs' else 1 + f xs'`. If `x` is a leaky boolean of the form `if [b] then true else false`, the conditional will be distributed into both branches, resulting in `if [b] then f xs' else 1 + f xs'`. But `f xs'` also produces a similar `if`, in both branches! For example, one more unrolling of the recursive calls may result in:

```
if [b]
then (if [b'] then f xs'' else 1 + f xs'')
else (if [b'] then 1 + f xs'' else 2 + f xs'')
```

Consequently, compiling f results in a $\widehat{\text{if}}$ tree that is exponential in the length of the list, causing exponential blowup in both memory and running time.

To ameliorate this problem, we have implemented an optimization that `tapes` leaky expressions earlier (called “early-tape” optimization), by wrapping the whole body of this recursive function with a composition of section and retraction: $\widehat{\mathbb{Z}}\#r$ (`tape` ($\widehat{\mathbb{Z}}\#s$ (f x s))). This transformation does not change the semantics of f , and the combination of `tape` and $\widehat{\mathbb{Z}}\#s$ forces the $\widehat{\text{if}}$ of each recursive call to be securely reduced to a single oblivious integer to avoid building up the $\widehat{\text{if}}$ tree. This simple strategy is quite effective for a class of programs which return combinations of primitive types and tuples.

Another deficiency of a naive implementation is that the sizes of oblivious data may be repeatedly computed. For example, the retraction function of the oblivious list with public view k requires computing $\widehat{\text{list}}$ k as a size, but its recursive calls on $k-1$ also requires computing $\widehat{\text{list}}$ ($k-1$), while this computation has already been done in the previous iteration. This redundant computation attaches a quadratic term to the complexity of the retraction function. As a result, linear scan of an oblivious list becomes quadratic.

To remove the redundant size computation, our compiler implements a classic technique called *tupling* [53, 54], which merges some recursive functions into a tuple to avoid multiple traversals over the same data. Consider the previous example: the oblivious list retraction function $\widehat{\text{list}}\#r$ has type $\mathbb{Z} \rightarrow \mathcal{A} \rightarrow \widetilde{\text{list}}$ in OIL, while the oblivious list function $\widehat{\text{list}}$, i.e., the size function in OIL, has type $\mathbb{Z} \rightarrow \mathbb{N}$. We create a new recursive function $\widehat{\text{list}}\#r_tupled$ of type $\mathbb{Z} \rightarrow \mathbb{N} \times (\mathcal{A} \rightarrow \widetilde{\text{list}})$, whose definition should be equivalent to the tuple of those two functions:

$$\widehat{\text{list}}\#r_tupled\ k = (\widehat{\text{list}}\ k, \widehat{\text{list}}\#r\ k)$$

The actual recursive definition of this tupled function is obtained by a transformation based on the unfold-simplify-fold rules [55, 56], which eventually replaces the calls to $\widehat{\text{list}}$ and $\widehat{\text{list}}\#r$ with recursive calls to this function itself. To see how this avoids repetitive size computation, observe that each recursive call to this tupled function on $k-1$ returns not only the result of $\widehat{\text{list}}\#r$ ($k-1$) but also $\widehat{\text{list}}$ ($k-1$), meaning that we can calculate $\widehat{\text{list}}$ k from that incrementally.

5.5 Experiments

Our experiments consist of a set of case studies that showcase the applicability of our approach, summarized in [Figure 5.22](#), and a set of microbenchmarks that examine the empirical benefits of being able to trade off security for performance.

5.5.1 Case Study: Medical Records

Our first collection of case studies are inspired by problems in the healthcare setting, where legal and privacy concerns keep parties from freely sharing their data. These benchmarks use a variety of data structures: patient data is represented as a record with fields for a patient’s ID, age, height and weight, a database is encoded as a list of patient records, and a classifier is implemented as a decision tree over health data. The oblivious types for these data structures admit interesting public views: a particular health record may choose to keep either its ID, or medical data (height and weight) secret, as in [Figure 3.5](#); a database adopts the privacy settings (i.e., revealing either ID or medical data) of its individual records, and a decision tree obscures the threshold that a feature is compared against at a given node, but not the overall structure. Using these representations, we have implemented a number of secure computations: biometric matching (minimum euclidean distance) between a single record and a database, calculating the percentage of healthy members of an age group according to the Body Mass Index (BMI), calculating statistics such as mean and variance over multiple private databases, and classifying a patient record using a private decision tree. [Figure 5.22](#) includes each of these programs. Notably, the computation in each of these benchmarks was written without a privacy policy in mind; instead, our compiler took care of enforcing each policy, as encoded by an oblivious type and section and retraction functions.

5.5.2 Case Study: Dating Application

Consider a functionality of matching potential soulmates. Each party owns their private profile with personal information, such as gender, income and education. They also have a private preference for their partner, encoded as predicates over profiles of *both* parties.

Computation	Alice's input	Bob's input	Description
is-taller	record	record	a variant of the millionaire problem, comparing height
is-obese-by-id	database	ID	whether the record of a given ID is obese (according to BMI)
healthy-rate-by-age	database	a range of ages	the percentage of healthy members of an age range (according to BMI)
min-euclidean-distance	database	record	the minimum euclidean distance between a database and a given record
database-analytics	database	database	calculate the mean and variance of the ages over these databases
mean-squared-error	database	database of BMIs	the mean squared error between the estimated BMIs (from Bob) and the actual BMIs (from Alice); the two databases may not contain the same records, and are matched with IDs (similar to joining tables)
decision-tree	decision tree	record	classify a medical record via a private decision tree
dating	profile and predicate	profile and predicate	match the dating preferences of two parties; each party provides a private profile and a preference encoded as a predicate over profiles
secure-calculator	expression and assignment	expression and assignment	a 2-round arithmetic expression evaluation; each party provides a private arithmetic expression and some private variable assignment depending on the round
voting	tabulated votes	tabulated votes	return the candidate with the most votes
k-means	list of vectors	list of vectors	partition vectors using the k-means clustering algorithm

Figure 5.22. Summary of programs used in our case studies

These predicates are expressions with boolean connectives, integer arithmetic and numeric comparisons. For example, one user may stipulate that the sum of both parties' income exceeds a particular amount. The peer matching function takes these private profiles and predicates, and returns a boolean indicating whether they are a good match, by evaluating the predicate expression on the profiles. The private predicates have many potential policies. As predicates are essentially ASTs, participating parties may agree on disclosing only the depth of the predicates, or revealing the AST nodes but not the operands, or even revealing only the boolean connectives but keep the integer expressions secret. In TAYPE, the peer matching function, its auxiliary functions, and the data types they depend on can be implemented in the conventional way, without knowing the policies. The private matching function can be obtained by composing with the desired policy. Updating policy or updating the matching algorithm can be done independently.

5.5.3 Case Study: Secure Calculator

To showcase our support for computations involving richly structured data, we have implemented a secure interpreter for a simple arithmetic expression language. In this case study, each user provides a private expression and an assignment to some variables. The result is securely computed by evaluating the first party's expression using the second party's assignment; the result of this expression is then used to evaluate the second party's expression, along with

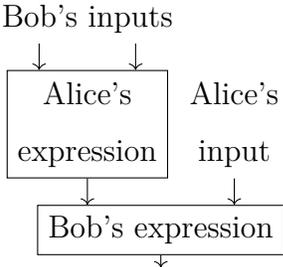


Figure 5.23. Workflow of the secure calculator

the first party's private value, as shown in [Figure 5.23](#). Not only does this case study use a rich data structure for expression, it also shows that we can readily compose the generated library functions to implement a more complex workflow, such as a multi-round computation.

5.5.4 Discussion

As mentioned in [Chapter 1](#), in existing frameworks, it is the program’s responsibility to enforce the privacy policy. In contrast, our medical records, dating application and secure calculator case studies demonstrate that, in our framework secure functionality can be written in a conventional functional language, agnostic to a particular privacy policy. On the other hand, implementing oblivious types, section and retraction functions is analogous to other common programming tasks: an oblivious type is essentially a different representation of the underlying data type, while section and retraction functions are effectively conversion functions between oblivious and public data types. Importantly, our abstraction allows programmers to write all these “boilerplate” functions once and for all, regardless of a particular target computation.

5.5.5 Microbenchmarks

To evaluate the performance of our compiler,⁷ we have built a number of microbenchmarks that showcase the performance tradeoffs between privacy and performance. Our first microbenchmark is a standard classification scenario, where one party wants to classify their private data using a decision tree belonging to another party [[57–59](#)]. The data being classified is given as a tuple with eight private integers as features. This experiment considers 4 public views for the decision tree: maximum height, the spine, spine including the feature index of each node, and the whole tree. Note that this last view is not unrealistic: in outsourced secure computation, such as FHE [[7](#)], the decision tree owner may perform all computation, independent of the other party. In this scenario, the whole tree can be revealed because the computing party owns it, but the computation should not reveal any information about the other party’s data. The definition of the decision tree is shown in [Figure 5.24a](#), together with an “oblivious” version that simply reveals the whole tree. [Figure 5.24](#) also includes views for three other policies. The section and retraction functions for each view are analogous to those in [Figure 5.3](#). For each public view, we test on a small tree of depth 1, and other

⁷↑ All results are averaged across 10 runs, on an M1 MacBook Pro with 16 GB memory. All parties run on the same host with local network communication.

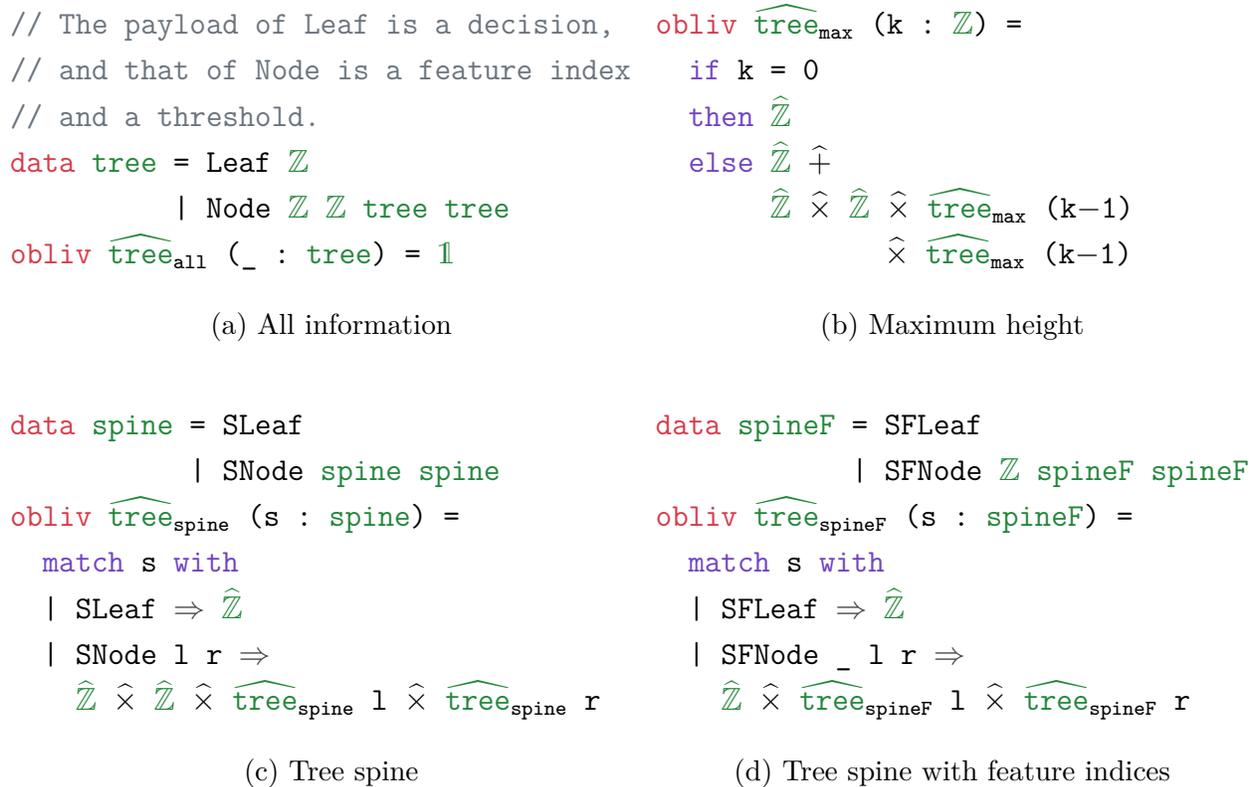


Figure 5.24. Definitions of oblivious decision trees with different public views

trees of depth 16. A *full tree* has exponentially many nodes, while an *eighth sparse tree* has roughly 1/8 of the nodes in a full tree, and a *very sparse tree* has only 16 nodes.

Figure 5.25 reports the performance impact of each view on the total run time. The results are as expected: revealing the whole tree results in the best performance, while sharing only the maximum height is quite slow. In the case of maximum height, the number of nodes in the actual decision tree does not affect the performance, as the structure of the tree is kept secret. Knowing both the spine and the feature index of each node improves performance, compared to knowing

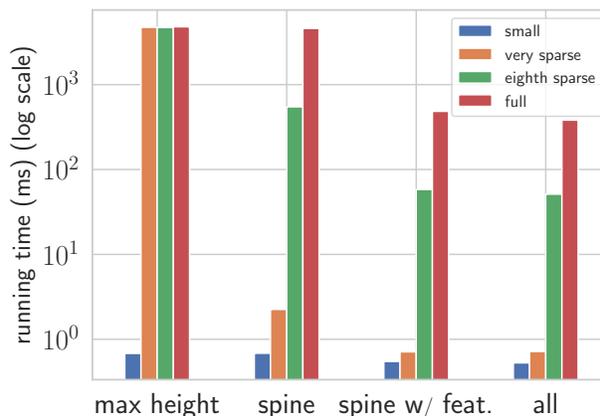
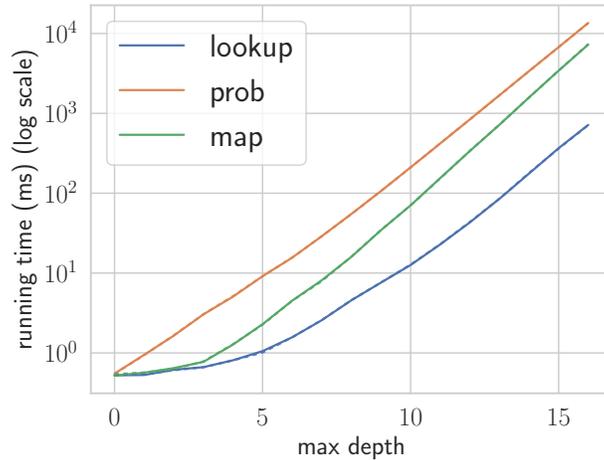
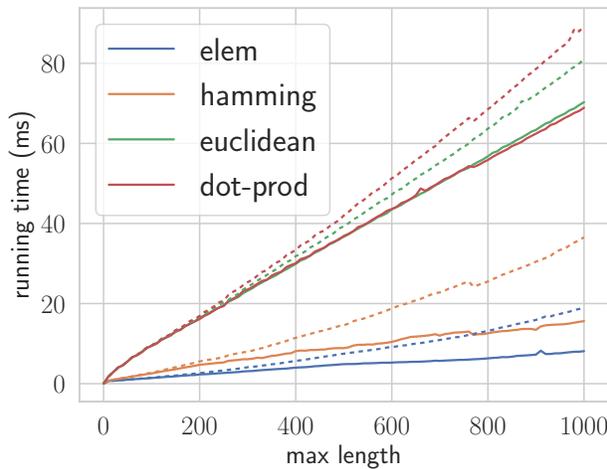


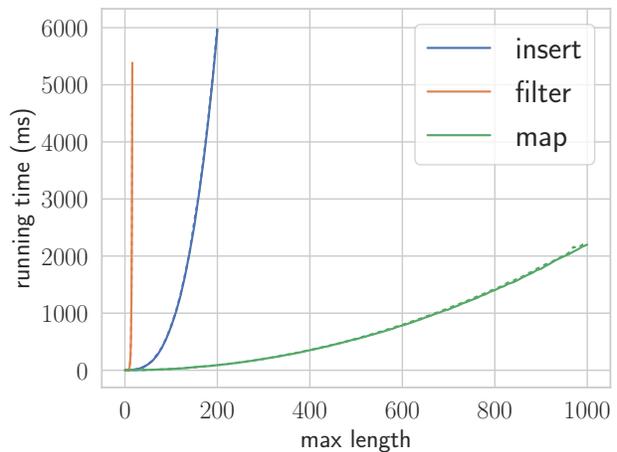
Figure 5.25. Decision tree



(a) Tree example



(b) List examples (primitive)



(c) List examples (complex)

Figure 5.26. Microbenchmarks

only the spine, as the computation does not need to obscure which feature is used at each decision point. When the underlying decision tree is relatively full, leaking more information about its structure does not improve performance, as the program does not need to perform wasted computation to ensure a constant time algorithm. Indeed, the fuller the private tree is, the less is gained by a more permissive public view. Of course, the owner of the tree must ultimately decide if they are willing to reveal how the tree is close to this worst case scenario. Again, the decision algorithm is agnostic of the actual public views, allowing for swapping privacy policies without any changes to the program logic.

We also evaluate the performance of a set of standard operations on trees, using its maximum height as the public view. These benchmarks consist of a membership test, computing the probability of an event given a probability tree diagram, and a map function that adds a private integer to each node in a tree. [Figure 5.26a](#) presents the performance results for these benchmarks. Despite the inherently expensive cryptography required by the conservative public view, all the benchmarks finish in under 15 seconds.

Finally, we have implemented a similar set of microbenchmarks for oblivious lists, using the length of the list as the public view. We subdivide these benchmarks into those that return a primitive value (i.e., an integer or boolean), and those that return an oblivious list. Even though the returned lists often reveal too much about the private input, they could be useful as an intermediate result of a bigger computation or as an input to the next round of computation. [Figure 5.26b](#) presents the performance for the first category, which includes a membership check, computing of the hamming and euclidean distances between two lists, and taking the dot product of two lists. All of these examples are amenable to the optimization mentioned in the previous section, resulting in reasonable running times. We use a dotted line for results without our tupling optimization, and a solid line for when the optimization is enabled. With tupling is used, their performance is linear in the size of the input list (as it is in the insecure setting). The second category includes insertion into a sorted list, and two higher-order examples: mapping a function that adds a private integer to all the elements of a list, and a filter function that drops all the elements greater than a private integer. Since these examples do not return primitive values, the early-tape optimization does not apply, resulting in slower performance, as [Figure 5.26c](#) shows. The tupling optimization does not have much impact, as its gains are overshadowed by the complexity of having to delay repairs to the leaky result values.

5.6 Conclusion

Secure multiparty computation enables different parties to compute functions over private data without leaking extra information, but writing these applications remains challenging. Existing high-level MPC languages require programs to explicitly enforce privacy policies,

making it difficult to update policies and to explore tradeoffs between privacy guarantees and performance. This chapter presented TAYPE, a language for secure multiparty applications that decouples these concerns. Our experiments feature a diverse set of benchmarks that were written without security policies in mind, and a wide range of security policies that went beyond whether a particular field is “secret or not”. Our results demonstrate the performance benefits that can result from being able to easily trade off privacy for performance.

6. TAYPSI: STATIC ENFORCEMENT OF POLICIES

TAYPE (Chapter 5) decouples privacy policies from programmatic concerns, allowing users to write applications over structured data that are agnostic to any particular privacy policy. To do so, TAYPE implements a novel form of the *tape semantics* (Chapter 4). This semantics allows insecure operations whose evaluation *could* violate a policy to appear in a program, as long as the results of these operations are *eventually* protected. Under tape semantics, such operations are lazily deferred until it is safe to execute them, effectively *dynamically* “repairing” potential leaks at runtime. Using TAYPE, programmers can thus build a privacy-preserving version of a standard functional program by composing it with a policy, specified as a *dependent type* equipped with security labels, relying on tape semantics to enforce the policy during execution. Unfortunately, while this enforcement strategy disentangles privacy concerns from program logic, it also introduces considerable overhead for applications that construct or manipulate structured data with complex privacy requirements. Thus, this strategy does not scale to the sorts of complex applications that could greatly benefit from this separation of concerns.

This chapter presents TAYPSI, a policy-agnostic language for writing MPC applications that eliminates this overhead by instead transforming a non-secure function into a version that *statically* enforces a user-provided privacy policy. TAYPSI extends TAYPE with a form of dependent sums, which we call Ψ -types, that package together the public and private components of an algebraic data type (ADT). Each Ψ -type is equipped with a set of Ψ -structures which play an important role in our translation, enabling it to, e.g., efficiently combine subcomputations that produce ADTs with different privacy policies. Our experimental evaluation demonstrates that this strategy yields considerable performance improvements over the enforcement strategy used by TAYPE, yielding exponential improvements on the most complex benchmarks in our evaluation suite.

To summarize, the contributions of this chapter are as follows:

- We present TAYPSI, a version of TAYPE extended with Ψ -types, a form of dependent sums that enables modular translation of non-secure programs into efficient, secure versions. This language is equipped with a security type system that offers the same

guarantees as TAYPE: after jointly computing a well-typed function, neither party can learn more about the other’s private data than what can be gleaned from their own data and the output of the function.

- We develop an algorithm that combines a program written in the standard fragment of TAYPSI with a privacy policy to produce a secure version that statically enforces the desired policy. We prove that this algorithm generates well-typed (and hence secure) target programs that are guaranteed to preserve the semantics of the source programs.
- We evaluate our approach on a range of case studies and microbenchmarks. Our experimental results demonstrate exponential performance improvements over the previous state-of-the-art (TAYPE) on several complicated benchmarks, while simultaneously showing no performance regression on the remaining benchmarks.

An artifact containing the Coq mechanization of the core calculus $\lambda_{\text{OADT}\Psi}$, the implementation of TAYPSI, its source code, and the source for all the benchmarks in our experiments with instructions is publicly available [60].

6.1 Overview

Before presenting the full details of our approach, we begin with an overview of TAYPSI’s strategy for building privacy-preserving applications. Consider the simple `filter` function in Figure 6.1, which drops all the elements in a list above a certain bound.⁸ Suppose Alice owns some integers, and wants to know which of those integers are less than some threshold integer belonging to Bob, but neither party wants to share their data with the other. Using oblivious algebraic data types (OADTs) and the oblivious language from

```

data list = Nil | Cons ℤ list
fn filter : list → ℤ → list =
  λxs y ⇒
    match xs with
    | Nil ⇒ Nil
    | Cons x xs' ⇒
      if x ≤ y
      then Cons x (filter xs' y)
      else filter xs' y

```

Figure 6.1. Filtering a list

⁸↑TAYPSI supports higher-order functions, but our overview will use this specialized version for presentation purposes.

Chapter 3, we can encode a secure version of this `filter` function, allowing Alice and Bob to encrypt their data and then jointly compute `filter`, without leaking information about the encrypted data beyond what they can infer from the final disclosed output.

The particular policy (i.e., the chosen public view) that a secure application enforces can greatly impact the performance of that application, since the control flow of an application cannot depend on private data. In the case of our example, this means that the number of recursive calls to `filter` depends on the public information Alice is willing to share. If Alice only wants to share the maximum length of her list, for example, its encrypted version must be padded with dummy encrypted values, and a secure version of `filter` must recurse over these dummy elements, in order to avoid leaking information to Bob through its control flow. On the other hand, if Alice does not mind sharing the exact number of integers she owns, the joint computation will not have to go over these values, allowing a secure version of `filter` to be computed more efficiently.

TAYPSI allows Alice and Bob to encode their private data and policies as OADTs. Figure 6.2 shows two OADTs for the type `list`: $\widehat{\text{list}}_{\leq}$, whose public view is the maximum length of a list, and $\widehat{\text{list}}_{=}$, whose public view is the exact length.

Similar to TAYPE, in the implementation of TAYPSI, oblivious values are represented using arrays of secure values. To ensure that attackers cannot learn anything from the “memory layout” of an OADT value, the size of this array is the same for all values of a particular OADT. As an example, the encoding of the list `Cons 10 (Cons 20 Nil)` as an oblivious list of type $\widehat{\text{list}}_{\leq} 2$ is $\widehat{\text{inr}} ([10], \widehat{\text{inr}} ([20], ()))$, where $\widehat{\text{inr}} (\widehat{\text{inl}})$ is the oblivious counterpart of standard sum injection `inr (inl)`. Under the hood, this oblivious value is represented as an array holding four secure values; in the remainder of this section, we will informally write this value as `[Cons, 10, Cons, 20]`, where `[Cons]` is a synonym of the tag `[inr]` for readability. As another example, the empty list `Nil` is encoded as $\widehat{\text{inl}} ()$; it is also represented using an

```

obliv  $\widehat{\text{list}}_{\leq}$  (k : N) =
  if k = 0 then 1
  else 1 +  $\widehat{\mathbb{Z}}$  ×  $\widehat{\text{list}}_{\leq}$  (k-1)

obliv  $\widehat{\text{list}}_{=}$  (k : N) =
  if k = 0 then 1
  else  $\widehat{\mathbb{Z}}$  ×  $\widehat{\text{list}}_{=}$  (k-1)

```

Figure 6.2. Oblivious lists with maximum and exact length public views

array with four elements, $[\text{Nil}, -, -, -]$, where the last three elements are dummy encrypted values (denoted by $-$). Our compiler uses the type of $\widehat{\text{inl}}$ to automatically pad this array with these values, in order to ensure that it is indistinguishable from other private values of $\widehat{\text{list}}_{\leq} 2$.

Users can directly implement privacy-preserving applications in TAYPSI using OADTs and secure operations, but this requires manually instrumenting programs so that their control flow only depends on public information. Under this discipline, the implementation of a secure function intertwines program logic and privacy policies: the secure version of `filter` requires a different implementation depending on whether Alice is willing to share the exact length of her list, or an upper bound on that length. TAYPE (Chapter 5) decouples these concerns by allowing programs to include unsafe computations and repairing unsafe computations at runtime, using a novel form of semantics called *tape semantics*. As an example of this approach, in TAYPE, a secure implementation of `filter` that allows Alice to only share an upper bound on the size of her list can be written as:

```
fn  $\widehat{\text{filter}}_{\leq} : (k : \mathbb{N}) \rightarrow \widehat{\text{list}}_{\leq} k \rightarrow \widehat{\mathbb{Z}} \rightarrow \widehat{\text{list}}_{\leq} k =$ 
   $\lambda k \widehat{x}s \widehat{y} \Rightarrow \widehat{\text{list}}_{\leq} \#s k (\text{filter } (\widehat{\text{list}}_{\leq} \#r k \widehat{x}s) (\widehat{\mathbb{Z}} \#r \widehat{y}))$ 
```

The type signature of $\widehat{\text{filter}}_{\leq}$ specifies the policy it must follow. Intuitively, its implementation first “decrypts” the private inputs, applying the standard `filter` function to those values, and then “re-encrypts” the filtered list. In this example, the retractions of the private inputs $\widehat{x}s$ and \widehat{y} are unsafe computations that would violate the desired policy if they were computed naively. Fortunately, using the tape semantics prevents this from occurring by deferring these computations until it is safe to do so. Less fortunately, the runtime overhead of dynamic policy enforcement makes it hard to scale private applications manipulating structured data. As one data point, the secure version of `filter` produced by TAYPE takes more than 5 seconds to run with an oblivious list $\widehat{\text{list}}_{\leq}$ with sixteen elements, and its performance grows exponentially worse as the number of elements increases.

To understand the source of this slowdown, consider a computation that filters a private list containing 10 and 20 with integer 15: $\widehat{\text{filter}}_{\leq} 2 [\text{Cons}, 10, \text{Cons}, 20] [15]$. The first step in evaluating this function is to compute $\widehat{\text{list}}_{\leq} \#r 2 [\text{Cons}, 10, \text{Cons}, 20]$. Completely

reducing this expression leaks information, so tape semantics instead stops evaluation after producing the following computation:

```
mux [false] Nil (Cons ( $\widehat{Z}\#r$  [10])) (mux [false] Nil (Cons ( $\widehat{Z}\#r$  [20]) Nil)))
```

The two [false]s are the results of securely checking if the two constructors in the input list are Nil. Observe that evaluating either `mux` or $\widehat{Z}\#r$ would reveal private information, so the evaluation of these operations is deferred. This delayed computation can be thought of as an “if-tree” whose internal nodes are the private conditions needed to compute the final results, and whose leaves hold the result of the computation along each corresponding control flow path. To make progress, tape semantics distributes the context surrounding a delayed computation, `filter` and then $\widehat{list}_{\leq}\#s$ in this example, into each of its leaves; having done so, those leaves can be further evaluated. Importantly, in our example, the leaves of this if-tree are eventually re-encrypted using $\widehat{list}_{\leq}\#s$. The tape semantics does so in a secure way, so that $\widehat{Z}\#r$ [10] becomes [10] again, and each result list is converted to a secure value of the expected OADT. Once the branches of a `mux` node have been reduced to oblivious values of the same type, the node itself can be securely reduced using the secure semantics of `mux`. Unfortunately, the if-tree produced by the tape semantics can grow exponentially large before its `mux` nodes can be reduced. For example, after applying `filter` to the if-tree produced by $\widehat{list}_{\leq}\#r$, the resulting if-tree has a leaf corresponding to every possible list that `filter` could produce; the number of these leaves is exponential in the maximum length of the input list. As any surrounding computation, i.e., $\widehat{list}_{\leq}\#s$ in our example, can be distributed to each of these leaves, an exponential number of computations may need to be performed before the if-tree can be collapsed.

To remedy these limitations, this chapter proposes to instead compile an insecure program into a secure version that *statically* enforces a specified policy. To do so, we extend TAYPE, the secure language from [Chapter 5](#) with Ψ -types, a form of *dependent sums* (or dependent pairs) that packs public views and the oblivious data into a uniform representation. For example, $\Psi\widehat{list}_{\leq}$ is the oblivious list \widehat{list}_{\leq} with its public view: $\langle 2, \widehat{inr} ([10], \widehat{inr} ([20], ())) \rangle$ and $\langle 2, \widehat{inl} () \rangle$ are elements of type $\Psi\widehat{list}_{\leq}$, corresponding to the examples in the previous section. The first component of this pair-like syntax is a public view and the second component

is an OADT whose public view is exactly the first component. This allows users to again derive a private filter function from its type signature:

```
fn  $\widehat{\text{filter}}_{\leq}$  :  $\Psi\widehat{\text{list}}_{\leq} \rightarrow \widehat{\mathbb{Z}} \rightarrow \Psi\widehat{\text{list}}_{\leq} = \%lift \text{ filter}$ 
```

Users no longer need to explicitly provide the public views for either the output or any intermediate subroutines: both are automatically inferred. As a result, the policy specification of $\widehat{\text{filter}}_{\leq}$ more directly corresponds to the type signature of `filter`. In addition, specifying policies using Ψ -types avoids mistakes in the supplied public views: using TAYPE, if the programmer mistakenly specifies the return type $\widehat{\text{list}}_{\leq} (k-1)$ for a secure version of `filter`, for example, the resulting implementation may truncate the last element of the result list. A keyword `%lift` is used to translate the standard non-secure function `filter` to a private version that respects the policy specification.

To understand how this translation works, consider a naive approach where each algebraic data type (ADT) is thought of as an abstract interface, whose operations correspond to the *introduction* and *elimination* forms of the algebraic data type⁹. An ADT, e.g., `list`, as well as any corresponding Ψ -type, e.g., $\Psi\widehat{\text{list}}_{\leq}$ and $\Psi\widehat{\text{list}}_{=}$, are implementations or *instances* of this interface. For example, an interface for list operations is:

```
ListLike t = {
  Nil : 1 → t;
  Cons :  $\widehat{\mathbb{Z}} \times t \rightarrow t$ ;
  match : t → (1 →  $\alpha$ ) → ( $\widehat{\mathbb{Z}} \times t \rightarrow \alpha$ ) →  $\alpha$ 
}
```

As long as $\Psi\widehat{\text{list}}_{\leq}$ and $\Psi\widehat{\text{list}}_{=}$ implement this interface, we could straightforwardly translate `filter` to a secure version, as shown in Figure 6.3. This strategy does not rely on unsafe retractions like $\widehat{\text{list}}_{\leq} \#r$, as private data always remains in its secure form, eliminating the need to defer unsafe computations, which is the source of exponential slowdowns in TAYPE. Unfortunately, there are several obstacles to directly implementing this strategy. First, an ADT and an OADT may not agree on the type signatures of the abstract interface. `ListLike` fixes the argument types of operations like `Cons` and `match`, meaning that `list` is

⁹As TAYPSI already supports general recursion, we use pattern matching instead of recursion schemes as our elimination forms.

```

fn  $\widehat{\text{filter}}_{\leq} : \Psi\widehat{\text{list}}_{\leq} \rightarrow \widehat{\mathbb{Z}} \rightarrow \Psi\widehat{\text{list}}_{\leq} = \lambda \text{xs } y \Rightarrow$ 
   $\widehat{\text{list}}_{\leq} \# \text{match } \text{xs}$ 
  ( $\lambda \_ \Rightarrow \widehat{\text{list}}_{\leq} \# \text{Nil } ()$ )
  ( $\lambda (\text{x}, \text{xs}') \Rightarrow$ 
    mux ( $\text{x} \widehat{\leq} \text{y}$ ) ( $\widehat{\text{list}}_{\leq} \# \text{Cons } \text{x } (\widehat{\text{filter}}_{\leq} \text{xs}' \text{y})$ )
      ( $\widehat{\text{filter}}_{\leq} \text{xs}' \text{y}$ ))

```

(a) Secure `filter` using OADT $\widehat{\text{list}}_{\leq}$

```

fn  $\widehat{\text{filter}}_{=} : \Psi\widehat{\text{list}}_{=} \rightarrow \widehat{\mathbb{Z}} \rightarrow \Psi\widehat{\text{list}}_{=} = \lambda \text{xs } y \Rightarrow$ 
   $\widehat{\text{list}}_{=} \# \text{match } \text{xs}$ 
  ( $\lambda \_ \Rightarrow \widehat{\text{list}}_{=} \# \text{Nil } ()$ )
  ( $\lambda (\text{x}, \text{xs}') \Rightarrow$ 
    mux ( $\text{x} \widehat{\leq} \text{y}$ ) ( $\widehat{\text{list}}_{=} \# \text{Cons } \text{x } (\widehat{\text{filter}}_{=} \text{xs}' \text{y})$ )
      ( $\widehat{\text{filter}}_{=} \text{xs}' \text{y}$ ))

```

(b) Secure `filter` using OADT $\widehat{\text{list}}_{=}$

Figure 6.3. Naive translation of `filter` to secure versions

not an instance of this abstract interface, despite `list` being a very reasonable (albeit very permissive) policy! In general, different OADTs may only be able to implement operations with specific signatures. Second, a private function may involve a mixture of oblivious types. Thus, some functions may need to coerce from one type to a “more” secure version. For example, if the policy of $\widehat{\text{filter}}_{\leq}$ is $\Psi\widehat{\text{list}}_{\leq} \rightarrow \mathbb{Z} \rightarrow \Psi\widehat{\text{list}}_{\leq}$, its second argument `y` will need to be converted to $\widehat{\mathbb{Z}}$ in order to evaluate $\text{x} \widehat{\leq} \text{y}$. A secure list that discloses its exact length may similarly need to be converted to one disclosing its maximum length. Third, this naive translation results in ill-typed programs, because the branches of a `mux` may have mismatched public views. In $\widehat{\text{filter}}_{\leq}$, for example, the branches of `mux` may evaluate to $\langle 2, [\text{Cons}, 10, \text{Cons}, 20] \rangle$ and $\langle 1, [\text{Cons}, 20] \rangle$, respectively. Thus, TAYPSI’s secure type system will (rightly) reject $\widehat{\text{filter}}_{\leq}$ as leaky. Lastly, the signatures that should be ascribed to any subsidiary function calls may not be obvious. Consider the following client of `filter`:

```

fn filter5 : list → list = λxs ⇒ filter xs 5

```

If `filter5` is given a signature $\Psi\widehat{\text{list}}_{\leq} \rightarrow \Psi\widehat{\text{list}}_{\leq}$, we would like to use a secure version of the `filter` function with the type $\Psi\widehat{\text{list}}_{\leq} \rightarrow \mathbb{Z} \rightarrow \Psi\widehat{\text{list}}_{\leq}$, as the threshold argument is publicly known. In general, a function may have many private versions, and we should infer which version to use at each callsite: a recursive function may even recursively call a different “version” of itself.

To solve these challenges, we generalize the abstract interface described above into a set of more flexible structures, which we collectively refer to as Ψ -structures (Section 6.3). Intuitively, each category of Ψ -structures solves one of the challenges described above. Our translation algorithm (Section 6.4) generates a set of typing *constraints* for the intermediate expressions in a program. These constraints are then solved using the set of available Ψ -structures, resulting in multiple private versions of the necessary functions and ruling out the infeasible ones, e.g., $\widehat{\text{filter}}_{=}$.

Figure 6.4 presents the methods of each category of Ψ -structures of $\widehat{\text{list}}_{\leq}$. The first two methods, $\widehat{\text{list}}_{\leq}\#\text{s}$ and $\widehat{\text{list}}_{\leq}\#\text{r}$, are its section and retraction functions, belonging to the OADT-structure category. Unlike TAYPE, these two functions are not directly used to derive secure implementations of functions. In fact, our type system guarantees that retraction functions are never used in a secure computation, because TAYPSI does not rely on tape semantics to repair unsafe computation (the `unsafe fn` keyword tells our type checker that $\widehat{\text{list}}_{\leq}\#\text{r}$ is potentially leaky). Our implementation of TAYPSI exposes section and retraction functions as part of the API of the secure library it generates, however, so that client programs can conceal their private input and reveal the output of secure computations. This structure also includes a `view` method, which our translation uses to select the public view needed to safely convert a `list` into a $\Psi\widehat{\text{list}}_{\leq}$. Figure 6.4 does not show coercion methods, but the programmers can define a coercion from $\Psi\widehat{\text{list}}_{=}$ to $\Psi\widehat{\text{list}}_{\leq}$, for example.

The next set of methods belong to the intro-structure and elim-structure category. These introduction ($\widehat{\text{list}}_{\leq}\#\text{Nil}$ and $\widehat{\text{list}}_{\leq}\#\text{Cons}$) and elimination ($\widehat{\text{list}}_{\leq}\#\text{match}$) methods construct and destruct private list, respectively. As we construct and manipulate data, these methods build the private version, calculate its public view, and record that view in Ψ -types. Their type signatures are specified by the programmers, as long as the signatures are *compatible* with $\mathbb{Z} \times \text{list}$ (Section 6.3).

OADT-STRUCTURE

```

fn  $\widehat{\text{list}}_{\leq} \#s : (k : \mathbb{N}) \rightarrow \text{list} \rightarrow \widehat{\text{list}}_{\leq} k = \lambda k \text{ xs} \Rightarrow$ 
  if  $k = 0$  then  $()$ 
  else match xs with
    | Nil  $\Rightarrow \widehat{\text{inl}} ()$ 
    | Cons x xs'  $\Rightarrow \widehat{\text{inr}} (\widehat{\mathbb{Z}} \#s \ x, \widehat{\text{list}}_{\leq} \#s (k-1) \text{ xs}')$ 

```

```

unsafe fn  $\widehat{\text{list}}_{\leq} \#r : (k : \mathbb{N}) \rightarrow \widehat{\text{list}}_{\leq} k \rightarrow \text{list} = \lambda k \Rightarrow$ 
  if  $k = 0$  then  $\lambda \_ \Rightarrow \text{Nil}$ 
  else  $\lambda \text{xs} \Rightarrow \widehat{\text{match}} \text{xs}$  with
    |  $\widehat{\text{inl}} \_ \Rightarrow \text{Nil}$ 
    |  $\widehat{\text{inr}} (x, \text{xs}') \Rightarrow \text{Cons} (\widehat{\mathbb{Z}} \#r \ x) (\widehat{\text{list}}_{\leq} \#r (k-1) \text{ xs}')$ 

```

```

fn  $\widehat{\text{list}}_{\leq} \#view : \text{list} \rightarrow \mathbb{N} = \text{length}$ 

```

INTRO/ELIM-STRUCTURE

```

fn  $\widehat{\text{list}}_{\leq} \#\text{Nil} : \mathbb{1} \rightarrow \Psi \widehat{\text{list}}_{\leq} = \lambda \_ \Rightarrow \langle 0, () \rangle$ 

```

```

fn  $\widehat{\text{list}}_{\leq} \#\text{Cons} : \widehat{\mathbb{Z}} \times \Psi \widehat{\text{list}}_{\leq} \rightarrow \Psi \widehat{\text{list}}_{\leq} = \lambda (x, \langle k, \text{xs} \rangle) \Rightarrow \langle k+1, \widehat{\text{inr}} (x, \text{xs}) \rangle$ 

```

```

fn  $\widehat{\text{list}}_{\leq} \#\text{match} : \Psi \widehat{\text{list}}_{\leq} \rightarrow (\mathbb{1} \rightarrow \alpha) \rightarrow (\widehat{\mathbb{Z}} \times \Psi \widehat{\text{list}}_{\leq} \rightarrow \alpha) \rightarrow \alpha =$ 
   $\lambda \langle k, \text{xs} \rangle \text{ f1 f2} \Rightarrow$ 
  (if  $k = 0$  then  $\lambda \_ \Rightarrow \text{f1} ()$ 
  else  $\lambda \text{xs} \Rightarrow \widehat{\text{match}} \text{xs}$  with
    |  $\widehat{\text{inl}} \_ \Rightarrow \text{f1} ()$ 
    |  $\widehat{\text{inr}} (x, \text{xs}') \Rightarrow \text{f2} (x, \langle k-1, \text{xs}' \rangle) : \widehat{\text{list}}_{\leq} k \rightarrow \alpha) \text{ xs}$ 

```

JOIN-STRUCTURE

```

fn  $\widehat{\text{list}}_{\leq} \#\text{join} : \mathbb{N} \rightarrow \mathbb{N} \rightarrow \mathbb{N} = \text{max}$ 

```

```

fn  $\widehat{\text{list}}_{\leq} \#\text{reshape} : (k : \mathbb{N}) \rightarrow (k' : \mathbb{N}) \rightarrow \widehat{\text{list}}_{\leq} k \rightarrow \widehat{\text{list}}_{\leq} k' = \lambda k \text{ k}' \Rightarrow$ 
  if  $k' = 0$  then  $\lambda \_ \Rightarrow ()$ 
  else if  $k = 0$  then  $\lambda \_ \Rightarrow \widehat{\text{inl}} ()$ 
  else  $\lambda \text{xs} \Rightarrow \widehat{\text{match}} \text{xs}$  with
    |  $\widehat{\text{inl}} \_ \Rightarrow \widehat{\text{inl}} ()$ 
    |  $\widehat{\text{inr}} (x, \text{xs}') \Rightarrow \widehat{\text{inr}} (x, \widehat{\text{list}}_{\leq} \#\text{reshape} (k-1) (k'-1) \text{ xs}')$ 

```

Figure 6.4. Ψ -structures of $\widehat{\text{list}}_{\leq}$

The `join` and `reshape` methods in the join-structure category enable translated programs to include private conditionals whose branches return OADT values with different public views. As an example, consider the following private conditional whose branches have Ψ -types:

```
mux [true] ⟨2, [Cons,10,Cons,20]⟩ ⟨1, [Cons,20]⟩
```

To build a version of this program that does not reveal `[true]`, TAYPSI uses `join` to calculate a common public view that “covers” both branches. In this example, $\widehat{\text{list}}_{\leq} \# \text{join}$ chooses a public view of 2, as a list with at most one element also has at most two elements. Our translation then uses the `reshape` method to convert both branches to use this common public view. In our example, `[Cons,20]`, an oblivious list of maximum length 1, is converted into the list `[Cons,20,Nil,-]`, which has maximum length 2. Since both branches in the resulting program have the same public view, it is safe to evaluate `mux`: the resulting list is equivalent to `⟨2, mux [true] [Cons,10,Cons,20] [Cons,20,Nil,-]⟩`. As we will see later, not all OADTs admit join structures, e.g., $\widehat{\text{list}}_{=}$, but our translation generates constraints that take advantage of any that are available, failing when these constraints cannot be resolved in a way that guarantees security. Note that these two methods are key to avoiding the slowdown exhibited by TAYPE’s enforcement strategy: they allow functions that may return different private representations to be *eagerly* evaluated, instead of being lazily deferred in a way that requires an exponential number of subcomputations to resolve.

In summary, to develop a secure application in TAYPSI, programmers first implement its desired functionality, e.g., `filter`, in the public fragment of TAYPSI, independently of any particular privacy policy. Policies are separately defined as oblivious algebraic data types, e.g., $\widehat{\text{list}}_{\leq}$, and their Ψ -structures. Users can then automatically derive a secure version of their application by providing the desired policy in the form of a type signature involving Ψ -types, relying on TAYPSI’s compiler to produce a privacy-preserving implementation. The type system of TAYPSI, like TAYPE’s, provides a strong security guarantee in the form of an obliviousness theorem ([Theorem 6.2.1](#)). This obliviousness theorem is a variant of noninterference [21], and ensures that well-typed programs in TAYPSI are *secure by construction*: no private information can be inferred even by an attacker capable of observing every state of a program’s execution. Our compilation algorithm is further guaranteed

to generate a secure implementation that preserves the behavior of the original program (Theorem 6.3.3). TAYPSI’s formal guarantees (Section 6.3.7) do not cover equi-termination of the source and target programs: when the public view lacks sufficient information to bound the computation of the original program, the secure version will not terminate, in order to avoid leaking information through its termination behavior.

The following three sections formally develop the language TAYPSI, the Ψ -structures, and our translation algorithm.

6.2 Taypsi, Formally

This section presents $\lambda_{\text{OADT}\Psi}$, the core calculus for secure computation that we will use to explain our translation. This calculus extends the existing λ_{OADT} calculus (Chapter 3) with Ψ -types, and uses ML-style ADTs in lieu of explicit `fold` and `unfold` operations. For simplicity, $\lambda_{\text{OADT}\Psi}$ does not include public sums and oblivious integers, which are straightforward to add.

6.2.1 Syntax

Figure 6.5 presents the syntax of $\lambda_{\text{OADT}\Psi}$. Types and expressions are in the same syntax class, as $\lambda_{\text{OADT}\Psi}$ is dependently typed, but we use e for expressions and τ for types when possible. A $\lambda_{\text{OADT}\Psi}$ program consists of a set of *global definitions* of data types, functions and oblivious types. Definitions in each of these classes are allowed to refer to themselves, permitting recursive types and general recursion in both function and oblivious type definitions. We use x for variable names, C for constructor names, T for type names, and \hat{T} for oblivious type names. Each constructor of an ADT definition takes exactly one argument, but this does not harm expressivity: this argument is `1` for constructors that take no arguments, e.g., `Nil`, and a tuple of types for constructors that have more than one argument, e.g., `Cons` takes an argument of type $\mathbb{Z} \times \text{list}$.

In addition to standard types and dependent function types (II), $\lambda_{\text{OADT}\Psi}$ includes oblivious booleans ($\hat{\mathbb{B}}$) and oblivious sum types ($\hat{+}$). The elimination forms of these types are oblivious conditionals `mux` and oblivious case analysis `match`, respectively. The branches of both expressions must be private and each branch has to be fully evaluated before the expression

$e, \tau ::=$ <ul style="list-style-type: none"> $\mathbb{1} \mid \mathbb{B} \mid \widehat{\mathbb{B}} \mid \tau \times \tau \mid \tau \widehat{+} \tau$ $\prod x:\tau, \tau$ $\Psi \widehat{T}$ $x \mid T$ $() \mid b$ $\lambda x:\tau \Rightarrow e$ $\text{let } x = e \text{ in } e$ $e \ e \mid C \ e \mid \widehat{T} \ e$ $\text{if } e \text{ then } e \text{ else } e$ $\text{mux } e \ e \ e$ (e, e) $\langle e, e \rangle$ $\pi_b \ e$ $\widehat{\iota}_b \langle \tau \rangle \ e$ $\text{match } e \ \text{with } x \Rightarrow e \mid x \Rightarrow e$ $\text{match } e \ \text{with } \overline{C} \ x \Rightarrow e$ $\widehat{\mathbb{B}} \#_s \ e$ $[b] \mid [\iota_b \langle \widehat{\omega} \rangle \ \widehat{v}]$ 	<p>EXPRESSIONS:</p> <ul style="list-style-type: none"> simple types dependent function type Ψ-type variable and type names unit and boolean constants function abstraction let binding applications conditional oblivious conditional pair dependent pair (Ψ-pair) product and Ψ-type projection oblivious sum injection oblivious sum elimination ADT elimination boolean section runtime boxed values
$D ::=$ <ul style="list-style-type: none"> $\text{data } T = \overline{C} \ \tau$ $\text{fn } x:\tau = e$ $\text{obliv } \widehat{T} \ (x:\tau) = \tau$ 	<p>GLOBAL DEFINITIONS</p> <ul style="list-style-type: none"> algebraic data type definition (recursive) function definition (recursive) oblivious type definition
$\widehat{\omega} ::= \mathbb{1} \mid \widehat{\mathbb{B}} \mid \widehat{\omega} \times \widehat{\omega} \mid \widehat{\omega} \widehat{+} \widehat{\omega}$	<p>OBLIVIOUS TYPE VALUES</p>
$\widehat{v} ::= () \mid [b] \mid (\widehat{v}, \widehat{v}) \mid [\iota_b \langle \widehat{\omega} \rangle \ \widehat{v}]$	<p>OBLIVIOUS VALUES</p>
$v ::= \widehat{v} \mid b \mid (v, v) \mid \langle v, v \rangle \mid \lambda x:\tau \Rightarrow e \mid C \ v$	<p>VALUES</p>

Figure 6.5. $\lambda_{\text{OADT}\Psi}$ syntax with extensions to λ_{OADT} highlighted

can take an atomic step to a final result. Boolean section $\widehat{\mathbb{B}}\#_s$ is a primitive operation that “encrypts” a boolean expression to an oblivious version. Oblivious injection $\widehat{\iota}_b$ (i.e., $\widehat{\text{inl}}$ and $\widehat{\text{inr}}$) are the oblivious counterparts of the standard constructors for sums. Other terms are mostly standard, although let bindings (`let`), conditionals (`if`) and pattern matching (`match`) are allowed to return a type, as $\lambda_{\text{OADT}\Psi}$ supports type-level computation.

The key addition over λ_{OADT} is the Ψ -type, $\Psi \widehat{T}$. It is constructed from a pair expression $\langle \cdot, \cdot \rangle$ that packs the public view and the oblivious data together, and has the same eliminators π_1

and π_2 as normal products. As an example, $\langle 3, \widehat{\text{list}}_{\leq} \#s \ 3 \ (\text{Cons } 1 \ \text{Nil}) \rangle$ creates a Ψ -pair of type $\Psi \widehat{\text{list}}_{\leq}$ with public view 3, using the section function from Figure 6.4. Projecting out the second component of a pair using π_2 produces a value of type $\widehat{\text{list}}_{\leq} \ 3$. A Ψ -type is essentially a dependent sum type $(\Sigma x:\tau, \widehat{T} \ x)$, with the restriction that τ is the public view of \widehat{T} , and that $\widehat{T} \ x$ is an oblivious type.

Since $\lambda_{\text{OADT}\Psi}$ has type-level computation, oblivious types have normal forms; oblivious type values $(\widehat{\omega})$ are essentially polynomials formed by primitive oblivious types. We also have the oblivious values of oblivious boolean and sum type. Note that these “boxed” values only appear at runtime, our semantics use these to model encrypted booleans and tagged sums.

6.2.2 Semantics

Figure 6.6 shows a selection of the small-step semantics rules of $\lambda_{\text{OADT}\Psi}$ (the omitted rules are identical to λ_{OADT}), with judgment $\Sigma \vdash e \longrightarrow e'$. The global context Σ is a map from names to a global definition, which is elided for brevity as it is fixed in these rules. The semantics of $\lambda_{\text{OADT}\Psi}$ is similar to λ_{OADT} , with the addition of S-PSIPROJ to handle the projection of dependent pairs, which is simply the same as normal projection. S-CTX reduces subterms according to the evaluation contexts defined in Figure 6.6. The first few contexts take care of the type-level reduction of product and oblivious sum type. The type annotation of oblivious injection ι_b is reduced to a type value first, before reducing the payload. The evaluation contexts for **mux** capture the intuition that all components of a private conditional have to be normalized to values first to avoid leaking the private condition through control flow channels.

S-OMATCH evaluates a pattern matching expression for oblivious sums. Similar to **mux**, oblivious pattern matching needs to ensure the reduction does not reveal private information about the discriminatee, e.g., whether it is the left injection or right injection. To do so, we reduce a **match** to an oblivious conditional that uses the private tag. The pattern variable in the “correct” branch is of course instantiated by the payload in the discriminatee, while the pattern variable in the “wrong” branch is an arbitrary value of the corresponding type, synthesized

$e \longrightarrow e'$

$$\begin{array}{c}
\text{S-CTX} \\
\frac{e \longrightarrow e'}{\mathcal{E}[e] \longrightarrow \mathcal{E}[e']} \\
\\
\text{S-FUN} \\
\frac{\text{fn } x:\tau = e \in \Sigma}{x \longrightarrow e} \\
\\
\text{S-OADT} \\
\frac{\text{obliv } \hat{T} \ (x:\tau) = \tau' \in \Sigma}{\hat{T} \ v \longrightarrow [v/x]\tau'} \\
\\
\text{S-APP} \\
\frac{}{(\lambda x:\tau \Rightarrow e) \ v \longrightarrow [v/x]e} \\
\\
\text{S-IF} \\
\frac{}{\text{if } b \ \text{then } e_1 \ \text{else } e_2 \longrightarrow \text{ite}(b, e_1, e_2)} \\
\\
\text{S-MUX} \\
\frac{}{\text{mux } [b] \ v_1 \ v_2 \longrightarrow \text{ite}(b, v_1, v_2)} \\
\\
\text{S-MATCH} \\
\frac{}{\text{match } C_i \ v \ \text{with } \bar{C} \ x \Rightarrow e \longrightarrow [v/x]e_i} \\
\\
\text{S-PROJ} \\
\frac{}{\pi_b \ (v_1, v_2) \longrightarrow \text{ite}(b, v_1, v_2)} \\
\\
\text{S-SEC} \\
\frac{}{\hat{\mathbb{B}}\#s \ b \longrightarrow [b]} \\
\\
\text{S-OINJ} \\
\frac{}{\hat{t}_b \langle \hat{\omega} \rangle \ \hat{v} \longrightarrow [t_b \langle \hat{\omega} \rangle \ \hat{v}]} \\
\\
\text{S-PSIPROJ} \\
\frac{}{\pi_b \ \langle v_1, v_2 \rangle \longrightarrow \text{ite}(b, v_1, v_2)} \\
\\
\text{S-OMATCH} \\
\frac{\hat{v}_1 \Leftarrow \hat{\omega}_1 \quad \hat{v}_2 \Leftarrow \hat{\omega}_2}{\widehat{\text{match}} \ [t_b \langle \hat{\omega}_1 \hat{+} \hat{\omega}_2 \rangle \ \hat{v}] \ \text{with } x \Rightarrow e_1 \mid x \Rightarrow e_2 \longrightarrow \text{mux } [b] \ \text{ite}(b, [\hat{v}/x]e_1, [\hat{v}_1/x]e_1) \\ \text{ite}(b, [\hat{v}_2/x]e_2, [\hat{v}/x]e_2)} \\
\\
\text{EVALUATION CONTEXTS} \\
\mathcal{E} ::= \square \times \tau \mid \hat{\omega} \times \square \mid \square \hat{+} \tau \mid \hat{\omega} \hat{+} \square \\
\mid \text{let } x = \square \ \text{in } e \mid e \ \square \mid \square \ v \mid C \ \square \mid \hat{T} \ \square \\
\mid \text{if } \square \ \text{then } e \ \text{else } e \mid \text{mux } \square \ e \ e \mid \text{mux } v \ \square \ e \mid \text{mux } v \ v \ \square \\
\mid (\square, e) \mid (v, \square) \mid \langle \square, e \rangle \mid \langle v, \square \rangle \mid \pi_b \ \square \\
\mid \hat{t}_b \langle \square \rangle \ e \mid \hat{t}_b \langle \hat{\omega} \rangle \ \square \mid \widehat{\text{match}} \ \square \ \text{with } x \Rightarrow e \mid x \Rightarrow e \\
\mid \text{match } \square \ \text{with } \bar{C} \ x \Rightarrow e \mid \hat{\mathbb{B}}\#s \ \square
\end{array}$$

Figure 6.6. Selected small-step semantics rules of $\lambda_{\text{OADT}\Psi}$

from the judgment $\hat{v} \Leftarrow \hat{\omega}$, whose definition has been presented in [Figure 3.9](#). When evaluating a $\widehat{\text{match}}$ statement whose discriminée is $[\text{inl} \langle \hat{\mathbb{B}} \hat{+} \hat{\mathbb{B}} \times \hat{\mathbb{B}} \rangle [\text{true}]]$, the pattern variable in the second branch can be substituted by $([\text{true}], [\text{true}])$, $([\text{false}], [\text{true}])$, or any other pair of oblivious booleans.

$$\begin{array}{c}
\boxed{\Gamma \vdash e : \tau} \\
\\
\text{T-CONV} \\
\frac{\Gamma \vdash e : \tau \quad \tau \equiv \tau' \quad \Gamma \vdash \tau' :: *}{\Gamma \vdash e : \tau'} \\
\\
\text{T-ABS} \\
\frac{x : \tau_1, \Gamma \vdash e : \tau_2 \quad \Gamma \vdash \tau_1 :: *}{\Gamma \vdash \lambda x : \tau_1 \Rightarrow e : \Pi x : \tau_1, \tau_2} \\
\\
\text{T-APP} \\
\frac{\Gamma \vdash e_2 : \Pi x : \tau_1, \tau_2 \quad \Gamma \vdash e_1 : \tau_1}{\Gamma \vdash e_2 \ e_1 : [e_1/x] \tau_2} \\
\\
\text{T-IF} \\
\frac{\Gamma \vdash e_0 : \mathbb{B} \quad \Gamma \vdash e_1 : [\text{true}/z] \tau \quad \Gamma \vdash e_2 : [\text{false}/z] \tau}{\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : [e_0/z] \tau} \\
\\
\text{T-CTOR} \\
\frac{\text{data } T = \overline{C} \ \tau \in \Sigma \quad \Gamma \vdash e : \tau_i}{\Gamma \vdash C_i \ e : T} \\
\\
\text{T-MATCH} \\
\frac{\text{data } T = \overline{C} \ \tau \in \Sigma \quad \Gamma \vdash e_0 : T \quad \forall i. x : \tau_i, \Gamma \vdash e_i : [C_i \ x/z] \tau'}{\Gamma \vdash \text{match } e_0 \text{ with } \overline{C} \ x \Rightarrow e : [e_0/z] \tau'} \\
\\
\text{T-MUX} \\
\frac{\Gamma \vdash e_0 : \widehat{\mathbb{B}} \quad \Gamma \vdash \tau :: *^0 \quad \Gamma \vdash e_1 : \tau \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash \text{mux } e_0 \ e_1 \ e_2 : \tau} \\
\\
\text{T-PSIPAIR} \\
\frac{\text{obliv } \widehat{T} \ (x : \tau) = \tau' \in \Sigma \quad \Gamma \vdash e_1 : \tau \quad \Gamma \vdash e_2 : \widehat{T} \ e_1}{\Gamma \vdash \langle e_1, e_2 \rangle : \Psi \widehat{T}} \\
\\
\text{T-PSIPROJ}_1 \\
\frac{\text{obliv } \widehat{T} \ (x : \tau) = \tau' \in \Sigma \quad \Gamma \vdash e : \Psi \widehat{T}}{\Gamma \vdash \pi_1 \ e : \tau} \\
\\
\text{T-PSIPROJ}_2 \\
\frac{\text{obliv } \widehat{T} \ (x : \tau) = \tau' \in \Sigma \quad \Gamma \vdash e : \Psi \widehat{T}}{\Gamma \vdash \pi_2 \ e : \widehat{T} \ (\pi_1 \ e)}
\end{array}$$

Figure 6.7. Selected typing rules of $\lambda_{\text{OADT}\Psi}$

6.2.3 Type System

Similar to λ_{OADT} , types in $\lambda_{\text{OADT}\Psi}$ are classified by *kinds* which specify how protected a type is, in addition to ensuring the types are well-formed. For example, an oblivious type, e.g., $\widehat{\mathbb{B}}$, kinded by $*^0$, can be used as branches of an oblivious conditional, but not as a public view, which can only be kinded by $*^P$. A mixed kind $*^M$ is used to classify function types and types that consist of both public and oblivious components, e.g., $\mathbb{B} \times \widehat{\mathbb{B}}$. A type with a mixed kind cannot be used as a public view or in private context.

The type system of $\lambda_{\text{OADT}\Psi}$ is defined by a pair of typing and kinding judgments, $\Sigma; \Gamma \vdash e : \tau$ and $\Sigma; \Gamma \vdash \tau :: \kappa$, with global context Σ (which is again elided for brevity) and the standard

$\Gamma \vdash \tau :: \kappa$

$$\begin{array}{c}
\text{K-SUB} \\
\frac{\Gamma \vdash \tau :: \kappa \quad \kappa \sqsubseteq \kappa'}{\Gamma \vdash \tau :: \kappa'} \\
\\
\text{K-OADT} \\
\frac{\text{obliv } \hat{T} (x:\tau) = \tau' \in \Sigma \quad \Gamma \vdash e : \tau}{\Gamma \vdash \hat{T} e :: *^0} \\
\\
\text{K-PI} \\
\frac{\Gamma \vdash \tau_1 :: * \quad x : \tau_1, \Gamma \vdash \tau_2 :: *}{\Gamma \vdash \Pi_{x:\tau_1, \tau_2} :: *^M} \\
\\
\text{K-OSUM} \\
\frac{\Gamma \vdash \tau_1 :: *^0 \quad \Gamma \vdash \tau_2 :: *^0}{\Gamma \vdash \tau_1 \hat{+} \tau_2 :: *^0} \\
\\
\text{K-PSI} \\
\frac{\text{obliv } \hat{T} (x:\tau) = \tau' \in \Sigma}{\Gamma \vdash \Psi \hat{T} :: *^M} \\
\\
\text{K-IF} \\
\frac{\Gamma \vdash e_0 : \mathbb{B} \quad \Gamma \vdash \tau_1 :: *^0 \quad \Gamma \vdash \tau_2 :: *^0}{\Gamma \vdash \text{if } e_0 \text{ then } \tau_1 \text{ else } \tau_2 :: *^0} \\
\\
\text{K-MATCH} \\
\frac{\text{data } T = \overline{C} \tau \in \Sigma \quad \Gamma \vdash e_0 : T \\
\forall i. x : \tau_i, \Gamma \vdash \tau'_i :: *^0}{\Gamma \vdash \text{match } e_0 \text{ with } \overline{C} x \Rightarrow \tau' :: *^0}
\end{array}$$

Figure 6.8. Selected kinding rules of $\lambda_{\text{OADT}\Psi}$

$\Sigma \vdash D$

$$\begin{array}{c}
\text{DT-FUN} \\
\frac{\cdot \vdash \tau :: * \quad \cdot \vdash e : \tau}{\Sigma \vdash \text{fn } x:\tau = e} \\
\\
\text{DT-ADT} \\
\frac{\forall i. \cdot \vdash \tau_i :: *^P}{\Sigma \vdash \text{data } T = \overline{C} \tau} \\
\\
\text{DT-OADT} \\
\frac{\cdot \vdash \tau :: *^P \quad x : \tau \vdash \tau' :: *^0}{\Sigma \vdash \text{obliv } \hat{T} (x:\tau) = \tau'}
\end{array}$$

Figure 6.9. $\lambda_{\text{OADT}\Psi}$ global definitions typing

typing context Γ . [Figure 6.7](#) and [Figure 6.8](#) presents a subset of our typing and kinding rules; the omitted rules are identical to the ones in λ_{OADT} .

The security type system [27] of $\lambda_{\text{OADT}\Psi}$ enforces a few key invariants. First, oblivious types can only be constructed from oblivious types, which is enforced by the kinding rules, such as K-OSUM. Otherwise, the attacker could infer the private tag of an oblivious sum, e.g., $\mathbb{B} \hat{+} \mathbb{1}$, by observing its public payload. Second, oblivious operations, e.g., `mux`, require their subterms to be oblivious, to avoid leaking private information via control flow channels. T-MUX, for example, requires both branches to be typed by an oblivious type, otherwise an attacker may infer the private condition by observing the result, as in `mux [true] true false`. Third,

type-level computation is only defined for oblivious types and cannot depend on private information. Thus, **K-IF** and **K-MATCH** requires all branches to have oblivious kinds, and the condition to be public. The type `mux [true] 1 $\hat{\mathbb{B}}$` is ill-typed, since the “shape” of the data reveals the private condition.

The typing rules for Ψ -types are defined similarly to the rules of standard dependent sums. **T-PSIPAIR** introduces a dependent pair, where the type of the second component depends on the first component. In contrast to standard dependent sum type, Ψ -type has the restriction that the first component must be public, and the second component must be oblivious. This condition is implicitly enforced by the side condition that $\hat{\mathbb{T}}$ is an OADT with public view type τ . [Figure 6.9](#) shows the typing rules for global definitions; **DT-OADT** prescribes exactly this restriction. The rules for the first and second projection of Ψ -type, **T-PSIPROJ₁** and **T-PSIPROJ₂**, are very similar to the corresponding rules for standard dependent sum types. Observe that a Ψ -type always has mixed kind, as in **K-PSI**, because it consists of both public and oblivious components.

T-CONV allows conversion between equivalent types, such as `if true then $\hat{\mathbb{B}}$ else 1` and $\hat{\mathbb{B}}$. The equivalence judgment $\tau \equiv \tau'$ is defined by a set of *parallel reduction* rules, which are mostly identical to the rules in λ_{OADT} . The converted type is nonetheless required to be well-kinded.

Note that these rules cannot be used to type check retraction functions, e.g., `list≤#r` from [Figure 6.4](#), and for good reason: these functions reveal private information. Nevertheless, we still want to check that these sorts of “leaky” functions have standard type safety properties, i.e., progress and preservation. To do so, we use a version of these rules that simply omit some security-related side-conditions about oblivious kinding: removing $\Gamma \vdash \tau :: *^0$ from **T-MUX** allows the branches of a `mux` to disclose the private condition, for example. The implementation of TAYPSI’s type checker uses a “mode” flag to indicate whether security-related side-conditions should be checked. Our implementation ensures that secure functions never use any leaky functions.

6.2.4 Metatheory

With our addition of Ψ -types, $\lambda_{\text{OADT}\Psi}$ enjoys the standard type safety properties (i.e., progress and preservation), and, more importantly, the same security guarantees as λ_{OADT} :

Theorem 6.2.1 (Obliviousness). *If $e_1 \approx e_2$ and $\cdot \vdash e_1 : \tau_1$ and $\cdot \vdash e_2 : \tau_2$, then*

1. $e_1 \longrightarrow^n e'_1$ if and only if $e_2 \longrightarrow^n e'_2$ for some e'_2 .
2. if $e_1 \longrightarrow^n e'_1$ and $e_2 \longrightarrow^n e'_2$, then $e'_1 \approx e'_2$.

Here, $e \approx e'$ means the two expressions are indistinguishable, i.e., they only differ in their unobservable oblivious values, and $e \longrightarrow^n e'$ means e reduces to e' in exactly n steps. This obliviousness theorem provides a strong security guarantee: well-typed programs that are indistinguishable produce traces that are pairwise indistinguishable. In other words, an attacker cannot infer any private information even by observing the execution trace of a program. All these results are mechanized in the Coq theorem prover, including the formalization of the core calculus and the proofs of soundness and obliviousness theorems.

6.3 Ψ -structures and Declarative Lifting

While our secure language makes it possible to encode structured data and privacy policies, and use them in a secure way, it does not quite achieve our main goal yet, i.e., to decouple privacy policies and programmatic concerns. To do so, we allow the programmers to implement the functionality of their secure application in a conventional way, that is using only the public, nondependent fragment of TAYPSI. We make this fragment explicit by requiring such programs to have *simple types*, denoted by η , defined in [Figure 6.10](#). For example, `filter` has simple type `list → ℤ → list`. Programs of simple types are the source programs to our lifting process that translates them to a private version against a policy, which stipulates the public information allowed to disclose in the program input and output. This policy on private functionality is specified by a *specification type*, denoted by θ , defined also in [Figure 6.10](#). For example, $\widehat{\text{filter}}_{\leq}$ has specification type $\Psi\widehat{\text{list}}_{\leq} \rightarrow \widehat{\mathbb{Z}} \rightarrow \Psi\widehat{\text{list}}_{\leq}$. Note that dependent types are not directly allowed in specifications, they are instead encapsulated

[θ]

<p>SIMPLE TYPES</p> $\eta ::= \mathbf{1} \mid \mathbb{B} \mid \mathbb{T} \mid \eta \times \eta \mid \eta \rightarrow \eta$	$[\mathbf{1}] = \mathbf{1} \qquad [\mathbb{B}] = [\widehat{\mathbb{B}}] = \mathbb{B}$	
<p>SPECIFICATION TYPES</p> $\theta ::= \mathbf{1} \mid \mathbb{B} \mid \widehat{\mathbb{B}} \mid \mathbb{T} \mid \Psi \widehat{\mathbb{T}}$ $\quad \mid \quad \theta \times \theta \mid \theta \rightarrow \theta$	$[\mathbb{T}] = \mathbb{T} \quad \text{where } \mathbb{T} \text{ is an ADT}$ $[\Psi \widehat{\mathbb{T}}] = \mathbb{T} \quad \text{where } \widehat{\mathbb{T}} \text{ is an OADT for } \mathbb{T}$	$[\theta \times \theta] = [\theta] \times [\theta] \qquad [\theta \rightarrow \theta] = [\theta] \rightarrow [\theta]$

Figure 6.10. Simple types, specification types and erasure

in Ψ -types. Simple types and specification types are additionally required to be well-kinded under empty local context, i.e., all ADTs and OADTs appear in them are defined.

However, not all specification types are valid with respect to a simple type. It is nonsensical to give `filter` the specification type $\widehat{\mathbb{Z}} \rightarrow \widehat{\mathbb{B}}$, for example. The specification types should still correspond to the simple types in some way: the specification type corresponding to `list` should at least be “list-like”. This correspondence is formally captured in the erasure function in Figure 6.10, which maps a specification type to the “underlying” simple type. For example, $\Psi \widehat{\text{list}}_{\leq}$ is erased to `list`. This function clearly induces an equivalence relation: the erasure $[\theta]$ is the representative of the equivalence class. We call this equivalence class a *compatibility class*, and say two types are *compatible* if they belong to the same compatibility class. For example, `list`, $\Psi \widehat{\text{list}}_{\leq}$ and $\Psi \widehat{\text{list}}_{=}$ are in the same compatibility class $[\text{list}]$. This erasure operation is straightforwardly extended to typing contexts, $[\Gamma]$, by erasing every specification type in Γ and leaving other types untouched.

Our translation transforms source programs with simple types into target programs with the desired (compatible) specification types. As mentioned in Section 6.1, this *lifting* process depends on a set of Ψ -structures which explain how to translate certain operations associated with an OADT.

6.3.1 OADT Structures

Every global OADT definition $\widehat{\mathbb{T}}$ must be equipped with an *OADT-structure*, defined below.

Definition 6.3.1 (OADT-structure). An OADT-structure of an OADT \widehat{T} , with public view type τ , consists of the following (TAYPSI) type and functions:

- A public type $T :: *^P$, which is the public counterpart of \widehat{T} . We say \widehat{T} is an OADT for T .
- A section function $\mathbf{s} : \prod \mathbf{k} : \tau, T \rightarrow \widehat{T} \mathbf{k}$, which converts a public type to its oblivious counterpart.
- A retraction function $\mathbf{r} : \prod \mathbf{k} : \tau, \widehat{T} \mathbf{k} \rightarrow T$, which converts an oblivious type to its public version.
- A public view function $\nu : T \rightarrow \tau$, which creates a valid view of the public type.
- A binary relation \preceq over values of types T and τ ; $\mathbf{v} \preceq \mathbf{k}$ reads as \mathbf{v} has public view \mathbf{k} , or \mathbf{k} is a valid public view of \mathbf{v} .

These operations are required to satisfy the following axioms:

- (A-O₁) \mathbf{s} and \mathbf{r} are a valid section and retraction, i.e., \mathbf{r} is a left-inverse for \mathbf{s} , given a valid public view: for any values $\mathbf{v} : T$, $\mathbf{k} : \tau$ and $\widehat{\mathbf{v}} : \widehat{T} \mathbf{k}$, if $\mathbf{v} \preceq \mathbf{k}$ and $\mathbf{s} \mathbf{k} \mathbf{v} \rightarrow^* \widehat{\mathbf{v}}$, then $\mathbf{r} \widehat{\mathbf{v}} \rightarrow^* \mathbf{v}$.
- (A-O₂) the result of \mathbf{r} always has valid public view: $\mathbf{r} \widehat{\mathbf{v}} \rightarrow^* \mathbf{v}$ implies $\mathbf{v} \preceq \mathbf{k}$ for all values $\mathbf{k} : \tau$, $\widehat{\mathbf{v}} : \widehat{T} \mathbf{k}$ and $\mathbf{v} : T$.
- (A-O₃) ν produces a valid public view: $\nu \mathbf{v} \rightarrow^* \mathbf{k}$ implies $\mathbf{v} \preceq \mathbf{k}$, given any values $\mathbf{v} : T$ and $\mathbf{k} : \tau$.

For example, $\widehat{\text{list}}_{\leq}$ is equipped with the OADT-structure with the public type `list`, section function $\widehat{\text{list}}_{\leq} \# \mathbf{s}$, retraction function $\widehat{\text{list}}_{\leq} \# \mathbf{r}$ and view function $\widehat{\text{list}}_{\leq} \# \text{view}$, all of which are shown in [Figure 6.4](#). TAYPSI users do not need to explicitly give the public type of an OADT-structure, as it can be inferred from the types of the other functions. The binary relation \preceq is only used in the proof of correctness of our translation, so TAYPSI users can also elide it. In the case of $\widehat{\text{list}}_{\leq}$, \preceq simply states the length of the list is no larger than the public view.

6.3.2 Join Structures

In order for Ψ -types to be flexibly used in the branches of secure control flow structures, our translation must be able to find a common public view for both branches, and to convert an OADT to use this view. To do so, an OADT can *optionally* be equipped with a *join-structure*.

Definition 6.3.2 (join-structure). A join-structure of an OADT \widehat{T} for \mathbb{T} , with public view type τ , consists of the following operations:

- A binary relation \sqsubseteq on τ , used to compare two public views.
- A join function $\sqcup : \tau \rightarrow \tau \rightarrow \tau$, which computes an upper bound of two public views¹⁰.
- A reshape function $\uparrow : \prod \mathbf{k} : \tau, \prod \mathbf{k}' : \tau, \widehat{T} \mathbf{k} \rightarrow \widehat{T} \mathbf{k}'$, which converts an OADT to one with a different public view.

such that:

- (A-R₁) \sqsubseteq is a partial order on τ .
- (A-R₂) the join function produces an upper bound: given values $\mathbf{k}_1, \mathbf{k}_2$ and \mathbf{k} of type τ , if $\mathbf{k}_1 \sqcup \mathbf{k}_2 \rightarrow^* \mathbf{k}$, then $\mathbf{k}_1 \sqsubseteq \mathbf{k}$ and $\mathbf{k}_2 \sqsubseteq \mathbf{k}$.
- (A-R₃) the validity of public views is monotone with respect to the binary relation \sqsubseteq : for any values $\mathbf{v} : \mathbb{T}, \mathbf{k} : \tau$ and $\mathbf{k}' : \tau$, if $\mathbf{v} \preceq \mathbf{k}$ and $\mathbf{k} \sqsubseteq \mathbf{k}'$, then $\mathbf{v} \preceq \mathbf{k}'$.
- (A-R₄) the reshape function produces equivalent value, as long as the new public view is valid: for any values $\mathbf{v} : \mathbb{T}, \mathbf{k} : \tau, \mathbf{k}' : \tau, \widehat{\mathbf{v}} : \widehat{T} \mathbf{k}$ and $\widehat{\mathbf{v}}' : \widehat{T} \mathbf{k}'$, if $\mathbf{r} \mathbf{k} \widehat{\mathbf{v}} \rightarrow^* \mathbf{v}$ and $\mathbf{v} \preceq \mathbf{k}'$ and $\uparrow \mathbf{k} \mathbf{k}' \widehat{\mathbf{v}} \rightarrow^* \widehat{\mathbf{v}}'$, then $\mathbf{r} \mathbf{k}' \widehat{\mathbf{v}}' \rightarrow^* \mathbf{v}$.

Figure 6.4 defines the join and reshape functions $\widehat{\text{list}}_{\leq} \# \text{join}$ and $\widehat{\text{list}}_{\leq} \# \text{reshape}$. The partial order for this join structure is simply the total order on integers, and the join is simply the maximum of the two numbers. Not all OADTs have a sensible join-structure: oblivious lists using their exact length as a public view cannot be combined if they have different

¹⁰↑It is a bit misleading to call the operation \sqcup “join”, as it only computes an upper bound, not necessarily the lowest one. However, it *should* compute a supremum for performance reasons: intuitively, larger public view means more padding.

$$\begin{array}{c}
\frac{\theta \in \{1, \widehat{\mathbb{B}}\}}{\lambda\theta \triangleright \lambda\widehat{\mathbf{b}} \ x \ y \Rightarrow \text{mux} \ \widehat{\mathbf{b}} \ x \ y} \\
\\
\frac{\lambda\theta_1 \triangleright \widehat{\text{ite}}_1 \quad \lambda\theta_2 \triangleright \widehat{\text{ite}}_2}{\lambda\theta_1 \times \theta_2 \triangleright \lambda\widehat{\mathbf{b}} \ x \ y \Rightarrow (\widehat{\text{ite}}_1 \ \widehat{\mathbf{b}} \ (\pi_1 \ x) \ (\pi_1 \ y), \widehat{\text{ite}}_2 \ \widehat{\mathbf{b}} \ (\pi_2 \ x) \ (\pi_2 \ y))} \\
\\
\frac{\lambda\theta_2 \triangleright \widehat{\text{ite}}_2}{\lambda\theta_1 \rightarrow \theta_2 \triangleright \lambda\widehat{\mathbf{b}} \ x \ y \Rightarrow \lambda z \Rightarrow \widehat{\text{ite}}_2 \ \widehat{\mathbf{b}} \ (x \ z) \ (y \ z)} \\
\\
\frac{(\widehat{\mathbb{T}}, \sqcup, \uparrow) \in \mathcal{S}_{\sqcup}}{\lambda\Psi\widehat{\mathbb{T}} \triangleright \lambda\widehat{\mathbf{b}} \ x \ y \Rightarrow \text{let } k = \pi_1 \ x \sqcup \pi_1 \ y \text{ in } \langle k, \text{mux} \ \widehat{\mathbf{b}} \ (\uparrow (\pi_1 \ x) \ k \ (\pi_2 \ x)) \ (\uparrow (\pi_1 \ y) \ k \ (\pi_2 \ y)) \rangle}
\end{array}$$

Figure 6.11. Mergeability

lengths. If such lists are the branches of an oblivious conditional, lifting will either fail or coerce both to an OADT with a join-structure.

Join structures induce a *mergeability* relation, defined in Figure 6.11, that can be used to decide if a specification type can be used in oblivious conditionals. We say θ is *mergeable* if $\lambda\theta \triangleright \widehat{\text{ite}}$, with witness $\widehat{\text{ite}}$ of type $\widehat{\mathbb{B}} \rightarrow \theta \rightarrow \theta \rightarrow \theta$. We will write $\lambda\theta$ when we do not care about the witness. This witness can be thought of as a generalized, drop-in replacement of **mux**: we simply translate **mux** to the derived $\widehat{\text{ite}}$ if the result type is mergeable. The case of Ψ -type captures this intuition: we first join the public views, and reshape all branches to this common public view, before we select the correct one privately using **mux**. This rule looks up the necessary methods from the context of join structures \mathcal{S}_{\sqcup} . Other cases are straightforward: we simply fall back to **mux** for primitive types, and the derivation for product and function types are done congruently.

6.3.3 Introduction and Elimination Structures

An ADT is manipulated by its introduction and elimination forms. To successfully lift a public program using ADTs, we need structures to explain how the primitive operations

of its ADTs are handled in their OADT counterparts. Thus, an OADT \widehat{T} can *optionally* be equipped with an *introduction-structure* (intro-structure) and an *elimination-structure* (elim-structure), defined below. These structures are optional because some programs only consume ADTs, without constructing any new ADT values (and vice versa): a function that checks membership in a list only requires an elim-structure on lists, for example. Intuitively, the axioms of these structures require the introduction and elimination methods of an OADT to behave like those of the corresponding ADT. This is formalized using a pair of logical refinement relations on values ($\mathcal{V}_n[\cdot]$) and expressions ($\mathcal{E}_n[\cdot]$); these relations are formally defined in [Section 6.3.6](#).

Definition 6.3.3 (intro-structure). An intro-structure of an OADT \widehat{T} for ADT T , with global definition `data T = \overline{C} η` , consists of a set of functions \widehat{C}_i , each corresponding to a constructor C_i . The type of \widehat{C}_i is $\theta_i \rightarrow \Psi\widehat{T}$, where $[\theta_i] = \eta_i$ (note that DT-ADT guarantees that η_i is a simple type). The particular θ_i an intro-structure uses is determined by the author of that structure.

Each \widehat{C}_i is required to logically refine the corresponding constructor (A-I₁): given any values $v : [\theta]$ and $v' : \theta$, if $(v, v') \in \mathcal{V}_n[\theta]$, then $(C_i\ v, \widehat{C}_i\ v') \in \mathcal{E}_n[\Psi\widehat{T}]$.

Definition 6.3.4 (elim-structure). An elim-structure of an OADT \widehat{T} for ADT T , with global definition `data T = \overline{C} η` , consists of a family of functions $\widehat{\text{match}}_\alpha$, indexed by the possible return types. The type of $\widehat{\text{match}}_\alpha$ is $\Psi\widehat{T} \rightarrow (\overline{\theta} \rightarrow \alpha) \rightarrow \alpha$, where $[\theta_i] = \eta_i$ for each θ_i in the function arguments corresponding to alternatives.

Each $\widehat{\text{match}}_\alpha$ is required to logically refine the pattern matching expression, specialized with ADT T and return type α . The sole axiom of this structure (A-E₁) only considers return type α being a specification type: given values $v_i : \eta_i$, $\langle k, \widehat{v} \rangle : \widehat{T}\ k$, $\lambda x \Rightarrow e_i : [\theta_i] \rightarrow [\alpha]$ and $\lambda x \Rightarrow e'_i : \theta_i \rightarrow \alpha$, if $r\ k\ \widehat{v} \rightarrow^* C_i\ v_i$ and $(\lambda x \Rightarrow e_i, \lambda x \Rightarrow e'_i) \in \mathcal{V}_n[\theta_i \rightarrow \alpha]$ then $([v_i/x] e_i, \widehat{\text{match}}\ \langle k, \widehat{v} \rangle\ (\overline{\lambda x \Rightarrow e'_i})) \in \mathcal{E}_n[\alpha]$.

The types of the oblivious introduction and elimination forms in these structures are only required to be compatible with the public counterparts. The programmers can choose which specific OADTs to use according to their desired privacy policy. [Figure 6.4](#) shows the constructors and pattern matching functions for $\widehat{\text{list}}_{\leq}$.

$$\boxed{\theta \rightsquigarrow \theta' \triangleright \uparrow}$$

$$\begin{array}{c}
\frac{}{\theta \rightsquigarrow \theta \triangleright \lambda x \Rightarrow x} \quad \frac{}{\mathbb{B} \rightsquigarrow \widehat{\mathbb{B}} \triangleright \lambda x \Rightarrow \widehat{\mathbb{B}} \# s \ x} \quad \frac{(\widehat{\mathbb{T}}, \mathbb{T}, s, r, \nu, \preceq) \in \mathcal{S}_\omega}{\mathbb{T} \rightsquigarrow \Psi \widehat{\mathbb{T}} \triangleright \lambda x \Rightarrow \langle \nu \ x, s \ (\nu \ x) \ x \rangle} \\
\\
\frac{\uparrow : \Psi \widehat{\mathbb{T}} \rightarrow \Psi \widehat{\mathbb{T}}' \in \mathcal{S}_\uparrow}{\Psi \widehat{\mathbb{T}} \rightsquigarrow \Psi \widehat{\mathbb{T}}' \triangleright \uparrow} \quad \frac{\theta_1 \rightsquigarrow \theta'_1 \triangleright \uparrow_1 \quad \theta_2 \rightsquigarrow \theta'_2 \triangleright \uparrow_2}{\theta_1 \times \theta_2 \rightsquigarrow \theta'_1 \times \theta'_2 \triangleright \lambda x \Rightarrow (\uparrow_1(\pi_1 \ x), \uparrow_2(\pi_2 \ x))} \\
\\
\frac{\theta'_1 \rightsquigarrow \theta_1 \triangleright \uparrow_1 \quad \theta_2 \rightsquigarrow \theta'_2 \triangleright \uparrow_2}{\theta_1 \rightarrow \theta_2 \rightsquigarrow \theta'_1 \rightarrow \theta'_2 \triangleright \lambda x \Rightarrow \lambda y \Rightarrow \uparrow_2(x \ (\uparrow_1 y))}
\end{array}$$

Figure 6.12. Coercion

The elim-structure of an OADT consists of a family of destructors, whose return type α does not necessarily range over all types. For example, $\widehat{\text{match}}_\alpha$ of $\widehat{\text{list}}_\leq$, $\widehat{\text{list}}_\leq \# \text{match}$ in Figure 6.4, requires α to be a mergeable type, due to the use of $\widehat{\text{match}}$, which imposes a restriction similarly to mux . Such constraints on α are automatically inferred and enforced.

6.3.4 Coercion Structures

As discussed in Section 6.1, we may need to convert an oblivious type to another, either due to a mismatch from input to output, or due to its lack of certain structures. For example, $\widehat{\text{list}}_=$ does not have join structure, so if the branches of an oblivious conditional has type $\Psi \widehat{\text{list}}_=$, they should be coerced to $\Psi \widehat{\text{list}}_\leq$, when such a coercion is available.

Two compatible OADTs may form a *coercion-structure*, shown below.

Definition 6.3.5 (coercion-structure). A coercion-structure of a pair of compatible OADTs $\widehat{\mathbb{T}}$ and $\widehat{\mathbb{T}}'$ for \mathbb{T} , with public view type τ and τ' respectively, consists of a coercion function \uparrow of type $\Psi \widehat{\mathbb{T}} \rightarrow \Psi \widehat{\mathbb{T}}'$.

The coercion should produce an equivalent value (A-C₁): given values $v : \mathbb{T}$, $\langle k, \widehat{v} \rangle : \Psi \widehat{\mathbb{T}}$ and $\langle k', \widehat{v}' \rangle : \Psi \widehat{\mathbb{T}}'$, if $r \ k \ \widehat{v} \longrightarrow^* v$ and $\uparrow \langle k, \widehat{v} \rangle \longrightarrow^* \langle k', \widehat{v}' \rangle$, then $r \ k' \ \widehat{v}' \longrightarrow^* v$.

This structure only defines the coercion between two Ψ -types. Figure 6.12 generalizes the coercion relation to any (compatible) specification types. We say θ is *coercible* to θ' if $\theta \rightsquigarrow \theta' \triangleright \uparrow$, with witness \uparrow of type $\theta \rightarrow \theta'$. We may write $\theta \rightsquigarrow \theta'$ when we do not care about

the witness. The rules of this relation are straightforward. The context of coercion structures \mathcal{S}_\dagger and the context of OADT structures \mathcal{S}_ω are used to look up the necessary methods in the corresponding rules. The rule for coercing a function type is contravariant. Note that we can always coerce a public type to an OADT by running the section function, and the public view can be selected by the view function in the OADT structure.

6.3.5 Declarative Lifting

With these Ψ -structures, we define a declarative lifting relation, which describes what the lifting procedure is allowed to derive at a high level. This lifting relation is given by the judgment $\mathcal{S}; \mathcal{L}; \Sigma; \Gamma \vdash e : \theta \triangleright \dot{e}$. It is read as the expression e of type $[\theta]$ is lifted to the expression \dot{e} of target type θ , under various contexts. The Ψ -structure context \mathcal{S} consists of the set of OADT-structures (\mathcal{S}_ω), join-structures (\mathcal{S}_\sqcup), intro-structures (\mathcal{S}_I), elim-structures (\mathcal{S}_E) and coercion-structures (\mathcal{S}_\dagger), respectively. The global definition context Σ is the same as the one used in the typing relation. The local context Γ is also similar to the one in the typing relation, but it keeps track of the target types of local variables instead of source types. Finally, the lifting context \mathcal{L} consists of entries of the form $x : \theta \triangleright \dot{x}$, which associates the global function x of type $[\theta]$ with a generated function \dot{x} of the target type θ . A single global function may have multiple target types, i.e., multiple private versions, either specified by the users or by the callsites. For example, \mathcal{L} may contain $\text{filter} : \widehat{\Psi\text{list}}_{\leq} \rightarrow \widehat{\mathbb{Z}} \rightarrow \widehat{\Psi\text{list}}_{\leq} \triangleright \widehat{\text{filter}}_1$ and $\text{filter} : \widehat{\Psi\text{list}}_{=} \rightarrow \widehat{\mathbb{Z}} \rightarrow \widehat{\Psi\text{list}}_{\leq} \triangleright \widehat{\text{filter}}_2$.

Figure 6.13 shows a selection of rules of the declarative lifting relation (the full rules are in Appendix B.1). We elide most contexts as they are fixed, and simply write $\Gamma \vdash e : \theta \triangleright \dot{e}$ for brevity. Most rules are simply congruences and similar to typing rules. L-FUN outsources the lifting of a function call to the lifting context. L-IF₂ handles the case when the condition is lifted to an oblivious boolean by delegating the translation to the mergeability relation. Similarly, L-CTOR₂ and L-MATCH₂ query the contexts of the intro-structures and elim-structures, and use the corresponding instances as the drop-in replacement, when we are constructing or destructing Ψ -types. Lastly, L-COERCE coerces an expression nondeterministically using the coercion relation.

$$\boxed{\Gamma \vdash e : \theta \triangleright \dot{e}}$$

$$\begin{array}{c}
\text{L-LIT} \\
\frac{}{\Gamma \vdash b : \mathbb{B} \triangleright b}
\end{array}
\quad
\begin{array}{c}
\text{L-VAR} \\
\frac{x : \theta \in \Gamma}{\Gamma \vdash x : \theta \triangleright x}
\end{array}
\quad
\begin{array}{c}
\text{L-FUN} \\
\frac{x : \theta \triangleright \dot{x} \in \mathcal{L}}{\Gamma \vdash x : \theta \triangleright \dot{x}}
\end{array}
\quad
\begin{array}{c}
\text{L-ABS} \\
\frac{x : \theta_1, \Gamma \vdash e : \theta_2 \triangleright \dot{e}}{\Gamma \vdash \lambda x : [\theta_1] \Rightarrow e : \theta_1 \rightarrow \theta_2 \triangleright \lambda x : \theta_1 \Rightarrow \dot{e}}
\end{array}$$

$$\begin{array}{c}
\text{L-APP} \\
\frac{\Gamma \vdash e_2 : \theta_1 \rightarrow \theta_2 \triangleright \dot{e}_2 \quad \Gamma \vdash e_1 : \theta_1 \triangleright \dot{e}_1}{\Gamma \vdash e_2 \ e_1 : \theta_2 \triangleright \dot{e}_2 \ \dot{e}_1}
\end{array}$$

$$\begin{array}{c}
\text{L-LET} \\
\frac{\Gamma \vdash e_1 : \theta_1 \triangleright \dot{e}_1 \quad x : \theta_1, \Gamma \vdash e_2 : \theta_2 \triangleright \dot{e}_2}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \theta_2 \triangleright \text{let } x = \dot{e}_1 \text{ in } \dot{e}_2}
\end{array}$$

$$\begin{array}{c}
\text{L-IF}_1 \\
\frac{\Gamma \vdash e_0 : \mathbb{B} \triangleright \dot{e}_0 \quad \Gamma \vdash e_1 : \theta \triangleright \dot{e}_1 \quad \Gamma \vdash e_2 : \theta \triangleright \dot{e}_2}{\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : \theta \triangleright \text{if } \dot{e}_0 \text{ then } \dot{e}_1 \text{ else } \dot{e}_2}
\end{array}$$

$$\begin{array}{c}
\text{L-IF}_2 \\
\frac{\Gamma \vdash e_0 : \widehat{\mathbb{B}} \triangleright \dot{e}_0 \quad \Gamma \vdash e_1 : \theta \triangleright \dot{e}_1 \quad \Gamma \vdash e_2 : \theta \triangleright \dot{e}_2}{\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : \theta \triangleright \widehat{\text{ite}} \ \dot{e}_0 \ \dot{e}_1 \ \dot{e}_2}
\end{array}
\quad
\begin{array}{c}
\text{L-CTOR}_1 \\
\frac{\text{data } \mathbb{T} = \overline{\mathbb{C}} \ \overline{\eta} \in \Sigma \quad \Gamma \vdash e : \eta_i \triangleright \dot{e}}{\Gamma \vdash \mathbb{C}_i \ e : \mathbb{T} \triangleright \mathbb{C}_i \ \dot{e}}
\end{array}$$

$$\begin{array}{c}
\text{L-CTOR}_2 \\
\frac{\widehat{\mathbb{C}}_i : \theta_i \rightarrow \Psi \widehat{\mathbb{T}} \in \mathcal{S}_I \quad \Gamma \vdash e : \theta_i \triangleright \dot{e}}{\Gamma \vdash \mathbb{C}_i \ e : \Psi \widehat{\mathbb{T}} \triangleright \widehat{\mathbb{C}}_i \ \dot{e}}
\end{array}
\quad
\begin{array}{c}
\text{L-MATCH}_1 \\
\frac{\text{data } \mathbb{T} = \overline{\mathbb{C}} \ \overline{\eta} \in \Sigma \quad \Gamma \vdash e_0 : \mathbb{T} \triangleright \dot{e}_0 \quad \forall i. x : \eta_i, \Gamma \vdash e_i : \theta' \triangleright \dot{e}_i}{\Gamma \vdash \text{match } e_0 \text{ with } \overline{\mathbb{C}} \ x \Rightarrow e : \theta' \triangleright \text{match } \dot{e}_0 \text{ with } \overline{\mathbb{C}} \ x \Rightarrow \dot{e}_i}
\end{array}$$

$$\begin{array}{c}
\text{L-MATCH}_2 \\
\frac{\widehat{\text{match}} : \Psi \widehat{\mathbb{T}} \rightarrow (\theta \rightarrow \theta') \rightarrow \theta' \in \mathcal{S}_E \quad \Gamma \vdash e_0 : \Psi \widehat{\mathbb{T}} \triangleright \dot{e}_0 \quad \forall i. x : \theta_i, \Gamma \vdash e_i : \theta' \triangleright \dot{e}_i}{\Gamma \vdash \text{match } e_0 \text{ with } \overline{\mathbb{C}} \ x \Rightarrow e : \theta' \triangleright \widehat{\text{match}} \ \dot{e}_0 \ (\lambda x : \theta \Rightarrow \dot{e}_i)}
\end{array}$$

$$\begin{array}{c}
\text{L-COERCE} \\
\frac{\Gamma \vdash e : \theta \triangleright \dot{e} \quad \theta \mapsto \theta' \triangleright \uparrow}{\Gamma \vdash e : \theta' \triangleright \uparrow \dot{e}}
\end{array}$$

Figure 6.13. Selected declarative lifting rules

This lifting relation in [Figure 6.13](#) only considers one expression. In practice, the users specify a set of functions and their target types to lift. The result of our lifting procedure is a lifting context \mathcal{L} which maps these functions and target types to the corresponding generated functions, as well as any other functions and the inferred target types that these functions depend on. The global context Σ is also extended with the definitions of the generated functions. To make this more clear, we say a lifting context is *derivable*, denoted by $\vdash \mathcal{L}$, if and only if, for any $\mathbf{x} : \theta \triangleright \dot{\mathbf{x}} \in \mathcal{L}$, $\mathbf{fn} \ \mathbf{x} : [\theta] = \mathbf{e} \in \Sigma$ and $\mathbf{fn} \ \dot{\mathbf{x}} : \theta = \dot{\mathbf{e}} \in \Sigma$ for some \mathbf{e} and $\dot{\mathbf{e}}$, such that $\mathcal{S}; \mathcal{L}; \Sigma; \cdot \vdash \mathbf{e} : \theta \triangleright \dot{\mathbf{e}}$. In other words, any definitions of the lifted functions in \mathcal{L} can be derived from the lifting relation in [Figure 6.13](#). Note that the derivation of a function definition is under a lifting context with possibly an entry of this function itself. This is similar to the role of global context in type checking, as TAYPSI supports mutually recursive functions. The goal of our algorithm ([Section 6.4](#)) is then to find such a derivable lifting context that includes the user-specified liftings.

6.3.6 Logical Refinement

The correctness of the lifting procedure is framed as a *logical refinement* between expressions of specification types and those of simple types; this relationship is defined as a step-indexed logical relation [61]. As is common, this relation is defined via a pair of set-valued type denotations: a value interpretation $\mathcal{V}_n \llbracket \theta \rrbracket$ and an expression interpretation $\mathcal{E}_n \llbracket \theta \rrbracket$. We say an expression \mathbf{e}' of type θ refines \mathbf{e} of type $[\theta]$ (within n steps) if $(\mathbf{e}, \mathbf{e}') \in \mathcal{E}_n \llbracket \theta \rrbracket$. In other words, \mathbf{e}' preserves the behavior of \mathbf{e} , in that if \mathbf{e}' terminates at a value, \mathbf{e} must terminate at an equivalent value. The equivalence between values is dictated by $\mathcal{V}_n \llbracket \theta \rrbracket$.

[Figure 6.14](#) shows the complete definition of the logical relation. All pairs in the relations must be closed and well-typed, i.e., their interpretations have the forms:

$$\begin{aligned} \mathcal{V}_n \llbracket \theta \rrbracket &= \{ (\mathbf{v}, \mathbf{v}') \mid \cdot \vdash \mathbf{v} : [\theta] \wedge \cdot \vdash \mathbf{v}' : \theta \wedge \dots \} \\ \mathcal{E}_n \llbracket \theta \rrbracket &= \{ (\mathbf{e}, \mathbf{e}') \mid \cdot \vdash \mathbf{e} : [\theta] \wedge \cdot \vdash \mathbf{e}' : \theta \wedge \dots \} \end{aligned}$$

For brevity, we leave this requirement implicit in [Figure 6.14](#).

$\mathcal{V}_n[\theta]$

$$\begin{aligned}
\mathcal{V}_n[\mathbf{1}] &= \mathcal{V}_n[\mathbf{B}] = \mathcal{V}_n[\mathbf{T}] = \{ (v, v') \mid 0 < n \implies v = v' \} \\
\mathcal{V}_n[\widehat{\mathbf{B}}] &= \{ (b, [b']) \mid 0 < n \implies b = b' \} \\
\mathcal{V}_n[\widehat{\Psi\mathbf{T}}] &= \{ (v, \langle k, \widehat{v} \rangle) \mid 0 < n \implies r \ k \ \widehat{v} \longrightarrow^* v \} \\
\mathcal{V}_n[\theta_1 \times \theta_2] &= \{ ((v_1, v_2), (v'_1, v'_2)) \mid (v_1, v'_1) \in \mathcal{V}_n[\theta_1] \wedge (v_2, v'_2) \in \mathcal{V}_n[\theta_2] \} \\
\mathcal{V}_n[\theta_1 \rightarrow \theta_2] &= \left\{ (\lambda x : [\theta_1] \Rightarrow e, \lambda x : \theta_1 \Rightarrow e') \mid \begin{array}{l} \forall i < n. \forall (v, v') \in \mathcal{V}_i[\theta_1]. \\ ([v/x]e, [v'/x]e') \in \mathcal{E}_i[\theta_2] \end{array} \right\}
\end{aligned}$$

 $\mathcal{E}_n[\theta]$

$$\mathcal{E}_n[\theta] = \{ (e, e') \mid \forall i < n. \forall v'. e' \longrightarrow^i v' \implies \exists v. e \longrightarrow^* v \wedge (v, v') \in \mathcal{V}_{n-i}[\theta] \}$$

Figure 6.14. A logical relation for refinement

The definitions are mostly standard. The most interesting case is the value interpretation of Ψ -type: we say the pair of a public view and an oblivious value of an OADT is equivalent to a public value of the corresponding ADT when the oblivious value can be retracted to the public value. Intuitively, an encrypted value is equivalent to the value it decrypts to. The base cases of the value interpretation are also guarded by the condition that we still have steps left, i.e., greater than 0. This requirement maintains the pleasant property that the interpretations $\mathcal{V}_0[\theta]$ and $\mathcal{E}_0[\theta]$ are total relations on closed values and expressions, respectively, of type θ . The proof also uses a straightforward interpretation of typing context, $\mathcal{G}_n[\Gamma]$, whose definition is in [Appendix B.1](#).

This relation also gives rise to a semantic characterization of the lifting context. We say a lifting context is n -valid, denoted by $\models_n \mathcal{L}$, if and only if, for any $x : \theta \triangleright \dot{x} \in \mathcal{L}$, $(x, \dot{x}) \in \mathcal{E}_n[\theta]$. If $\models_n \mathcal{L}$ for any n , we say \mathcal{L} is valid, denoted by $\models \mathcal{L}$. The validity is essentially a semantic correctness of \mathcal{L} .

6.3.7 Metatheory of Lifting

The first key property of the lifting relation is well-typedness, which guarantees the security of translated programs, thanks to [Theorem 6.2.1](#).

Theorem 6.3.1 (Regularity of declarative lifting). *Suppose \mathcal{L} is well-typed and $\mathcal{S}; \mathcal{L}; \Sigma; \Gamma \vdash e : \theta \triangleright \dot{e}$. We have $\Sigma; [\Gamma] \vdash e : [\theta]$ and $\Sigma; \Gamma \vdash \dot{e} : \theta$.*

Our lifting relation ensures that lifted expressions refine source expressions in fewer than n steps, as long as every lifted program in \mathcal{L} is also semantically correct in fewer than n steps. As is common in logical relation proofs, this proof requires a more general theorem about open terms.

Theorem 6.3.2 (Correctness of declarative lifting of closed terms). *Suppose $\mathcal{S}; \mathcal{L}; \Sigma; \cdot \vdash e : \theta \triangleright \dot{e}$ and $\vDash_n \mathcal{L}$. We have $(e, \dot{e}) \in \mathcal{E}_n[[\theta]]$.*

Finally, [Theorem 6.3.3](#) provides a strong result of the correctness of our translation. Any lifting context that is derived using the rules of [Figure 6.13](#) is semantically correct. In other words, if every pair of source program and lifted program in \mathcal{L} are in our lifting relation, they also satisfy our refinement criteria.

Theorem 6.3.3 (Correctness of declarative lifting). *$\vdash \mathcal{L}$ implies $\vDash \mathcal{L}$.*

Our notion of logical refinement only provides partial correctness guarantees, as can be seen in the definition of $\mathcal{E}_n[[\cdot]]$. As a result, the lifting relation does not guarantee equi-termination: it is possible that a lifted program will diverge when the source program terminates. This can occur when an `if` is replaced by a `mux`: since the latter fully executes both branches, this effectively changes the semantics of a conditional from a lazy evaluation strategy to an eager strategy. Using a public value to bound the recursion depth in order to guarantee termination is a common practice in data-oblivious computation, for the reasons discussed in [Section 6.1](#). While the public view of an OADT naturally serves as a measure in many cases, including all of the case studies and benchmarks in our evaluation, in theory it is possible for a user to provide a policy to a function that results in a nonterminating lifted version. In

this situation, users must either specify a different policy, or rewrite the functions to recurse on a different argument, e.g., a fuel value.

All the proofs in this section are available in [Appendix B.3](#).

6.4 Algorithmic Lifting

[Figure 6.15](#) presents the overall workflow of our lifting algorithm. This algorithm starts with a set of *goals*, i.e., pairs of source functions tagged with the `%lift` keywords and their desired specification types. We then run our lifting algorithm on all the functions in these goals, as well as any functions they depend on, transforming each of these functions to an oblivious version parameterized by *typed macros* and type variables, along with a set of constraints over these type variables. After solving the constraints, we obtain a set of type assignments for each function. Note that a single function may have multiple type assignments, one for each occurrence in a goal and callsite. For example, `filter` may have the type assignment for the goal $\Psi\widehat{\text{list}}_{\leq} \rightarrow \widehat{\mathbb{Z}} \rightarrow \Psi\widehat{\text{list}}_{\leq}$ generated by `%lift`, and the assignment for $\Psi\widehat{\text{list}}_{\leq} \rightarrow \mathbb{Z} \rightarrow \Psi\widehat{\text{list}}_{\leq}$ generated by the call in `filter5` from [Section 6.1](#). Finally, we generate the private versions of all the lifted functions by instantiating their type variables and expanding away any macros. The lifting context from the last section is simply these lifted functions and their generated private versions.

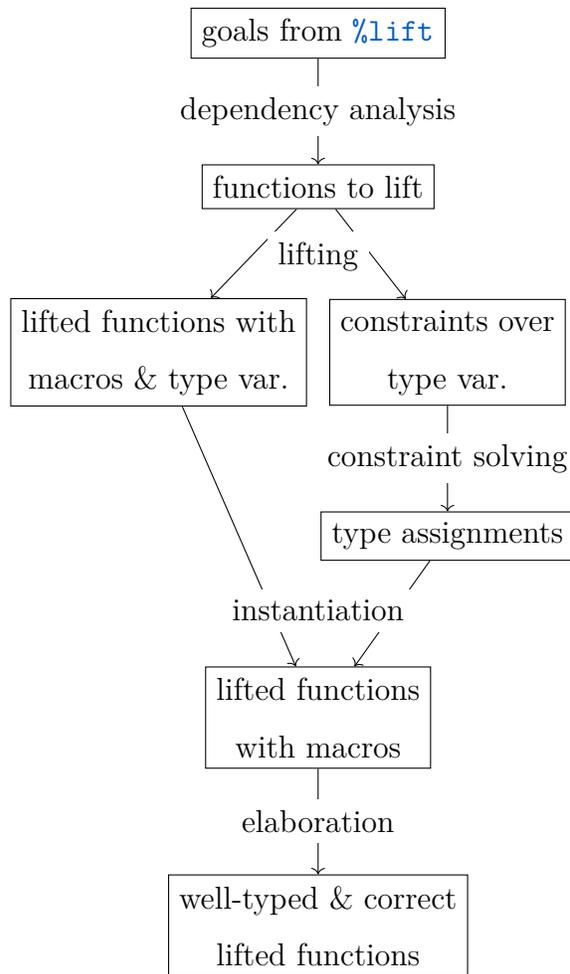


Figure 6.15. Translation pipeline

The lifting algorithm is defined using the judgment $\Sigma; \Gamma \vdash e : \eta \sim X \triangleright \dot{e} \mid \mathcal{C}$. It reads as the source expression e of type η is lifted to the target expression \dot{e} whose type is a *type variable* X as a placeholder for the specification type, and generates constraints \mathcal{C} . The source expression e is required to be in *administrative normal form* (ANF) [51], which is guaranteed by our type checker. In particular, type annotations are added to `let`-bindings, and the body of every `let` is either another `let` or a variable. Importantly, this means the last expression of a sequence of `let` must be a variable. The output of this algorithm is an expression \dot{e} containing macros (which will be discussed shortly), and the constraints \mathcal{C} . Unlike the declarative rules, this algorithm keeps track of the source type η , which is used to restrict the range of the type variables. Consequently, every entry of the typing context Γ has the form $x : \eta \sim X$, meaning that local variable x has type η in the source program and type X in the target program. For example, after the lifting algorithm has processed the function arguments of `filter` in Figure 6.1, the typing context contains entries $xs : \text{list} \sim X_1$ and $y : \mathbb{Z} \sim X_2$, with freshly generated type variables X_1 and X_2 .

The typed macros, defined in Figure 6.16, are an essential part of the output of the lifting algorithm, and permit a form of ad-hoc polymorphism, that allows the algorithm to cleanly separate constraint solving from program generation. These macros take types as parameters and elaborate to expressions, under the contexts \mathcal{S} , \mathcal{L} and Σ implicitly. These macros are effectively thin “wrappers” of their corresponding language constructs and the previously defined relations. The conditional macro `%ite`, for example, corresponds to the `if` expression, but the condition may be oblivious. The constructor macro `%C` is a “smart” constructor that may construct a Ψ -type. The pattern matching macro `%match` is similar to `%C` but for eliminating a type compatible with an ADT. Lastly, `%↑` and `%x` is simply a direct wrapper of the mergeable relation and the lifting context \mathcal{L} , respectively. Note that the derivation of these macro are completely determined by the type parameters.

Figure 6.17 defines the constraints used in the algorithm, where θ^+ is the specification types extended with type variables. The constraint $X \in [\eta]$ means type variable X belongs to the compatibility class of η . In other words, $[X] = \eta$. Each macro is accompanied by a constraint on its type parameters. These constraints mean that the corresponding macros are resolvable. More formally, this means they can elaborate to some expressions according to the

$$\begin{array}{c}
\boxed{\%ite(\theta_0, \theta; e_0, e_1, e_2) \triangleright \dot{e}} \\
\frac{}{\%ite(\mathbb{B}, \theta; e_0, e_1, e_2) \triangleright \text{if } e_0 \text{ then } e_1 \text{ else } e_2} \quad \frac{\uparrow \theta \triangleright \widehat{ite}}{\%ite(\widehat{\mathbb{B}}, \theta; e_0, e_1, e_2) \triangleright \widehat{ite} \ e_0 \ e_1 \ e_2} \\
\boxed{\%C(\theta, \theta'; e) \triangleright \dot{e}} \\
\frac{\text{data } T = \overline{C} \ \eta \in \Sigma}{\%C_i(\eta_i, T; e) \triangleright C_i \ e} \quad \frac{\widehat{C}_i : \theta_i \rightarrow \Psi \widehat{T} \in \mathcal{S}_I}{\%C_i(\theta_i, \Psi \widehat{T}; e) \triangleright \widehat{C}_i \ e} \\
\boxed{\%match(\theta_0, \bar{\theta}, \theta'; e_0, \bar{e}) \triangleright \dot{e}} \\
\frac{\text{data } T = \overline{C} \ \eta \in \Sigma}{\%match(T, \bar{\eta}, \theta'; e_0, \bar{e}) \triangleright \text{match } e_0 \ \text{with } \overline{C} \ x \Rightarrow e} \\
\frac{\widehat{match} : \Psi \widehat{T} \rightarrow (\overline{\theta \rightarrow \theta'}) \rightarrow \theta' \in \mathcal{S}_E}{\%match(\Psi \widehat{T}, \bar{\theta}, \theta'; e_0, \bar{e}) \triangleright \widehat{match} \ e_0 \ (\lambda x : \theta \Rightarrow e)} \\
\boxed{\%\uparrow(\theta, \theta'; e) \triangleright \dot{e}} \quad \boxed{\%x(\theta) \triangleright \dot{e}} \\
\frac{\theta \mapsto \theta' \triangleright \uparrow}{\%\uparrow(\theta, \theta'; e) \triangleright \uparrow e} \quad \frac{x : \theta \triangleright \dot{x} \in \mathcal{L}}{\%x(\theta) \triangleright \dot{x}}
\end{array}$$

Figure 6.16. Typed macros

CONSTRAINTS

$$\begin{array}{l}
c ::= x \in [\eta] \mid \theta^+ = \theta^+ \\
\mid \ \%ite(\theta^+, \theta^+) \mid \%C(\theta^+, \theta^+) \mid \%match(\theta^+, \bar{\theta}^+, \theta^+) \mid \%\uparrow(\theta^+, \theta^+) \mid \%x(\theta^+)
\end{array}$$

Figure 6.17. Constraints

rules in [Figure 6.16](#) for any expression arguments. As a result, after solving all constraints and concretizing the type variables, all macros in the lifted expression \dot{e} can be fully elaborated away.

[Figure 6.18](#) shows a selection of lifting algorithm rules (the full rules are in [Appendix B.2](#)). Coercions only happen when we lift variables, as in A-VAR. This works because the source program is in ANF, so each expression is bound to a variable which has the opportunity to get coerced. For example, the argument to a function or constructor, in A-APP and A-CTOR, is always a variable in ANF, and recursively lifting it allows the application of

$$\boxed{\Gamma \vdash e : \eta \sim X \triangleright \dot{e} \mid \mathcal{C}}$$

$$\text{A-LIT} \quad \frac{}{\Gamma \vdash b : \mathbb{B} \sim X \triangleright b \mid X = \mathbb{B}}$$

$$\text{A-VAR} \quad \frac{x : \eta \sim X \in \Gamma}{\Gamma \vdash x : \eta \sim X' \triangleright \% \uparrow(X, X'; x) \mid \% \uparrow(X, X')}$$

$$\text{A-FUN} \quad \frac{\text{fn } x : \eta = e \in \Sigma}{\Gamma \vdash x : \eta \sim X \triangleright \% x(X) \mid \% x(X)}$$

$$\text{A-ABS} \quad \frac{X_1, X_2 \text{ fresh} \quad x : \eta_1 \sim X_1, \Gamma \vdash e : \eta_2 \sim X_2 \triangleright \dot{e} \mid \mathcal{C}}{\Gamma \vdash \lambda x : \eta_1 \Rightarrow e : \eta_1 \rightarrow \eta_2 \sim X \triangleright \lambda x : X_1 \Rightarrow \dot{e} \mid X_1 \in [\eta_1], X_2 \in [\eta_2], X = X_1 \rightarrow X_2, \mathcal{C}}$$

$$\text{A-APP} \quad \frac{X_1 \text{ fresh} \quad x_2 : \eta_1 \rightarrow \eta_2 \sim X \in \Gamma \quad \Gamma \vdash x_1 : \eta_1 \sim X_1 \triangleright \dot{e}_1 \mid \mathcal{C}}{\Gamma \vdash x_2 \ x_1 : \eta_2 \sim X_2 \triangleright x_2 \ \dot{e}_1 \mid X_1 \in [\eta_1], X = X_1 \rightarrow X_2, \mathcal{C}}$$

$$\text{A-LET} \quad \frac{X_1 \text{ fresh} \quad \Gamma \vdash e_1 : \eta_1 \sim X_1 \triangleright \dot{e}_1 \mid \mathcal{C}_1 \quad x : \eta_1 \sim X_1, \Gamma \vdash e_2 : \eta_2 \sim X_2 \triangleright \dot{e}_2 \mid \mathcal{C}_2}{\Gamma \vdash \text{let } x : \eta_1 = e_1 \text{ in } e_2 : \eta_2 \sim X_2 \triangleright \text{let } x : X_1 = \dot{e}_1 \text{ in } \dot{e}_2 \mid X_1 \in [\eta_1], \mathcal{C}_1, \mathcal{C}_2}$$

$$\text{A-IF} \quad \frac{x_0 : \mathbb{B} \sim X_0 \in \Gamma \quad \Gamma \vdash e_1 : \eta \sim X \triangleright \dot{e}_1 \mid \mathcal{C}_1 \quad \Gamma \vdash e_2 : \eta \sim X \triangleright \dot{e}_2 \mid \mathcal{C}_2}{\Gamma \vdash \text{if } x_0 \text{ then } e_1 \text{ else } e_2 : \eta \sim X \triangleright \% \text{ite}(X_0, X; x_0, \dot{e}_1, \dot{e}_2) \mid \% \text{ite}(X_0, X), \mathcal{C}_1, \mathcal{C}_2}$$

$$\text{A-CTOR} \quad \frac{\text{data } T = \overline{C} \ \overline{\eta} \in \Sigma \quad X_i \text{ fresh} \quad \Gamma \vdash x : \eta_i \sim X_i \triangleright \dot{e} \mid \mathcal{C}}{\Gamma \vdash C_i \ x : T \sim X \triangleright \% C_i(X_i, X; \dot{e}) \mid X_i \in [\eta_i], \% C_i(X_i, X), \mathcal{C}}$$

$$\text{A-MATCH} \quad \frac{\overline{X} \text{ fresh} \quad x_0 : T \sim X_0 \in \Gamma \quad \text{data } T = \overline{C} \ \overline{\eta} \in \Sigma \quad \forall i. x : \eta_i \sim X_i, \Gamma \vdash e_i : \eta' \sim X' \triangleright \dot{e}_i \mid \mathcal{C}_i}{\Gamma \vdash \text{match } x_0 \text{ with } \overline{C} \ x \Rightarrow e : \eta' \sim X' \triangleright \% \text{match}(X_0, \overline{X}, X'; x_0, \overline{e}) \mid \overline{X} \in [\eta], \% \text{match}(X_0, \overline{X}, X'), \overline{C}}$$

Figure 6.18. Selected algorithmic lifting rules

A-VAR. On the other hand, the top-level program is always in `let`-binding form, whose last expression is always a variable too, allowing coercion of the whole program. However, not all variables are subject to coercions: the function x_2 in A-APP, the condition x_0 in A-IF and the discriminée x_0 in A-MATCH are kept as they are, for example. Coercing these variables would be unnecessary and undesirable. For example, coercing the condition in a conditional only makes the generated program more expensive: there is no reason to coerce from \mathbb{B} to $\widehat{\mathbb{B}}$, and use `mux` instead of `if`. Another key invariant we enforce in our algorithmic rules is that every fresh variable is “guarded” by a compatibility class constraint. For example, in A-ABS, the freshly generated variables X_1 and X_2 belong to the classes η_1 and η_2 , respectively. This constraint ensures that every type variable can be finitely enumerated, as every compatibility class is a finite set, bounded by the number of available OADTs. As a result, constraint solving in our context is decidable. Finally, if an expression is translated to a macro, a corresponding constraint is added to ensure this macro is resolvable.

We use the judgment $\mathcal{S}; \mathcal{L}; \Sigma; \sigma \models \mathcal{C}$ to mean the assignment σ satisfies a set of constraints \mathcal{C} , under the context of Ψ -structure, lifting context and global definition context. The constraints generated by our lifting algorithm use type variables X as placeholders for the target type of the function being lifted. To solve a goal with a particular target type θ , we add a constraint to \mathcal{C} that equates the placeholder with the stipulated type, i.e., $X = \theta$. Our constraint solver then attempts to find type assignments that satisfy the constraints in \mathcal{C} ; the resulting assignment is used to generate private versions of all the functions in the set of goals, as well as the accompanying lifting context.

6.4.1 Constraint Solving

At a high level, our solver reduces all constraints, except for function call constraints (`%x`), to quantifier-free formulas in a finite domain theory, which can be efficiently solved using an off-the-shelf solver. Function call constraints are recursively solved once their type arguments have been concretized by discharging the other constraints. When a function call constraint is unsatisfiable, we add a new refutation constraint and invoke the solver again to find a new instantiation of type parameters. As an example of this process, in

order to ascribe `filter` the type $\widehat{\Psi}_{\text{list}_=} \rightarrow \widehat{\mathbb{Z}} \rightarrow \widehat{\Psi}_{\text{list}_\leq}$, we first add the constraint $X = \widehat{\Psi}_{\text{list}_=} \rightarrow \widehat{\mathbb{Z}} \rightarrow \widehat{\Psi}_{\text{list}_\leq}$ to the constraints generated by the lifting algorithm $\cdot \vdash \dots : \text{list} \rightarrow \mathbb{Z} \rightarrow \text{list} \sim X \triangleright \dot{e} \mid \mathcal{C}$. Solving the other constraints may concretize the type variable of function call constraint `%filter`(X), i.e., the type of the recursive call to `filter`, to `%filter`($\widehat{\Psi}_{\text{list}_=} \rightarrow \widehat{\mathbb{Z}} \rightarrow \widehat{\Psi}_{\text{list}_\leq}$). Recursively solving this subgoal assuming the original goal is solved, i.e., extending the lifting context with the original goal, results in immediate success, as the subgoal is simply in the lifting context. On the other hand, if the type of the recursive call is instantiated as `%filter`($\widehat{\Psi}_{\text{list}_=} \rightarrow \widehat{\mathbb{Z}} \rightarrow \widehat{\Psi}_{\text{list}_=}$), the same constraints generated by lifting `filter` are solved, with an additional constraint $X = \widehat{\Psi}_{\text{list}_=} \rightarrow \widehat{\mathbb{Z}} \rightarrow \widehat{\Psi}_{\text{list}_=}$. However, this set of constraints is unsatisfiable, as $\widehat{\text{list}}_=$ has no join structure, so we add a refutation constraint to the context that forces the solver to not generate this assignment again. In general, the type of the recursive call to `filter` may be concretized to any types compatible with $\text{list} \rightarrow \mathbb{Z} \rightarrow \text{list}$. The number of such compatible types is bounded, as the number of arguments of this function and the number of OADTs are themselves bounded. The function `filter` has $3 \times 2 \times 3 = 18$ possible type assignments. In the worst case scenario, the algorithm eventually terminates after exhausting all 18 combinations.

More formally, the constraint solving algorithm uses two maps. \mathcal{F} maps function names to the generated constraints with a type variable X as the placeholder for the potential specification type, resulting from the lifting algorithm: every entry of \mathcal{F} has the form $f \mapsto (X, \mathcal{C}, \mathcal{C}')$. The generated constraints are partitioned into \mathcal{C} , which consists of all constraints except for function call constraints, and \mathcal{C}' , which consists of the functional call constraints (`%x`). Initially, \mathcal{F} consists of all functions collected from the keyword `%lift` and the functions they depend on. \mathcal{M} maps a pair of function name f and its target type θ , called a *goal*, to a type assignment σ : each entry has the form $(f, \theta) \mapsto \sigma$. Initially \mathcal{M} is empty.

The constraint solving algorithm takes an initial \mathcal{F} and \mathcal{M} , and a goal, and returns an updated \mathcal{M} that consists of the type assignments for the goal and all the subgoals this goal depends on. The constraint solver is applied to all goals from `%lift`, and the final \mathcal{M} is the union of all returned \mathcal{M} . With \mathcal{M} , we can generate functions from the type assignments for each goal. The global context is subsequently extended by these generated functions, and

Inputs: Constraint map \mathcal{F} , type assignment map \mathcal{M} , function name \mathbf{f} and target type θ

Output: Updated type assignment map \mathcal{M}'

```

function SOLVE( $\mathcal{F}$ ,  $\mathcal{M}$ ,  $\mathbf{f}$ ,  $\theta$ )
  if  $(\mathbf{f}, \theta) \mapsto \sigma \in \mathcal{M}$  then return  $\mathcal{M}$ 
   $(\mathbf{X}, \mathcal{C}, \mathcal{C}') \leftarrow \mathcal{F}(\mathbf{f})$ 
   $\phi \leftarrow \text{LOWER}(\mathbf{X} = \theta, \mathcal{C})$ 
  if QF-FD-SOLVE( $\phi$ ) returns unsat then fail
  else if QF-FD-SOLVE( $\phi$ ) returns sat with  $\sigma$  then
    if there is a  $\%g(\theta') \in \mathcal{C}'$  s.t. SOLVE( $\mathcal{F}$ ,  $\mathcal{M}[(\mathbf{f}, \theta) \mapsto \sigma]$ ,  $\mathbf{g}$ ,  $\sigma(\theta')$ ) fails then
      | SOLVE( $\mathcal{F}[\mathbf{f} \mapsto (\mathbf{X}, \mathcal{C} \cup \{\theta' \neq \sigma(\theta')\}]$ ,  $\mathcal{C}'$ ),  $\mathcal{M}$ ,  $\mathbf{f}$ ,  $\theta$ )
    else
      |  $\{(\mathbf{f}, \theta) \mapsto \sigma\} \cup \cup \text{SOLVE}(\mathcal{F}, \mathcal{M}[(\mathbf{f}, \theta) \mapsto \sigma], \mathbf{g}, \sigma(\theta'))$  for all  $\%g(\theta') \in \mathcal{C}'$ 

```

Figure 6.19. Constraint solving algorithm

the final lifting context is constructed by pairing each goal with its corresponding generated function.

Figure 6.19 presents a naive algorithm for constraint solving. The subroutine LOWER is used to reduce all constraints except for function call constraints to formulas in the *quantifier-free finite domain theory* (QF_FD), then an off-the-shelf solver (Z3 [40]) is used to solve them. LOWER first decomposes all compatibility class constraints, i.e., $\mathbf{X} \in [\eta]$, into *atomic* classes. For example, $\mathbf{X} \in [\text{list} \rightarrow \mathbb{Z}]$ is decomposed into $\mathbf{X}_1 \in [\text{list}]$ and $\mathbf{X}_2 \in [\mathbb{Z}]$, with \mathbf{X} substituted by $\mathbf{X}_1 \rightarrow \mathbf{X}_2$ in all other constraints. Then the newly generated compatibility class constraints can be reduced to a disjunction of all possible specification types in this class. For example, $\mathbf{X}_1 \in [\text{list}]$ reduces to $\mathbf{X}_1 = \text{list} \vee \mathbf{X}_1 = \widehat{\Psi \text{list}}_{\leq} \vee \mathbf{X}_1 = \widehat{\Psi \text{list}}_{=}$, if $\widehat{\text{list}}_{\leq}$ and $\widehat{\text{list}}_{=}$ are the only OADTs for `list`. After this step, all type variables are compatible with an atomic simple type, so we can also decompose other constraints into a set of base cases according to the rules of their relations. Note that while mergeability is not one of the constraints generated by the lifting algorithm, such constraints still arise from `%ite` and `%match`. When multiple assignments are valid for a particular type variable, we prefer the “cheaper”, i.e., more permissive, solution. For example, we prefer public type `list` over OADTs if possible, and prefer $\widehat{\text{list}}_{=}$ over $\widehat{\text{list}}_{\leq}$. We encode these preferences as *soft constraints*, and assign a bigger penalty to more restrictive types. The penalty is inferred by

analyzing the coercion relations: a more permissive type can be coerced to a more restrictive type but not the other way around, because the more restrictive type hides more information.

Once other constraints are solved, the type parameters of function call constraints are concretized. The algorithm is then recursively applied to this subgoal, i.e., the functional call paired with the concretized type, assuming the original goal is solved by extending \mathcal{M} with the original goal and its type assignments. This handles potential (mutual) recursion. If a subgoal fails, its target type will be added as a refutation to the corresponding constraints, and backtrack.

It is easy to see that this algorithm terminates: every recursive call to SOLVE either adds a refutation that reduces the search space of the non-function-call constraints, or extends \mathcal{M} which is finitely bounded by the number of functions and the number of OADTs, hence reducing the search space of function-call constraints.

Our implementation maintains a more sophisticated state of \mathcal{M} and applies several optimizations to reduce the number of calls to the external solver. We also exploit the incremental solver of Z3 to help with performance.

6.4.2 Metatheory of Algorithmic Lifting

The lifting algorithm enjoys a soundness theorem with respect to the declarative lifting relation. As a result, our algorithm inherits the well-typedness and correctness properties of the declarative version. The statement of this theorem follows how the algorithm is used: if the generated constraints, equating the function type variable with the specification type, are satisfiable by the type assignment σ , instantiating the lifted expression with σ and elaborating the macros results in a target expression that is valid under the declarative lifting relation:

Theorem 6.4.1 (Soundness of algorithmic lifting). *Suppose $\Sigma; \cdot \vdash e : \eta \sim \mathbf{x} \triangleright \dot{e} \mid \mathcal{C}$. Given a specification type θ , if $\mathcal{S}; \mathcal{L}; \Sigma; \sigma \vDash \mathbf{x} = \theta, \mathcal{C}$, then $\sigma(\dot{e})$ elaborates to an expression \dot{e}' , such that $\mathcal{S}; \mathcal{L}; \Sigma; \cdot \vdash e : \theta \triangleright \dot{e}'$.*

The proof of this theorem is available in [Appendix B.3](#).

6.5 Implementation

Our compilation pipeline takes as input a source program, including any OADTs, Ψ -structures, and macros (e.g., `%lift`), in the public fragment of TAYPSI and privacy policies (i.e., security-type signatures) for all target functions. After typing the source program using a bidirectional type checker, our lifting pass generates secure versions of the specified functions and their dependencies, using Z3 [40] as its constraint solver. The resulting TAYPSI functions are translated into OIL (Section 5.3), an ML-style functional language equipped with oblivious arrays and secure array operations: OADTs are converted to serialized versions which are stored in secure arrays, and all oblivious operations are translated into secure array operations. After applying some optimizations, our pipeline outputs an OCaml library providing secure implementations of all the specified functions, including section and retraction functions for encrypting private data and decrypting the results of a joint computation. After linking this library to a *driver* that provides the necessary cryptographic primitives (i.e., secure integer arithmetic), programmers can build secure MPC applications on top of this API. The evaluation in Section 6.6 uses a driver implemented using the popular open-source EMP toolkit [52].

6.5.1 Optimizations

Our implementation of TAYPSI implements three optimizations which further improve the performance of the programs it generates.

The *smart array* optimization supports zero-cost array slicing and concatenation, and eliminates redundant operations over the serialized representation of oblivious data. To reduce the overhead of constructing and destructing oblivious data, our implementation does not create new arrays when performing array slicing and concatenation. It allows the results to follow the original structure in TAYPSI for as long as possible, until a `mux` forces them to be flattened. Conceptually, the smart arrays delay these operations, performing them all at once when flattening is required. On the other hand, to eliminate the redundant cryptographic operations, one observation underlying this optimization is that evaluating a `mux` whose branches are encrypted versions of publicly-known values is

unnecessary: `mux [b] ($\widehat{\mathbb{B}}\#s$ true) ($\widehat{\mathbb{B}}\#s$ false)` is equivalent to `[b]`, for example. This situation frequently occurs in map-like functions, where the constructor used in each branch of a function is publicly known. Under the hood, the serialized encoding of the result of `map` uses a boolean tag to indicate which constructor was used to build it, i.e., `Nil` or `Cons`; this boolean is determined by the tag of the input list, e.g., `mux [tag] [true] [false]`. Of course, the tag used in each branch is publicly known: `map` always returns `Nil` if the input list is empty, and returns a `Cons` otherwise. Thus, we can safely reuse the `[tag]` of the input list to label the result of `map`, for similar reasons as the previous example. The smart array optimization exploits this observation by marking when section functions are applied to public values instead of, for example, immediately evaluating `$\widehat{\mathbb{B}}\#s$ true` to the encrypted value `[true]`. Then, when performing a `mux`, the smart array first checks if both branches are “fake” private values, safely reducing the `mux` to its private condition if so, without actually performing any cryptographic operations. In addition, we keep track of whether a value is “arbitrary”, e.g., used for padding. We can simply return a branch of a `mux` if the other branch is an arbitrary value.

The *reshape guard* optimization instruments `reshape` instances to first check if the public views of two private values are identical, omitting the `reshape` operation if so. Reshaping OADTs to the same public view is a common scenario, especially when the partial order defined on a public view type is a total order. For example, the join of the public views of `$\widehat{\text{list}}_{\leq}$` is the maximum one, which is always one of these public views. Therefore, reshaping the private lists of this maximum length should not require any additional work.

The *memoization* optimization caches the sizes of the private representation of data in order to avoid recalculating this information, which is needed to create and slice oblivious arrays. Similar to TAYPE, invoking these size computations in a recursive function can potentially introduce asymptotic slowdown. TAYPE eliminates these repeated computations using a tupling optimization that merges the size functions with the section and retraction functions. However, in TAYPSI, this size calculation can happen in any arbitrary private functions, in addition to section and retraction functions, making tupling optimization brittle to apply. Instead, TAYPSI memoizes the map from public views to the sizes of the private representation to avoid recalculation. If an OADT uses integer public view, then the

memoization technique is standard using a hash table for retrieving the calculated sizes. On the other hand, it is not efficient to use ADT public views as keys of a hash table. We instead automatically embed the sizes within a public view type itself, and rewrite its introduction and elimination forms and all TAYPSI programs that use them accordingly. For example, a Peano number public view `peano` is augmented as follows.

```
data peano = peano_memo × ℕ
data peano_memo = Zero | Succ peano
```

Whenever we need to calculate the size of a private representation using this public view `peano`, we simply project out its size (of type `ℕ`), without recalculating the size function.

6.6 Evaluation

Our evaluation considers the following research questions:

RQ1 How does the performance of TAYPSI’s transformation-based approach compare to the dynamic enforcement strategy of TAYPE?

RQ2 What is the compilation overhead of TAYPSI’s translation strategy?

6.6.1 Microbenchmark Performance

To answer **RQ1**, we have evaluated the performance of a set of microbenchmarks compiled with both TAYPSI and TAYPE. Both approaches are equipped with optimizations that are unique to their enforcement strategies: TAYPSI’s reshape guard optimization is not applicable to TAYPE, and TAYPE features an *early tape* optimization that does not make sense for TAYPSI.¹¹ Our evaluation also includes a version of TAYPE that implements TAYPSI’s smart array optimization (TAYPE-SA), in order to provide a comparison of the two approaches at their full potential.

Our benchmarks are a superset of the benchmarks from [Section 5.5](#). [Figure 6.20](#) presents the experimental results.¹² These experiments fix the public views of private lists and trees

¹¹↑TAYPE also implements a tupling optimization, but this is analogous to TAYPSI’s memoization optimization.

¹²↑All results are averaged across 5 runs, on an M1 MacBook Pro with 16 GB memory. All parties run on the same host with local network communication.

Benchmark	TAYPE (ms)	TAYPE-SA (ms)	TAYPSI (ms)
elem_1000 [†]	8.15	8.11	8.02 (98.47%, 98.89%)
hamming_1000 [†]	15.09	15.21	14.46 (95.79%, 95.04%)
euclidean_1000 [†]	67.43	67.55	67.32 (99.84%, 99.66%)
dot_prod_1000 [†]	66.12	66.19	66.41 (100.43%, 100.33%)
nth_1000 [†]	11.98	12.05	12.04 (100.54%, 99.93%)
map_1000	2139.55	5.07	5.14 (0.24%, 101.44%)
filter_200	failed	failed	86.86 (N/A, N/A)
insert_200	5796.69	88.92	88.07 (1.52%, 99.04%)
insert_list_100	failed	failed	4667.66 (N/A, N/A)
append_100	4274.7	45.09	44.18 (1.03%, 97.99%)
take_200	169.07	3.05	3.09 (1.83%, 101.15%)
flat_map_200	failed	failed	7.3 (N/A, N/A)
span_200	13529.34	124.79	91.22 (0.67%, 73.09%)
partition_200	failed	failed	176.49 (N/A, N/A)
<hr/>			
elem_16 [†]	446.81	459.1	404.9 (90.62%, 88.19%)
prob_16 [†]	13082.52	12761.7	12735.16 (97.34%, 99.79%)
map_16	4414.69	262.14	215.67 (4.89%, 82.27%)
filter_16	8644.14	452.04	433.7 (5.02%, 95.94%)
swap_16	failed	failed	4251.36 (N/A, N/A)
path_16	failed	6657.07	894.88 (N/A, 13.44%)
insert_16	83135.81	8093.81	1438.87 (1.73%, 17.78%)
bind_8	21885.65	494.98	532.86 (2.43%, 107.65%)
collect_8	failed	failed	143.38 (N/A, N/A)

Figure 6.20. Running times for each benchmark in milliseconds. The TAYPSI column also reports the percentage of running time relative to TAYPE and TAYPE-SA. A **failed** entry indicates the benchmark either timed out after 5 minutes or exceeded the memory bound of 8 GB. List and tree benchmarks appear above and below the double line, respectively.

to be their maximum length and maximum depth, respectively; the suffix of each benchmark name indicates the public view used. The benchmarks annotated with [†] simply traverse the data type in order to produce a primitive value, e.g., an integer; these include membership (`elem`), hamming distance (`hamming`), minimum euclidean distance (`euclidean`), dot product (`dot_prod`), secure index look up (`nth`) and computing the probability of an event given a probability tree diagram (`prob`). The programs generated by TAYPE, TAYPE-SA and TAYPSI all exhibit similar performance on these benchmarks. The remaining benchmarks all construct structured data values, i.e., the sort of application on which TAYPSI is expected to shine.

In addition to standard list operations, the list benchmarks include insertion into a sorted list (`insert`) and insertion of a list of elements into a sorted list (`insert_list`) (both lists have public view 100). The tree examples include a filter function that removes all nodes (including any subtrees) greater than a given private integer (`filter`), swapping subtrees if the node matches a private integer (`swap`), computing a subtree reached following a list of “going left” and “going right” directions (`path`), insertion into a binary search tree (`insert`), replacing the leaves of a tree with a given tree (`bind`), and collecting all nodes smaller than a private integer into a list (`collect`).

Dynamic policy enforcement either fails to finish within 5 minutes or exceeds an 8 GB memory bound on almost half of the last set of benchmarks, due to the exponential blowup discussed in [Section 6.1](#). For those benchmarks that do finish, TAYPSI’s enforcement strategy results in a fraction of the total execution time compared to TAYPE. Compared to the version of TAYPE using smart arrays, TAYPSI still performs comparably or better, although the gap is somewhat narrowed: functions like `map` do not suffer from exponential blowup, so these benchmarks benefit mostly from the smart array optimization. In summary, these results demonstrate that a static enforcement strategy performs considerably better than a dynamic one on many benchmarks, and works roughly as well on the remainder.

6.6.2 Impact of Optimization

To evaluate the performance impact of TAYPSI’s three optimizations, we conducted an ablation study on their effect. The results, shown in [Figure 6.21](#), indicate that our smart array optimization is the most important, providing up to almost 800x speedup in the best case. As suggested by [Figure 6.20](#), this optimization also helps significantly with the performance of TAYPE, although not enough to outweigh the exponential blowup innate in its dynamic approach. The other optimizations also improve performance, albeit not as significantly. As our memoization pass caches public views of arbitrary type, we have also conducted an ablation study for these examples using ADT public views instead: the list examples use Peano number to encode the maximum length of a list, and the tree examples use the upper

Benchmark	No SA (ms)		No RG (ms)		No Memo (ms)	
elem_1000	18.37	(2.29x)	8.06	(1.0x)	17.76	(2.21x)
hamming_1000	51.73	(3.58x)	14.53	(1.01x)	35.5	(2.46x)
euclidean_1000	79.07	(1.17x)	67.31	(1.0x)	76.36	(1.13x)
dot_prod_1000	87.77	(1.32x)	66.15	(1.0x)	77.33	(1.16x)
nth_1000	22.69	(1.88x)	12.18	(1.01x)	20.53	(1.7x)
map_1000	2106.43	(409.89x)	139.91	(27.23x)	37.71	(7.34x)
filter_200	5757.28	(66.29x)	93.93	(1.08x)	114.7	(1.32x)
insert_200	255.43	(2.9x)	94.61	(1.07x)	89.32	(1.01x)
insert_list_100	22806.87	(4.89x)	5186.07	(1.11x)	4771.28	(1.02x)
append_100	4226.32	(95.66x)	50.79	(1.15x)	61.77	(1.4x)
take_200	169.45	(54.91x)	12.92	(4.19x)	4.68	(1.52x)
flat_map_200	5762.63	(789.08x)	16.99	(2.33x)	60.03	(8.22x)
span_200	5924.1	(64.95x)	99.83	(1.09x)	120.09	(1.32x)
partition_200	11528.0	(65.32x)	185.16	(1.05x)	231.06	(1.31x)
<hr/>						
elem_16	433.73	(1.07x)	404.05	(1.0x)	402.15	(0.99x)
prob_16	13019.56	(1.02x)	12746.24	(1.0x)	12731.89	(1.0x)
map_16	4410.84	(20.45x)	635.18	(2.95x)	213.96	(0.99x)
filter_16	8674.71	(20.0x)	1131.02	(2.61x)	440.16	(1.01x)
swap_16	8671.52	(2.04x)	5471.4	(1.29x)	4246.39	(1.0x)
path_16	9108.54	(10.18x)	1083.21	(1.21x)	888.95	(0.99x)
insert_16	19101.36	(13.28x)	2151.83	(1.5x)	1432.92	(1.0x)
bind_8	19647.83	(36.87x)	870.93	(1.63x)	534.3	(1.0x)
collect_8	11830.6	(82.51x)	152.29	(1.06x)	186.92	(1.3x)

Figure 6.21. Impact of turning off the smart array (No SA), reshape guard (No RG), and public view memoization (No Memo) optimizations. Each column presents running time in milliseconds and the slowdown relative to that of the fully optimized version reported in [Figure 6.20](#).

bound of the spines. In this study ([Figure 6.22](#)), we observe up to 9 times speed up in the list examples, with minimal regression in tree examples.

6.6.3 Compilation Overhead

To measure the overhead of TAYPSI’s use of an external solver to resolve constraints, we have profiled the compilation of a set of larger programs drawn from TAYPE’s benchmark suite. The first two benchmark suites (List and Tree) in [Figure 6.23](#) include all the microbenchmarks from previous section. The next benchmark, List (stress), consists of the

Benchmark	Base (ms)	No Memoization (ms)	
elem_1000	13.45	18.2	(1.35x)
hamming_1000	25.98	36.35	(1.4x)
euclidean_1000	73.22	77.07	(1.05x)
dot_prod_1000	77.65	77.92	(1.0x)
nth_1000	17.8	20.89	(1.17x)
map_1000	19.95	42.53	(2.13x)
filter_200	87.22	118.54	(1.36x)
insert_200	89.01	89.77	(1.01x)
insert_list_100	4719.13	4809.16	(1.02x)
append_100	45.05	63.83	(1.42x)
take_200	4.08	5.17	(1.27x)
flat_map_200	7.32	69.1	(9.45x)
span_200	92.73	124.06	(1.34x)
partition_200	176.68	238.27	(1.35x)
elem_16	425.41	416.13	(0.98x)
prob_16	12772.3	12762.84	(1.0x)
map_16	268.07	255.66	(0.95x)
filter_16	494.38	485.96	(0.98x)
swap_16	4407.25	4329.39	(0.98x)
path_16	940.83	944.69	(1.0x)
insert_16	1603.88	1770.76	(1.1x)
bind_8	553.0	585.06	(1.06x)
collect_8	143.79	187.25	(1.3x)

Figure 6.22. Impact of turning off the public view memoization (No Memoization) optimization for the examples using ADT public views. The No Memoization column also reports the slowdown relative to Base, the fully optimized version.

same microbenchmarks as List with 5 additional list OADTs. The purpose of this synthetic suite is to examine the impact of the number of OADTs on the search space. The remaining benchmarks represent larger, more realistic applications which demonstrate the expressivity and usability of TAYPSI.

The last three columns of [Figure 6.23](#) report the results of these experiments: total compilation time (Total), time spent on constraint solving (Solver) and the number of solver queries (#Queries). The group of columns in the middle of the table describes features that can impact the performance of our constraint-based approach: the number

Suite	#Functions	#Types	#Atoms	#Queries	Total (s)	Solver (s)
List	20	7	70	84	0.47	0.081
Tree	14	9	44	31	0.47	0.024
List (stress)	20	12	70	295	3.45	2.8
Dating	4	13	16	10	0.58	0.019
Medical Records	20	19	58	51	0.48	0.072
Secure Calculator	2	9	6	5	1.34	0.013
Decision Tree	2	13	6	16	0.28	0.016
K-means	16	11	68	86	1.62	0.95
Miscellaneous	11	7	42	47	0.26	0.065

Figure 6.23. Impact of constraint solving on compilation speed

of functions (`#Functions`) being translated, the number of atomic types (`#Types`), and the total number of atomic types used in function types (`#Atoms`). For example, the `List` benchmark features 7 atomic types: `public` and `oblivious` booleans, integers and lists, as well as an unsigned integer type (i.e., natural numbers). The number of atomic types in the function `filter : list → ℤ → list` is 3. In the worst case scenario, our constraint solving algorithm will explore every combination of types that are compatible with this signature, resulting in the constraints associated with `filter` being solved $2 * 2 * 2 = 8$ times. Exactly how many compatible types the constraint solving algorithm explores depends on many factors: the user-specified policies, the complexity of the functions, the calls to other functions and so on. We chose these 3 metrics as a coarse approximation of the solution space. Our results show that the solver overhead is quite minimal for most benchmarks, and in general solving time per query is low thanks to our encoding of constraints in an efficiently decidable logic.

6.7 Conclusion

Secure multiparty computation allows joint computation over private data from multiple parties, while keeping that data secure. TAYPE has considered how to make languages for MPC more accessible by allowing privacy requirements to be decoupled from functionality, relying on dynamic enforcement of policies. Unfortunately, the resulting overhead of this strategy made it difficult to scale applications manipulating structured data. This chapter

presents TAYPSI, a policy-agnostic language for oblivious computation that transforms programs to instead *statically* enforce a user-provided privacy policy. The resulting programs are guaranteed to be both well-typed, and hence secure, and equivalent to the source program. Our experimental results show this strategy yields considerable performance improvements over prior approaches, while maintaining a clean separation between privacy and programmatic concerns.

7. RELATED WORK

Secure computation was first formally introduced by Yao [1] alongside his proposed solution, Garbled Circuits. In secure computation, an untrusted party may observe the whole execution of the secure program, or infer some private information from other side-channels. Enabling secure computations that use algebraic data types that also hide their structures is a key motivation of this work. Secure computation techniques can be broadly divided into those using multiparty computation and those relying on outsourced computation [2, 62]. Those in the former category typically use protocols based on either Garbled Circuits or secret-sharing schemes [5, 6, 63]. In the realm of outsourced computation, solutions are typically based on fully homomorphic encryption [7, 8], but can also be supported by virtualization [9, 10] or secure processors [11]. Our implementation uses the EMP toolkit [52] for its secure backend, but is compatible with other solutions under the mild requirement that they implement primitives for secure integer operations.

Many high-level programming languages have been proposed that support some form of secure computation [4]. Their goals are similar to ours in that they provide high-level language support for writing secure programs. However, most do not support (recursive) data structures at all, or assume the structural information is always public. To the best of our knowledge, none of these languages decouple security policies and program logic, as TAYPE and TAYPSI do. Obliv-C [13] is a C-like oblivious language. Algebraic data types can be encoded with the C-style `struct` keyword with pointers. Since their oblivious types are restricted to base C types, however, the structure of the defined ADT is public. It would be possible to implement oblivious ADT in Obliv-C by manually padding and using the data types according to their public views. The language provides a `~obliv` keyword that can be used to dynamically track the maximum bound of a data type, at the cost of some additional user effort. Moreover, if the programmers decide to use a different public view, they have to fix every place where this data type is used. PICCO [14, 64] is also a C-like language for secure computation, which supports C pointers to private data (possibly at private locations) and dynamic memory allocation [65]. ObliVM [15] is a Java-like language which also has a `struct` keyword to define data types, but only supports public structures,

much like Obliv-C. Wysteria and Wys* [16, 17] are functional languages that focus on mixed-mode computation. While they do not support recursive data types, both languages include simple polynomial types and primitive arrays. In contrast, our languages do not consider mixed-mode computation. Symphony [20, 66] is a successor of Wysteria which permits more reactive applications through a combination of first-class support for coordinating parties and primitives for secret-sharing and -recombination. Symphony also supports recursive data types which may contain private data, e.g., a tree whose leaves contain oblivious payloads, but does not obfuscate the structure of those datatypes. λ_{obliv} [18] is a functional programming language for oblivious computation that focuses on probabilistic programs, making it suitable for implementing some oblivious cryptographic algorithms, such as ORAM, although it does not include algebraic data types. In contrast, our work does not consider probabilistic programs, though it could be an interesting future direction. Our approach and theirs share similar threat models and guarantees of obliviousness. Other secure computation toolchains include FairplayMP [67], Sharemind [68], CBMC-GC [69], SCVM [70], TinyGarble [71], and Frigate [72].

Several prior works have considered how to compile secure programs into more efficient secure versions. Viaduct [19, 73] is a compiler that transforms high-level programs into secure distributed versions by intelligently selecting an efficient combination of protocols for subcomputations. The HyCC toolchain [74] similarly transforms a C program into a version that combines different MPC protocols to optimize performance. TASTY [75], ABY [76], EzPC [77] and MOTION [78] are similar frameworks for enabling *mixed-protocol computation*. The HACCLE toolchain [79] uses staging to generate efficient garbled circuits from a high-level language. Compiler techniques, e.g., vectorization, have been studied for optimizing fully homomorphic encryption (FHE) applications [59, 80–83].

Constant-time languages protect programs from inadvertently leaking private information through timing channels by providing atomic constructs and carefully tracking information control. This is also a goal of our system, and our solution to this problem is similar. The first formal study of constant time algorithms was in the context of cache-based attacks [9]. Barthe *et al.* [84] extended the formally verified CompCert compiler [85] to ensure constant time execution. Our obliviousness theorem provides a formal guarantee of a constant-time

property. Jasmin [86, 87] is a framework for implementing high-performance cryptography. It achieves constant-time security by embedding Jasmin programs into Dafny [88] which enables automated proofs of this property (and memory safety). FaCT [89] is a high-level language for writing constant-time computation using (non-recursive) data types. One of its unique features is a front-end compiler to transform a well-typed (but potentially not constant-time) FaCT program into a constant-time FaCT program. In TAYPE and TAYPSI, the programmers can simply encode programs in the conventional fragment and then convert them to oblivious programs (that are constant-time) by composing privacy policies and the standard programs.

Another popular cryptographic technique for hiding private information of data structures is oblivious RAM [26, 90, 91] (ORAM). ORAM provides primitives to access an encrypted memory buffer without revealing the access pattern, except for the number of accesses. There have been proposals for generically constructing oblivious data structures using ORAM [92]. Oblivious data structures constructed this way hide the access patterns of a sequence of data structure operations. This line of work in general does not consider leakage through side-channels. While our solution also naturally hides access patterns, we also assume a much stronger adversary who can observe the whole computation. On the other hand, certain data structures may yield asymptotically better performance if encoded using ORAM. Integrating ORAM into our systems for performance gains while maintaining strong security guarantees is a promising future direction.

Our approach of type-based information flow control to enforce obliviousness, a form of *noninterference*, follows a body of work in *security-type systems* [27, 29]. To the best of our knowledge, our system is the first to combine a dependent type system with large elimination and a security-type system. While most security-type systems tag types with labels classifying the sensitivity of data, our dependent type system tags kinds instead to keep track of whether a term is oblivious. On the other hand, our leakage labels are similar to these security labels, and used to track if a term is leaky. Our notion of retraction bears some resemblance to delimited information release [93]. In a system with delimited information release, the programmers may choose to reveal some private information, similar to retraction functions in λ_{OADT^+} . However, our semantics guarantees retraction never releases any private information. Another difference from a standard security-type system is that we use explicit

coercion via section functions instead of implicit subtyping to convert public types to secure types. On the one hand, our typing rules and semantics for oblivious types and non-oblivious types are quite different. On the other hand, implicit subtyping does not make sense in the case of ADTs. To convert a public ADT to an oblivious one, we not only need to know how the oblivious ADT is represented, but also to infer the public views.

Our obliviousness guarantee is a strengthened variant of *memory trace obliviousness* (MTO) [28], which itself provides stronger guarantees than most information flow type systems. Under MTO, the patterns of memory access generated by a program are required to be indistinguishable, in addition to its result. This work also proposed a language based on *Oblivious RAM* [26, 90, 91] and transformation techniques to ensure this property. However, this threat model is weaker than that of this dissertation. On the one hand, it does not consider timing channel: while memory access traces include instruction fetches, which ensures the branches of a secure conditional always run the same number of instructions, the instructions themselves can still exhibit different timing behaviors. For example, the program `if s > 0 then s := p + p else s := p * p` is secure in their model, as both branches produce the same memory access pattern (including instruction fetches), but the second branch is slower, assuming multiplication is slower than addition in the CPU. On the other hand, under MTO, adversaries cannot observe the instructions executed by the CPU. This is not the case in the MPC setting (especially in the secret-sharing-based schemes), as every party is a potential adversary that can observe instructions: `if s > 0 then s := p + 1 else s := p + 2` is accepted in their model, but an adversary in our model is able to discern if the program is computing `p+1` or `p+2`, even if they have the same timing behavior. In contrast, the traces we consider include every program state under a small-step semantics, which rules out these two examples.

Jeeves [94] and our work have a shared goal of decoupling security policies from program logic. While they both employ a similar high-level strategy of relying on the language to automatically enforce policies, their different settings result in very different solutions. In Jeeves’ programming model, each piece of data is equipped with a pair of high- and low-level views: a username, for example, may have a high confidentiality view of “Alice”, but a low view of “Anonymous”. The language then uses the view stipulated by the privacy policy and

current execution context, ensuring that information is only visible to observers with the proper authority. In the MPC setting, however, no party is allowed to observe the private data of other parties. Thus, no party can view all the data necessary for the computation, making it impossible to compute a correct result by simply replacing data with some predetermined value, like “Anonymous”.

Our oblivious types can also be viewed as a kind of refinement types [36–38]: the oblivious tree in Chapter 3 can be understood as trees with a maximum depth stipulated by the type index, for example. However, this declarative specification does not explain how to represent such an oblivious tree. Nonetheless, this view of subset types suggests a future direction of integrating refinement typing into our system to ensure the correct use of public indices and to enable simpler policy specifications. Dependent type systems with large elimination can be found in many theorem provers, such as Coq [33] and Agda [35]. These languages are designed more towards theorem proving and thus only admit total functions, while our languages allow general recursion and hence nontermination. A notable dependently typed language with nontermination is Zombie [41–43], though our goals are drastically different.

Nanevski *et al.* [95] show how Relational Hoare Type Theory can be used to encode and verify a variety of security policies in a theorem prover using dependent types. While capable of specifying security policies like noninterference, their encoding does not address termination behavior and only characterizes the final output value, and thus does not protect against control flow leaks. In addition, users have to manually verify these properties in the proof assistant. In contrast, we consider a much stronger threat model, and all of our oblivious calculi protect against a larger class of leaks. The calculi additionally provide a fixed security guarantee in the form of our obliviousness theorem, which any well-typed program enjoys “for free”, without any additional user effort.

In our mechanized formalization, the correctness and security guarantees provided by the underlying cryptographic primitives are baked into our semantics and notion of indistinguishability. There is a body of work about formally verified cryptography [96–99], which could be integrated into our work in the future to provide a stronger formal guarantee. Some of these solutions have focused on verifying multiparty computation [100, 101].

8. CONCLUSIONS AND FUTURE WORK

Writing secure applications that do not leak private information presents challenges in several dimensions. The limited support for rich data structures in existing oblivious languages discourages users from implementing general applications. The inability to express complex policies makes it difficult to comply with real-world privacy requirements. The intermixing of application logic and privacy policies forces programmers to write specialized versions of the same program for each set of privacy requirements. This dissertation introduced various programming language techniques that tackle these problems: private structured data and complex policies can be encoded using oblivious algebraic data types, and a form of modularity that decouples privacy and programmatic concerns can be achieved using tape semantics or static program transformations. This thesis has formally developed these techniques via a family of core calculi, which it used to establish key metatheoretic results, including a strong security guarantee called obliviousness and the standard type safety property. These techniques have been implemented in the languages and compilation pipelines, TAYPE and TAYPSI, providing an end-to-end programming environment for developing secure applications and evaluating performance and usability.

The work presented in this dissertation is a step towards wider adoption of privacy-preserving techniques, but there is still much room to make these policy-agnostic languages more practical. First, to enable a wider range of applications, we should increase the expressivity of these languages: adding support for mutable data, probabilistic computation and reactive programs that allow for interaction between participating entities in a policy-agnostic way are all important directions. Another limitation of the current approach is that OADTs can only encode tree-like structures, similar to ADTs. Supporting graph data structures will enable more secure applications. Second, while the separation of policies and application logic reduces user burden, many potential improvements to usability remain. As one example, TAYPE and TAYPSI do not support type polymorphism, and as a result, users need to define, e.g., distinct types for lists of integers and lists of booleans. Writing OADTs and their associated methods requires programmers to be proficient in dependent types. At the same time, many of these definitions are straightforward but tedious boilerplate.

A better approach would be to derive OADTs from a higher-level predicate that specifies the publicly shared information and synthesize the associated methods. The current implement of TAYPSI can generate nonterminating programs if a specified policy does not provide enough public information to bound the recursion depth. Guaranteeing equi-termination for the generated secure programs would free users from manually reasoning about termination. Third, good performance is necessary for encouraging adoption of these language techniques. The abstractions developed for the modularity described in this dissertation often come with performance penalty. Users may sometimes need to manually optimize a program to avoid some efficiency pitfalls caused by the secure semantics of `mux`, for example. Thus, additional optimization techniques are needed to generate secure programs that are as efficient as hand-crafted versions. One promising direction is to rewrite program control flow in a way that exploits publicly available information to remove secure operations.

REFERENCES

- [1] A. C. Yao, “Protocols for secure computations,” in *23rd Annual Symposium on Foundations of Computer Science (Sfcs 1982)*, Nov. 1982, pp. 160–164. DOI: [10.1109/SFCS.1982.38](https://doi.org/10.1109/SFCS.1982.38).
- [2] D. Evans, V. Kolesnikov, and M. Rosulek, “A Pragmatic Introduction to Secure Multi-Party Computation,” *Foundations and Trends® in Privacy and Security*, vol. 2, no. 2-3, pp. 70–246, 2018, ISSN: 2474-1558, 2474-1566. DOI: [10.1561/33000000019](https://doi.org/10.1561/33000000019). [Online]. Available: <http://www.nowpublishers.com/article/Details/SEC-019>.
- [3] P. Laud and L. Kamm, Eds., *Applications of Secure Multiparty Computation* (Cryptography and Information Security Series volume 13). Amsterdam, Netherlands: IOS Press, 2015, 253 pp., ISBN: 978-1-61499-532-6 978-1-61499-531-9.
- [4] M. Hastings, B. Hemenway, D. Noble, and S. Zdancewic, “SoK: General Purpose Compilers for Secure Multi-Party Computation,” in *2019 2019 IEEE Symposium on Security and Privacy (SP)*, Los Alamitos, CA, USA: IEEE Computer Society, May 2019, pp. 479–496. DOI: [10.1109/SP.2019.00028](https://doi.org/10.1109/SP.2019.00028). [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/SP.2019.00028>.
- [5] O. Goldreich, “Towards a theory of software protection and simulation by oblivious RAMs,” in *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, ser. STOC ’87, New York, NY, USA: Association for Computing Machinery, Jan. 1, 1987, pp. 182–194, ISBN: 978-0-89791-221-1. DOI: [10.1145/28395.28416](https://doi.org/10.1145/28395.28416). [Online]. Available: <http://doi.org/10.1145/28395.28416>.
- [6] A. Beimel, “Secret-Sharing Schemes: A Survey,” in *Coding and Cryptology*, Y. M. Chee *et al.*, Eds., ser. Lecture Notes in Computer Science, Berlin, Heidelberg: Springer, 2011, pp. 11–46, ISBN: 978-3-642-20901-7. DOI: [10.1007/978-3-642-20901-7_2](https://doi.org/10.1007/978-3-642-20901-7_2).
- [7] C. Gentry, “Fully homomorphic encryption using ideal lattices,” in *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, ser. STOC ’09, New York, NY, USA: Association for Computing Machinery, May 31, 2009, pp. 169–178, ISBN: 978-1-60558-506-2. DOI: [10.1145/1536414.1536440](https://doi.org/10.1145/1536414.1536440). [Online]. Available: <http://doi.org/10.1145/1536414.1536440>.
- [8] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti, “A Survey on Homomorphic Encryption Schemes: Theory and Implementation,” *ACM Computing Surveys (CSUR)*, vol. 51, no. 4, 79:1–79:35, Jul. 25, 2018, ISSN: 0360-0300. DOI: [10.1145/3214303](https://doi.org/10.1145/3214303). [Online]. Available: <http://doi.org/10.1145/3214303>.

- [9] G. Barthe, G. Betarte, J. Campo, C. Luna, and D. Pichardie, “System-level Non-interference for Constant-time Cryptography,” in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’14, Scottsdale, Arizona, USA: Association for Computing Machinery, Nov. 3, 2014, pp. 1267–1279, ISBN: 978-1-4503-2957-6. DOI: [10.1145/2660267.2660283](https://doi.org/10.1145/2660267.2660283). [Online]. Available: <http://doi.org/10.1145/2660267.2660283>.
- [10] G. Barthe, G. Betarte, J. D. Campo, and C. Luna, “System-Level Non-interference of Constant-Time Cryptography. Part I: Model,” *Journal of Automated Reasoning*, vol. 63, no. 1, pp. 1–51, Jun. 1, 2019, ISSN: 1573-0670. DOI: [10.1007/s10817-017-9441-5](https://doi.org/10.1007/s10817-017-9441-5). [Online]. Available: <https://doi.org/10.1007/s10817-017-9441-5>.
- [11] M. E. Hoekstra. “Intel SGX for Dummies (Intel SGX Design Objectives).” (Nov. 30, 2015), [Online]. Available: <https://www.intel.com/content/www/us/en/develop/blogs/protecting-application-secrets-with-intel-sgx.html>.
- [12] D. Malkhi, N. Nisan, B. Pinkas, and Y. Sella, “Fairplay - a secure two-party computation system,” in *Proceedings of the 13th Conference on USENIX Security Symposium - Volume 13*, ser. SSYM’04, USA: USENIX Association, Aug. 13, 2004, p. 20.
- [13] S. Zahur and D. Evans, “Obliv-C: A Language for Extensible Data-Oblivious Computation,” 1153, 2015. [Online]. Available: <https://eprint.iacr.org/2015/1153>.
- [14] Y. Zhang, A. Steele, and M. Blanton, “PICCO: A general-purpose compiler for private distributed computation,” in *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*, ser. CCS ’13, New York, NY, USA: Association for Computing Machinery, Nov. 4, 2013, pp. 813–826, ISBN: 978-1-4503-2477-9. DOI: [10.1145/2508859.2516752](https://doi.org/10.1145/2508859.2516752). [Online]. Available: <https://dl.acm.org/doi/10.1145/2508859.2516752>.
- [15] C. Liu, X. S. Wang, K. Nayak, Y. Huang, and E. Shi, “ObliVM: A Programming Framework for Secure Computation,” in *2015 IEEE Symposium on Security and Privacy*, May 2015, pp. 359–376. DOI: [10.1109/SP.2015.29](https://doi.org/10.1109/SP.2015.29).
- [16] A. Rastogi, M. A. Hammer, and M. Hicks, “Wysteria: A Programming Language for Generic, Mixed-Mode Multiparty Computations,” in *2014 IEEE Symposium on Security and Privacy*, May 2014, pp. 655–670. DOI: [10.1109/SP.2014.48](https://doi.org/10.1109/SP.2014.48).

- [17] A. Rastogi, N. Swamy, and M. Hicks, “Wys*: A DSL for Verified Secure Multi-party Computations,” in *Principles of Security and Trust*, F. Nielson and D. Sands, Eds., ser. Lecture Notes in Computer Science, Springer International Publishing, 2019, pp. 99–122, ISBN: 978-3-030-17138-4. [Online]. Available: https://doi.org/10.1007/978-3-030-17138-4_5.
- [18] D. Darais, I. Sweet, C. Liu, and M. Hicks, “A language for probabilistically oblivious computation,” *Proceedings of the ACM on Programming Languages*, vol. 4, pp. 1–31, POPL Jan. 2020, ISSN: 2475-1421, 2475-1421. DOI: [10.1145/3371118](https://doi.org/10.1145/3371118). arXiv: [1711.09305](https://arxiv.org/abs/1711.09305). [Online]. Available: <http://dl.acm.org/doi/10.1145/3371118>.
- [19] C. Acay, R. Recto, J. Gancher, A. C. Myers, and E. Shi, “Viaduct: An extensible, optimizing compiler for secure distributed programs,” in *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, ser. PLDI 2021, New York, NY, USA: Association for Computing Machinery, Jun. 19, 2021, pp. 740–755, ISBN: 978-1-4503-8391-2. DOI: [10.1145/3453483.3454074](https://doi.org/10.1145/3453483.3454074). [Online]. Available: <https://doi.org/10.1145/3453483.3454074>.
- [20] I. Sweet, D. Darais, D. Heath, W. Harris, R. Estes, and M. Hicks, “Symphony: Expressive Secure Multiparty Computation with Coordination,” *The Art, Science, and Engineering of Programming*, vol. 7, no. 3, 14:1–14:55, Feb. 15, 2023, ISSN: 2473-7321. DOI: [10.22152/programming-journal.org/2023/7/14](https://doi.org/10.22152/programming-journal.org/2023/7/14). [Online]. Available: <https://programming-journal.org/2023/7/14/>.
- [21] J. A. Goguen and J. Meseguer, “Security Policies and Security Models,” in *1982 IEEE Symposium on Security and Privacy*, Apr. 1982, pp. 11–11. DOI: [10.1109/SP.1982.10014](https://doi.org/10.1109/SP.1982.10014).
- [22] Q. Ye and B. Delaware, “Oblivious algebraic data types,” *Proceedings of the ACM on Programming Languages*, vol. 6, no. POPL, 51:1–51:29, Jan. 2022. DOI: [10.1145/3498713](https://doi.org/10.1145/3498713). [Online]. Available: <https://doi.org/10.1145/3498713>.
- [23] Q. Ye and B. Delaware, “Taype: A Policy-Agnostic Language for Oblivious Computation,” *Proceedings of the ACM on Programming Languages*, vol. 7, no. PLDI, 147:1001–147:1025, Jun. 2023. DOI: [10.1145/3591261](https://doi.org/10.1145/3591261). [Online]. Available: <https://dl.acm.org/doi/10.1145/3591261>.
- [24] Q. Ye and B. Delaware, “Taypsi: Static Enforcement of Privacy Policies for Policy-Agnostic Oblivious Computation,” *Proceedings of the ACM on Programming Languages*, vol. 8, no. OOPSLA1, Apr. 2024. DOI: [10.1145/3649861](https://doi.org/10.1145/3649861). [Online]. Available: <https://dl.acm.org/doi/10.1145/3649861>.

- [25] R. Canetti, “Security and Composition of Multiparty Cryptographic Protocols,” *Journal of Cryptology*, vol. 13, no. 1, pp. 143–202, Jan. 1, 2000, ISSN: 1432-1378. DOI: [10.1007/s001459910006](https://doi.org/10.1007/s001459910006). [Online]. Available: <https://doi.org/10.1007/s001459910006>.
- [26] O. Goldreich, S. Micali, and A. Wigderson, “How to play ANY mental game,” in *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, ser. STOC '87, New York, New York, USA: Association for Computing Machinery, Jan. 1, 1987, pp. 218–229, ISBN: 978-0-89791-221-1. DOI: [10.1145/28395.28420](https://doi.org/10.1145/28395.28420). [Online]. Available: <http://doi.org/10.1145/28395.28420>.
- [27] A. Sabelfeld and A. Myers, “Language-based information-flow security,” *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 1, pp. 5–19, Jan. 2003, ISSN: 1558-0008. DOI: [10.1109/JSAC.2002.806121](https://doi.org/10.1109/JSAC.2002.806121).
- [28] C. Liu, M. Hicks, and E. Shi, “Memory Trace Oblivious Program Execution,” in *2013 IEEE 26th Computer Security Foundations Symposium*, Jun. 2013, pp. 51–65. DOI: [10.1109/CSF.2013.11](https://doi.org/10.1109/CSF.2013.11).
- [29] S. A. Zdancewic, “Programming languages for information security,” Ph.D. dissertation, Cornell University, USA, 2002, 327 pp.
- [30] W. W. Tait, “Intensional interpretations of functionals of finite type I,” *The Journal of Symbolic Logic*, vol. 32, no. 2, pp. 198–212, Aug. 1967, ISSN: 0022-4812, 1943-5886. DOI: [10.2307/2271658](https://doi.org/10.2307/2271658). [Online]. Available: <https://www.cambridge.org/core/journals/journal-of-symbolic-logic/article/intensional-interpretations-of-functionals-of-finite-type-i/9F30EA199783BD797DF6FA44525F114E>.
- [31] G. D. Plotkin, “Lambda Definability and Logical Relations,” University of Edinburgh, Memorandum SAI-RM-4, 1973. [Online]. Available: https://homepages.inf.ed.ac.uk/gdp/publications/logical_relations_1973.pdf.
- [32] G. Barthe and T. Coquand, “An Introduction to Dependent Type Theory,” in *Applied Semantics*, G. Barthe, P. Dybjer, L. Pinto, and J. Saraiva, Eds., Berlin, Heidelberg: Springer, 2002, pp. 1–41, ISBN: 978-3-540-45699-5. DOI: [10.1007/3-540-45699-6_1](https://doi.org/10.1007/3-540-45699-6_1).
- [33] The Coq Development Team, *The Coq Proof Assistant*, Zenodo, Jun. 2023. DOI: [10.5281/zenodo.8161141](https://doi.org/10.5281/zenodo.8161141). [Online]. Available: <https://zenodo.org/records/8161141>.

- [34] L. de Moura and S. Ullrich, “The Lean 4 Theorem Prover and Programming Language,” in *Automated Deduction – CADE 28*, A. Platzer and G. Sutcliffe, Eds., Cham: Springer International Publishing, 2021, pp. 625–635, ISBN: 978-3-030-79876-5. DOI: [10.1007/978-3-030-79876-5_37](https://doi.org/10.1007/978-3-030-79876-5_37).
- [35] Agda Developers, *Agda*, Mar. 2024. [Online]. Available: <https://agda.readthedocs.io/>.
- [36] H. Xi and F. Pfenning, “Dependent types in practical programming,” in *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ser. POPL ’99, New York, NY, USA: Association for Computing Machinery, Jan. 1999, pp. 214–227, ISBN: 978-1-58113-095-9. DOI: [10.1145/292540.292560](https://doi.org/10.1145/292540.292560). [Online]. Available: <http://doi.org/10.1145/292540.292560>.
- [37] P. M. Rondon, M. Kawaguchi, and R. Jhala, “Liquid Types,” in *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI ’08, New York, NY, USA: ACM, 2008, pp. 159–169, ISBN: 978-1-59593-860-2. DOI: [10.1145/1375581.1375602](https://doi.org/10.1145/1375581.1375602). [Online]. Available: <http://doi.acm.org/10.1145/1375581.1375602>.
- [38] M. Kawaguchi, P. Rondon, and R. Jhala, “Type-based data structure verification,” in *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI ’09, Dublin, Ireland: Association for Computing Machinery, Jun. 2009, pp. 304–315, ISBN: 978-1-60558-392-1. DOI: [10.1145/1542476.1542510](https://doi.org/10.1145/1542476.1542510). [Online]. Available: <http://doi.org/10.1145/1542476.1542510>.
- [39] N. Swamy *et al.*, “Dependent types and multi-monadic effects in F*,” in *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ser. POPL ’16, St. Petersburg, FL, USA: Association for Computing Machinery, Jan. 2016, pp. 256–270, ISBN: 978-1-4503-3549-2. DOI: [10.1145/2837614.2837655](https://doi.org/10.1145/2837614.2837655). [Online]. Available: <http://doi.org/10.1145/2837614.2837655>.
- [40] L. de Moura and N. Bjørner, “Z3: An Efficient SMT Solver,” in *Tools and Algorithms for the Construction and Analysis of Systems*, C. R. Ramakrishnan and J. Rehof, Eds., ser. Lecture Notes in Computer Science, Berlin, Heidelberg: Springer, 2008, pp. 337–340, ISBN: 978-3-540-78800-3. DOI: [10.1007/978-3-540-78800-3_24](https://doi.org/10.1007/978-3-540-78800-3_24).
- [41] V. Sjöberg *et al.*, “Irrelevance, Heterogeneous Equality, and Call-by-value Dependent Type Systems,” *Electronic Proceedings in Theoretical Computer Science*, vol. 76, pp. 112–162, Feb. 2012, ISSN: 2075-2180. DOI: [10.4204/EPTCS.76.9](https://doi.org/10.4204/EPTCS.76.9). arXiv: [1202.2923](https://arxiv.org/abs/1202.2923). [Online]. Available: <http://arxiv.org/abs/1202.2923>.

- [42] V. Sjöberg and S. Weirich, “Programming up to Congruence,” in *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ser. POPL ’15, New York, NY, USA: Association for Computing Machinery, Jan. 2015, pp. 369–382, ISBN: 978-1-4503-3300-9. DOI: [10.1145/2676726.2676974](https://doi.org/10.1145/2676726.2676974). [Online]. Available: <http://doi.org/10.1145/2676726.2676974>.
- [43] V. Sjöberg, “A Dependently Typed Language with Nontermination,” *Publicly Accessible Penn Dissertations*, Jan. 2015. [Online]. Available: <https://repository.upenn.edu/edissertations/1137>.
- [44] H. Barendregt, “Introduction to generalized type systems,” *Journal of Functional Programming*, vol. 1, no. 2, pp. 125–154, Apr. 1991, ISSN: 0956-7968, 1469-7653. DOI: [10.1017/S0956796800020025](https://doi.org/10.1017/S0956796800020025). [Online]. Available: <https://www.cambridge.org/core/journals/journal-of-functional-programming/article/introduction-to-generalized-type-systems/869991BA6A99180BF96A616894C6D710>.
- [45] H. P. Barendregt, “Lambda Calculi with Types,” in *Handbook of Logic in Computer Science*, J. Spurr, S. Abramsky, and D. M. Gabbay, Eds., Oxford University Press, Dec. 1992, ISBN: 978-0-19-853761-8. DOI: [10.1093/oso/9780198537618.003.0002](https://doi.org/10.1093/oso/9780198537618.003.0002). [Online]. Available: <https://doi.org/10.1093/oso/9780198537618.003.0002>.
- [46] B. C. Pierce, Ed., *Advanced Topics in Types and Programming Languages*. Cambridge, Mass: MIT Press, 2005, ISBN: 978-0-262-16228-9.
- [47] H. P. Barendregt, *The Lambda Calculus: Its Syntax and Semantics* (Studies in Logic and the Foundations of Mathematics v. 103), Rev. ed. Amsterdam ; New York : New York, N.Y: North-Holland ; Sole distributors for the U.S.A. and Canada, Elsevier Science Pub. Co, 1984, ISBN: 978-0-444-86748-3 978-0-444-87508-2.
- [48] Q. Ye and B. Delaware, *Oblivious Algebraic Data Types: POPL22 Artifact*, Zenodo, Oct. 2021. DOI: [10.5281/zenodo.5652106](https://doi.org/10.5281/zenodo.5652106). [Online]. Available: <https://zenodo.org/record/5652106>.
- [49] Q. Ye and B. Delaware, *Taype: A Policy-Agnostic Language for Oblivious Computation: PLDI23 Artifact*, Zenodo, Apr. 2023. DOI: [10.5281/zenodo.7806981](https://doi.org/10.5281/zenodo.7806981). [Online]. Available: <https://zenodo.org/record/7806981>.
- [50] J. Dunfield and N. Krishnaswami, “Bidirectional Typing,” *ACM Computing Surveys*, vol. 54, no. 5, 98:1–98:38, May 2021, ISSN: 0360-0300. DOI: [10.1145/3450952](https://doi.org/10.1145/3450952). [Online]. Available: <https://doi.org/10.1145/3450952>.

- [51] C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen, “The essence of compiling with continuations,” *ACM SIGPLAN Notices*, vol. 28, no. 6, pp. 237–247, Jun. 1993, ISSN: 0362-1340. DOI: [10.1145/173262.155113](https://doi.org/10.1145/173262.155113). [Online]. Available: <https://doi.org/10.1145/173262.155113>.
- [52] X. Wang, A. J. Malozemoff, and J. Katz, *EMP-toolkit: Efficient MultiParty computation toolkit*, <https://github.com/emp-toolkit>, 2016.
- [53] R. S. Bird, “Using circular programs to eliminate multiple traversals of data,” *Acta Informatica*, vol. 21, no. 3, pp. 239–250, Oct. 1984, ISSN: 1432-0525. DOI: [10.1007/BF00264249](https://doi.org/10.1007/BF00264249). [Online]. Available: <https://doi.org/10.1007/BF00264249>.
- [54] W.-N. Chin, “Towards an automated tupling strategy,” in *Proceedings of the 1993 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, ser. PEPM '93, New York, NY, USA: Association for Computing Machinery, Aug. 1993, pp. 119–132, ISBN: 978-0-89791-594-6. DOI: [10.1145/154630.154643](https://doi.org/10.1145/154630.154643). [Online]. Available: <https://doi.org/10.1145/154630.154643>.
- [55] R. M. Burstall and J. Darlington, “A Transformation System for Developing Recursive Programs,” *Journal of the ACM*, vol. 24, no. 1, pp. 44–67, Jan. 1977, ISSN: 0004-5411. DOI: [10.1145/321992.321996](https://doi.org/10.1145/321992.321996). [Online]. Available: <https://doi.org/10.1145/321992.321996>.
- [56] V. F. Turchin, R. M. Nirenberg, and D. V. Turchin, “Experiments with a super-compiler,” in *Proceedings of the 1982 ACM Symposium on LISP and Functional Programming*, ser. LFP '82, New York, NY, USA: Association for Computing Machinery, Aug. 1982, pp. 47–55, ISBN: 978-0-89791-082-8. DOI: [10.1145/800068.802134](https://doi.org/10.1145/800068.802134). [Online]. Available: <https://doi.org/10.1145/800068.802134>.
- [57] D. J. Wu, T. Feng, M. Naehrig, and K. Lauter, “Privately Evaluating Decision Trees and Random Forests,” *Proceedings on Privacy Enhancing Technologies*, vol. 2016, no. 4, pp. 335–355, Oct. 1, 2016, ISSN: 2299-0984. DOI: [10.1515/popets-2016-0043](https://doi.org/10.1515/popets-2016-0043). [Online]. Available: <https://petsymposium.org/popets/2016/popets-2016-0043.php>.
- [58] Á. Kiss, M. Naderpour, J. Liu, N. Asokan, and T. Schneider, “SoK: Modular and Efficient Private Decision Tree Evaluation,” *Proceedings on Privacy Enhancing Technologies*, vol. 2019, no. 2, pp. 187–208, Apr. 1, 2019, ISSN: 2299-0984. DOI: [10.2478/popets-2019-0026](https://doi.org/10.2478/popets-2019-0026). [Online]. Available: <https://petsymposium.org/popets/2019/popets-2019-0026.php>.

- [59] R. Malik, V. Singhal, B. Gottfried, and M. Kulkarni, “Vectorized secure evaluation of decision forests,” in *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, ser. PLDI 2021, New York, NY, USA: Association for Computing Machinery, Jun. 18, 2021, pp. 1049–1063, ISBN: 978-1-4503-8391-2. DOI: [10.1145/3453483.3454094](https://doi.org/10.1145/3453483.3454094). [Online]. Available: <https://dl.acm.org/doi/10.1145/3453483.3454094>.
- [60] Q. Ye and B. Delaware, *Taypsi: Static Enforcement of Privacy Policies for Policy-Agnostic Oblivious Computation: OOPSLA24 Artifact*, Zenodo, Mar. 2024. DOI: [10.5281/zenodo.10701642](https://doi.org/10.5281/zenodo.10701642). [Online]. Available: <https://zenodo.org/records/10701642>.
- [61] A. Ahmed, “Step-Indexed Syntactic Logical Relations for Recursive and Quantified Types,” in *Programming Languages and Systems*, P. Sestoft, Ed., ser. Lecture Notes in Computer Science, Berlin, Heidelberg: Springer, 2006, pp. 69–83, ISBN: 978-3-540-33096-7. DOI: [10.1007/11693024_6](https://doi.org/10.1007/11693024_6).
- [62] C. Hazay and Y. Lindell, *Efficient Secure Two-Party Protocols: Techniques and Constructions* (Information Security and Cryptography). Berlin ; London: Springer, 2010, 263 pp., ISBN: 978-3-642-14302-1.
- [63] U. Maurer, “Secure multi-party computation made simple,” *Discrete Applied Mathematics*, Coding and Cryptography, vol. 154, no. 2, pp. 370–381, Feb. 1, 2006, ISSN: 0166-218X. DOI: [10.1016/j.dam.2005.03.020](https://doi.org/10.1016/j.dam.2005.03.020). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0166218X05002428>.
- [64] A. Rathore, M. Blanton, M. Gaboardi, and L. Ziarek, *A Formal Model for Secure Multiparty Computation*, May 2023. DOI: [10.48550/arXiv.2306.00308](https://doi.org/10.48550/arXiv.2306.00308). arXiv: [2306.00308](https://arxiv.org/abs/2306.00308) [cs]. [Online]. Available: <http://arxiv.org/abs/2306.00308>.
- [65] Y. Zhang, M. Blanton, and G. Almashaqbeh, “Implementing Support for Pointers to Private Data in a General-Purpose Secure Multi-Party Compiler,” *ACM Transactions on Privacy and Security*, vol. 21, no. 2, 6:1–6:34, Dec. 2017, ISSN: 2471-2566. DOI: [10.1145/3154600](https://doi.org/10.1145/3154600). [Online]. Available: <https://dl.acm.org/doi/10.1145/3154600>.
- [66] I. Sweet, D. Darais, D. Heath, R. Estes, W. Harris, and M. Hicks, “Symphony: A concise language model for MPC,” in *Informal Proceedings of the Workshop on Foundations on Computer Security (FCS)*, Jun. 2021.

- [67] A. Ben-David, N. Nisan, and B. Pinkas, “FairplayMP: A system for secure multi-party computation,” in *Proceedings of the 15th ACM Conference on Computer and Communications Security*, ser. CCS ’08, New York, NY, USA: Association for Computing Machinery, Oct. 2008, pp. 257–266, ISBN: 978-1-59593-810-7. DOI: [10.1145/1455770.1455804](https://doi.org/10.1145/1455770.1455804). [Online]. Available: <https://dl.acm.org/doi/10.1145/1455770.1455804>.
- [68] D. Bogdanov, S. Laur, and J. Willemson, “Sharemind: A Framework for Fast Privacy-Preserving Computations,” in *Computer Security - ESORICS 2008*, S. Jajodia and J. Lopez, Eds., Berlin, Heidelberg: Springer, 2008, pp. 192–206, ISBN: 978-3-540-88313-5. DOI: [10.1007/978-3-540-88313-5_13](https://doi.org/10.1007/978-3-540-88313-5_13).
- [69] A. Holzer, M. Franz, S. Katzenbeisser, and H. Veith, “Secure two-party computations in ANSI C,” in *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, ser. CCS ’12, New York, NY, USA: Association for Computing Machinery, Oct. 2012, pp. 772–783, ISBN: 978-1-4503-1651-4. DOI: [10.1145/2382196.2382278](https://doi.org/10.1145/2382196.2382278). [Online]. Available: <https://dl.acm.org/doi/10.1145/2382196.2382278>.
- [70] C. Liu, Y. Huang, E. Shi, J. Katz, and M. Hicks, “Automating Efficient RAM-Model Secure Computation,” in *2014 IEEE Symposium on Security and Privacy*, May 2014, pp. 623–638. DOI: [10.1109/SP.2014.46](https://doi.org/10.1109/SP.2014.46). [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/6956591>.
- [71] E. M. Songhori, S. U. Hussain, A.-R. Sadeghi, T. Schneider, and F. Koushanfar, “TinyGarble: Highly Compressed and Scalable Sequential Garbled Circuits,” in *2015 IEEE Symposium on Security and Privacy*, May 2015, pp. 411–428. DOI: [10.1109/SP.2015.32](https://doi.org/10.1109/SP.2015.32). [Online]. Available: <https://ieeexplore.ieee.org/document/7163039?denied=>.
- [72] B. Mood, D. Gupta, H. Carter, K. Butler, and P. Traynor, “Frigate: A Validated, Extensible, and Efficient Compiler and Interpreter for Secure Computation,” in *2016 IEEE European Symposium on Security and Privacy (EuroSec’16)*, Mar. 2016, pp. 112–127. DOI: [10.1109/EuroSP.2016.20](https://doi.org/10.1109/EuroSP.2016.20). [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/7467350>.
- [73] C. Acay, J. Gancher, R. Recto, and A. C. Myers. “Secure Synthesis of Distributed Cryptographic Applications (Technical Report).” arXiv: [2401.04131](https://arxiv.org/abs/2401.04131) [cs]. (Jan. 5, 2024), [Online]. Available: <http://arxiv.org/abs/2401.04131>, preprint.

- [74] N. Büscher, D. Demmler, S. Katzenbeisser, D. Kretzmer, and T. Schneider, “HyCC: Compilation of Hybrid Protocols for Practical Secure Computation,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, (Toronto, Canada), ser. CCS ’18, New York, NY, USA: ACM, 2018, pp. 847–861, ISBN: 978-1-4503-5693-0. DOI: [10.1145/3243734.3243786](https://doi.org/10.1145/3243734.3243786). [Online]. Available: <http://doi.acm.org/10.1145/3243734.3243786>.
- [75] W. Henecka, S. Kögl, A.-R. Sadeghi, T. Schneider, and I. Wehrenberg, “TASTY: Tool for automating secure two-party computations,” in *Proceedings of the 17th ACM Conference on Computer and Communications Security*, ser. CCS ’10, Chicago, Illinois, USA: Association for Computing Machinery, Oct. 4, 2010, pp. 451–462, ISBN: 978-1-4503-0245-6. DOI: [10.1145/1866307.1866358](https://doi.org/10.1145/1866307.1866358). [Online]. Available: <http://doi.org/10.1145/1866307.1866358>.
- [76] D. Demmler, T. Schneider, and M. Zohner, “ABY - A Framework for Efficient Mixed-Protocol Secure Two-Party Computation,” in *Proceedings 2015 Network and Distributed System Security Symposium*, San Diego, CA: Internet Society, 2015, ISBN: 978-1-891562-38-9. DOI: [10.14722/ndss.2015.23113](https://doi.org/10.14722/ndss.2015.23113). [Online]. Available: <https://www.ndss-symposium.org/ndss2015/ndss-2015-programme/aby---framework-efficient-mixed-protocol-secure-two-party-computation/>.
- [77] N. Chandran, D. Gupta, A. Rastogi, R. Sharma, and S. Tripathi, “EzPC: Programmable and Efficient Secure Two-Party Computation for Machine Learning,” in *2019 IEEE European Symposium on Security and Privacy (EuroSP)*, Jun. 2019, pp. 496–511. DOI: [10.1109/EuroSP.2019.00043](https://doi.org/10.1109/EuroSP.2019.00043). [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8806756>.
- [78] L. Braun, D. Demmler, T. Schneider, and O. Tkachenko, “MOTION - A Framework for Mixed-Protocol Multi-Party Computation,” *ACM Transactions on Privacy and Security*, vol. 25, no. 2, 8:1–8:35, Mar. 4, 2022, ISSN: 2471-2566. DOI: [10.1145/3490390](https://doi.org/10.1145/3490390). [Online]. Available: <http://doi.org/10.1145/3490390>.
- [79] Y. Bao *et al.*, “HACCLE: Metaprogramming for secure multi-party computation,” in *Proceedings of the 20th ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences*, ser. GPCE 2021, New York, NY, USA: Association for Computing Machinery, Oct. 17, 2021, pp. 130–143, ISBN: 978-1-4503-9112-2. DOI: [10.1145/3486609.3487205](https://doi.org/10.1145/3486609.3487205). [Online]. Available: <https://doi.org/10.1145/3486609.3487205>.

- [80] R. Dathathri, B. Kostova, O. Saarikivi, W. Dai, K. Laine, and M. Musuvathi, “EVA: An encrypted vector arithmetic language and compiler for efficient homomorphic computation,” in *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 2020, New York, NY, USA: Association for Computing Machinery, Jun. 11, 2020, pp. 546–561, ISBN: 978-1-4503-7613-6. DOI: [10.1145/3385412.3386023](https://doi.org/10.1145/3385412.3386023). [Online]. Available: <https://dl.acm.org/doi/10.1145/3385412.3386023>.
- [81] M. Cowan, D. Dangwal, A. Alaghi, C. Trippel, V. T. Lee, and B. Reagen, “Porcupine: A synthesizing compiler for vectorized homomorphic encryption,” in *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, ser. PLDI 2021, New York, NY, USA: Association for Computing Machinery, Jun. 18, 2021, pp. 375–389, ISBN: 978-1-4503-8391-2. DOI: [10.1145/3453483.3454050](https://doi.org/10.1145/3453483.3454050). [Online]. Available: <https://dl.acm.org/doi/10.1145/3453483.3454050>.
- [82] R. Malik, K. Sheth, and M. Kulkarni, “Coyote: A Compiler for Vectorizing Encrypted Arithmetic Circuits,” in *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, ser. ASPLOS 2023, New York, NY, USA: Association for Computing Machinery, Mar. 25, 2023, pp. 118–133, ISBN: 978-1-4503-9918-0. DOI: [10.1145/3582016.3582057](https://doi.org/10.1145/3582016.3582057). [Online]. Available: <https://dl.acm.org/doi/10.1145/3582016.3582057>.
- [83] A. Viand, P. Jattke, M. Haller, and A. Hithnawi, “HECO: Fully Homomorphic Encryption Compiler,” presented at the 32nd USENIX Security Symposium (USENIX Security 23), Anaheim, CA, USA: USENIX Association, Aug. 2023, pp. 4715–4732, ISBN: 978-1-939133-37-3. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity23/presentation/viand>.
- [84] G. Barthe *et al.*, “Formal verification of a constant-time preserving C compiler,” *Proceedings of the ACM on Programming Languages*, vol. 4, 7:1–7:30, POPL Dec. 20, 2019. DOI: [10.1145/3371075](https://doi.org/10.1145/3371075). [Online]. Available: <http://doi.org/10.1145/3371075>.
- [85] X. Leroy, “Formal verification of a realistic compiler,” *Communications of the ACM*, vol. 52, no. 7, p. 107, Jul. 1, 2009, ISSN: 00010782. DOI: [10.1145/1538788.1538814](https://doi.org/10.1145/1538788.1538814). [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1538788.1538814>.
- [86] J. B. Almeida *et al.*, “Jasmin: High-Assurance and High-Speed Cryptography,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’17, New York, NY, USA: Association for Computing Machinery, Oct. 30, 2017, pp. 1807–1823, ISBN: 978-1-4503-4946-8. DOI: [10.1145/3133956.3134078](https://doi.org/10.1145/3133956.3134078). [Online]. Available: <https://doi.org/10.1145/3133956.3134078>.

- [87] J. B. Almeida *et al.*, “The Last Mile: High-Assurance and High-Speed Cryptographic Implementations,” in *2020 IEEE Symposium on Security and Privacy (SP)*, May 2020, pp. 965–982. DOI: [10.1109/SP40000.2020.00028](https://doi.org/10.1109/SP40000.2020.00028).
- [88] K. R. M. Leino, “Dafny: An Automatic Program Verifier for Functional Correctness,” in *Logic for Programming, Artificial Intelligence, and Reasoning*, E. M. Clarke and A. Voronkov, Eds., Berlin, Heidelberg: Springer, 2010, pp. 348–370, ISBN: 978-3-642-17511-4. DOI: [10.1007/978-3-642-17511-4_20](https://doi.org/10.1007/978-3-642-17511-4_20).
- [89] S. Cauligi *et al.*, “FaCT: A DSL for timing-sensitive computation,” in *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 2019, Phoenix, AZ, USA: Association for Computing Machinery, Jun. 8, 2019, pp. 174–189, ISBN: 978-1-4503-6712-7. DOI: [10.1145/3314221.3314605](https://doi.org/10.1145/3314221.3314605). [Online]. Available: <http://doi.org/10.1145/3314221.3314605>.
- [90] O. Goldreich and R. Ostrovsky, “Software protection and simulation on oblivious RAMs,” *Journal of the ACM*, vol. 43, no. 3, pp. 431–473, May 1, 1996, ISSN: 0004-5411. DOI: [10.1145/233551.233553](https://doi.org/10.1145/233551.233553). [Online]. Available: <http://doi.org/10.1145/233551.233553>.
- [91] E. Stefanov *et al.*, “Path ORAM: An extremely simple oblivious RAM protocol,” in *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*, ser. CCS ’13, New York, NY, USA: Association for Computing Machinery, Nov. 4, 2013, pp. 299–310, ISBN: 978-1-4503-2477-9. DOI: [10.1145/2508859.2516660](https://doi.org/10.1145/2508859.2516660). [Online]. Available: <http://doi.org/10.1145/2508859.2516660>.
- [92] X. S. Wang *et al.*, “Oblivious Data Structures,” in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’14, Scottsdale, Arizona, USA: Association for Computing Machinery, Nov. 3, 2014, pp. 215–226, ISBN: 978-1-4503-2957-6. DOI: [10.1145/2660267.2660314](https://doi.org/10.1145/2660267.2660314). [Online]. Available: <http://doi.org/10.1145/2660267.2660314>.
- [93] A. Sabelfeld and A. C. Myers, “A Model for Delimited Information Release,” in *Software Security - Theories and Systems*, K. Futatsugi, F. Mizoguchi, and N. Yonezaki, Eds., ser. Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2004, pp. 174–191, ISBN: 978-3-540-37621-7. [Online]. Available: https://doi.org/10.1007/978-3-540-37621-7_9.

- [94] J. Yang, K. Yessenov, and A. Solar-Lezama, “A language for automatically enforcing privacy policies,” in *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages - POPL '12*, Philadelphia, PA, USA: ACM Press, 2012, p. 85, ISBN: 978-1-4503-1083-3. DOI: [10.1145/2103656.2103669](https://doi.org/10.1145/2103656.2103669). [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2103656.2103669>.
- [95] A. Nanevski, A. Banerjee, and D. Garg, “Dependent Type Theory for Verification of Information Flow and Access Control Policies,” *ACM Transactions on Programming Languages and Systems*, vol. 35, no. 2, 6:1–6:41, Jul. 1, 2013, ISSN: 0164-0925. DOI: [10.1145/2491522.2491523](https://doi.org/10.1145/2491522.2491523). [Online]. Available: <http://doi.org/10.1145/2491522.2491523>.
- [96] G. Barthe, B. Grégoire, and S. Zanella Béguelin, “Formal Certification of Code-based Cryptographic Proofs,” in *Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, (Savannah, GA, USA), ser. POPL '09, New York, NY, USA: ACM, 2009, pp. 90–101, ISBN: 978-1-60558-379-2. DOI: [10.1145/1480881.1480894](https://doi.org/10.1145/1480881.1480894). [Online]. Available: <http://doi.acm.org/10.1145/1480881.1480894>.
- [97] G. Barthe, B. Grégoire, S. Héraud, and S. Z. Béguelin, “Computer-Aided Security Proofs for the Working Cryptographer,” in *Advances in Cryptology – CRYPTO 2011*, P. Rogaway, Ed., ser. Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2011, pp. 71–90, ISBN: 978-3-642-22792-9.
- [98] C. Abate *et al.*, “SSProve: A Foundational Framework for Modular Cryptographic Proofs in Coq,” in *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*, Jun. 2021, pp. 1–15. DOI: [10.1109/CSF51468.2021.00048](https://doi.org/10.1109/CSF51468.2021.00048).
- [99] M. Barbosa *et al.*, “SoK: Computer-Aided Cryptography,” in *2021 IEEE Symposium on Security and Privacy (SP)*, May 2021, pp. 777–795. DOI: [10.1109/SP40001.2021.00008](https://doi.org/10.1109/SP40001.2021.00008). [Online]. Available: <https://ieeexplore.ieee.org/document/9519449>.
- [100] M. Backes, M. Maffei, and E. Mohammadi, “Computationally Sound Abstraction and Verification of Secure Multi-Party Computations,” in *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2010)*, K. Lodaya and M. Mahajan, Eds., ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 8, Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2010, pp. 352–363, ISBN: 978-3-939897-23-1. DOI: [10.4230/LIPIcs.FSTTCS.2010.352](https://doi.org/10.4230/LIPIcs.FSTTCS.2010.352). [Online]. Available: <http://drops.dagstuhl.de/opus/volltexte/2010/2877>.

- [101] H. Haagh, A. Karbyshev, S. Oechsner, B. Spitters, and P.-Y. Strub, “Computer-Aided Proofs for Multiparty Computation with Active Security,” in *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, Jul. 2018, pp. 119–131. DOI: [10.1109/CSF.2018.00016](https://doi.org/10.1109/CSF.2018.00016).

A. OMITTED RULES AND PROOFS FOR TAYPE

This chapter documents the omitted rules from [Chapter 5](#): full semantics rules ([Appendix A.1](#)), typing and kinding rules ([Appendix A.2](#)), and TAYPE-to-OIL translation rules ([Appendix A.3](#)). A proof of the totality of the translation algorithm is also included ([Appendix A.4](#)).

A.1 Semantics

[Figure A.2](#) and [Figure A.3](#) show the small-step operational semantics relation of core TAYPE: $\Sigma \vdash e \longrightarrow e'$. The global context Σ is fixed in the rules, so we elide it for brevity. We also elide type annotations in conditionals and pattern matching expressions as they are not relevant in these rules. The auxiliary relation used in S-OMATCH is exactly the same as the one in $\lambda_{\text{O}ADT}$ ([Figure 3.9](#)). S-CTX refers to evaluation contexts \mathcal{E} shown in [Figure A.1](#).

EVALUATION CONTEXTS

$$\begin{aligned}
 \mathcal{E} ::= & \square \hat{\times} \tau \mid \hat{\omega} \hat{\times} \square \mid \square \hat{+} \tau \mid \hat{\omega} \hat{+} \square \\
 & \mid \text{let } x : \tau = \square \text{ in } e \mid e \square \mid \square v \mid C \square \mid \hat{T} \square \\
 & \mid \text{if } \square \text{ then } e \text{ else } e \mid \text{mux } \square e e \mid \text{mux } v \square e \mid \text{mux } v v \square \\
 & \mid \widehat{\text{if}} \square \text{ then } e \text{ else } e \mid \widehat{\text{if}} v \text{ then } \square \text{ else } e \mid \widehat{\text{if}} v \text{ then } v \text{ else } \square \\
 & \mid (\square, e) \mid (v, \square) \mid \langle \square, e \rangle \mid \langle v, \square \rangle \mid \hat{t}_b \langle \square \rangle e \mid \hat{t}_b \langle \hat{\omega} \rangle \square \\
 & \mid \square \oplus e \mid v \oplus \square \mid \square \hat{\oplus} e \mid v \hat{\oplus} \square \\
 & \mid \text{match } \square \text{ with } \overline{C} x \Rightarrow e \mid \text{match } \square \text{ with } (x_1, x_2) \Rightarrow e \\
 & \mid \widehat{\text{match}} \square \text{ with } x \Rightarrow e \mid x \Rightarrow e \mid \widehat{\text{match}} \square \text{ with } [x_1, x_2] \Rightarrow e \\
 & \mid \hat{B}\#s \square \mid \hat{Z}\#s \square \mid \hat{Z}\#r \square \mid \text{tape } \square \mid \uparrow \square
 \end{aligned}$$

Figure A.1. Core TAYPE evaluation context

S-ADD and S-OADD simply evaluate the integer operations according to the denotational domain. Most rules involving promotions (e.g., S-SECINTPROM, S-APPPROM and S-ADDPROM) require the resulting expressions also get promoted, because the \top leakage label should be preserved in each step. However, this is not explicitly necessary for S-IFPROM, as both branches of a conditional with promoted discriminée already have \top labels, which is enforced by the typing rules (e.g., the side condition in T-IFNODEP). S-PMATCHPROM and S-MATCHPROM are similar: the expressions these rules step to do not need to be promoted again, although their pattern variables are substituted by promoted values.

$$\boxed{e \longrightarrow e'}$$

$$\begin{array}{c}
\text{S-CTX} \\
\frac{e \longrightarrow e'}{\mathcal{E}[e] \longrightarrow \mathcal{E}[e']} \\
\\
\text{S-APP} \quad \frac{}{(\lambda x: \iota \tau \Rightarrow e) v \longrightarrow [v/x]e} \quad \text{S-LET} \quad \frac{}{\text{let } x: \iota \tau = v \text{ in } e \longrightarrow [v/x]e} \\
\\
\text{S-IF} \quad \frac{}{\text{if } b \text{ then } e_1 \text{ else } e_2 \longrightarrow \text{ite}(b, e_1, e_2)} \\
\\
\text{S-PMATCH} \quad \frac{}{\text{match } (v_1, v_2) \text{ with } (x_1, x_2) \Rightarrow e \longrightarrow [v_2/x_2] [v_1/x_1]e} \\
\\
\text{S-MATCH} \quad \frac{}{\text{match } C_i v \text{ with } \bar{C} x \Rightarrow e \longrightarrow [v/x]e_i} \quad \text{S-OADT} \quad \frac{\text{obliv } \hat{T} (x: \tau) = \tau' \in \Sigma}{\hat{T} v \longrightarrow [v/x]\tau'} \quad \text{S-FUN} \quad \frac{\text{fn } x: \iota \tau = e \in \Sigma}{x \longrightarrow e} \\
\\
\text{S-SECBOOL} \quad \frac{}{\widehat{\mathbb{B}}\#s b \longrightarrow [b]} \quad \text{S-SECINT} \quad \frac{}{\widehat{\mathbb{Z}}\#s n \longrightarrow [n]} \quad \text{S-ADD} \quad \frac{}{n_1 \oplus n_2 \longrightarrow [[n_1 \oplus n_2]]} \quad \text{S-OADD} \quad \frac{}{[n_1] \hat{\oplus} [n_2] \longrightarrow [[n_1 \oplus n_2]]} \\
\\
\text{S-OINJ} \quad \frac{}{\widehat{\iota}_b \langle \widehat{\omega} \rangle \hat{v} \longrightarrow [\iota_b \langle \widehat{\omega} \rangle \hat{v}]} \quad \text{S-MUX} \quad \frac{}{\text{mux } [b] v_1 v_2 \longrightarrow \text{ite}(b, v_1, v_2)} \\
\\
\text{S-OPMATCH} \quad \frac{}{\widehat{\text{match}} [v_1, v_2] \text{ with } [x_1, x_2] \Rightarrow e \longrightarrow [v_2/x_2] [v_1/x_1]e} \\
\\
\text{S-OIF} \quad \frac{}{\widehat{\mathcal{E}}[\widehat{\text{if}} [b] \text{ then } v_1 \text{ else } v_2] \longrightarrow \widehat{\text{if}} [b] \text{ then } \widehat{\mathcal{E}}[v_1] \text{ else } \widehat{\mathcal{E}}[v_2]} \\
\\
\text{S-OMATCH} \quad \frac{\hat{v}_1 \Leftarrow \widehat{\omega}_1 \quad \hat{v}_2 \Leftarrow \widehat{\omega}_2}{\widehat{\text{match}} [\iota_b \langle \widehat{\omega}_1 + \widehat{\omega}_2 \rangle \hat{v}] \text{ with } x \Rightarrow e_1 \mid x \Rightarrow e_2 \longrightarrow \widehat{\text{if}} [b] \text{ then } \text{ite}(b, [\hat{v}/x]e_1, [\hat{v}_1/x]e_1) \text{ else } \text{ite}(b, [\hat{v}_2/x]e_2, [\hat{v}/x]e_2)}
\end{array}$$

Figure A.2. Core TAYPE semantics rules

$$\boxed{e \longrightarrow e'}$$

S-TAPEOIF

$$\frac{}{\text{tape } (\widehat{\text{if}} [b] \text{ then } v_1 \text{ else } v_2) \longrightarrow \text{mux } [b] (\text{tape } v_1) (\text{tape } v_2)}$$

S-TAPEPROM

$$\frac{}{\text{tape } (\uparrow v) \longrightarrow v}$$

S-RETADD₁

$$\frac{}{(\widehat{\mathbb{Z}}\#r [n_1]) \oplus (\widehat{\mathbb{Z}}\#r [n_2]) \longrightarrow \widehat{\mathbb{Z}}\#r ([n_1] \hat{\oplus} [n_2])}$$

S-RETADD₂

$$\frac{}{(\widehat{\mathbb{Z}}\#r [n_1]) \oplus (\uparrow n_2) \longrightarrow \widehat{\mathbb{Z}}\#r ([n_1] \hat{\oplus} (\widehat{\mathbb{Z}}\#s n_2))}$$

S-RETADD₃

$$\frac{}{(\uparrow n_1) \oplus (\widehat{\mathbb{Z}}\#r [n_2]) \longrightarrow \widehat{\mathbb{Z}}\#r ((\widehat{\mathbb{Z}}\#s n_1) \hat{\oplus} [n_2])}$$

S-SECRETINT

$$\frac{}{\widehat{\mathbb{Z}}\#s (\widehat{\mathbb{Z}}\#r [n]) \longrightarrow \uparrow [n]}$$

S-SECINTPROM

$$\frac{}{\widehat{\mathbb{Z}}\#s (\uparrow n) \longrightarrow \uparrow (\widehat{\mathbb{Z}}\#s n)}$$

S-SECBOOLPROM

$$\frac{}{\widehat{\mathbb{B}}\#s (\uparrow b) \longrightarrow \uparrow (\widehat{\mathbb{B}}\#s b)}$$

S-ADDPROM

$$\frac{}{(\uparrow n_1) \oplus (\uparrow n_2) \longrightarrow \uparrow (n_1 \oplus n_2)}$$

S-APPPROM

$$\frac{}{(\uparrow (\lambda x : \iota \tau \Rightarrow e)) v \longrightarrow \uparrow ([v/x] e)}$$

S-IFPROM

$$\frac{}{\text{if } \uparrow b \text{ then } e_1 \text{ else } e_2 \longrightarrow \text{ite}(b, e_1, e_2)}$$

S-PMATCHPROM

$$\frac{}{\text{match } \uparrow (v_1, v_2) \text{ with } (x_1, x_2) \Rightarrow e \longrightarrow [\uparrow v_2/x_2] [\uparrow v_1/x_1] e}$$

S-MATCHPROM

$$\frac{}{\text{match } \uparrow C_i v \text{ with } \overline{C} x \Rightarrow e \longrightarrow [\uparrow v/x] e_i}$$

Figure A.3. Core TAYPE semantics rules (cont.)

A.2 Type System

Programs in TAYPE are typed using a pair of typing and kinding judgments: $\Sigma; \Gamma \vdash e :_l \tau$ and $\Sigma; \Gamma \vdash \tau :: \kappa$. The global context Σ is elided for brevity, as it is fixed in all typing and kinding rules. [Figure A.4](#) shows the kinding rules, while [Figure A.5](#) and [Figure A.6](#) present the typing rules.

$$\boxed{\Gamma \vdash \tau :: \kappa}$$

$$\begin{array}{c}
 \text{K-SUB} \\
 \frac{\Gamma \vdash \tau :: \kappa \quad \kappa \sqsubseteq \kappa'}{\Gamma \vdash \tau :: \kappa'} \\
 \\
 \text{K-ADT} \\
 \frac{\text{data } T = \overline{C} \tau \in \Sigma}{\Gamma \vdash T :: *^P} \\
 \\
 \text{K-OADT} \\
 \frac{\text{obliv } \hat{T} (x:\tau) = \tau' \in \Sigma \quad \Gamma \vdash e :_{\perp} \tau}{\Gamma \vdash \hat{T} e :: *^0} \\
 \\
 \text{K-UNIT} \quad \text{K-BOOL} \quad \text{K-OBOOL} \quad \text{K-INT} \quad \text{K-OINT} \\
 \frac{}{\Gamma \vdash \mathbf{1} :: *^A} \quad \frac{}{\Gamma \vdash \mathbf{B} :: *^P} \quad \frac{}{\Gamma \vdash \hat{\mathbf{B}} :: *^0} \quad \frac{}{\Gamma \vdash \mathbf{Z} :: *^P} \quad \frac{}{\Gamma \vdash \hat{\mathbf{Z}} :: *^0} \\
 \\
 \text{K-PI} \quad \text{K-PROD} \\
 \frac{\Gamma \vdash \tau_1 :: * \quad x :_l \tau_1, \Gamma \vdash \tau_2 :: *}{\Gamma \vdash \prod x :_l \tau_1, \tau_2 :: *^M} \quad \frac{\Gamma \vdash \tau_1 :: \kappa \quad \Gamma \vdash \tau_2 :: \kappa}{\Gamma \vdash \tau_1 \times \tau_2 :: \kappa \sqcup *^P} \\
 \\
 \text{K-OPROD} \quad \text{K-OSUM} \\
 \frac{\Gamma \vdash \tau_1 :: *^0 \quad \Gamma \vdash \tau_2 :: *^0}{\Gamma \vdash \tau_1 \hat{\times} \tau_2 :: *^0} \quad \frac{\Gamma \vdash \tau_1 :: *^0 \quad \Gamma \vdash \tau_2 :: *^0}{\Gamma \vdash \tau_1 \hat{+} \tau_2 :: *^0} \\
 \\
 \text{K-LET} \quad \text{K-IF} \\
 \frac{\Gamma \vdash e :_{\perp} \tau \quad x :_{\perp} \tau, \Gamma \vdash \tau' :: *^0}{\Gamma \vdash \text{let } x :_{\perp} \tau = e \text{ in } \tau' :: *^0} \quad \frac{\Gamma \vdash e_0 :_{\perp} \mathbf{B} \quad \Gamma \vdash \tau_1 :: *^0 \quad \Gamma \vdash \tau_2 :: *^0}{\Gamma \vdash \text{if } e_0 \text{ then } \tau_1 \text{ else } \tau_2 :: *^0} \\
 \\
 \text{K-PMATCH} \\
 \frac{\Gamma \vdash e_0 :_{\perp} \tau_1 \times \tau_2 \quad x_1 :_{\perp} \tau_1, x_2 :_{\perp} \tau_2, \Gamma \vdash \tau :: *^0}{\Gamma \vdash \text{match } e_0 \text{ with } (x_1, x_2) \Rightarrow \tau :: *^0} \\
 \\
 \text{K-MATCH} \\
 \frac{\text{data } T = \overline{C} \tau \in \Sigma \quad \Gamma \vdash e_0 :_{\perp} T \quad \forall i. x :_{\perp} \tau_i, \Gamma \vdash \tau'_i :: *^0}{\Gamma \vdash \text{match } e_0 \text{ with } \overline{C} x \Rightarrow \tau' :: *^0}
 \end{array}$$

Figure A.4. Core TAYPE kinding rules

$\Gamma \vdash e :_l \tau$				
$\frac{\text{T-CONV} \quad \Gamma \vdash e :_l \tau \quad \tau \equiv \tau' \quad \Gamma \vdash \tau' :: *}{\Gamma \vdash e :_l \tau'}$	$\frac{\text{T-VAR} \quad x :_l \tau \in \Gamma}{\Gamma \vdash x :_l \tau}$	$\frac{\text{T-UNIT}}{\Gamma \vdash () :_{\perp} \mathbf{1}}$	$\frac{\text{T-LITBOOL}}{\Gamma \vdash b :_{\perp} \mathbb{B}}$	$\frac{\text{T-LITINT}}{\Gamma \vdash n :_{\perp} \mathbb{Z}}$
$\frac{\text{T-ABS} \quad x :_{l_1} \tau_1, \Gamma \vdash e :_{l_2} \tau_2 \quad \Gamma \vdash \tau_1 :: *}{\Gamma \vdash \lambda x :_{l_1} \tau_1 \Rightarrow e :_{l_2} \Pi x :_{l_1} \tau_1, \tau_2}$	$\frac{\text{T-APP} \quad \Gamma \vdash e_2 :_{l_2} \Pi x :_{l_1} \tau_1, \tau_2 \quad \Gamma \vdash e_1 :_{l_1} \tau_1}{\Gamma \vdash e_2 \ e_1 :_{l_2} [e_1/x] \tau_2}$			
$\frac{\text{T-LET} \quad \Gamma \vdash e_1 :_{l_1} \tau_1 \quad x :_{l_1} \tau_1, \Gamma \vdash e_2 :_{l_2} \tau_2}{\Gamma \vdash \text{let } x :_{l_1} \tau_1 = e_1 \text{ in } e_2 :_{l_2} [e_1/x] \tau_2}$	$\frac{\text{T-FUN} \quad \text{fn } x :_l \tau = e \in \Sigma}{\Gamma \vdash x :_l \tau}$			
$\frac{\text{T-IF} \quad \Gamma \vdash e_0 :_{\perp} \mathbb{B} \quad \Gamma \vdash e_1 :_l [\text{true}/z] \tau \quad \Gamma \vdash e_2 :_l [\text{false}/z] \tau}{\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 :_l [e_0/z] \tau}$	$\frac{\text{T-IFNODEP} \quad \Gamma \vdash e_0 :_{l_0} \mathbb{B} \quad l_0 \sqsubseteq l \quad \Gamma \vdash e_1 :_l \tau \quad \Gamma \vdash e_2 :_l \tau}{\Gamma \vdash \text{if}_{\tau} e_0 \text{ then } e_1 \text{ else } e_2 :_l \tau}$			
$\frac{\text{T-PAIR} \quad \Gamma \vdash e_1 :_l \tau_1 \quad \Gamma \vdash e_2 :_l \tau_2}{\Gamma \vdash (e_1, e_2) :_l \tau_1 \times \tau_2}$	$\frac{\text{T-PMATCH} \quad \Gamma \vdash e_0 :_{\perp} \tau_1 \times \tau_2 \quad x_1 :_{\perp} \tau_1, x_2 :_{\perp} \tau_2, \Gamma \vdash e :_l [(x_1, x_2)/z] \tau}{\Gamma \vdash \text{match } e_0 \text{ with } (x_1, x_2) \Rightarrow e :_l [e_0/z] \tau}$			
$\frac{\text{T-PMATCHNODEP} \quad \Gamma \vdash e_0 :_{l_0} \tau_1 \times \tau_2 \quad l_0 \sqsubseteq l \quad x_1 :_{l_0} \tau_1, x_2 :_{l_0} \tau_2, \Gamma \vdash e :_l \tau}{\Gamma \vdash \text{match}_{\tau} e_0 \text{ with } (x_1, x_2) \Rightarrow e :_l \tau}$	$\frac{\text{T-ADD} \quad \Gamma \vdash e_1 :_l \mathbb{Z} \quad \Gamma \vdash e_2 :_l \mathbb{Z}}{\Gamma \vdash e_1 \oplus e_2 :_l \mathbb{Z}}$	$\frac{\text{T-CTOR} \quad \text{data } T = \overline{C} \ \tau \in \Sigma \quad \Gamma \vdash e :_l \tau_i}{\Gamma \vdash C_i \ e :_l T}$		
$\frac{\text{T-MATCH} \quad \text{data } T = \overline{C} \ \tau \in \Sigma \quad \Gamma \vdash e_0 :_{\perp} T \quad \forall i. x :_{\perp} \tau_i, \Gamma \vdash e_i :_l [C_i \ x/z] \tau'}{\Gamma \vdash \text{match } e_0 \text{ with } \overline{C} \ x \Rightarrow e :_l [e_0/z] \tau'}$	$\frac{\text{T-MATCHNODEP} \quad \text{data } T = \overline{C} \ \tau \in \Sigma \quad \Gamma \vdash e_0 :_{l_0} T \quad l_0 \sqsubseteq l \quad \forall i. x :_{l_0} \tau_i, \Gamma \vdash e_i :_l \tau'}{\Gamma \vdash \text{match}_{\tau'} e_0 \text{ with } \overline{C} \ x \Rightarrow e :_l \tau'}$			

Figure A.5. Core TAYPE typing rules

$\Gamma \vdash e :_l \tau$

$$\begin{array}{c}
\text{T-MUX} \\
\frac{\Gamma \vdash e_0 :_{\perp} \widehat{\mathbb{B}} \quad \Gamma \vdash \tau :: *^0 \quad \Gamma \vdash e_1 :_{\perp} \tau \quad \Gamma \vdash e_2 :_{\perp} \tau}{\Gamma \vdash \text{mux } e_0 \ e_1 \ e_2 :_{\perp} \tau}
\end{array}
\qquad
\begin{array}{c}
\text{T-OPAIR} \\
\frac{\Gamma \vdash e_1 :_{\perp} \tau_1 \quad \Gamma \vdash e_2 :_{\perp} \tau_2 \quad \Gamma \vdash \tau_1 :: *^0 \quad \Gamma \vdash \tau_2 :: *^0}{\Gamma \vdash [e_1, e_2] :_{\perp} \tau_1 \widehat{\times} \tau_2}
\end{array}$$

$$\begin{array}{c}
\text{T-OPMATCH} \\
\frac{\Gamma \vdash e_0 :_{\perp} \tau_1 \widehat{\times} \tau_2 \quad x_1 :_{\perp} \tau_1, x_2 :_{\perp} \tau_2, \Gamma \vdash e :_l \tau}{\Gamma \vdash \widehat{\text{match}}_{e_0 : \tau_1 \widehat{\times} \tau_2} \text{ with } [x_1, x_2] \Rightarrow e :_l \tau}
\end{array}
\qquad
\begin{array}{c}
\text{T-OADD} \\
\frac{\Gamma \vdash e_1 :_{\perp} \widehat{\mathbb{Z}} \quad \Gamma \vdash e_2 :_{\perp} \widehat{\mathbb{Z}}}{\Gamma \vdash e_1 \widehat{\oplus} e_2 :_{\perp} \widehat{\mathbb{Z}}}
\end{array}$$

$$\begin{array}{c}
\text{T-OINJ} \\
\frac{\Gamma \vdash e :_{\perp} \text{ite}(b, \tau_1, \tau_2) \quad \Gamma \vdash \tau_1 \widehat{+} \tau_2 :: *^0}{\Gamma \vdash \widehat{!}_b \langle \tau_1 \widehat{+} \tau_2 \rangle e :_{\perp} \tau_1 \widehat{+} \tau_2}
\end{array}
\qquad
\begin{array}{c}
\text{T-OMATCH} \\
\frac{\Gamma \vdash e_0 :_{\perp} \tau_1 \widehat{+} \tau_2 \quad x :_{\perp} \tau_1, \Gamma \vdash e_1 :_{\top} \tau \quad x :_{\perp} \tau_2, \Gamma \vdash e_2 :_{\top} \tau}{\Gamma \vdash \widehat{\text{match}}_{\tau} e_0 : \tau_1 \widehat{+} \tau_2 \text{ with } x \Rightarrow e_1 \mid x \Rightarrow e_2 :_{\top} \tau}
\end{array}$$

$$\begin{array}{c}
\text{T-BOXEDLITBOOL} \\
\frac{}{\Gamma \vdash [b] :_{\perp} \widehat{\mathbb{B}}}
\end{array}
\qquad
\begin{array}{c}
\text{T-BOXEDLITINT} \\
\frac{}{\Gamma \vdash [n] :_{\perp} \widehat{\mathbb{Z}}}
\end{array}
\qquad
\begin{array}{c}
\text{T-BOXEDINJ} \\
\frac{[!_b \langle \widehat{\omega} \rangle \widehat{v}] \Leftarrow \widehat{\omega}}{\Gamma \vdash [!_b \langle \widehat{\omega} \rangle \widehat{v}] :_{\perp} \widehat{\omega}}
\end{array}
\qquad
\begin{array}{c}
\text{T-SECBOOL} \\
\frac{\Gamma \vdash e :_l \mathbb{B}}{\Gamma \vdash \widehat{\mathbb{B}} \#_s e :_l \widehat{\mathbb{B}}}
\end{array}$$

$$\begin{array}{c}
\text{T-SECINT} \\
\frac{\Gamma \vdash e :_l \mathbb{Z}}{\Gamma \vdash \widehat{\mathbb{Z}} \#_s e :_l \widehat{\mathbb{Z}}}
\end{array}
\qquad
\begin{array}{c}
\text{T-RETINT} \\
\frac{\Gamma \vdash e :_{\perp} \widehat{\mathbb{Z}}}{\Gamma \vdash \widehat{\mathbb{Z}} \#_r e :_{\top} \mathbb{Z}}
\end{array}
\qquad
\begin{array}{c}
\text{T-OIF} \\
\frac{\Gamma \vdash e_0 :_{\perp} \widehat{\mathbb{B}} \quad \Gamma \vdash e_1 :_{\top} \tau \quad \Gamma \vdash e_2 :_{\top} \tau}{\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 :_{\top} \tau}
\end{array}$$

$$\begin{array}{c}
\text{T-PROMOTE} \\
\frac{\Gamma \vdash e :_{\perp} \tau}{\Gamma \vdash \uparrow e :_{\top} \tau}
\end{array}
\qquad
\begin{array}{c}
\text{T-TAPE} \\
\frac{\Gamma \vdash e :_{\top} \tau \quad \Gamma \vdash \tau :: *^0}{\Gamma \vdash \text{tape } e :_{\perp} \tau}
\end{array}$$

Figure A.6. Core TAYPE typing rules (cont.)

A.3 Translation Algorithm

Figure A.7 translates TAYPE oblivious types to OIL expressions of size type. Figure A.8 and Figure A.9 translate TAYPE expressions to OIL expressions.

$$\boxed{\Gamma \vdash \tau \rightsquigarrow \mathbf{s}}$$

$$\begin{array}{c}
 \text{TR-UNITT} \\
 \hline
 \Gamma \vdash \mathbf{1} \rightsquigarrow \mathbf{0}
 \end{array}
 \quad
 \begin{array}{c}
 \text{TR-OB00L} \\
 \hline
 \Gamma \vdash \widehat{\mathbf{B}} \rightsquigarrow \mathbf{1}
 \end{array}
 \quad
 \begin{array}{c}
 \text{TR-OINT} \\
 \hline
 \Gamma \vdash \widehat{\mathbf{Z}} \rightsquigarrow \mathbf{1}
 \end{array}
 \quad
 \begin{array}{c}
 \text{TR-OPROD} \\
 \hline
 \Gamma \vdash \tau_1 \rightsquigarrow \mathbf{s}_1 \quad \Gamma \vdash \tau_2 \rightsquigarrow \mathbf{s}_2 \\
 \Gamma \vdash \tau_1 \widehat{\times} \tau_2 \rightsquigarrow \mathbf{s}_1 + \mathbf{s}_2
 \end{array}$$

$$\begin{array}{c}
 \text{TR-OSUM} \\
 \hline
 \Gamma \vdash \tau_1 \rightsquigarrow \mathbf{s}_1 \quad \Gamma \vdash \tau_2 \rightsquigarrow \mathbf{s}_2 \\
 \Gamma \vdash \tau_1 \widehat{+} \tau_2 \rightsquigarrow \mathbf{1} + \max \mathbf{s}_1 \ \mathbf{s}_2
 \end{array}
 \quad
 \begin{array}{c}
 \text{TR-TAPP} \\
 \hline
 \Gamma \vdash \widehat{\mathbf{T}} \ \mathbf{x} \rightsquigarrow \widehat{\mathbf{T}} \ \mathbf{x}
 \end{array}$$

$$\begin{array}{c}
 \text{TR-TLET} \\
 \hline
 \Gamma \vdash \mathbf{e} \rightsquigarrow_{\perp} \dot{\mathbf{e}} \quad \mathbf{x} :_{\perp} \tau_1, \Gamma \vdash \tau \rightsquigarrow \mathbf{s} \\
 \Gamma \vdash \mathbf{let} \ \mathbf{x} :_{\perp} \tau_1 = \mathbf{e} \ \mathbf{in} \ \tau \rightsquigarrow \mathbf{let} \ \mathbf{x} = \dot{\mathbf{e}} \ \mathbf{in} \ \mathbf{s}
 \end{array}$$

$$\begin{array}{c}
 \text{TR-TIF} \\
 \hline
 \Gamma \vdash \tau_1 \rightsquigarrow \mathbf{s}_1 \quad \Gamma \vdash \tau_2 \rightsquigarrow \mathbf{s}_2 \\
 \Gamma \vdash \mathbf{if} \ \mathbf{x}_0 \ \mathbf{then} \ \tau_1 \ \mathbf{else} \ \tau_2 \rightsquigarrow \mathbf{if} \ \mathbf{x}_0 \ \mathbf{then} \ \mathbf{s}_1 \ \mathbf{else} \ \mathbf{s}_2
 \end{array}$$

$$\begin{array}{c}
 \text{TR-TPMATCH} \\
 \hline
 \mathbf{x}_0 :_{\perp} \tau_1 \times \tau_2 \in \Gamma \quad \mathbf{x}_1 :_{\perp} \tau_1, \mathbf{x}_2 :_{\perp} \tau_2, \Gamma \vdash \tau \rightsquigarrow \mathbf{s} \\
 \Gamma \vdash \mathbf{match} \ \mathbf{x}_0 \ \mathbf{with} \ (\mathbf{x}_1, \mathbf{x}_2) \Rightarrow \tau \rightsquigarrow \mathbf{match} \ \mathbf{x}_0 \ \mathbf{with} \ (\mathbf{x}_1, \mathbf{x}_2) \Rightarrow \mathbf{s}
 \end{array}$$

$$\begin{array}{c}
 \text{TR-TMATCH} \\
 \hline
 \mathbf{data} \ \mathbf{T} = \overline{\mathbf{C}} \ \tau \in \Sigma \quad \mathbf{x}_0 :_{\perp} \mathbf{T} \in \Gamma \quad \forall i. \mathbf{x} :_{\perp} \tau_i, \Gamma \vdash \tau'_i \rightsquigarrow \mathbf{s}_i \\
 \Gamma \vdash \mathbf{match} \ \mathbf{x}_0 \ \mathbf{with} \ \overline{\mathbf{C}} \ \mathbf{x} \Rightarrow \tau' \rightsquigarrow \mathbf{match} \ \mathbf{x}_0 \ \mathbf{with} \ \overline{\mathbf{C}} \ \mathbf{x} \Rightarrow \mathbf{s}
 \end{array}$$

Figure A.7. Translating core TAYPE oblivious types to OIL sizes

$$\boxed{\Gamma \vdash e \rightsquigarrow_l \dot{e}}$$

$\frac{}{\Gamma \vdash \mathbf{b} \rightsquigarrow_{\perp} \mathbf{b}}$	$\frac{}{\Gamma \vdash \mathbf{n} \rightsquigarrow_{\perp} \mathbf{n}}$	$\frac{}{\Gamma \vdash \mathbf{x} \rightsquigarrow_l \mathbf{x}}$	$\frac{\text{TR-ABS } \mathbf{x} :_{l_1} \tau_1, \Gamma \vdash e \rightsquigarrow_l \dot{e}}{\Gamma \vdash \lambda \mathbf{x} :_{l_1} \tau_1 \Rightarrow e \rightsquigarrow_l \lambda \mathbf{x} \Rightarrow \dot{e}}$
$\frac{}{\Gamma \vdash \mathbf{x}_2 \ \mathbf{x}_1 \rightsquigarrow_l \mathbf{x}_2 \ \mathbf{x}_1}$	$\frac{}{\Gamma \vdash (\mathbf{x}_1, \mathbf{x}_2) \rightsquigarrow_l \begin{cases} (\mathbf{x}_1, \mathbf{x}_2) & \text{if } l = \perp \\ \widetilde{\text{pair}} \ \mathbf{x}_1 \ \mathbf{x}_2 & \text{if } l = \top \end{cases}}$		
$\frac{}{\Gamma \vdash \mathbf{C} \ \mathbf{x} \rightsquigarrow_l \begin{cases} \mathbf{C} \ \mathbf{x} & \text{if } l_0 = \perp \\ \widetilde{\mathbf{C}} \ \mathbf{x} & \text{if } l_0 = \top \end{cases}}$	$\frac{}{\Gamma \vdash \mathbf{x}_1 \oplus \mathbf{x}_2 \rightsquigarrow_l \begin{cases} \mathbf{x}_1 \oplus \mathbf{x}_2 & \text{if } l_0 = \perp \\ \mathbf{x}_1 \tilde{\oplus} \mathbf{x}_2 & \text{if } l_0 = \top \end{cases}}$		
$\frac{}{\Gamma \vdash \widehat{\mathbb{B}}\#\mathbf{s} \ \mathbf{x} \rightsquigarrow_l \begin{cases} \widehat{\mathbb{B}}\#\mathbf{s} \ \mathbf{x} & \text{if } l = \perp \\ \widetilde{\mathbb{B}}\#\mathbf{s} \ \mathbf{x} & \text{if } l = \top \end{cases}}$	$\frac{}{\Gamma \vdash \widehat{\mathbb{Z}}\#\mathbf{s} \ \mathbf{x} \rightsquigarrow_l \begin{cases} \widehat{\mathbb{Z}}\#\mathbf{s} \ \mathbf{x} & \text{if } l = \perp \\ \widetilde{\mathbb{Z}}\#\mathbf{s} \ \mathbf{x} & \text{if } l = \top \end{cases}}$		
$\frac{\text{TR-LET } \Gamma \vdash e_1 \rightsquigarrow_{l_1} \dot{e}_1 \quad \mathbf{x} :_{l_1} \tau_1, \Gamma \vdash e_2 \rightsquigarrow_l \dot{e}_2}{\Gamma \vdash \text{let } \mathbf{x} :_{l_1} \tau_1 = e_1 \text{ in } e_2 \rightsquigarrow_l \text{let } \mathbf{x} = \dot{e}_1 \text{ in } \dot{e}_2}$			
$\frac{\text{TR-IF } \mathbf{x}_0 :_{l_0} \mathbb{B} \in \Gamma \quad \Gamma \vdash e_1 \rightsquigarrow_l \dot{e}_1 \quad \Gamma \vdash e_2 \rightsquigarrow_l \dot{e}_2}{\Gamma \vdash \text{if}_{\tau} \ \mathbf{x}_0 \text{ then } e_1 \text{ else } e_2 \rightsquigarrow_l \begin{cases} \text{if } \mathbf{x}_0 \text{ then } \dot{e}_1 \text{ else } \dot{e}_2 & \text{if } l_0 = \perp \\ \widetilde{\text{if}} \ \widehat{\text{if}}(\tau) \ \mathbf{x}_0 \ \dot{e}_1 \ \dot{e}_2 & \text{if } l_0 = \top \end{cases}}$			
$\frac{\text{TR-PMATCH } \mathbf{x}_0 :_{l_0} \tau_1 \times \tau_2 \in \Gamma \quad \mathbf{x}_1 :_{l_0} \tau_1, \mathbf{x}_2 :_{l_0} \tau_2, \Gamma \vdash e \rightsquigarrow_l \dot{e}}{\Gamma \vdash \text{match}_{\tau} \ \mathbf{x}_0 \text{ with } (\mathbf{x}_1, \mathbf{x}_2) \Rightarrow e \rightsquigarrow_l \begin{cases} \text{match } \mathbf{x}_0 \text{ with } (\mathbf{x}_1, \mathbf{x}_2) \Rightarrow \dot{e} & \text{if } l_0 = \perp \\ \widetilde{\text{match}}_{\times} \ \text{prom}(\tau_1) \ \text{prom}(\tau_2) \ \widehat{\text{if}}(\tau) \ \mathbf{x}_0 \ (\lambda \mathbf{x}_1 \lambda \mathbf{x}_2 \Rightarrow \dot{e}) & \text{if } l_0 = \top \end{cases}}$			
$\frac{\text{TR-MATCH } \text{data } \mathbb{T} = \overline{\mathbf{C}} \ \tau \in \Sigma \quad \mathbf{x}_0 :_{l_0} \mathbb{T} \in \Gamma \quad \forall i. \mathbf{x} :_{l_0} \tau_i, \Gamma \vdash e_i \rightsquigarrow_l \dot{e}_i}{\Gamma \vdash \text{match}_{\tau} \ \mathbf{x}_0 \text{ with } \overline{\mathbf{C}} \ \mathbf{x} \Rightarrow e \rightsquigarrow_l \begin{cases} \text{match } \mathbf{x}_0 \text{ with } \overline{\mathbf{C}} \ \mathbf{x} \Rightarrow \dot{e} & \text{if } l_0 = \perp \\ \widetilde{\text{match}}_{\tau} \ \widehat{\text{if}}(\tau) \ \mathbf{x}_0 \ (\lambda \mathbf{x} \Rightarrow \dot{e}) & \text{if } l_0 = \top \end{cases}}$			

Figure A.8. Translating core TAYPE expressions to OIL expressions

A.4 Totality of the Translation Algorithm

The totality theorem relies on a few lemmas that we state as follows.

Lemma A.4.1 (Regularity). *If $\Gamma \vdash e :_l \tau$, then $\Gamma \vdash \tau :: \kappa$ for some kind κ .*

Lemma A.4.2 (Conversion of typing context). *If $x :_{l'} \tau_1, \Gamma \vdash e :_l \tau$ and $\tau_1 \equiv \tau_2$ with $\Gamma \vdash \tau_2 :: *$, then $x :_{l'} \tau_2, \Gamma \vdash e :_l \tau$.*

*If $x :_{l'} \tau_1, \Gamma \vdash \tau :: \kappa$ and $\tau_1 \equiv \tau_2$ with $\Gamma \vdash \tau_2 :: *$, then $x :_{l'} \tau_2, \Gamma \vdash \tau :: \kappa$.*

Lemma A.4.3 (Equality of equivalent public types).

1. *If $\Gamma \vdash \tau :: *$ and $\tau \equiv \mathbb{B}$, then $\tau = \mathbb{B}$.*
2. *If $\Gamma \vdash \tau :: *$, $\Gamma \vdash \mathbb{T} :: *$ and $\tau \equiv \mathbb{T}$, then $\tau = \mathbb{T}$.*
3. *If $\Gamma \vdash \tau :: *$, $\Gamma \vdash \tau_1 \times \tau_2 :: *$ and $\tau \equiv \tau_1 \times \tau_2$, then $\tau = \tau'_1 \times \tau'_2$ for some τ'_1 and τ'_2 such that $\tau_1 \equiv \tau'_1$ and $\tau_2 \equiv \tau'_2$.*

We will also use the expected typing and kinding inversion lemmas. The proofs of these lemmas are omitted here, which are available in the Coq formalization [49]. The last lemma essentially says if a well-kinded type is equivalent to a well-kinded public type, then they are equal. The case of product also falls into this interpretation if we consider the type former \times itself public, even though \times can be used to connect non-public components. In other words, the “heads” of the public type formers are equal.

Now we prove the following totality theorem.

Theorem A.4.4 (Totality of translation). *If $\Gamma \vdash e :_l \tau$ and e is in ANF, then $\Gamma \vdash e \rightsquigarrow_l \dot{e}$ for some OIL expression \dot{e} .*

*If $\Gamma \vdash \tau :: *^0$ and τ is in ANF, then $\Gamma \vdash \tau \rightsquigarrow \mathfrak{s}$ for some OIL expression \mathfrak{s} .*

Proof. By induction on the ANF structure of e and τ , we prove these two statements simultaneously. Every subterm of an ANF term is also in ANF, so we can always apply the induction hypotheses to subterms.

It is trivial to prove the cases on base oblivious types, literals, variable, constructor, (function and type) application, public and oblivious pairs, integer operations, primitive

sections and retractions, atomic conditional `mux` and the `tape` operation. We simply apply their corresponding translation rules, which do not have any assumptions.

The cases on public types (booleans, integers and public products) and function type vacuously hold, since they are not obviously kinded.

Case on lambda abstraction $\lambda x :_{l_1} \tau_1 \Rightarrow e$: By the well-typedness assumption and the inversion lemma, we have $x :_{l_1} \tau_1, \Gamma \vdash e :_l \tau'$ for some τ' . It then follows that $x :_{l_1} \tau_1, \Gamma \vdash e \rightsquigarrow_l \dot{e}$ for some \dot{e} , by the induction hypothesis. Applying TR-ABS concludes this case.

Case on let binding `let` $x :_{l_1} \tau_1 = e_1$ `in` e_2 : The proof for the first part (when this expression is well-typed) is similar to the case on lambda abstraction using rule TR-LET. On the other hand, if this is obviously kinded, we have $\Gamma \vdash e_1 :_{\perp} \tau_1$ and $x :_{\perp} \tau_1, \Gamma \vdash e_2 :: *^0$. By the induction hypotheses, $\Gamma \vdash e_1 \rightsquigarrow_{\perp} \dot{e}_1$ and $x :_{\perp} \tau_1, \Gamma \vdash e_2 \rightsquigarrow s$ for some \dot{e}_1 and s . We then discharge this case by TR-TLET.

Case on oblivious injection $\widehat{l}_b \langle \tau \rangle x$: By the well-typedness assumption and the inversion lemma, we know $l = \perp$ and τ is some $\tau_1 \widehat{+} \tau_2$ such that $\Gamma \vdash \tau_1 \widehat{+} \tau_2 :: *^0$, which implies that τ_1 and τ_2 are also obviously kinded by the kinding inversion lemma. The conclusion follows by TR-OINJ.

Case on promotion $\uparrow x$: We have $\Gamma \vdash x :_{\perp} \tau$ from the assumption. By the inversion lemma for variables, we get $x :_{\perp} \tau' \in \Gamma$ for some equivalent type τ' . This case then follows by TR-PROMOTE.

Cases on oblivious product and sum types, their pattern matching expressions and leaky conditional are similar.

The remaining cases on conditional, product and ADT pattern matching are the trickiest ones. We show the proofs for conditional and product pattern matching here; the case of of ADT pattern matching is analogous.

In the case of well-typed conditional, `if` _{τ} x_0 `then` e_1 `else` e_2 , we know e_1 and e_2 are well-typed with some label l , which discharges the two corresponding conditions in TR-IF by the induction hypotheses. By the inversion lemma for variables and the fact that $\Gamma \vdash x_0 :_{l_0} \mathbb{B}$, we have $x_0 :_{l_0} \tau' \in \Gamma$, for some τ' such that $\tau' \equiv \mathbb{B}$ and τ' is well-kinded. However, τ' must equal to \mathbb{B} by Lemma A.4.3, thus $x_0 :_{l_0} \mathbb{B} \in \Gamma$, required by TR-IF. The case of obviously kinded conditional is similar using TR-TIF.

In the case of well-typed product pattern matching, $\text{match}_\tau \mathbf{x}_0 \text{ with } (\mathbf{x}_1, \mathbf{x}_2) \Rightarrow \mathbf{e}$, we know $\Gamma \vdash \mathbf{x}_0 :_{l_0} \tau'_1 \times \tau'_2$, and \mathbf{e} is well-typed under the context $\mathbf{x}_1 :_{l_0} \tau'_1, \mathbf{x}_2 :_{l_0} \tau'_2, \Gamma$. By the inversion lemma for variables, $\mathbf{x}_0 :_{l_0} \tau' \in \Gamma$ for some well-kinded τ' and $\tau' \equiv \tau'_1 \times \tau'_2$. As $\tau'_1 \times \tau'_2$ is well-kinded by [Lemma A.4.1](#), it follows from [Lemma A.4.3](#) that $\tau' = \tau_1 \times \tau_2$ for some τ_1 and τ_2 , such that $\tau'_1 \equiv \tau_1$ and $\tau'_2 \equiv \tau_2$. We then conclude this case by applying TR-PMATCH, whose translation premise is obtained by [Lemma A.4.2](#) and the induction hypothesis. The case for obviously kinded product pattern matching is similar. \square

B. OMITTED RULES AND PROOFS FOR TAYPSI

This chapter documents the omitted rules from [Chapter 6](#): full logical refinement definitions and declarative lifting rules ([Appendix B.1](#)), as well as full algorithmic lifting rules ([Appendix B.2](#)). [Appendix B.3](#) proves the correctness of declarative lifting and the soundness of algorithmic lifting.

B.1 Declarative Lifting

[Figure B.1](#) gives the full definitions of the logical refinement from [Section 6.3.6](#). In particular, [Figure B.1](#) defines an interpretation for typing contexts that range over specification types, $\mathcal{G}_n[\Gamma]$. The codomain of this interpretation is substitutions (σ) of related value pairs.

$$\begin{array}{l}
 \boxed{\mathcal{V}_n[\theta]} \\
 \mathcal{V}_n[\mathbf{1}] = \mathcal{V}_n[\mathbb{B}] = \mathcal{V}_n[\mathbb{T}] = \{ (v, v') \mid 0 < n \implies v = v' \} \\
 \mathcal{V}_n[\widehat{\mathbb{B}}] = \{ (b, [b']) \mid 0 < n \implies b = b' \} \\
 \mathcal{V}_n[\Psi\widehat{\mathbb{T}}] = \{ (v, \langle k, \widehat{v} \rangle) \mid 0 < n \implies r \ k \ \widehat{v} \longrightarrow^* v \} \\
 \mathcal{V}_n[\theta_1 \times \theta_2] = \{ ((v_1, v_2), (v'_1, v'_2)) \mid (v_1, v'_1) \in \mathcal{V}_n[\theta_1] \wedge (v_2, v'_2) \in \mathcal{V}_n[\theta_2] \} \\
 \mathcal{V}_n[\theta_1 \rightarrow \theta_2] = \left\{ (\lambda x : [\theta_1] \Rightarrow e, \lambda x : \theta_1 \Rightarrow e') \mid \begin{array}{l} \forall i < n. \forall (v, v') \in \mathcal{V}_i[\theta_1]. \\ ([v/x] e, [v'/x] e') \in \mathcal{E}_i[\theta_2] \end{array} \right\} \\
 \boxed{\mathcal{E}_n[\theta]} \\
 \mathcal{E}_n[\theta] = \{ (e, e') \mid \forall i < n. \forall v'. e' \longrightarrow^i v' \implies \exists v. e \longrightarrow^* v \wedge (v, v') \in \mathcal{V}_{n-i}[\theta] \} \\
 \boxed{\mathcal{G}_n[\Gamma]} \\
 \mathcal{G}_n[\cdot] = \{ \emptyset \} \quad \mathcal{G}_n[x : \theta, \Gamma] = \{ \sigma[x \mapsto (v, v')] \mid \sigma \in \mathcal{G}_n[\Gamma] \wedge (v, v') \in \mathcal{V}_n[\theta] \}
 \end{array}$$

Figure B.1. Refinement as logical relation

[Figure B.2](#) presents the full rules of the declarative lifting relation, $\mathcal{S}; \mathcal{L}; \Sigma; \Gamma \vdash e : \theta \triangleright \dot{e}$. We elide most contexts as they are fixed in these rules for brevity.

$$\boxed{\Gamma \vdash e : \theta \triangleright \dot{e}}$$

$$\begin{array}{c}
\text{L-UNIT} \\
\hline
\Gamma \vdash () : \mathbf{1} \triangleright ()
\end{array}
\quad
\begin{array}{c}
\text{L-LIT} \\
\hline
\Gamma \vdash b : \mathbb{B} \triangleright b
\end{array}
\quad
\begin{array}{c}
\text{L-VAR} \\
x : \theta \in \Gamma \\
\hline
\Gamma \vdash x : \theta \triangleright x
\end{array}
\quad
\begin{array}{c}
\text{L-FUN} \\
x : \theta \triangleright \dot{x} \in \mathcal{L} \\
\hline
\Gamma \vdash x : \theta \triangleright \dot{x}
\end{array}$$

$$\begin{array}{c}
\text{L-ABS} \\
x : \theta_1, \Gamma \vdash e : \theta_2 \triangleright \dot{e} \\
\hline
\Gamma \vdash \lambda x : [\theta_1] \Rightarrow e : \theta_1 \rightarrow \theta_2 \triangleright \lambda x : \theta_1 \Rightarrow \dot{e}
\end{array}
\quad
\begin{array}{c}
\text{L-APP} \\
\Gamma \vdash e_2 : \theta_1 \rightarrow \theta_2 \triangleright \dot{e}_2 \quad \Gamma \vdash e_1 : \theta_1 \triangleright \dot{e}_1 \\
\hline
\Gamma \vdash e_2 \ e_1 : \theta_2 \triangleright \dot{e}_2 \ \dot{e}_1
\end{array}$$

$$\begin{array}{c}
\text{L-PAIR} \\
\Gamma \vdash e_1 : \theta_1 \triangleright \dot{e}_1 \quad \Gamma \vdash e_2 : \theta_2 \triangleright \dot{e}_2 \\
\hline
\Gamma \vdash (e_1, e_2) : \theta_1 \times \theta_2 \triangleright (\dot{e}_1, \dot{e}_2)
\end{array}
\quad
\begin{array}{c}
\text{L-PROJ} \\
\Gamma \vdash e : \theta_1 \times \theta_2 \triangleright \dot{e} \\
\hline
\Gamma \vdash \pi_b \ e : \text{ite}(b, \theta_1, \theta_2) \triangleright \pi_b \ \dot{e}
\end{array}$$

$$\begin{array}{c}
\text{L-LET} \\
\Gamma \vdash e_1 : \theta_1 \triangleright \dot{e}_1 \quad x : \theta_1, \Gamma \vdash e_2 : \theta_2 \triangleright \dot{e}_2 \\
\hline
\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \theta_2 \triangleright \text{let } x = \dot{e}_1 \text{ in } \dot{e}_2
\end{array}
\quad
\begin{array}{c}
\text{L-COERCE} \\
\Gamma \vdash e : \theta \triangleright \dot{e} \quad \theta \mapsto \theta' \triangleright \uparrow \\
\hline
\Gamma \vdash e : \theta' \triangleright \uparrow \dot{e}
\end{array}$$

$$\begin{array}{c}
\text{L-CTOR}_1 \\
\text{data } T = \overline{C} \ \eta \in \Sigma \quad \Gamma \vdash e : \eta_i \triangleright \dot{e} \\
\hline
\Gamma \vdash C_i \ e : T \triangleright C_i \ \dot{e}
\end{array}
\quad
\begin{array}{c}
\text{L-CTOR}_2 \\
\widehat{C}_i : \theta_i \rightarrow \Psi \widehat{T} \in \mathcal{S}_I \quad \Gamma \vdash e : \theta_i \triangleright \dot{e} \\
\hline
\Gamma \vdash C_i \ e : \Psi \widehat{T} \triangleright \widehat{C}_i \ \dot{e}
\end{array}$$

$$\begin{array}{c}
\text{L-IF}_1 \\
\Gamma \vdash e_0 : \mathbb{B} \triangleright \dot{e}_0 \quad \Gamma \vdash e_1 : \theta \triangleright \dot{e}_1 \quad \Gamma \vdash e_2 : \theta \triangleright \dot{e}_2 \\
\hline
\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : \theta \triangleright \text{if } \dot{e}_0 \text{ then } \dot{e}_1 \text{ else } \dot{e}_2
\end{array}$$

$$\begin{array}{c}
\text{L-IF}_2 \\
\Gamma \vdash e_0 : \widehat{\mathbb{B}} \triangleright \dot{e}_0 \quad \uparrow \theta \triangleright \widehat{\text{ite}} \quad \Gamma \vdash e_1 : \theta \triangleright \dot{e}_1 \quad \Gamma \vdash e_2 : \theta \triangleright \dot{e}_2 \\
\hline
\Gamma \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : \theta \triangleright \widehat{\text{ite}} \ \dot{e}_0 \ \dot{e}_1 \ \dot{e}_2
\end{array}$$

$$\begin{array}{c}
\text{L-MATCH}_1 \\
\text{data } T = \overline{C} \ \eta \in \Sigma \quad \Gamma \vdash e_0 : T \triangleright \dot{e}_0 \quad \forall i. x : \eta_i, \Gamma \vdash e_i : \theta' \triangleright \dot{e}_i \\
\hline
\Gamma \vdash \text{match } e_0 \text{ with } \overline{C} \ x \Rightarrow e : \theta' \triangleright \text{match } \dot{e}_0 \text{ with } \overline{C} \ x \Rightarrow \dot{e}
\end{array}$$

$$\begin{array}{c}
\text{L-MATCH}_2 \\
\widehat{\text{match}} : \Psi \widehat{T} \rightarrow (\overline{\theta} \rightarrow \theta') \rightarrow \theta' \in \mathcal{S}_E \quad \Gamma \vdash e_0 : \Psi \widehat{T} \triangleright \dot{e}_0 \quad \forall i. x : \theta_i, \Gamma \vdash e_i : \theta' \triangleright \dot{e}_i \\
\hline
\Gamma \vdash \text{match } e_0 \text{ with } \overline{C} \ x \Rightarrow e : \theta' \triangleright \widehat{\text{match}} \ \dot{e}_0 \ (\lambda x : \theta \Rightarrow \dot{e})
\end{array}$$

Figure B.2. Declarative lifting rules

B.2 Algorithmic Lifting

Figure B.3 and Figure B.4 presents the full rules of the lifting algorithm, $\Sigma; \Gamma \vdash e : \eta \sim X \triangleright \dot{e} \mid \mathcal{C}$.

$$\boxed{\Gamma \vdash e : \eta \sim X \triangleright \dot{e} \mid \mathcal{C}}$$

$$\begin{array}{c}
 \text{A-UNIT} \\
 \hline
 \Gamma \vdash () : \mathbf{1} \sim X \triangleright () \mid X = \mathbf{1} \\
 \\
 \text{A-LIT} \\
 \hline
 \Gamma \vdash b : \mathbb{B} \sim X \triangleright b \mid X = \mathbb{B} \\
 \\
 \text{A-VAR} \qquad \qquad \qquad \text{A-FUN} \\
 \frac{x : \eta \sim X \in \Gamma}{\Gamma \vdash x : \eta \sim X' \triangleright \% \uparrow(X, X'; x) \mid \% \uparrow(X, X')} \qquad \frac{\text{fn } x : \eta = e \in \Sigma}{\Gamma \vdash x : \eta \sim X \triangleright \% x(X) \mid \% x(X)} \\
 \\
 \text{A-ABS} \\
 \frac{X_1, X_2 \text{ fresh} \quad x : \eta_1 \sim X_1, \Gamma \vdash e : \eta_2 \sim X_2 \triangleright \dot{e} \mid \mathcal{C}}{\Gamma \vdash \lambda x : \eta_1 \Rightarrow e : \eta_1 \rightarrow \eta_2 \sim X \triangleright \lambda x : X_1 \Rightarrow \dot{e} \mid X_1 \in [\eta_1], X_2 \in [\eta_2], X = X_1 \rightarrow X_2, \mathcal{C}} \\
 \\
 \text{A-APP} \\
 \frac{X_1 \text{ fresh} \quad x_2 : \eta_1 \rightarrow \eta_2 \sim X \in \Gamma \quad \Gamma \vdash x_1 : \eta_1 \sim X_1 \triangleright \dot{e}_1 \mid \mathcal{C}}{\Gamma \vdash x_2 \ x_1 : \eta_2 \sim X_2 \triangleright x_2 \ \dot{e}_1 \mid X_1 \in [\eta_1], X = X_1 \rightarrow X_2, \mathcal{C}} \\
 \\
 \text{A-LET} \\
 \frac{X_1 \text{ fresh} \quad \Gamma \vdash e_1 : \eta_1 \sim X_1 \triangleright \dot{e}_1 \mid \mathcal{C}_1 \quad x : \eta_1 \sim X_1, \Gamma \vdash e_2 : \eta_2 \sim X_2 \triangleright \dot{e}_2 \mid \mathcal{C}_2}{\Gamma \vdash \text{let } x : \eta_1 = e_1 \text{ in } e_2 : \eta_2 \sim X_2 \triangleright \text{let } x : X_1 = \dot{e}_1 \text{ in } \dot{e}_2 \mid X_1 \in [\eta_1], \mathcal{C}_1, \mathcal{C}_2} \\
 \\
 \text{A-PAIR} \\
 \frac{x_1 : \eta_1 \sim X_1 \in \Gamma \quad x_2 : \eta_2 \sim X_2 \in \Gamma}{\Gamma \vdash (x_1, x_2) : \eta_1 \times \eta_2 \sim X \triangleright (x_1, x_2) \mid X = X_1 \times X_2} \\
 \\
 \text{A-PROJ} \\
 \frac{\text{ite}(b, X_2, X_1) \text{ fresh} \quad x : \eta_1 \times \eta_2 \sim X \in \Gamma}{\Gamma \vdash \pi_b \ x : \text{ite}(b, \eta_1, \eta_2) \sim \text{ite}(b, X_1, X_2) \triangleright \pi_b \ x \mid \text{ite}(b, X_2, X_1) \in [\text{ite}(b, \eta_2, \eta_1)], X = X_1 \times X_2}
 \end{array}$$

Figure B.3. Algorithmic lifting rules

$$\boxed{\Gamma \vdash e : \eta \sim X \triangleright \dot{e} \mid \mathcal{C}}$$

$$\text{A-IF} \quad \frac{x_0 : \mathbb{B} \sim X_0 \in \Gamma \quad \Gamma \vdash e_1 : \eta \sim X \triangleright \dot{e}_1 \mid \mathcal{C}_1 \quad \Gamma \vdash e_2 : \eta \sim X \triangleright \dot{e}_2 \mid \mathcal{C}_2}{\Gamma \vdash \text{if } x_0 \text{ then } e_1 \text{ else } e_2 : \eta \sim X \triangleright \%ite(X_0, X; x_0, \dot{e}_1, \dot{e}_2) \mid \%ite(X_0, X), \mathcal{C}_1, \mathcal{C}_2}$$

$$\text{A-CTOR} \quad \frac{\text{data } T = \overline{C} \overline{\eta} \in \Sigma \quad X_i \text{ fresh} \quad \Gamma \vdash x : \eta_i \sim X_i \triangleright \dot{e} \mid \mathcal{C}}{\Gamma \vdash C_i \ x : T \sim X \triangleright \%C_i(X_i, X; \dot{e}) \mid X_i \in [\eta_i], \%C_i(X_i, X), \mathcal{C}}$$

$$\text{A-MATCH} \quad \frac{\overline{X} \text{ fresh} \quad \text{data } T = \overline{C} \overline{\eta} \in \Sigma \quad x_0 : T \sim X_0 \in \Gamma \quad \forall i. x : \eta_i \sim X_i, \Gamma \vdash e_i : \eta' \sim X' \triangleright \dot{e}_i \mid \mathcal{C}_i}{\Gamma \vdash \text{match } x_0 \text{ with } \overline{C} \ x \Rightarrow e : \eta' \sim X' \triangleright \%match(X_0, \overline{X}, X'; x_0, \overline{e}) \mid \overline{X} \in [\overline{\eta}], \%match(X_0, \overline{X}, X'), \overline{C}}$$

Figure B.4. Algorithmic lifting rules (cont.)

B.3 Metatheory of Lifting

We say a lifting context \mathcal{L} is well-typed (under Σ) if and only if, for any $x : \theta \triangleright \dot{x} \in \mathcal{L}$, $\Sigma; \cdot \vdash x : [\theta]$ and $\Sigma; \cdot \vdash \dot{x} : \theta$. A lifting context is derivable, denoted by $\vdash \mathcal{L}$, if and only if, for any $x : \theta \triangleright \dot{x} \in \mathcal{L}$, $\text{fn } x : [\theta] = e \in \Sigma$ and $\text{fn } \dot{x} : \theta = \dot{e} \in \Sigma$ for some e and \dot{e} , such that $\mathcal{S}; \mathcal{L}; \Sigma; \cdot \vdash e : \theta \triangleright \dot{e}$. A lifting context is n -valid, denoted by $\vDash_n \mathcal{L}$, if and only if, for any $x : \theta \triangleright \dot{x} \in \mathcal{L}$, $(x, \dot{x}) \in \mathcal{E}_n[[\theta]]$. If $\vDash_n \mathcal{L}$ for any n , we say \mathcal{L} is valid, denoted by $\vDash \mathcal{L}$. Obviously, derivability or validity implies well-typedness.

The lemmas and theorems in this section assume a well-typed global context Σ . We also (explicitly or implicitly) use some standard results about $\lambda_{\text{OADT}\Psi}$ that are proved in the Coq formalization [60]: weakening lemmas, substitution lemmas, preservation theorem and canonical forms of values. We do not state these lemmas formally here, as they are all standard.

Lemma B.3.1. $[\eta] = \eta$.

Proof. By routine induction on η . □

Lemma B.3.2 (Regularity of mergeability). $\uparrow \theta \triangleright \widehat{\text{ite}}$ implies $\cdot \vdash \widehat{\text{ite}} : \widehat{\mathbb{B}} \rightarrow \theta \rightarrow \theta \rightarrow \theta$.

Proof. By routine induction on the derivation of $\lambda\theta \triangleright \widehat{\text{ite}}$ and applying typing rules as needed. \square

Lemma B.3.3 (Regularity of coercibility). $\theta \rightsquigarrow \theta' \triangleright \uparrow$ implies $[\theta] = [\theta']$ and $\cdot \vdash \uparrow : \theta \rightarrow \theta'$.

Proof. By routine induction on the derivation of $\theta \rightsquigarrow \theta' \triangleright \uparrow$ and applying typing rules as needed. \square

The following regularity theorem ensures that all lifted expressions are well-typed, which in turn provides the security guarantees, as well-typed programs are oblivious by the obliviousness theorem.

Theorem B.3.4 (Regularity of declarative lifting). *Suppose \mathcal{L} is well-typed and $\mathcal{S}; \mathcal{L}; \Sigma; \Gamma \vdash e : \theta \triangleright \dot{e}$. We have $\Sigma; [\Gamma] \vdash e : [\theta]$ and $\Sigma; \Gamma \vdash \dot{e} : \theta$.*

Proof. By induction on the derivation of the declarative lifting judgment.

The cases on L-UNIT, L-LIT, L-VAR and L-FUN are trivial.

The cases on L-ABS, L-APP, L-LET, L-PAIR, L-PROJ, L-IF₁, L-CTOR₁ and L-MATCH₁ are straightforward to prove, as they are simply congruence cases. L-CTOR₁ and L-MATCH₁ also rely on [Lemma B.3.1](#). Here we show only the proof on L-ABS. Other cases are similar.

To prove $\Gamma \vdash \lambda x : \theta_1 \Rightarrow \dot{e} : \theta_1 \rightarrow \theta_2$, by T-ABS, we need to show $x : \theta_1, \Gamma \vdash \dot{e} : \theta_2$, which follows immediately by the induction hypothesis. The side condition of θ_1 being well-kinded under Γ is immediate from the fact that specification types are well-kinded and the weakening lemma. The other part of this case, $\Gamma \vdash \lambda x : [\theta_1] \Rightarrow e : [\theta_1 \rightarrow \theta_2] = [\theta_1] \rightarrow [\theta_2]$ proceeds similarly.

Case L-IF₂: By [Lemma B.3.2](#), we know $\cdot \vdash \widehat{\text{ite}} : \widehat{\mathbb{B}} \rightarrow \theta \rightarrow \theta \rightarrow \theta$. Applying T-APP and the induction hypothesis, we have $\Gamma \vdash \widehat{\text{ite}} \dot{e}_0 \dot{e}_1 \dot{e}_2 : \theta$, as desired. The other half is straightforward.

Cases L-CTOR₂ and L-MATCH₂ rely on the properties of structures \mathcal{S}_I and \mathcal{S}_E , but otherwise proceed similarly to other cases.

Case L-COERCE: By the induction hypothesis, we have $[\Gamma] \vdash e : [\theta]$ and $\Gamma \vdash \dot{e} : \theta$. Since $[\theta] = [\theta']$ by [Lemma B.3.3](#), $[\Gamma] \vdash e : [\theta']$ as required. On the other hand, $\cdot \vdash \uparrow : \theta \rightarrow \theta'$ by [Lemma B.3.3](#). It follows immediately that $\Gamma \vdash \uparrow \dot{e} : \theta'$. \square

Recall that a substitution in the denotation of typing contexts maps names to pairs of related values. In the next lemma (and the rest of this section), we use the notations σ_1 and σ_2 for the substitution projections, i.e., substitutions that only use the first or the second component of the pairs.

Lemma B.3.5 (Multi-substitution). *Let $\sigma \in \mathcal{G}_n[\Gamma]$. If $\Gamma \vdash \mathbf{e} : \theta$, then $\cdot \vdash \sigma_2(\mathbf{e}) : \theta$, and if $[\Gamma] \vdash \mathbf{e} : \eta$, then $\cdot \vdash \sigma_1(\mathbf{e}) : \eta$.*

Proof. By routine induction on the structure of Γ , and applying substitution lemma when needed. Note that θ and η do not contain any local variables, so substitution on these types does nothing. \square

Lemma B.3.6. *Suppose $\cdot \vdash \eta :: *^P$. We have $(\mathbf{v}, \mathbf{v}') \in \mathcal{V}_n[\eta]$ for any $\mathbf{v} : \eta$. Conversely, if $(\mathbf{v}, \mathbf{v}') \in \mathcal{V}_n[\eta]$ for $n > 0$, then $\mathbf{v} = \mathbf{v}'$.*

Proof. By routine induction on the structure of η . \square

Lemma B.3.7 (Anti-monotonicity). *If $m \leq n$, then $\mathcal{V}_n[\theta] \subseteq \mathcal{V}_m[\theta]$, and $\mathcal{E}_n[\theta] \subseteq \mathcal{E}_m[\theta]$, and $\mathcal{G}_n[\Gamma] \subseteq \mathcal{G}_m[\Gamma]$, and $\vDash_n \mathcal{L} \implies \vDash_m \mathcal{L}$.*

Proof. To prove the case of value interpretation, proceed by routine induction on θ . The cases of other interpretations are straightforward. \square

Lemma B.3.8 (Correctness of mergeability). *Suppose $\lambda\theta \triangleright \widehat{\text{ite}}$. If $(\mathbf{v}_1, \dot{\mathbf{v}}_1) \in \mathcal{V}_n[\theta]$ and $(\mathbf{v}_2, \dot{\mathbf{v}}_2) \in \mathcal{V}_n[\theta]$, and $\widehat{\text{ite}}[\mathbf{b}] \dot{\mathbf{v}}_1 \dot{\mathbf{v}}_2 \longrightarrow^* \dot{\mathbf{v}}$, then $(\text{ite}(\mathbf{b}, \mathbf{v}_1, \mathbf{v}_2), \dot{\mathbf{v}}) \in \mathcal{V}_n[\theta]$.*

Proof. By induction on the derivation of mergeability. We assume $n > 0$ because otherwise it is trivial. The cases when θ is $\mathbf{1}$ or $\widehat{\mathbb{B}}$ are trivial. The case of product type is routine, similar to the case of function type.

Case on function type $\theta_1 \rightarrow \theta_2$: By assumption,

- $\mathbf{v}_1 = \lambda \mathbf{z} \Rightarrow \mathbf{e}_1$ for some \mathbf{e}_1
- $\dot{\mathbf{v}}_1 = \lambda \mathbf{z} \Rightarrow \dot{\mathbf{e}}_1$ for some $\dot{\mathbf{e}}_1$
- $\mathbf{v}_2 = \lambda \mathbf{z} \Rightarrow \mathbf{e}_2$ for some \mathbf{e}_2

- $\dot{v}_2 = \lambda z \Rightarrow \dot{e}_2$ for some \dot{e}_2

Suppose $\widehat{\text{ite}} [b] \dot{v}_1 \dot{v}_2 \longrightarrow^* \dot{v}$, i.e., $\dot{v} = \lambda z \Rightarrow \widehat{\text{ite}}_2 [b] (\dot{v}_1 z) (\dot{v}_2 z)$. Suppose $i < n$, and $(u, \dot{u}) \in \mathcal{V}_i[\theta_1]$. We want to show $(\text{ite}(b, [u/z]e_1, [u/z]e_2), \widehat{\text{ite}}_2 [b] (\dot{v}_1 \dot{u}) (\dot{v}_2 \dot{u})) \in \mathcal{E}_i[\theta_2]$. Suppose $j < i$. We have the following trace:

$$\begin{aligned}
\widehat{\text{ite}}_2 [b] (\dot{v}_1 \dot{u}) (\dot{v}_2 \dot{u}) &\longrightarrow^1 \widehat{\text{ite}}_2 [b] (\dot{v}_1 \dot{u}) ([\dot{u}/z]\dot{e}_2) \\
&\longrightarrow^{j_2} \widehat{\text{ite}}_2 [b] (\dot{v}_1 \dot{u}) \dot{w}_2 \\
&\longrightarrow^1 \widehat{\text{ite}}_2 [b] ([\dot{u}/z]\dot{e}_1) \dot{w}_2 \\
&\longrightarrow^{j_1} \widehat{\text{ite}}_2 [b] \dot{w}_1 \dot{w}_2 \\
&\longrightarrow^{j_3} \dot{w}
\end{aligned}$$

where $j = j_1 + j_2 + j_3 + 2$, with $[\dot{u}/z]\dot{e}_2 \longrightarrow^{j_2} \dot{w}_2$ and $[\dot{u}/z]\dot{e}_1 \longrightarrow^{j_1} \dot{w}_1$. We instantiate the assumptions with i and u and \dot{u} to get:

- $([u/z]e_1, [\dot{u}/z]\dot{e}_1) \in \mathcal{E}_i[\theta_2]$
- $([u/z]e_2, [\dot{u}/z]\dot{e}_2) \in \mathcal{E}_i[\theta_2]$

It then follows that:

- $[u/z]e_1 \longrightarrow^* w_1$ for some value w_1
- $(w_1, \dot{w}_1) \in \mathcal{V}_{i-j_1}[\theta_2]$
- $[u/z]e_2 \longrightarrow^* w_2$ for some value w_2
- $(w_2, \dot{w}_2) \in \mathcal{V}_{i-j_2}[\theta_2]$

Thus $\text{ite}(b, [u/z]e_1, [u/z]e_2) \longrightarrow^* \text{ite}(b, w_1, w_2)$. It remains to show $(\text{ite}(b, w_1, w_2), \dot{w}) \in \mathcal{V}_{i-j}[\theta_2]$, but it follows immediately by the induction hypothesis and [Lemma B.3.7](#).

Case on Ψ -type $\Psi\hat{\text{T}}$: By assumption,

- $\dot{v}_1 = \langle k_1, \hat{v}_1 \rangle$ for some values k_1 and \hat{v}_1
- $r k_1 \hat{v}_1 \longrightarrow^* v_1$

- $\hat{v}_2 = \langle k_2, \hat{v}_2 \rangle$ for some values k_2 and \hat{v}_2
- $r \ k_2 \ \hat{v}_2 \longrightarrow^* v_2$

Suppose $\widehat{\text{ite}} \ [b] \ \hat{v}_1 \ \hat{v}_2 \longrightarrow^* \dot{v}$. We get the following trace:

$$\begin{aligned}
& \widehat{\text{ite}} \ [b] \ \hat{v}_1 \ \hat{v}_2 \longrightarrow^* \langle k, \text{mux} \ [b] \ (\uparrow (\pi_1 \ \hat{v}_1) \ k \ (\pi_2 \ \hat{v}_1)) \\
& \qquad \qquad \qquad (\uparrow (\pi_1 \ \hat{v}_2) \ k \ (\pi_2 \ \hat{v}_2)) \rangle \\
& \longrightarrow^* \text{let } k = k_1 \sqcup k_2 \text{ in } \dots \\
& \longrightarrow^* \langle k, \text{mux} \ [b] \ (\uparrow (\pi_1 \ \hat{v}_1) \ k \ (\pi_2 \ \hat{v}_1)) \ (\dots) \rangle \\
& \longrightarrow^* \langle k, \text{mux} \ [b] \ (\uparrow k_1 \ k \ \hat{v}_1) \ (\dots) \rangle \\
& \longrightarrow^* \langle k, \text{mux} \ [b] \ \hat{v}'_1 \ (\uparrow (\pi_1 \ \hat{v}_2) \ k \ (\pi_2 \ \hat{v}_2)) \rangle \\
& \longrightarrow^* \langle k, \text{mux} \ [b] \ \hat{v}'_1 \ (\uparrow k_2 \ k \ \hat{v}_2) \rangle \\
& \longrightarrow^* \langle k, \text{mux} \ [b] \ \hat{v}'_1 \ \hat{v}'_2 \rangle \\
& \longrightarrow^* \langle k, \text{ite}(b, \hat{v}'_1, \hat{v}'_2) \rangle = \dot{v}
\end{aligned}$$

with $k_1 \sqcup k_2 \longrightarrow^* k$, $\uparrow k_1 \ k \ \hat{v}_1 \longrightarrow^* \hat{v}'_1$ and $\uparrow k_2 \ k \ \hat{v}_2 \longrightarrow^* \hat{v}'_2$. Because $v_1 \preceq k_1$ by A-O₂, and $k_1 \sqsubseteq k$ by A-R₂, we have $v_1 \preceq k$ by A-R₃. It follows that $r \ k \ \hat{v}'_1 \longrightarrow^* v_1$ by A-R₄. Similarly, we get $r \ k \ \hat{v}'_2 \longrightarrow^* v_2$. Hence, $r \ k \ \text{ite}(b, \hat{v}'_1, \hat{v}'_2) \longrightarrow^* \text{ite}(b, v_1, v_2)$. That is to say $(\text{ite}(b, v_1, v_2), \dot{v}) \in \mathcal{V}_n[\Psi\hat{\text{T}}]$, as desired. \square

Lemma B.3.9 (Correctness of coercibility). *Suppose $\theta \rightsquigarrow \theta' \triangleright \uparrow$. If $(v, \dot{v}) \in \mathcal{V}_n[\theta]$ and $\uparrow \dot{v} \longrightarrow^* \dot{v}'$, then $(v, \dot{v}') \in \mathcal{V}_n[\theta']$.*

Proof. By induction on the derivation of coercibility. We only consider $n > 0$ since it is otherwise trivial. The cases on identity coercion and boolean coercion are trivial. The case on $\Psi\hat{\text{T}} \rightsquigarrow \Psi\hat{\text{T}}'$ is immediate from A-C₁. The case on product type is routine, similar to the case of function type.

Case on $\mathsf{T} \rightsquigarrow \Psi\widehat{\mathsf{T}}$: Suppose $(\mathbf{v}, \dot{\mathbf{v}}) \in \mathcal{V}_n[\llbracket \mathsf{T} \rrbracket]$, i.e., $\mathbf{v} = \dot{\mathbf{v}}$ by definition, and $\uparrow \dot{\mathbf{v}} = \uparrow \mathbf{v} \longrightarrow^* \dot{\mathbf{v}}'$. We have the following trace:

$$\begin{aligned} \uparrow \mathbf{v} &\longrightarrow \langle \nu \ \mathbf{v}, \mathbf{s} \ (\nu \ \mathbf{v}) \ \mathbf{v} \rangle \\ &\longrightarrow^* \langle \mathbf{k}, \mathbf{s} \ \mathbf{k} \ \mathbf{v} \rangle \\ &\longrightarrow^* \langle \mathbf{k}, \widehat{\mathbf{v}} \rangle = \dot{\mathbf{v}}' \end{aligned}$$

with $\nu \ \mathbf{v} \longrightarrow^* \mathbf{k}$ and $\mathbf{s} \ \mathbf{k} \ \mathbf{v} \longrightarrow^* \widehat{\mathbf{v}}$. Knowing $\mathbf{v} \preceq \mathbf{k}$ by A-O₃, we have $\mathbf{r} \ \mathbf{k} \ \widehat{\mathbf{v}} \longrightarrow^* \mathbf{v}$ by A-O₁. But that is $(\mathbf{v}, \langle \mathbf{k}, \widehat{\mathbf{v}} \rangle) \in \mathcal{V}_n[\llbracket \Psi\widehat{\mathsf{T}} \rrbracket]$, as desired.

Case on $\theta_1 \rightarrow \theta_2 \rightsquigarrow \theta'_1 \rightarrow \theta'_2$: By assumption,

- $\mathbf{v} = \lambda \mathbf{y} \Rightarrow \mathbf{e}$ for some \mathbf{e}
- $\dot{\mathbf{v}} = \lambda \mathbf{y} \Rightarrow \dot{\mathbf{e}}$ for some $\dot{\mathbf{e}}$

Suppose $\uparrow \dot{\mathbf{v}} \longrightarrow^* \dot{\mathbf{v}}'$, i.e., $\dot{\mathbf{v}}' = \lambda \mathbf{y} \Rightarrow \uparrow_2(\dot{\mathbf{v}} \ (\uparrow_1 \mathbf{y}))$. We want to show $(\mathbf{v}, \dot{\mathbf{v}}') \in \mathcal{V}_n[\llbracket \theta'_1 \rightarrow \theta'_2 \rrbracket]$. Towards this goal fix $i < n$ and $(\mathbf{v}_1, \dot{\mathbf{v}}_1) \in \mathcal{V}_i[\llbracket \theta'_1 \rrbracket]$. It suffices to show $([\mathbf{v}_1/\mathbf{y}] \mathbf{e}, \uparrow_2(\dot{\mathbf{v}} \ (\uparrow_1 \dot{\mathbf{v}}_1))) \in \mathcal{E}_i[\llbracket \theta'_2 \rrbracket]$. Suppose $j < i$. We have the following trace:

$$\begin{aligned} \uparrow_2(\dot{\mathbf{v}} \ (\uparrow_1 \dot{\mathbf{v}}_1)) &\longrightarrow^{j_1} \uparrow_2(\dot{\mathbf{v}} \ \dot{\mathbf{v}}_1) \\ &\longrightarrow^1 \uparrow_2([\dot{\mathbf{v}}_1/\mathbf{y}] \dot{\mathbf{e}}) \\ &\longrightarrow^{j_2} \uparrow_2 \dot{\mathbf{v}}_2 \\ &\longrightarrow^{j_3} \dot{\mathbf{v}}'_2 \end{aligned}$$

where $j = j_1 + j_2 + j_3 + 1$, with $\uparrow_1 \dot{\mathbf{v}}_1 \longrightarrow^{j_1} \dot{\mathbf{v}}_1$ and $[\dot{\mathbf{v}}_1/\mathbf{y}] \dot{\mathbf{e}} \longrightarrow^{j_2} \dot{\mathbf{v}}_2$. By the induction hypothesis, we get $(\mathbf{v}_1, \dot{\mathbf{v}}_1) \in \mathcal{V}_i[\llbracket \theta_1 \rrbracket]$. It follows, by assumption, that $([\mathbf{v}_1/\mathbf{y}] \mathbf{e}, [\dot{\mathbf{v}}_1/\mathbf{y}] \dot{\mathbf{e}}) \in \mathcal{E}_i[\llbracket \theta_2 \rrbracket]$. Hence $[\mathbf{v}_1/\mathbf{y}] \mathbf{e} \longrightarrow^* \mathbf{v}_2$ for some \mathbf{v}_2 such that $(\mathbf{v}_2, \dot{\mathbf{v}}_2) \in \mathcal{V}_{i-j_2}[\llbracket \theta_2 \rrbracket]$. It then follows by the induction hypothesis that $(\mathbf{v}_2, \dot{\mathbf{v}}_2) \in \mathcal{V}_{i-j_2}[\llbracket \theta'_2 \rrbracket]$, which concludes the proof by [Lemma B.3.7](#). \square

Theorem B.3.10 (Correctness of declarative lifting of expressions). *Suppose $\mathcal{S}; \mathcal{L}; \Sigma; \Gamma \vdash \mathbf{e} : \theta \triangleright \dot{\mathbf{e}}$ and $\vDash_n \mathcal{L}$. Given a substitution $\sigma \in \mathcal{G}_n[\llbracket \Gamma \rrbracket]$, we have $(\sigma_1(\mathbf{e}), \sigma_2(\dot{\mathbf{e}})) \in \mathcal{E}_n[\llbracket \theta \rrbracket]$.*

Proof. By induction on the derivation of the declarative lifting judgment. Note that we do not fix step n in the induction. That means we can instantiate it when applying the induction hypothesis, but we also need to discharge $\vDash_n \mathcal{L}$ for the n we pick. We will not explicitly show this side-condition for brevity, because it is simply a consequence of [Lemma B.3.7](#) as long as we pick the same or a smaller step. The well-typedness side-conditions in logical relations are omitted, because they are trivial from [Theorem B.3.4](#) and preservation theorem. In addition, we only consider step $n > 0$ because the cases when $n = 0$ are always trivial.

The cases L-UNIT, L-LIT, L-VAR and L-FUN are trivial.

Case T-ABS: We need to show $(\lambda x: [\theta] \Rightarrow \sigma_1(e), \lambda x: \theta \Rightarrow \sigma_2(\dot{e})) \in \mathcal{E}_n[[\theta_1 \rightarrow \theta_2]]$. Suppose $i < n$. Because lambda abstraction can not take step, $i = 0$, and it suffices to show $(\lambda x: [\theta] \Rightarrow \sigma_1(e), \lambda x: \theta \Rightarrow \sigma_2(\dot{e})) \in \mathcal{V}_i[[\theta_1 \rightarrow \theta_2]]$. To this end fix $i < n$ and suppose that $(v, \dot{v}) \in \mathcal{V}_i[[\theta_1]]$ for some v and \dot{v} . We know $\sigma[x \mapsto (v, \dot{v})] \in \mathcal{G}_i[[x : \theta_1, \Gamma]]$, by assumption and [Lemma B.3.7](#). Now we specialize the induction hypothesis with $n = i$ and substitution $\sigma[x \mapsto (v, \dot{v})]$ to get $(\sigma_1[x \mapsto v](e), \sigma_2[x \mapsto \dot{v}](\dot{e})) = ([v/x]\sigma_1(e), [\dot{v}/x]\sigma_2(\dot{e})) \in \mathcal{E}_i[[\theta_2]]$, as required.

Case T-APP: We need to show $(\sigma_1(e_2) \sigma_1(e_1), \sigma_2(\dot{e}_2) \sigma_2(\dot{e}_1)) \in \mathcal{E}_n[[\theta_2]]$. Suppose that $i < n$, and $\sigma_2(\dot{e}_2) \sigma_2(\dot{e}_1) \longrightarrow^i \dot{v}$. By the semantics definition, we have the following reduction trace:

$$\begin{aligned} \sigma_2(\dot{e}_2) \sigma_2(\dot{e}_1) &\longrightarrow^{i_1} \sigma_2(\dot{e}_2) \dot{v}_1 \\ &\longrightarrow^{i_2} \dot{v}_2 \dot{v}_1 \\ &\longrightarrow^1 [\dot{v}_1/x] \dot{e}'_2 \\ &\longrightarrow^{i_3} \dot{v} \end{aligned}$$

where \dot{v}_2 is $\lambda x \Rightarrow \dot{e}'_2$ for some \dot{e}'_2 , and $i = i_1 + i_2 + i_3 + 1$, with $\sigma_2(\dot{e}_1) \longrightarrow^{i_1} \dot{v}_1$ and $\sigma_2(\dot{e}_2) \longrightarrow^{i_2} \dot{v}_2$.

Instantiating the two induction hypotheses with step n and σ , we get:

- $(\sigma_1(e_2), \sigma_2(\dot{e}_2)) \in \mathcal{E}_n[[\theta_1 \rightarrow \theta_2]]$
- $(\sigma_1(e_1), \sigma_2(\dot{e}_1)) \in \mathcal{E}_n[[\theta_1]]$

It follows that:

- $\sigma_1(\mathbf{e}_2) \longrightarrow^* \mathbf{v}_2$ for some \mathbf{v}_2
- $(\mathbf{v}_2, \dot{\mathbf{v}}_2) \in \mathcal{V}_{n-i_2}[\![\theta_1 \rightarrow \theta_2]\!]$
- $\sigma_1(\mathbf{e}_1) \longrightarrow^* \mathbf{v}_1$ for some \mathbf{v}_1
- $(\mathbf{v}_1, \dot{\mathbf{v}}_1) \in \mathcal{V}_{n-i_1}[\![\theta_1]\!]$

where \mathbf{v}_2 is $\lambda \mathbf{x} \Rightarrow \mathbf{e}'_2$ for some \mathbf{e}'_2 . It follows that $([\mathbf{v}_1/\mathbf{x}] \mathbf{e}'_2, [\dot{\mathbf{v}}_1/\mathbf{x}] \dot{\mathbf{e}}'_2) \in \mathcal{E}_{n-i_1-i_2-1}[\![\theta_2]\!]$, by specializing the value interpretation of function type to $n - i_1 - i_2 - 1$. As $i_3 < n - i_1 - i_2 - 1$, we know $[\mathbf{v}_1/\mathbf{x}] \mathbf{e}'_2 \longrightarrow^* \mathbf{v}$ for some \mathbf{v} , such that $(\mathbf{v}, \dot{\mathbf{v}}) \in \mathcal{V}_{n-i_1-i_2-i_3-1}[\![\theta_2]\!]$. Therefore,

$$\begin{aligned}
\sigma_1(\mathbf{e}_2) \ \sigma_1(\mathbf{e}_1) &\longrightarrow^* \sigma_1(\mathbf{e}_2) \ \mathbf{v}_1 \\
&\longrightarrow^* \mathbf{v}_2 \ \mathbf{v}_1 \\
&\longrightarrow [\mathbf{v}_1/\mathbf{x}] \mathbf{e}'_2 \\
&\longrightarrow^* \mathbf{v}
\end{aligned}$$

and $(\mathbf{v}, \dot{\mathbf{v}}) \in \mathcal{V}_{n-i}[\![\theta_2]\!]$, as desired.

Case L-LET: Observe the trace of the lifted expression:

$$\begin{aligned}
\text{let } \mathbf{x} = \sigma_2(\dot{\mathbf{e}}_1) \text{ in } \sigma_2(\dot{\mathbf{e}}_2) &\longrightarrow^{i_1} \text{let } \mathbf{x} = \dot{\mathbf{v}}_1 \text{ in } \sigma_2(\dot{\mathbf{e}}_2) \\
&\longrightarrow^1 [\dot{\mathbf{v}}_1/\mathbf{x}] \sigma_2(\dot{\mathbf{e}}_2) \\
&\longrightarrow^{i_2} \dot{\mathbf{v}}
\end{aligned}$$

The rest of the proof is similar to L-APP and L-ABS, with the induction hypothesis of the `let`-body specialized to step $n - i_1$.

Case L-PAIR: Similarly with the trace:

$$\begin{aligned}
(\sigma_2(\dot{\mathbf{e}}_1), \sigma_2(\dot{\mathbf{e}}_2)) &\longrightarrow^{i_1} (\dot{\mathbf{v}}_1, \sigma_2(\dot{\mathbf{e}}_2)) \\
&\longrightarrow^{i_2} (\dot{\mathbf{v}}_1, \dot{\mathbf{v}}_2)
\end{aligned}$$

Case L-PROJ: Similarly with the trace:

$$\begin{aligned} \pi_b \sigma_2(\dot{e}) &\longrightarrow^{i_1} \pi_b (\dot{v}_1, \dot{v}_2) \\ &\longrightarrow^1 \mathbf{ite}(b, \dot{v}_1, \dot{v}_2) \end{aligned}$$

Case L-IF₁: Similarly with the trace:

$$\begin{aligned} \mathbf{if} \sigma_2(\dot{e}_0) \mathbf{then} \sigma_2(\dot{e}_1) \mathbf{else} \sigma_2(\dot{e}_2) &\longrightarrow^{i_0} \mathbf{if} b \mathbf{then} \sigma_2(\dot{e}_1) \mathbf{else} \sigma_2(\dot{e}_2) \\ &\longrightarrow^1 \mathbf{ite}(b, \sigma_2(\dot{e}_1), \sigma_2(\dot{e}_2)) \\ &\longrightarrow^j \dot{v} \end{aligned}$$

Case L-CTOR₁: Similarly with the trace $C_i \sigma_2(\dot{e}) \longrightarrow^i C_i \dot{v}$. We also need to apply [Lemma B.3.6](#) to complete the proof.

Case L-MATCH₁: Similarly with the trace:

$$\begin{aligned} \mathbf{match} \sigma_2(\dot{e}_0) \mathbf{with} \overline{C \ x \Rightarrow \sigma_2(\dot{e})} &\longrightarrow^{j_0} \mathbf{match} C_i \ \dot{v}_0 \mathbf{with} \overline{C \ x \Rightarrow \sigma_2(\dot{e})} \\ &\longrightarrow^1 [\dot{v}_0/x] \sigma_2(\dot{e}_i) \\ &\longrightarrow^{j_1} \dot{v} \end{aligned}$$

Similar to L-CTOR₁, [Lemma B.3.6](#) is used.

Case L-IF₂: Suppose $i < n$. We have the following trace:

$$\begin{aligned} \widehat{\mathbf{ite}} \sigma_2(\dot{e}_0) \sigma_2(\dot{e}_1) \sigma_2(\dot{e}_2) &\longrightarrow^{i_2} \widehat{\mathbf{ite}} \sigma_2(\dot{e}_0) \sigma_2(\dot{e}_1) \dot{v}_2 \\ &\longrightarrow^{i_1} \widehat{\mathbf{ite}} \sigma_2(\dot{e}_0) \dot{v}_1 \dot{v}_2 \\ &\longrightarrow^{i_0} \widehat{\mathbf{ite}} [b] \dot{v}_1 \dot{v}_2 \\ &\longrightarrow^{i_3} \dot{v} \end{aligned}$$

where $i = i_0 + i_1 + i_2 + i_3$, with $\sigma_2(\dot{e}_0) \longrightarrow^{i_0} [b]$, $\sigma_2(\dot{e}_1) \longrightarrow^{i_1} \dot{v}_1$ and $\sigma_2(\dot{e}_2) \longrightarrow^{i_2} \dot{v}_2$. By the induction hypothesis, it follows that:

- $\sigma_1(\mathbf{e}_0) \longrightarrow^* b$

- $\sigma_1(\mathbf{e}_1) \longrightarrow^* \mathbf{v}_1$ for some \mathbf{v}_1
- $(\mathbf{v}_1, \dot{\mathbf{v}}_1) \in \mathcal{V}_{n-i_1}[\![\theta]\!]$
- $\sigma_1(\mathbf{e}_2) \longrightarrow^* \mathbf{v}_2$ for some \mathbf{v}_2
- $(\mathbf{v}_2, \dot{\mathbf{v}}_2) \in \mathcal{V}_{n-i_2}[\![\theta]\!]$

Hence we have:

$$\begin{aligned}
& \text{if } \sigma_1(\mathbf{e}_0) \text{ then } \sigma_1(\mathbf{e}_1) \text{ else } \sigma_1(\mathbf{e}_2) \longrightarrow^* \text{if } \mathbf{b} \text{ then } \sigma_1(\mathbf{e}_1) \text{ else } \sigma_1(\mathbf{e}_2) \\
& \longrightarrow \text{ite}(\mathbf{b}, \sigma_1(\mathbf{e}_1), \sigma_1(\mathbf{e}_2)) \\
& \longrightarrow^* \text{ite}(\mathbf{b}, \mathbf{v}_1, \mathbf{v}_2)
\end{aligned}$$

and $(\text{ite}(\mathbf{b}, \mathbf{v}_1, \mathbf{v}_2), \dot{\mathbf{v}}) \in \mathcal{V}_{n-i}[\![\theta]\!]$ by [Lemma B.3.8](#), as required.

Case L-CTOR₂: Similarly with the trace:

$$\begin{aligned}
\widehat{\mathbf{C}}_i \sigma_2(\dot{\mathbf{e}}) & \longrightarrow^{j_1} \widehat{\mathbf{C}}_i \dot{\mathbf{v}}_i \\
& \longrightarrow^{j_2} \langle \mathbf{k}, \widehat{\mathbf{v}} \rangle = \dot{\mathbf{v}}
\end{aligned}$$

We complete the proof by A-I₁.

Case L-MATCH₂: Similarly with the trace:

$$\begin{aligned}
\widehat{\text{match}} \sigma_2(\dot{\mathbf{e}}_0) \overline{(\lambda \mathbf{x} \Rightarrow \sigma_2(\dot{\mathbf{e}}))} & \longrightarrow^{j_0} \widehat{\text{match}} \dot{\mathbf{v}}_0 \overline{(\lambda \mathbf{x} \Rightarrow \sigma_2(\dot{\mathbf{e}}))} \\
& \longrightarrow^{j_1} \dot{\mathbf{v}}
\end{aligned}$$

We complete the proof by A-E₁, whose assumptions are discharged by the induction hypothesis.

Case L-COERCE: The goal is $(\sigma_1(\mathbf{e}), \uparrow \sigma_2(\dot{\mathbf{e}})) \in \mathcal{E}_n[\![\theta']]\!]$. Suppose $i < n$. We have the trace:

$$\begin{aligned}
\uparrow \sigma_2(\dot{\mathbf{e}}) & \longrightarrow^{i_1} \uparrow \dot{\mathbf{v}} \\
& \longrightarrow^{i_2} \dot{\mathbf{v}}'
\end{aligned}$$

where $i = i_1 + i_2$, with $\sigma_2(\dot{e}) \longrightarrow^{i_1} \dot{v}$. It follows by the induction hypothesis that $\sigma_1(e) \longrightarrow^* v$ for some v such that $(v, \dot{v}) \in \mathcal{V}_{n-i_1}[\llbracket \theta \rrbracket]$. Then, by [Lemma B.3.9](#), $(v, \dot{v}') \in \mathcal{V}_{n-i}[\llbracket \theta' \rrbracket]$, as desired. \square

Corollary B.3.11 (Correctness of declarative lifting of closed terms). *Suppose $\mathcal{S}; \mathcal{L}; \Sigma; \cdot \vdash e : \theta \triangleright \dot{e}$ and $\vDash_n \mathcal{L}$. We have $(e, \dot{e}) \in \mathcal{E}_n[\llbracket \theta \rrbracket]$.*

Theorem B.3.12 (Correctness of declarative lifting). $\vdash \mathcal{L}$ implies $\vDash \mathcal{L}$.

Proof. We need to show $\vDash_n \mathcal{L}$ for any n , i.e., $(x, \dot{x}) \in \mathcal{E}_n[\llbracket \theta \rrbracket]$ for any $x : \theta \triangleright \dot{x}$. The proof is a straightforward induction on n .

The base case is trivial, as $\mathcal{E}_0[\llbracket \theta \rrbracket]$ is a total relation as long as x and \dot{x} are well-typed, which is immediate from $\vdash \mathcal{L}$.

Suppose $\vDash_n \mathcal{L}$. We need to show $(x, \dot{x}) \in \mathcal{E}_{n+1}[\llbracket \theta \rrbracket]$. From $\vdash \mathcal{L}$, we know x and \dot{x} are defined by e and \dot{e} in Σ , respectively. It is easy to see it suffices to prove $(e, \dot{e}) \in \mathcal{E}_n[\llbracket \theta \rrbracket]$, because x and \dot{x} take exactly one step to e and \dot{e} . But that is immediate by [Corollary B.3.11](#). \square

In the following theorem, we write $\sigma(\Gamma_2)$ to obtain the second projection of the typing context with the type variables substituted, i.e., each $x : \eta \sim X$ in Γ is mapped to $x : \sigma(X)$. It may be counterintuitive that we do not require θ to be compatible with η or each pair in the typing context to be compatible, but these side conditions are indeed not needed; if the generated constraints (under empty context) are satisfiable, the compatibility conditions should hold by construction.

Theorem B.3.13 (Soundness of algorithmic lifting of open terms). *Suppose $\Sigma; \Gamma \vdash e : \eta \sim X \triangleright \dot{e} \mid \mathcal{C}$. Given a specification type θ , if $\mathcal{S}; \mathcal{L}; \Sigma; \sigma \vDash X = \theta, \mathcal{C}$, then $\sigma(\dot{e})$ elaborates to an expression \dot{e}' , such that $\mathcal{S}; \mathcal{L}; \Sigma; \sigma(\Gamma_2) \vdash e : \theta \triangleright \dot{e}'$.*

Proof. By induction on the derivation of the algorithmic lifting judgment. The cases of A-UNIT, A-LIT and A-FUN are trivial. The cases of A-LET, A-PAIR and A-PROJ are similar to A-ABS and A-APP, so we only show the proofs of these two cases. The cases of A-CTOR and A-MATCH are similar to A-IF (and other cases), so we only show the proof of A-IF.

Case A-VAR: Suppose $\sigma \vDash X' = \theta, \% \uparrow(X, X')$. It follows that:

- $\theta = \sigma(\mathbf{X}')$
- $\sigma(\mathbf{X}) \mapsto \sigma(\mathbf{X}') \triangleright \uparrow$
- $\% \uparrow(\sigma(\mathbf{X}), \sigma(\mathbf{X}'); \mathbf{x}) \triangleright \uparrow \mathbf{x}$

Therefore, $\sigma(\% \uparrow(\mathbf{X}, \mathbf{X}'; \mathbf{x})) = \% \uparrow(\sigma(\mathbf{X}), \sigma(\mathbf{X}'); \mathbf{x})$ elaborates to $\uparrow \mathbf{x}$. Since $\mathbf{x} : \eta \sim \mathbf{X} \in \Gamma$ by assumption, $\mathbf{x} : \sigma(\mathbf{X}) \in \sigma(\Gamma_2)$. We can then apply L-COERCE and L-VAR to derive $\sigma(\Gamma_2) \vdash \mathbf{x} : \sigma(\mathbf{X}') \triangleright \uparrow \mathbf{x}$.

Case A-ABS: Suppose $\sigma \vDash \mathbf{X} = \theta, \mathbf{X}_1 \in [\eta_1], \mathbf{X}_2 \in [\eta_2], \mathbf{X} = \mathbf{X}_1 \rightarrow \mathbf{X}_2, \mathcal{C}$. It follows that:

- $\theta = \sigma(\mathbf{X}) = \sigma(\mathbf{X}_1) \rightarrow \sigma(\mathbf{X}_2)$
- $[\sigma(\mathbf{X}_1)] = \eta_1$
- $[\sigma(\mathbf{X}_2)] = \eta_2$

Because $\sigma \vDash \mathbf{X}_2 = \sigma(\mathbf{X}_2), \mathcal{C}$, we have by the induction hypothesis:

- $\sigma(\dot{\mathbf{e}})$ elaborates to some $\dot{\mathbf{e}}'$
- $\mathbf{x} : \sigma(\mathbf{X}_1), \sigma(\Gamma_2) \vdash \mathbf{e} : \sigma(\mathbf{X}_2) \triangleright \dot{\mathbf{e}}'$

Thus $\sigma(\lambda \mathbf{x} : \mathbf{X}_1 \Rightarrow \dot{\mathbf{e}}) = \lambda \mathbf{x} : \sigma(\mathbf{X}_1) \Rightarrow \sigma(\dot{\mathbf{e}})$ elaborates to $\lambda \mathbf{x} : \sigma(\mathbf{X}_1) \Rightarrow \dot{\mathbf{e}}'$, and $\sigma(\Gamma_2) \vdash \lambda \mathbf{x} : \eta_1 \Rightarrow \mathbf{e} : \sigma(\mathbf{X}_1) \rightarrow \sigma(\mathbf{X}_2) \triangleright \lambda \mathbf{x} : \sigma(\mathbf{X}_1) \Rightarrow \dot{\mathbf{e}}'$, by L-ABS.

Case A-APP: Suppose $\sigma \vDash \mathbf{X}_2 = \theta, \mathbf{X}_1 \in [\eta_1], \mathbf{X} = \mathbf{X}_1 \rightarrow \mathbf{X}_2, \mathcal{C}$. It follows that:

- $\theta = \sigma(\mathbf{X}_2)$
- $[\sigma(\mathbf{X}_1)] = \eta_1$
- $\sigma(\mathbf{X}) = \sigma(\mathbf{X}_1) \rightarrow \sigma(\mathbf{X}_2)$

Because $\sigma \vDash \mathbf{X}_1 = \sigma(\mathbf{X}_1), \mathcal{C}$, we have by the induction hypothesis:

- $\sigma(\dot{\mathbf{e}}_1)$ elaborates to some $\dot{\mathbf{e}}'_1$
- $\sigma(\Gamma_2) \vdash \mathbf{x}_1 : \sigma(\mathbf{X}_1) \triangleright \dot{\mathbf{e}}'_1$

Therefore, $\sigma(\mathbf{x}_2 \dot{e}_1) = \mathbf{x}_2 \sigma(\dot{e}_1)$ elaborates to $\mathbf{x}_2 \dot{e}'_1$, and $\sigma(\Gamma_2) \vdash \mathbf{x}_2 \mathbf{x}_1 : \sigma(\mathbf{X}_2) \triangleright \mathbf{x}_2 \dot{e}'_1$, by L-APP and L-VAR.

Case A-IF: Suppose $\sigma \vDash \mathbf{X} = \theta, \%ite(\mathbf{X}_0, \mathbf{X}), \mathcal{C}_1, \mathcal{C}_2$. We know $\sigma(\mathbf{X}) = \theta$. Because $\sigma \vDash \mathbf{X} = \sigma(\mathbf{X}), \mathcal{C}_1$ and $\sigma \vDash \mathbf{X} = \sigma(\mathbf{X}), \mathcal{C}_2$, we have by the induction hypothesis:

- $\sigma(\dot{e}_1)$ elaborates to some \dot{e}'_1
- $\sigma(\Gamma_2) \vdash e_1 : \sigma(\mathbf{X}) \triangleright \dot{e}'_1$
- $\sigma(\dot{e}_2)$ elaborates to some \dot{e}'_2
- $\sigma(\Gamma_2) \vdash e_2 : \sigma(\mathbf{X}) \triangleright \dot{e}'_2$

Since $\sigma \vDash \%ite(\mathbf{X}_0, \mathbf{X}), \%ite(\sigma(\mathbf{X}_0), \sigma(\mathbf{X}); \mathbf{x}_0, \dot{e}'_1, \dot{e}'_2) \triangleright \dot{e}$ for some \dot{e} . By inverting this judgment, we consider two cases. In the first case, the condition has type \mathbb{B} :

- $\%ite(\sigma(\mathbf{X}_0), \sigma(\mathbf{X}); \mathbf{x}_0, \dot{e}'_1, \dot{e}'_2) \triangleright \mathbf{if} \ \mathbf{x}_0 \ \mathbf{then} \ \dot{e}'_1 \ \mathbf{else} \ \dot{e}'_2$
- $\sigma(\mathbf{X}_0) = \mathbb{B}$

In this case, $\%ite(\sigma(\mathbf{X}_0), \sigma(\mathbf{X}); \mathbf{x}_0, \sigma(\dot{e}_1), \sigma(\dot{e}_2))$ elaborates to $\mathbf{if} \ \mathbf{x}_0 \ \mathbf{then} \ \dot{e}'_1 \ \mathbf{else} \ \dot{e}'_2$, and, by L-IF₁ and L-VAR, $\sigma(\Gamma_2) \vdash \mathbf{if} \ \mathbf{x}_0 \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2 : \sigma(\mathbf{X}) \triangleright \mathbf{if} \ \mathbf{x}_0 \ \mathbf{then} \ \dot{e}'_1 \ \mathbf{else} \ \dot{e}'_2$.

The second case has condition type $\widehat{\mathbb{B}}$:

- $\lambda \sigma(\mathbf{X}) \triangleright \widehat{\mathbf{ite}}$
- $\%ite(\sigma(\mathbf{X}_0), \sigma(\mathbf{X}); \mathbf{x}_0, \dot{e}'_1, \dot{e}'_2) \triangleright \widehat{\mathbf{ite}} \ \mathbf{x}_0 \ \dot{e}'_1 \ \dot{e}'_2$
- $\sigma(\mathbf{X}_0) = \widehat{\mathbb{B}}$

In this case, $\%ite(\sigma(\mathbf{X}_0), \sigma(\mathbf{X}); \mathbf{x}_0, \sigma(\dot{e}_1), \sigma(\dot{e}_2))$ elaborates to $\widehat{\mathbf{ite}} \ \mathbf{x}_0 \ \dot{e}'_1 \ \dot{e}'_2$, and, by L-IF₂ and L-VAR, $\sigma(\Gamma_2) \vdash \mathbf{if} \ \mathbf{x}_0 \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2 : \sigma(\mathbf{X}) \triangleright \widehat{\mathbf{ite}} \ \mathbf{x}_0 \ \dot{e}'_1 \ \dot{e}'_2$. \square

Corollary B.3.14 (Soundness of algorithmic lifting). *Suppose $\Sigma; \cdot \vdash e : \eta \sim \mathbf{X} \triangleright \dot{e} \mid \mathcal{C}$. Given a specification type θ , if $\mathcal{S}; \mathcal{L}; \Sigma; \sigma \vDash \mathbf{X} = \theta, \mathcal{C}$, then $\sigma(\dot{e})$ elaborates to an expression \dot{e}' , such that $\mathcal{S}; \mathcal{L}; \Sigma; \cdot \vdash e : \theta \triangleright \dot{e}'$.*