

Certifying compilation for Standard ML in a type analysis framework

Leaf Eames Petersen

May, 2005
CMU-CS-05-135

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:

Robert Harper, Chair
Karl Crary
Peter Lee
Greg Morrisett (Harvard University)

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy*

© 2005 Leaf Eames Petersen

This research was sponsored by the National Science Foundation under grant nos. CCR-0121633 and EIA-9706572, and the US Air Force under contract no. F-19628-95-C-0050. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government or any other entity.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE MAY 2005		2. REPORT TYPE		3. DATES COVERED 00-00-2005 to 00-00-2005	
4. TITLE AND SUBTITLE Certifying compilation for Standard ML in a type analysis framework				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University,School of Computer Science,Pittsburgh,PA,15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 248	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Keywords: Types, Compilers, Standard ML, Typed Assembly Language, Certifying Compilation, Typed Register Allocation, Type Analysis

Abstract

Certified code is native machine code that is annotated with an automatically checkable certificate attesting to the conformance of the program to a specified safety policy. Certified code frameworks have been built based on first-order logic (PCC) and on types (TAL). Compilers generating certified code have been built for safe subsets of C and for Java(tm).

Type-preserving compilers such as the TILT/ML compiler implement compilation as transformations on *typed* internal languages. Types are used by the compiler for internal verification, and for optimization purposes. Type analysis can be used to implement optimizations such as non-uniform data representation and tag-free garbage collection. However, none of the existing type-preserving compilers for full-scale languages maintain type information all the way to the machine-code level, and hence are not yet able to generate certified code.

In this thesis, I demonstrate that certified compilation is possible in a type analysis framework by extending the TILT/ML compiler to generate certified code in the form of typed assembly language without compromising the existing optimizations of the compiler. *This work demonstrates that a compiler can use types to generate certified binaries for a full modern language even in the presence of aggressive type-analysis based optimization.*

To my wife, Adriana

Acknowledgments

This thesis would not have come to be without the help of many people.

First among these are my mother and father, who first started me on the path of learning and provided me the freedom and the means to follow it as far as I desired. They gave me my education, in every sense.

All of the members of my committee provided excellent guidance and feedback, as well as many interesting conversations that informed this work. Bob Harper must be thanked especially: for teaching and mentoring me through this long process, for encouraging me to explore my ideas, and for seeing me through to the other side. Bob should also be thanked for relentlessly evicting the “wicked whiches” and other strange beasts that lurk in the dark corners of my prose. Karl Crary contributed a great deal to the framework underlying this work: that it was possible at all is a credit to the extent to which he foresaw the issues and prepared the ground. His insight and assistance were invaluable.

The system described in this thesis builds on the efforts of many other people. Without the work done on the **TILT** compiler by Perry Cheng, Chris Stone, and others, none of this would have been possible. I should especially like to thank David Swasey for doing a great deal of thankless work to make the **TALx86** backend (and the compiler in general) much more robust, usable, and correct. This thesis also benefited hugely from the excellent work done by Greg Morrisett and the rest of the **TALx86** team.

Many other people have offered me their help, support, and friendship while I have been at CMU, including Sharon Burks, Catherine Copetas, Frank Pfenning, Rose Hoberman, Derek Dreyer, Rowan Davies, Dan Spoonhower, Tara Cheesman, Chris Palmer, Carrie Sparks, and many others. Thanks for making my time in graduate school so much more enjoyable!

Leaf Petersen
May, 2005

Contents

1	Introduction	1
1.1	Types in compilation	1
1.2	The non-certifying TILT compiler	3
1.3	The certifying TILT compiler	4
1.3.1	The theoretical compiler	4
1.3.2	Compiler implementation	5
1.4	Overview	8
2	MIL	9
2.1	Kinds and constructors	9
2.2	Proper types	11
2.3	Terms	11
3	LIL	13
3.1	The LIL syntax and static semantics	13
3.1.1	Typing contexts	15
3.1.2	Constructors and Kinds	15
3.1.3	Terms	18
3.2	Useful properties of the LIL	20
4	Example: Floating point array flattening	25
4.1	An optimized array strategy	25
4.2	LIL implementation of optimized arrays	25
4.2.1	Static encodings	26
4.2.2	Dynamic encodings	27
4.2.3	Type analyzing terms	28
5	The MIL to LIL translation	31
5.1	Translation overview	31
5.2	Static encodings of constructors	32
5.2.1	The kind translation	33
5.2.2	The constructor translation	34
5.2.3	Translation of typing contexts	37
5.2.4	Proofs of soundness for the constructor and kind translation	37
5.2.5	Additional properties of the constructor translation	40
5.3	Dynamic encoding of constructors	41

5.3.1	Dynamic encoding types	41
5.3.2	Notational issues	42
5.3.3	Dynamic encodings	44
5.3.4	Dynamic encoding of type contexts	45
5.3.5	Proofs for the dynamic encoding translations	45
5.4	Type translations	50
5.4.1	Proper MIL types	50
5.4.2	The term typing context translation	50
5.4.3	Proofs of soundness for the proper type translations	50
5.4.4	The type translation respects equivalence	52
5.5	The term translation	52
5.5.1	Definitions for the term translation	52
5.5.2	The term level translation	55
5.5.3	Proof that the term level translation preserves typing	57
5.6	The complete term level translation	71
6	Closure converted LIL	75
6.1	Introduction	75
6.2	The closure conversion translation	75
6.2.1	Constructors	76
6.2.2	Typing contexts	76
6.2.3	Terms	77
6.2.4	Heap Values and Programs	78
6.3	Soundness of closure conversion	80
6.3.1	Constructors	80
6.3.2	Typing contexts	83
6.3.3	Terms	83
7	TILTAL	91
7.1	TILTAL overview	91
7.1.1	Stacks	91
7.1.2	Values	94
7.1.3	Operands: Registers and stack slots	95
7.1.4	Instructions and instruction sequences	95
7.1.5	Heap values, heaps, and register files	96
7.1.6	Programs	96
7.1.7	Typing contexts	96
7.2	TILTAL derived forms	96
7.2.1	Partial instruction sequences	96
7.2.2	Heap fragments	97
8	The LIL to TILTAL translation	99
8.1	The constructor level	99
8.1.1	Kinds	99
8.1.2	The constructor translation	99
8.1.3	Soundness of the type translation	103

8.2	The term level: preliminaries	103
8.2.1	Register and stack slot allocation	103
8.2.2	Derived instructions	107
8.3	The term translation	111
8.3.1	Values	112
8.3.2	Operations	117
8.3.3	Soundness of the operation translations	119
8.3.4	Expressions	127
8.3.5	Soundness of the expression translation	131
8.3.6	Heap values, heaps, and programs	147
8.4	The complete translation rules	153
9	Implementation	163
9.1	Singleton Elimination	163
9.1.1	Preserving sharing of types	165
9.1.2	Optimizations for singleton elimination	165
9.1.3	Runtime behavior	166
9.2	MIL to LIL	167
9.3	Closure Conversion	168
9.3.1	Closure conversion strategies	168
9.3.2	Recursive code closure conversion	168
9.4	General Optimization	169
9.5	Lil To Tal	170
9.5.1	Code generation	170
9.5.2	Register allocation	171
9.5.3	Additional TILT constructs	172
9.6	Engineering	174
9.6.1	Type representation	174
9.6.2	Traversing programs	175
9.6.3	Fast substitution	175
9.6.4	Fast alpha-equivalence	175
9.6.5	Weak Head-Normalization	176
9.6.6	Engineering the LIL type checker	176
9.7	Compilation units	180
9.7.1	Compiler generated files	180
9.8	Measurements	181
9.8.1	Benchmarks	181
9.8.2	Type size	181
9.8.3	Benchmark unit sizes	184
9.8.4	Libraries and linking	185
9.8.5	Run time	191
10	Conclusions and future work	195
10.1	Summary	195
10.1.1	Theory	195
10.1.2	Practice	195

10.1.3	Conclusions	196
10.1.4	Compiler engineering	196
10.2	Future work	197
10.2.1	Optimization	197
10.2.2	Garbage collection	198
10.2.3	Infrastructure improvements	198
10.2.4	Reducing the trusted computing base	198
10.3	Conclusions	199
A	MIL static semantics	201
B	LIL static semantics	209
C	TILTAL static and dynamic semantics	219
C.1	TILTAL Static semantics	219
C.2	TILTAL dynamic semantics	226
C.2.1	Definitions	226
C.2.2	Transitions	227

List of Figures

1.1	TILT Architecture	3
1.2	Structure of the theoretical compiler	5
1.3	Structure of the certifying TILT backend	6
2.1	MIL syntax	10
3.1	LIL syntax	14
3.2	Constructor constants and their kinds	16
5.1	The constructor translation (static encoding)	35
5.2	Derived non-named-form expressions	43
5.3	Typing rules for derived (non-named-form) expressions	43
5.4	The dynamic encoding translation	44
6.1	Closure converting recursive function definitions	79
7.1	TILTAL syntax	92
7.2	Judgements defining well-typedness of TILTAL programs	93
8.1	The constructor translation	100
8.2	LIL to TILTAL translation judgements	111
8.3	Exception case translation in TILTAL	130
8.4	Exception handler implementation in TILTAL	132
9.1	Structure of the certifying TILT backend	164
9.2	Benchmark timings with singleton elimination (normalized to timings without singleton elimination).	166
9.3	Benchmarks	182
9.4	Size breakdown of individual benchmark units, in kilobytes.	183
9.5	Size breakdown of individual benchmark units, relative.	184
9.6	Total sizes of compilation groups (kilobytes).	186
9.7	Total sizes of compilation groups (relative).	187
9.8	Breakdown of certificate overhead across individual Basis units (kilobytes)	187
9.9	Assembly file size vs assembled output size (kilobytes)	189
9.10	Absolute sizes of large units (kilobytes).	190
9.11	Relative sizes of large units.	191
9.12	Normalized runtime (MLton = 1)	192

C.1	TILTAL transitions (part I)	228
C.2	TILTAL transitions (part II)	229
C.3	TILTAL transitions (part III)	230

Chapter 1

Introduction

1.1 Types in compilation

It is an unfortunate fact about the state of programming that even good programs sometimes go wrong. It is not uncommon for programs to crash or misbehave, whether accidentally, or with malicious intent. This problem has been greatly compounded in recent years by the proliferation of *mobile* code. More and more of the code that is run is downloaded in bits and pieces from various sources. Examples of this include Java-script and Java(TM) programs downloaded into a web browser, applications downloaded directly over the internet, and code run on behalf of others (such as the SETI@HOME project, which uses donated spare processor cycles to further the search for extra-terrestrial life). The proliferation of mobile code is expected only to increase as networked technology becomes more a part of everyday life.

However, it is particularly hard to trust the behavior of this sort of code. The code producer may be unknown to the code consumer, or the identity of the producer may be spoofed. Moreover, even trusted producers occasionally produce programs that go wrong.

It would certainly seem desirable to be able to rule out programs that are unsafe. Unfortunately, determining the safety of arbitrary machine code is undecidable. A compromise solution that has been in existence for several decades, is to restrict ourselves to programming in a language which has the property that all programs are safe. The resulting compiled code is therefore safe by construction, *modulo the correctness of the compiler*. Languages such as LISP, ML and Java all have this property: that all programs written in these languages are guaranteed with some level of certainty not to produce undefined and unsafe behavior.

Unfortunately for the purposes of mobile code, the safety properties enjoyed by these languages are only guaranteed at the source level: once the source code has been compiled to machine instructions, the safety guarantee lies only in the implicit property of being in the image of the safe language under compilation. One solution to this is to ship around instead something tantamount to source code, allowing the consumer to validate the code independently. This is in essence the Java(TM) byte-code solution. This places much of the burden of compilation on the code consumer, who must still in turn trust that their own compiler is correct.

The two most commonly used solutions then, are either to accept arbitrary machine code based on trust in the producer of the code, or to accept only annotated high-level code and to trust the local compiler. Both of these solutions leave much to be desired.

To better this, several systems have been proposed for annotating machine code in such a way

that safety remains *checkable*. Along with the code, the code consumer receives a certificate that can be used to *check* that the code conforms to the correctness assertions that it claims. The only software that the code consumer must still trust is the checker itself. Two notable examples of this include Proof Carrying Code (PCC) from CMU [Nec98] and Typed Assembly Language (TAL) from Cornell [MWCG98]. The PCC system provides for machine code to be annotated with proofs of safety properties done in first-order logic. The code consumer simply checks that the proofs are indeed correct – a relatively simple procedure. This system is very general, in that it can be used to certify any property which can be expressed in the logic. TAL on the other hand specializes to the particular property of type safety. TAL provides for a machine code which is annotated with types, which can then be type checked by the code consumer.

A certifying compiler is one which produces, along with its normal output, a *certificate* which can be used to check that the generated code is safe according to some policy. Certifying compilers have been written translating safe subsets of C to both PCC and TAL [MCG⁺99, NL98]. More ambitiously, a full scale Java(TM) compiler has been written targeting PCC [CLN⁺00].

The **TILT** (TIL Two) compiler is an optimizing compiler developed at CMU that implements the full Standard ML '98 definition and includes support for separate compilation. Important ideas pioneered in **TILT** and its predecessor TIL include using intensional polymorphism [HM95] to reduce the cost of implementing polymorphism and garbage collection. Compilation proceeds as a series of typed transformations into successively lower level typed languages. Type information is used to allow for optimized data representations and to do “almost tag-free” garbage collection. Prior to this work however, type information was mostly erased in **TILT** well before the transformation to machine code was made, and hence safety properties of the resulting code could only be asserted - not checked.

TILT uses types during compilation for optimization purposes, and consequently requires an intermediate language with a very expressive type theory. Previous work on typed assembly languages has primarily focused on preserving type information for certification purposes. I claim that these two uses of low level typed languages are compatible. *A compiler can use types to generate certified binaries while retaining the ability to perform complicated type based optimizations on a full modern language.*

I have demonstrated this by extending the TILT-ML compiler to maintain type information used in the intermediate passes of the compiler all the way through code generation, producing certified binaries without sacrificing the ability to perform type analysis optimizations. This dissertation gives a careful theoretical description of the key elements of the compilation process, and proves soundness theorems for the translations between the major intermediate languages. A description is also given of the actual implementation, including some empirical results.

It is becoming increasingly clear that type preserving compilation, in addition to serving as a mechanism for enabling safe mobile code, also provides great benefits in its own right as a compiler engineering technique. Just as type safe languages allow programmers to write correct code more quickly, type preserving compilation allows compiler implementers to write correct compilers more quickly. Type checking the intermediate representations of programs within the compiler allows many or most compiler bugs to be caught during compilation and to be localized to the particular point of failure in the compiler. Whole classes of pernicious compiler bugs (such as those that involve stack or memory corruption) can be eliminated by producing checkably type safe code. One of the results of this thesis is to provide further evidence of the efficacy of this technique, and to add to the body of experience in the engineering of type preserving compilers. This aspect of




Figure 1.1: TILT Architecture

the implementation experience is discussed further in section 10.1.4.

1.2 The non-certifying TILT compiler

The past years have seen a great deal of interest in the idea of “typed compilers” that maintain type information deep into the compilation process. Such type information can be exploited by the compiler internally to allow for optimized data representations and to do tag-free garbage collection, as well as providing the compiler with a basis for internal correctness checks. This work was pioneered in the TIL compiler at CMU [TMC⁺96], and has been adopted by numerous other compilers, including the Glasgow Haskell Compiler. Other recent work has also suggested the possibility of maintaining type information through to the machine code as a form of certification [MWCG97].

The TIL compiler clearly demonstrated that typed compilation was both feasible and desirable. However, TIL compiled only the core language of Standard ML: the powerful modular features that are one of the most important elements of SML were not dealt with. The TIL Two (**TILT**) compiler was aimed at addressing this shortcoming.

Figure 1.1 depicts the structure of the non-certifying **TILT** compiler. Its architecture is based around two typed intermediate languages. The initial elaboration from SML source targets a structures calculus called the HIL (High Intermediate Language). This language is relatively close to SML, and among other things provides the interface language used for separate compilation. After elaboration (and hence typechecking), the HIL is translated to a second typed language called the **MIL** (Middle Intermediate Language) through a process called phase splitting [HMM90]. The phase splitting process maps each SML structure into separate type and term level records, representing the static and dynamic portions of the structure. Similarly, SML functors are mapped to type and term level functions. In this fashion, modular programs are translated into programs containing only lambda calculus terms.

The **MIL** is the language in which almost all of the optimization passes implemented in **TILT**

are done. This constrains the design of the **MIL**, since it must be possible to express the results of all of the desired optimizations in a typed fashion. In particular, it is important that primitives for data representation optimizations be present at this level. By “hiding” type analysis inside of a few primitives, the **MIL** avoids the need for a general typecase construct as used in the λ_i^{ML} calculus. Nonetheless, the fact that some **MIL** primitives do indeed analyze their types mandates a type passing interpretation for the **MIL** operational semantics.

All of the intermediate languages of the **TILT** compiler up to and including the **MIL** are typed, and all of the compiler passes on these languages are type-preserving in the sense that they map well-typed programs to well-typed programs. Unfortunately for certification purposes, the subsequent languages from the **MIL** on down are not typed, and hence the generated code cannot in general be proven safe. This dissertation replaces this un-typed backend with a new type-preserving backend that produces certified code.

1.3 The certifying TILT compiler

One of the major goals of certifying compilation is to ensure that the certifying compiler is type-preserving: that is, that it maps well-typed programs in the source language to well-typed programs in the target language. In order to show that this is the case, I spend the first part of this dissertation presenting idealized versions of the compiler intermediate languages, and proving the soundness of the translations between them. The next two sections describe the idealized compiler and its relation to the implementation.

1.3.1 The theoretical compiler

The theoretical portion of this dissertation describes the framework for a translation mapping the original **TILT** internal language (the **MIL**, described in chapter 2) down to a typed assembly language that I call **TILTAL** (Typed Assembly Language for **TILT**). This translation uses as an intermediate stage a new internal language called the **LIL** (Low Level Language). The **LIL** is an impredicatively typed lambda calculus based on Crary and Weirich’s LX [CW99]. Figure 1.2 describes the structure of the theoretical compiler.

The **LIL**, described in chapter 3, provides a very rich type system in which the type analysis of **TILT** can be represented using term level constructs. In addition to various engineering benefits, this fact allows us to take a type erasure interpretation, instead of the type passing interpretation apparently mandated by the **MIL** primitives. The fact that we can embed type analysis into the term level reflects in a typed fashion exactly the techniques already used in an untyped fashion for implementing type passing languages. Chapter 4 gives a brief introduction to the type analysis methodology of the **LIL** via a worked example.

The translation from **MIL** to **LIL** serves to make type analysis and type representations explicit in the term language. This translation is described in detail in chapter 5. A subsequent pass mapping **LIL** terms to **LIL** terms performs closure conversion. The closure conversion translation is described in chapter 6.

Finally, I give a translation from the **LIL** into **TILTAL**. **TILTAL** code is essentially machine code with type annotations added allowing it to be typechecked. Currently, there are in addition to the standard assembly language instructions several typed primitives corresponding to assembler macros. These primitives handle memory allocation (and hence the interaction with the garbage




Figure 1.2: Structure of the theoretical compiler


collector) and array bounds checking. The **TILTAL** language is introduced in chapter 7, and the translation from **LIL** programs to **TILTAL** programs is given in chapter 8.

The theoretical **LIL** is actually very close to the **LIL** as implemented in the compiler, mostly lacking only an extended set of basic primitive operations. The presentation of the **TILTAL** target language is intended to suggest how the ideas used to implement type analysis in the **LIL** translate down to the assembly code level. The actual **TALx86** implementation departs significantly from that given here. However, for the most part the theoretical translations between the various languages are designed to capture all or most of the essential details of the implementation. In particular, the **LIL** to **TILTAL** translation of chapter 8 attempts as much as possible to account for the actual code generation techniques used by the implementation.

The main result of the theoretical portion of this dissertation is a proof that the compilation of **LIL** programs into **TILTAL** programs preserves types, in the sense that each of the individual translations is proved to map well-typed terms to well-typed terms. Proofs of soundness for the various translation phases are given in the chapters in which they are defined.

1.3.2 Compiler implementation

The second half of this thesis is an actual implementation of a certifying compiler for Standard ML. The theoretical compiler discussed in the previous section is intended as a model for the implementation. Where the original architecture (see figure 1.1) switches to an untyped language, the new implementation must instead take the typed output and continue the compilation process in a typed setting, as shown in figure 1.3. The certifying **TILT** implementation is discussed in chapter 9.



The implementation of certifying **TILT** required the addition of essentially five additional compiler stages.

- **MIL** singleton elimination
- Translation to **LIL**
- Closure conversion
- Optimization
- Translation to **TALx86**

To simplify the meta-theory of the **LIL**, I have chosen not to support the singleton kinds of the original **MIL** as implemented in the compiler. Therefore, before (or concurrent with) the translation to **LIL**, singletons must be eliminated from **MIL** code. A proof that this is possible, and a simple algorithm for doing so has been given by Crary [Cra00]. The implementation of this algorithm and its effects on compiled code are discussed in section 9.1.

The translation to **LIL** is described in detail in chapter 5. The most interesting part of this translation is the process of making type analysis explicit. The **MIL** has no general notion of type analysis, instead using type analyzing primitives to implement specific optimizations such as floating point array unboxing and flattening function arguments into records. The translation to **LIL** compiles these primitives into uses of a general type analysis mechanism.

The **MIL** implements all of the optimizations that **TILT** supports, so extended optimization of **LIL** code is mostly unnecessary. However, the translations tend to produce code that can be improved significantly by a simple optimizer: for example, the closure converter frequently introduces dead code and projections from known records. Rather than making other phases do significantly more work to avoid this, it was simpler and cleaner to implement a basic optimization pass for the **LIL**. Note that the optimizer does not have a counterpart in the theoretical model.

It would be appealing to rely as well on the **MIL** closure converter, and only translate closure converted code. Unfortunately, the **MIL** notion of closure conversion is not easily compatible with the **LIL** notion, since the **MIL** must closure-convert types as data. Translating closure converted **MIL** code would require “de-closure converting” types, which is not appealing. For this reason, a separate closure conversion pass for the **LIL** was implemented, closely modeled on the formal closure conversion translation from chapter 6.

In order to allow intermediate code to be validated, it was also useful to implement a type checker for the **LIL** language. This allows the output of the various phases of the compiler to be checked for errors internally. The ability to check type correctness of internal program representations has proved valuable in the development of **TILT**.

Lastly, a translation is given mapping **LIL** programs into **TALx86** programs. Note that here, the implementation diverges significantly from the formal model described in chapter 8 in that the **TILTAL** and **TALx86** languages are quite different. However, the theoretical model was carefully designed to capture many of the interesting algorithmic approaches used in the actual implementation.

The final portion of the dissertation gives a basic quantitative evaluation of the implementation, including measurements of the size of the generated code and certificates to get a feeling for the overhead of certification. Some measurements of run time behavior of the compiled programs are also given.

It is important to state here that while I measure these quantities to gain some understanding of the cost of my approach, optimizing for small certificates and fast certification time is beyond the scope of this dissertation. The purpose of this work is to demonstrate the feasibility of using types to certify type analyzing code generated from a full-scale, general purpose language. Engineering the representations, while certainly important for making the compiler practical, was not a primary goal.

1.4 Overview

In the next chapter, I provide a brief introduction to a simplified version of the **MIL** intermediate language which is the original source language for the compilation process described in this dissertation. Chapter 3 describes the **LIL** language which serves as the main intermediate language of the new backend. The complete static semantics of the **MIL** and the **LIL** are given in appendixes A and B, respectively. Chapter 4 gives an introduction to the style of type-analysis used in the **LIL** via an extended example. The translation from **MIL** to **LIL** is described in detail in chapter 5, and a proof of soundness is given. Closure conversion of **LIL** terms is described and proved sound in chapter 6. Chapter 7 introduces the **TILTAL** typed assembly language which serves as the target of the theoretical compiler. A translation from **LIL** programs to **TILTAL** programs is described and proved sound in chapter 8. Finally, the implementation and its relation to the theoretical presentation is discussed in chapter 9, and some empirical results about the implementation are presented.

Chapter 2

MIL

I describe the **MIL** here only briefly, as the theoretical and practical aspects of this calculus have been discussed in detail elsewhere [PCHS00, SH99, VDP⁺03]. For the purposes of the translation, I assume that singletons have already been eliminated [Cra00], although in practice it may be desirable to do this concurrently with the translation. The syntax of the singleton free **MIL** is given in figure 2.1. The **MIL** is a predicative lambda calculus based on Girard’s F_ω , extended with primitives for type analysis. The intention is that these primitives are definable in terms of the λ_i^{ML} calculus of Harper and Morrisett [HM95]. In the untyped **TILT** back-ends, these primitives are only compiled directly after moving to an untyped calculus. The complete static semantics for the **MIL** are given in appendix A and are fairly straightforward. The remainder of this section gives a high-level overview of the structure of the language and describes the type-analysis methodology it employs.

2.1 Kinds and constructors

The kind structure of the **MIL** is relatively simple in the absence of singletons, with function and product kinds to classify constructor functions and tuples, and the base kind T_{32} to classify constructors which classify terms. The notation T_{32} indicates that the terms classified by constructors of this kind are intended to be represented by a 32 bit quantity after compilation. I use the term “constructor” in preference to the term “type” when referring to constructors of arbitrary or unspecified kind. I reserve the standard terminology “type” for constructors of kind T_{32} .

Most of the the base constructors are completely standard, with the exception of the treatment of sums and the constructs for type analysis. The sum type encodes the number of non value-carrying fields directly, instead of simply using unit as the carried value. In addition, the **MIL** has a known sum type, corresponding to the type of a sum for which the branch inhabited is known. A special projection operation projects the carried value out of a known sum. This allows the case construct to avoid destructing its value which may be unnecessary if the arm doesn’t actually use the carried value.

Type analysis is present at the type level in the form of the **Vararg** construct and implicitly in the **Array** type. The **Vararg** type classifies the term-level **vararg** construct, which is used to implement non-standard calling conventions for functions which take tuples as arguments. In **TILT** tuple arguments to functions are always flattened into registers if the number of fields in the tuple is small. In order to make this work with polymorphic functions, it is necessary to use type analysis.

$$\begin{aligned}
\kappa &::= \mathbf{T}_{32} \mid \kappa_1 \rightarrow \kappa_2 \mid \kappa_1 \times \kappa_2 \\
c &::= \alpha \mid \mathbf{Int} \mid \mathbf{Boxf} \mid (\vec{c}) \rightarrow c \mid c_1 \times \dots \times c_n \\
&\quad \mid \mu(\alpha, \beta).(c_1, c_2) \mid \langle c_1, c_2 \rangle \mid \pi_1 c \mid \pi_2 c \mid \lambda(\alpha::\kappa).c \mid c_1 c_2 \\
&\quad \mid \mathbf{Vararg}_{c_1 \rightarrow c_2} \mid \mathbf{Vararg}_{c_1 \rightarrow^t c_2} \mid \mathbf{Sum}_i(\vec{c}) \mid \mathbf{Sum}_i^j(\vec{c}) \\
&\quad \mid \mathbf{Array}_c \mid \mathbf{Farray} \mid \mathbf{Exn} \mid \mathbf{Dyntag}_c \\
\tau &::= T(c) \mid \forall[\alpha::\kappa](\vec{\tau})(i) \rightarrow \tau \mid \forall[\alpha::\kappa](\vec{\tau})(i) \rightarrow^t \tau \mid \tau_1 \times \dots \times \tau_n \\
fv &::= \mathbf{r} \mid x_f \\
sv &::= x \mid i \mid \mathbf{roll}_c sv \mid \mathbf{unroll}_c sv \\
&\quad \mid \mathbf{inj_tag}_{\mathbf{Sum}_i(\vec{c})}^j \\
opr &::= sv \mid \mathbf{vararg}_{c_1 \rightarrow c_2} sv \mid \mathbf{onearg}_{c_1 \rightarrow c_2} sv \\
&\quad \mid \mathbf{boxf} fv \mid \mathbf{unboxf} sv \mid fv \\
&\quad \mid \mathbf{proj}^i sv \mid \mathbf{inj}_{\mathbf{Sum}_i(\vec{c})}^j sv \\
&\quad \mid \langle \vec{sv} \rangle \mid \mathbf{select}^i sv \\
&\quad \mid \mathbf{case}(sv)(x_1.e_1, \dots, x_n.e_n) \mid \mathbf{handle}_\tau(e_1, x.e_2) \\
&\quad \mid sv[\vec{c}](\vec{sv})(\vec{fv}) \mid \mathbf{inj_dyn}_c(sv_1, sv_2) \\
&\quad \mid \mathbf{raise}_\tau sv \mid \mathbf{mkexntag}_c \\
&\quad \mid \mathbf{exncase}_\tau(sv)(sv_1 \Rightarrow x_1.e_1, - \Rightarrow e_2) \\
&\quad \mid \mathbf{sub}_c(sv_1, sv_2) \mid \mathbf{fsub}(sv_1, sv_2) \\
&\quad \mid \mathbf{upd}_c(sv_1, sv_2, sv_3) \mid \mathbf{fupd}(sv_1, sv_2, fv) \\
&\quad \mid \mathbf{array}_c(sv_1, sv_2) \mid \mathbf{farray}(sv, fv) \\
e &::= sv \mid \mathbf{let}_\tau \mathbf{rec}_\tau f[\alpha::\kappa](x::\tau)(\vec{x}_f).e \mathbf{in} e \\
&\quad \mid \mathbf{let}_\tau x = opr \mathbf{in} e \mid \mathbf{let}_\tau x_f = opr \mathbf{in} e \\
\Delta &::= \bullet \mid \Delta, \alpha:\kappa \\
\Gamma &::= \bullet \mid \Gamma, x:\tau \mid \Gamma, x_{64}
\end{aligned}$$

Figure 2.1: MIL syntax

The **Vararg** type can be thought of as simply a type level typecase in a λ_i^{ML} like language which branches on the argument type of a function type. If the argument type is a small tuple, it returns a multi-argument function type where the fields of the tuple have been flattened. Otherwise, it leaves the function type unchanged.

The **Array** type implicitly distinguishes between arrays of boxed floating point numbers and other arrays, in order to flatten the boxed float arrays. This is discussed in more detail below. The rest of the constructor level is a standard typed lambda calculus, classified by the function and product kinds.

2.2 Proper types

Also described in figure 2.1 is the syntax for the type level. Unlike the constructor level which corresponds to the notion of *types as data*, the type level in a predicative system corresponds to the notion of *types as classifiers*. The constructor level is included into the type level via an explicit inclusion $T(c)$. The type level also contains classifiers for polymorphic functions and tuples of terms. The duplication of the tuple type at the type level indicates the possibility of constructing pairs containing polymorphic functions which is not provided for by the constructor level.

2.3 Terms

For the most part, the term-level **MIL** is a standard lambda calculus with a few non-standard primitives. The presentation here restricts the syntax to a named form [FSDF93, Mog89] to reflect the form used internally in the **TILT** compiler. Named form forces incremental calculations to be given names via variable binding (hence *named* form) which is important for various compiler passes and code generation. Unlike most lambda calculi, the MIL also includes low level data representation primitives (such as float boxing and unboxing primitives). This allows data representation optimizations to be expressed at the level of the **MIL**.

A key optimization that **TILT** implements is the use of non-uniform data representation. Many implementations of languages with polymorphism require that all values fit into a word. In particular, array elements must always be word-sized, which means that arrays of 64 bit floats (for example) must actually be arrays of pointers to floats. This is unfortunate because of the wasted space, the extra indirections to access the data, and because of the consequent loss of data locality. **TILT** avoids this by incorporating type analysis into the array primitives. By passing types at runtime and allowing code to dispatch on them, unboxed floating point arrays can be used with the appropriate subscript stride chosen at runtime.

The MIL calculus differs from the λ_i^{ML} calculus of [HM95] in that it does not contain an explicit type analysis construct such as *typerec* or *typecase*. This does not mean however that the idea of intensional type analysis has been abandoned: rather, the type analysis has been hidden inside the primitives that need to use it.

For example, **TILT** deals with floating point numbers by syntactically distinguishing between boxed and unboxed floats, with appropriate term level conversions between them. This allows the optimizer to deal directly with data representation optimizations, even at the relatively high level of the MIL. The syntactic restriction on unboxed floats prevents polymorphic functions from being instantiated with the unboxed float type, so that all polymorphic values take up 32 bits. All unboxed floating point arguments are segregated so that they may be passed in float registers.

A special **Farray** type is provided corresponding to the type of arrays of unboxed floats, with corresponding term-level operations **farray**, **fsub** and **fupd**.

One obvious problem with this is that arrays of values whose type is not statically known to be **Float** would seem to have to use a boxed representation. By using type analysis in the array constructor as well as the subscript operator however, at least some of the difficulty can be avoided. Essentially, the **Array** constructor incorporates a typecase which selects the appropriate array representation based on the type of the carried values. Similarly, the term-level subscript and update operations must dispatch on the type to select the appropriate operation. In the case that the type turns out to be **Boxf**, the subscript operation will also be forced to re-box the float before returning it, since subscripting into an array of boxed floats returns a value of type **Boxf**.

Type analysis is also encoded into the **vararg** and **onearg** primitives which implement special calling conventions for single argument functions whose argument type is a small record type. For example, the term $\text{vararg}_{c_1 \rightarrow c_2} sv$ corresponds roughly to the following code using an explicit λ_i^{ML} style typecase:

```

typecase c1 of
  Record[] => lambda[] . sv <>
| Record[t] => lambda[x;t] . sv <x>
| Record[t1,t2] => lambda[x1;t1,x2;t2] . sv <x1,x2>
| _ => f

```

The **onearg** construct is the inverse of **vararg**, turning a variable argument function back to a normal function. The maximum number of fields that can be flattened (here shown as two) is a machine dependent parameter of the type theory.

The **MIL** hides type analysis by replacing certain stylized uses of typecase with primitives that analyze their type arguments. At some point however, it becomes necessary to make this analysis explicit. Currently in the compiler, this happens when the **MIL** is translated to a low-level untyped language. One of the major challenges in pushing type information down to the machine level in **TILT** is making this analysis explicit in a typed language that is amenable to translation to a typed assembly language. The next chapter discusses a strategy for doing this by reflecting type analysis into the term level of a more powerful lambda calculus.

Chapter 3

LIL

There is an apparent disparity between the type-passing model of type analysis and the constructs available at the machine level. Evaluation of terms depends on the type arguments as well as the term arguments, but real machine processors are untyped. The ad-hoc solution originally used in **TILT** is to translate to an untyped setting, choosing term-level representatives for the type data in the process. Type analysis then becomes simple branches on the values chosen to represent various types. This untyped approach was taken because of the absence of a well-understood type system for typing such code.

In order to make this ad-hoc solution viable for producing typed certified code, Cray, Weirich, and Morrisett describe a type theory that permits term-level representation of types and type analysis in a typed setting [CWM98] using primitive terms to represent types. Cray and Weirich subsequently extend this notion [CW99] to a type theory in which representation of types is definable using only the ordinary lambda calculus term constructs with an enriched type system. The **LIL** adopts this idea for its treatment of type analysis, and also extends the **MIL** with constructs for representing closures.

3.1 The LIL syntax and static semantics

Figure 3.1 describes the complete syntax of the **LIL**. As with the **MIL**, the language is syntactically restricted to named form. This is not particularly necessary for theoretical purposes, but matches more closely the implementation.

Kinds form the top of the syntactic hierarchy, and are generally written using the meta-variable κ . Kinds classify the constructors, for which the meta-variables c is generally used. In addition the meta-variables τ and ϕ are used to distinguish constructors of kind T_{32} , T_{64} respectively. This is purely a presentational distinction however and does not correspond to an actual syntactic distinction.

At the term level, the classes of small 32 bit and 64 values are notated as sv and fv respectively; 32 bit and 64 bit operations are notated as opr and $fopr$ respectively; and expressions are notated using the meta-variable e . Note that the distinction between 32 and 64 bit values is a syntactic distinction and not merely a notational convention, and similarly for operations.

LIL programs (p) consist of a mutually recursive set (d) of heap bindings ($hval$), and an executable expression e . In the theoretical presentation here, code functions are the only heap values

$$\begin{aligned}
\kappa &::= \mathsf{T}_{32} \mid \mathsf{T}_{64} \mid 1 \mid \kappa_1 \rightarrow \kappa_2 \mid \kappa_1 \times \kappa_2 \\
&\quad \mid +[\kappa_1, \dots, \kappa_n] \mid j \mid \mu j. \kappa \mid \forall j. \kappa \\
c, \tau, \phi &::= * \mid \alpha \mid \lambda \alpha : k. c \mid c_1 c_2 \\
&\quad \mid \langle c_1, c_2 \rangle \mid \pi_1 c \mid \pi_2 c \\
&\quad \mid \mathsf{inj}_i^{+[k_i]} c \mid \mathsf{case}(c, [\alpha_1.c_1, \dots, \alpha_n.c_n]) \\
&\quad \mid \mathsf{fold}_{\mu j. k} c \mid \mathsf{pr}(j, \alpha : \kappa, \rho : (j \rightarrow \kappa'), \iota : \mathsf{in} c) \\
&\quad \mid \mathsf{Float} \mid \mathsf{Int} \mid \mathsf{Boxed} \mid \mathsf{Void} \mid \times \mid \rightarrow \mid \mathsf{Code} \\
&\quad \mid c[\kappa] \mid \Lambda j. c \mid \exists \mid \forall \mid \mathsf{Rec} \mid \vee \\
&\quad \mid \mathsf{Array}_{32} \mid \mathsf{Array}_{64} \mid \mathsf{Dyntag} \mid \mathsf{Dyn} \mid \mathsf{Tag} \\
fv &::= x_{64} \mid \mathsf{t} \\
sv &::= x \mid \ell \mid i \\
&\quad \mid \mathsf{inj_union}_c sv \mid \mathsf{inj_dyn}_\tau(sv_1, sv_2) \\
&\quad \mid \mathsf{unroll}_\tau sv \mid \mathsf{roll}_\tau sv \\
&\quad \mid \mathsf{pack} sv \mathsf{as} \exists \alpha : \kappa. \tau \mathsf{ hiding } c \mid sv[c] \mid \mathsf{tag}_i \\
fopr &::= fv \mid \mathsf{unbox} sv \mid \mathsf{sub}_\phi(sv_1, sv_2) \\
opr &::= sv \mid \mathsf{select}^i sv \\
&\quad \mid \mathsf{case}(sv)(x.e_1, \dots, x.e_n) \\
&\quad \mid \mathsf{dyncase}(sv)(sv_1 \Rightarrow x_1.e_1, - \Rightarrow e) \\
&\quad \mid \mathsf{dyntag}_\tau \mid \mathsf{raise}_\tau sv \mid \mathsf{handle}_\tau(e_1, x.e_2) \\
&\quad \mid \mathsf{box} fv \mid \langle \vec{sv} \rangle \\
&\quad \mid sv(\vec{sv})(\vec{fv}) \mid \mathsf{call} sv(\vec{sv})(\vec{fv}) \\
&\quad \mid \mathsf{array}_\tau(sv_1, sv_2) \mid \mathsf{array}_\phi(sv, fv) \mid \mathsf{sub}_\tau(sv_1, sv_2) \\
&\quad \mid \mathsf{upd}_\phi(sv_1, sv_2, fv) \mid \mathsf{upd}_\tau(sv_1, sv_2, sv_3) \\
e &::= \mathsf{let} \mathsf{rec}_\tau f[\alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n](x_1 : \tau_1, \dots, x_m : \tau_m)(z_1 : \phi_1, \dots, z_k : \phi_k). e \mathsf{ in } e \\
&\quad \mid sv \mid \mathsf{let}_\tau x = opr \mathsf{ in } e \mid \mathsf{let} x_{64} = fopr \mathsf{ in } e \\
&\quad \mid \mathsf{let} [\alpha, x] = \mathsf{unpack} sv \mathsf{ in } e \mid \mathsf{let} \langle \beta, \gamma \rangle = c \mathsf{ in } e \\
&\quad \mid \mathsf{let} (\mathsf{fold} \beta) = c \mathsf{ in } e \\
&\quad \mid \mathsf{let} (\mathsf{inj}_i \beta) = (c, sv) \mathsf{ in } e \\
hval &::= \mathsf{code}_\tau[\alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n](x_1 : \tau_1, \dots, x_m : \tau_m)(z_1 : \phi_1, \dots, z_k : \phi_k). e \\
d &::= \epsilon \mid d, \ell : \tau \mapsto hval \\
p &::= \mathsf{letrec} d \mathsf{ in } e \\
\Psi &::= \bullet \mid \Psi, \ell : \tau \\
\Delta &::= \bullet \mid \Delta, j \mid \Delta, \alpha : \kappa \\
\Gamma &::= \bullet \mid \Gamma, x : \tau \mid \Gamma, x_{64} : \phi
\end{aligned}$$

Figure 3.1: LIL syntax

permitted: however, in practice additional statically allocated data is placed in the heap section.¹

The next several sections discuss some of the interesting syntactic aspects of the **LIL** and introduce the relevant typing judgements. The complete static semantics for the **LIL** can be found in appendix B. Many of the constructs described here are also discussed in detail by Cray and Weirich [CW99].

3.1.1 Typing contexts

Context judgements	
Heap contexts	$\vdash \Psi \text{ ok}$
Constructor contexts	$\vdash \Delta \text{ ok}$
Term contexts	$\Delta \vdash \Gamma \text{ ok}$

There are three sorts of typing contexts for the **LIL**: heap contexts Ψ , constructor contexts Δ , and term contexts Γ . Heap contexts associate closed types with labels and are derived from the top level heap of a program. Constructor contexts actually serve two purposes: they bind free kind variables, and they bind free constructor variables at kinds (which may refer to previously bound kind variables). In principle these two contexts could be separated, but since kind variables are uni-typed it seems unnecessary. Term contexts bind 32 and 64 bit term variables at types, which may refer to previously bound constructor variables. I assume that variables for each of the different syntactic classes (constructors, kinds, 32 bit terms and 64 bit terms) are drawn from mutually disjoint infinite sets.

A heap context is well-formed if all of the labels in its domain are distinct, and if each of the types in its range is well-formed in an empty constructor context. This latter constraint reflects the fact that heap values are bound at the top-level of programs.

A constructor context is well-formed if all of the kind variables in its domain are distinct; and if all of the constructor variables in its domain are distinct, and if each type in its range is well-formed in the context preceding its binding.

A term context is well-formed if all of the 32 bit variables in its range are distinct, each type at which a 32 bit variable is bound is well-formed at kind T_{32} ; and if all of the 64 bit variables in its range are distinct, and if each type at which a 64 bit variable is bound is well-formed at kind T_{64} .

3.1.2 Constructors and Kinds

Constructor and kind judgements	
Kinds	$\Delta \vdash \kappa \text{ ok}$
Constructors	$\Delta \vdash c : \kappa$
Equivalence	$\Delta \vdash c \equiv c' : \kappa$

The constructor and kind typing judgements are unsurprising. The judgement $\Delta \vdash \kappa \text{ ok}$ defines what it means for a kind κ to be well-formed in a constructor context Δ . The judgement $\Delta \vdash c : \kappa$

¹The practice of referring to the static data segment of **LIL** programs as a “heap” reflects the standard terminology of the literature in this area: however, this usage is slightly misleading since in practice a **LIL** “heap” is understood to correspond more closely to the code and data segments of an executable image, rather than to dynamically allocated heap space. Nonetheless, since this distinction is not apparent in the dynamic semantics of **TILTAL**, and since the usage has become standard in the literature, I will continue to refer to this portion of a **LIL** program as the heap. Sorry Bob!

Constants and their kinds		
$\text{Float}:\text{T}_{64}$	$\text{Int}:\text{T}_{32}$	$\text{Void}:\text{T}_{32}$
$\text{Array}_{32}:\text{T}_{32} \rightarrow \text{T}_{32}$	$\text{Array}_{64}:\text{T}_{64} \rightarrow \text{T}_{32}$	$\text{Boxed}:\text{T}_{64} \rightarrow \text{T}_{32}$
$\text{Tag}:\text{nat} \rightarrow \text{T}_{32}$	$\text{DynTag}:\text{T}_{32} \rightarrow \text{T}_{32}$	$\text{Dyn}:\text{T}_{32}$
$\times:\text{T}_{32}\text{list} \rightarrow \text{T}_{32}$	$\rightarrow:\text{T}_{32}\text{list} \rightarrow \text{T}_{64}\text{list} \rightarrow \text{T}_{32} \rightarrow \text{T}_{32}$	
$\bigvee:\text{T}_{32}\text{list} \rightarrow \text{T}_{32}$	$\text{Code}:\text{T}_{32}\text{list} \rightarrow \text{T}_{64}\text{list} \rightarrow \text{T}_{32} \rightarrow \text{T}_{32}$	
$\forall:\forall j.(j \rightarrow \text{T}_{32}) \rightarrow \text{T}_{32}$	$\exists:\forall j.(j \rightarrow \text{T}_{32}) \rightarrow \text{T}_{32}$	
$\text{Rec}:\forall j.((j \rightarrow \text{T}_{32}) \rightarrow (j \rightarrow \text{T}_{32})) \rightarrow j \rightarrow \text{T}_{32}$		

Figure 3.2: Constructor constants and their kinds

defines what it means for a constructor c to be well-formed at a kind κ in a context Δ . And the judgement $\Delta \vdash c \equiv c' : \kappa$ defines when two constructors are equivalent. Note that the kind and the context in the equivalence judgement are present for technical reasons: they should not change the set of equivalences on well-formed constructors (though they do rule out equivalences on ill-formed constructors). The complete definitions of these judgements are given in appendix B.

The most important change from the **MIL** is the enrichment of the kind structure: most importantly, the addition of sum kinds with corresponding introduction and elimination forms at the constructor level. Universal kinds $(\forall j.\kappa)$ and inductive kinds $(\mu j.\kappa)$ are also provided. Kind variables bound by inductive kinds are restricted to occur only positively. In order to provide for the possibility of more general 64 bit types, the **LIL** uses an explicit kind distinction, providing kinds T_{32} and T_{64} corresponding to the kind of the types of 32 and 64 bit expressions respectively.

The meta-variables c , τ and ϕ are used to represent constructors of arbitrary kind, kind T_{32} and kind T_{64} respectively. The lambda calculus portion of the constructor level contains the usual introduction and elimination forms for sums, pairs, lambdas, unit, and kind abstraction. These constructs are entirely standard. Inductive kinds are introduced with a fold construct, $\text{fold}_{\mu j.\kappa} c$, which injects a constructor c of kind $\kappa[\mu j.\kappa/j]$ into the kind $\mu j.\kappa$.

The elimination form for inductive kinds is one of the most complex constructs in the **LIL** and requires some explanation. The essential idea is to provide a form of primitive recursion over inductive kinds at the constructor level. A primitive recursive constructor $\text{pr}(j, \alpha:\kappa, \rho:(j \rightarrow \kappa').c)$ recursively defines a function from $\mu j.\kappa$ to $\kappa'[\mu j.\kappa/j]$, with the body of the function give by c . The variable α is the argument to the function, and stands for the unfolded argument (nominally of kind $\kappa[\mu j.\kappa/j]$). However, in order to ensure that the function is only recursively called on a sub-component of the argument (and hence is guaranteed to terminate), α is bound at kind κ with occurrences of j left abstract, and the recursive variable ρ always has j as its domain kind.

The constructors classifying the term level are presented in the **LIL** as constants of higher kind. These constructors, given in figure 3.2, include constants for impredicative universal and existential types, general parameterized recursive types, arrays, tagged values, sums, integers, floating point numbers, boxed 64 bit values, pairs, and functions. Formulating these constructors as higher order constants makes the interaction with type analysis easier to deal with. The kinds of several of the constants refer to the kind of lists of constructors. This is defined as $\text{list}[k] = \mu j'.1 + k \times j'$. Throughout this dissertation I will frequently use the usual ML list syntax for constructor lists. The kind nat describing the encoding of natural numbers is definable directly in the **LIL** in the

usual fashion, and is used in the typing rules as well. When n is a natural number, I write $\bar{n}$ for the constructor of kind **nat** representing n .

For the most part, the constructor constants are fairly straightforward. For example, $\times c$ corresponds to the type of a tuple, the types of the fields of which are given by the elements of c . If $c = [\tau_1, \tau_2]$, then this corresponds to $\tau_1 \times \tau_2$ in a more standard notation. Similarly, the arguments to the $\rightarrow$ constructor correspond to the types of the 32 bit and 64 bit arguments of the function, and its return type. Hence, $\rightarrow([\tau])([\phi])(\tau)$ corresponds to $(\tau; \phi) \rightarrow \tau$.

This higher-order abstract syntax methodology is very convenient for both the theory and the implementation but can interfere with readability. For example: the type of the polymorphic pairing function written in this style appears as

$$\forall[T_{32}](\lambda(\alpha:T_{32}).\forall[T_{32}](\lambda(\beta:T_{32}).\rightarrow[a][\rightarrow[b]](\times[a,b])))$$

(Note that even here I have used derived notation for lists). To enhance readability, I will frequently make use of more standard notation. So in this example, I would write the type as:

$$\forall[\alpha:T_{32}, \beta:T_{32}].\alpha \rightarrow \beta \rightarrow (\alpha \times \beta)$$

This practice is also done extensively in the actual implementation via libraries which provide defined forms implementing the more standard type constructors in terms of the HOAS style constructors.

The **Tag** constructor classifies sum tags, and takes as its argument a constructor of kind **nat** indicating which branch of the sum the value inhabits.

Sum types are dealt with in the **LIL** using union types. In principle we allow injection into any union type. However, the case construct limits its arguments to have union types composed solely of tags or tagged records for which the tags are disjoint, and cover the full range of tags from the zero to the largest tag. So for example, the type $\vee[\text{Tag}(0), \times[\text{Tag}(1), \tau]]$ is inhabited by terms which are either **tag**₀ or a pair containing a τ in its second field, tagged with **tag**₁ in the first field. This union type therefore describes a valid argument to case. In point of fact, the tag on the pair is not necessary and could be elided. More generally, when τ is a pointer, then both the tag and the pointer indirection can be eliminated. Note though that in the presence of unknown types such an optimization again requires type analysis. The actual implementation of sums in **TILT** in fact performs this optimization. However, since it adds nothing new to the problem, the theoretical discussion here assumes a simple treatment of sums in which all value-carrying arms are tagged.

The universal and existential constructors take as arguments a kind indicating the domain of quantification and a constructor function giving the body. The standard $\forall(\alpha:\kappa).c$ becomes $\forall[\kappa](\lambda(\alpha::\kappa).c)$, and similarly for the existential. Finally, parameterized recursive types are defined with the **Rec** constructor. **Rec** $[\kappa](\lambda(\rho:\kappa \rightarrow T_{32}).\lambda(\alpha:\kappa).c)(c')$ (where c' has kind κ) corresponds to recursively defining a constructor of kind $\kappa \rightarrow T_{32}$ and applying it to c' . Within the body of c , ρ stands for the recursively defined constructor, and α stands for the parameter. Therefore, the unfolding of such a type is given by $c_r(\lambda(\alpha:\kappa).(\text{Rec}[\kappa](c_r)(\alpha)))(c')$, where $c_r = (\lambda(\rho:\kappa \rightarrow T_{32}).\lambda(\alpha:\kappa).c)$.

In order to allow closure conversion to be done within the **LIL** language, two “function” types are given. The $\rightarrow$ primitive type constructor when applied to arguments classifies functions in the usual sense. The **Code** primitive on the other hand, classifies “code” functions: that is, code that is closed.

3.1.3 Terms

Term judgements	
Small values	$\Psi; \Delta; \Gamma \vdash sv : \tau$
64 bit values	$\Psi; \Delta; \Gamma \vdash fv : \phi$
32 bit Operations	$\Psi; \Delta; \Gamma \vdash opr : \tau \text{ opr}_{32}$
64 bit Operations	$\Psi; \Delta; \Gamma \vdash fopr : \phi \text{ opr}_{64}$
Expressions	$\Psi; \Delta; \Gamma \vdash e : \tau \text{ exp}$
Heap values	$\Psi \vdash h : \tau \text{ hval}$
Heaps	$\Psi \vdash d \text{ ok}$
Programs	$\Psi \vdash p : \tau$

The expression level of the **LIL** is divided into five syntactic classes: 64 bit values (fv), 32 bit values (sv), 32 bit operations (opr), 64 bit operations ($fopr$) and expressions (e). Programs are syntactically restricted to a named form where all intermediate computations are named in the usual fashion. In addition to the five classes making up expressions, there are an additional three syntactic classes making up full **LIL** programs: heap values ($hval$), heaps (d), and programs (p).

The well-formedness judgements for the five expression classes are all of the same essential form, defining what it means for a term to be well-formed at a type τ (or ϕ) in heap context Ψ , constructor context Δ , and term context Γ . The other three judgements apply to heaps and programs which do not have free variables of any sort. Therefore, the only context present in these judgements is the heap context, which describes the free labels of a heap or a program. As usual, any heap value may refer to the label bound to any other heap value: that is, all the heap values are mutually recursively defined. The presence of a heap context in the program well-formedness judgement leaves open the possibility of compiling programs against externally defined labels.

Small values

Values include variables, constants, polymorphic instantiation, and existential introduction. Sum tags tag_i are made explicit, and a coercion is introduced to inject terms into the union type.

64 bit values

The only 64 bit values present in the **LIL** are 64 bit float constants and variables.

32 bit operations

Operations are computations that return values, and are bound to variables within expressions. For simplicity, I include values into the operations to unify let binding into a single mechanism. Other operations include unrolling of recursive types, tuple introduction and tuple selection, boxing, sum case, known sum projection, exception constructs, and boxing of 64 bit values. Array update and creation operations are provided for both 32 and 64 bit arrays, along with the array subscript operation for 32 bit arrays. Note that all memory allocation in the **LIL** is explicit and present in this level, whether through the tuple introduction, the box primitive or the array creation primitives. The only exception to this is function introduction which implicitly allocates a closure: this is dealt with via closure conversion which turns uses of functions into uses of tuples and code. This is described in detail in chapter 6 and section 9.3.

64 bit operations

The 64 bit operations include unboxing of 64 bit values and 64 bit array subscripts, as well as the inclusion of 64 bit values.

Expressions

Expressions in the **LIL** are either small values, or let bindings of any of several forms. Recursive function binding, 32 bit variable binding, and 64 bit variable binding are provided. Existential unpacking is also included at this level. The most interesting expressions however, are the type refinement bindings that support the technology upon which type analysis in the **LIL** is built.

The most important of these is the **vcase**, or “virtual case” construct. Crary and Weirich [CW99] describe a system for implementing type analysis as case analysis on constructor-level sums, essentially defining typecase in terms of more standard type theoretic constructs. They also describe a variant of their system which allows for a type erasure interpretation but still supports type analysis as a programming idiom. In this variant, term-level sums are used to stand for analyzable constructors. The **vcase** construct provides a mechanism whereby information about the identity of types encoded as terms can be reflected back into the type level.

To understand how this works, consider a value v of type $\text{case}(\alpha, \beta_1.c, \beta_2.\text{Void})$. According to the type given, v is either of type c or of type Void depending on whether α gets instantiated with a left or a right injection. But since there are no closed values of type Void , it is apparent that α can not be instantiated with a right injection, since to do so would imply that v has type Void . In essence, v serves as a witness to the fact that α will be instantiated with the left injection.

This fact can be propagated back into the type system by using **vcase**. When the argument to **vcase** is a variable, its arms are type-checked with the variable replaced with the appropriate injection in the context and in the body of the arm (that is, in the left arm, the variable is replaced with the left injection, and similarly with the right). Therefore, in the second branch of the expression $\text{vcase}(\alpha, \beta_1.e, \beta_2.\text{dead } v)$ the value v has type Void , which implies that this branch cannot be reached and is dead code. Within the body of e , α is known to be the left injection, and v is known to have type c . When types are erased, it is sound to erase the **vcase** as well, since one arm is known to be dead code. The **vcase** construct is the key to implementing type analysis, as the subsequent section will show.

In addition, two special binding forms exist for refining constructor paths into variables for analysis. When α has kind $\kappa_1 \times \kappa_2$, the pair refinement operator $\text{let}\langle\beta, \gamma\rangle = \alpha \text{ in } e$ replaces all occurrences of α in e with the pair $\langle\beta, \gamma\rangle$. This means that projections from α can be turned into variables. The **vcase** construct can only refine the type of its argument when the argument is a variable: this pair refinement construct allows this to be extended to paths as well. The $\text{let}(\text{fold } \beta) = \alpha \text{ in } e$ expression serves a similar purpose when α has a recursive kind.

Heap values

The only heap values currently supported in the theoretical treatment of the **LIL** are code functions, which are necessary for closure conversion. Each heap value is required to be closed with respect to variables, but may refer to any other heap value via its label.

Heaps

A heap is a collection of mutually recursively defined heap values. A heap is well-formed in a heap context if each of its constituent heap values is well-formed at the type at which its label is bound. Note that I require that all of the heap labels be present in the heap context: for non-recursive functions this is not strictly necessary, but it serves the additional purpose of enforcing the property that all labels in the heap are distinct (via well-formedness check on heap contexts).

$$\frac{\vdash \Psi \text{ ok}}{\Psi \vdash \epsilon \text{ ok}}$$

$$\frac{\Psi[\ell:\tau] \vdash hval:\tau \text{ hval} \quad \Psi[\ell:\tau] \vdash d \text{ ok}}{\Psi[\ell:\tau] \vdash d, \ell:\tau \mapsto hval \text{ ok}}$$

Programs

LIL programs consist of a heap, which binds labels to heap values, and an expression which computes the “value” of the program. Since the heap values are (potentially) mutually recursive, the heap is checked in a context including bindings for all of the labels in the heap. For notational purposes, I define an operation on heaps, $\Psi(d)$, that corresponds to the heap context produced by taking the label and type from each binding in a heap

$$\Psi(\epsilon) \stackrel{\text{def}}{=} \bullet$$

$$\Psi(d, \ell:\tau \mapsto hval) \stackrel{\text{def}}{=} \Psi(d), \ell:\tau$$

The well-formedness rule for program simply checks the heap and the expression portion of the program in the heap context extended with the bindings from the heap.

$$\frac{\Psi, \Psi(d) \vdash \text{ ok} \quad \Psi, \Psi(d) \vdash d \text{ ok} \quad \Psi, \Psi(d); \bullet; \bullet \vdash e:\tau \text{ exp}}{\Psi \vdash \text{ letrec } d \text{ in } e:\tau}$$

3.2 Useful properties of the LIL

Subsequent proofs rely on certain properties of the **LIL**: in particular, that well-typedness is preserved under substitution and that weakening for typing contexts is admissible.

Lemma 1 (LIL Substitution)

If $\Delta, \alpha:\kappa' \vdash c:\kappa$ and $\Delta \vdash c':\kappa'$ then $\Delta \vdash c[c'/\alpha]:\kappa$.

Proof: (By induction on c)

We proceed by cases on the last rule of the derivation.

1. If c is a constant or $*$, then $c[c'/\alpha] = c$. By assumption, $\Delta, \alpha:\kappa_1 \vdash c:\kappa$, so by strengthening, $\Delta \vdash c:\kappa$.

2. $\Delta[\alpha:\kappa'] \vdash \alpha' : \kappa$

- $\alpha' = \alpha$. Then $c[c'/\alpha] = c_1$. By assumption, $\Delta \vdash c' : \kappa'$ and $\Delta, \alpha:\kappa' \vdash \alpha : \kappa$. But by inspection of the typing rules, it is clear that the derivation must end in a use of the variable rule, so $\kappa = \kappa'$, and hence $\Delta \vdash c' : \kappa$
- $\alpha' \neq \alpha$. Then the result follows as in the constant case above.

3. $\Delta[\alpha:\kappa'] \vdash \lambda(\beta:\kappa_\beta).c : \kappa_\beta \rightarrow \kappa_r$. Then by assumption, $\Delta, \alpha:\kappa', \beta:\kappa_\beta \vdash c : \kappa_r$. By induction, $\Delta, \beta:\kappa_\beta \vdash c[c'/\alpha] : \kappa_r$. (Note that we assume that β is chosen appropriately to avoid capture.) Then by the lambda rule, $\Delta \vdash \lambda(\beta:\kappa_\beta)(c[c'/\alpha]) : \kappa$, and hence by definition of substitution, $\Delta \vdash (\lambda(\beta:\kappa_\beta)c)[c'/\alpha] : \kappa$.

4. $\Delta[\alpha:\kappa'] \vdash c_1 \ c_2 : \kappa$. Then by assumption, $\Delta, \alpha:\kappa' \vdash c_1 : \kappa_2 \rightarrow \kappa$ and $\Delta, \alpha:\kappa' \vdash c_2 : \kappa_2$. By induction, $\Delta \vdash c_1[c'/\alpha] : \kappa_2 \rightarrow \kappa$ and $\Delta \vdash c_1[c'/\alpha] : \kappa_2$. The result follows directly then from the application rule and the definition of substitution: $\Delta \vdash (c_1 \ c_2)[c'/\alpha] : \kappa$.

5. $\Delta[\alpha:\kappa'] \vdash \langle c_1, c_2 \rangle : \kappa_1 \times \kappa_2$. Then by assumption $\Delta, \alpha:\kappa' \vdash c_1 : \kappa_1$ and $\Delta, \alpha:\kappa' \vdash c_2 : \kappa_2$. By induction, $\Delta \vdash c_1[c'/\alpha] : \kappa_1$ and $\Delta \vdash c_2[c'/\alpha] : \kappa_2$. The result follows directly then from the pair rule and the definition of substitution: $\Delta \vdash \langle c_1, c_2 \rangle[c'/\alpha] : \kappa$.

6. $\Delta[\alpha:\kappa'] \vdash \pi_1 c_p : \kappa_1$. Then by assumption $\Delta, \alpha:\kappa' \vdash c_p : \kappa \times \kappa_2$. By induction, $\Delta \vdash c_p[c'/\alpha] : \kappa \times \kappa_2$. The result follows from the the projection rule and the definition of substitution.

7. $\Delta[\alpha:\kappa'] \vdash \pi_2 c_p : \kappa_2$. As for π_2 .

8. $\Delta[\alpha:\kappa'] \vdash \text{inj}_i c : +[\kappa_1, \dots, \kappa_n]$. Then by assumption $\Delta, \alpha:\kappa' \vdash c : \kappa_i$. By induction, $\Delta \vdash c[c'/\alpha] : \kappa_i$. Finally, by the injection rule and the definition of substitution:

$$\Delta, \alpha:\kappa' \vdash (\text{inj}_i c)[c'/\alpha] : +[\kappa_1, \dots, \kappa_n]$$

9. $\Delta[\alpha:\kappa'] \vdash \text{case } c[(\alpha_1.c_1, \dots, \alpha_n.c_n)] : \kappa$. By assumption:

- $\Delta, \alpha:\kappa' \vdash c : +[\kappa_1, \dots, \kappa_n]$
- $\Delta, \alpha:\kappa', \alpha_i:\kappa_i \vdash c_i : \kappa$

By induction:

- $\Delta \vdash c[c'/\alpha] : +[\kappa_1, \dots, \kappa_n]$
- $\Delta[\alpha_i:\kappa_i] \vdash c_i[c'/\alpha'] : \kappa$

(Note that the α_i may always be chosen to avoid capture). Then by the case rule and the definition of substitution, $\Delta \vdash (\text{case } c[(\alpha_1.c_1, \dots, \alpha_n.c_n)])[c'/\alpha] : \kappa$

10. $\Delta[\alpha:\kappa'] \vdash \text{fold}_{\mu j.\kappa} c : \mu j.\kappa$. Then by assumption, $\Delta[\alpha:\kappa'] \vdash c : \kappa[\mu j.\kappa/j]$ and $\Delta[\alpha:\kappa'] \vdash \mu j.\kappa \text{ ok}$. By induction, $\Delta[\alpha:\kappa'] \vdash c[c'/\alpha] : \kappa[\mu j.\kappa/j]$, and so by the fold rule and the definition of substitution, $\Delta[\alpha:\kappa'] \vdash (\text{fold}_{\mu j.\kappa} c)[c'/\alpha] : \kappa[\mu j.\kappa/j]$

11. $\Delta[\alpha:\kappa'] \vdash \mathbf{pr}(j, \beta:\kappa_1, \rho:(j \rightarrow \kappa_2), \mathbf{in} \ c) : \mu j. \kappa_1 \rightarrow \kappa_2$ Then by assumption,

$$\Delta[\alpha:\kappa'] \vdash \mathbf{pr}(j, \beta:\kappa_1, \rho:(j \rightarrow \kappa_2), \mathbf{in} \ c) : \mu j. \kappa_1 \rightarrow \kappa_2$$

By induction

$$\Delta, j, \beta:\kappa_1, \rho:(j \rightarrow \kappa_2) \vdash c[c'/\alpha] : \kappa_2$$

Note that we can choose variables appropriately to avoid capture, and hence by the primitive recursion rule and the definition of substitution

$$\Delta[\alpha:\kappa'] \vdash \mathbf{pr}(j, \beta:\kappa_1, \rho:(j \rightarrow \kappa_2), \mathbf{in} \ c)[c'/\alpha] : \mu j. \kappa_1 \rightarrow \kappa_2$$

12. $\Delta[\alpha:\kappa'] \vdash c[\kappa] : \kappa_2[\kappa/j]$.

Then by assumption:

$$\Delta[\alpha:\kappa'] \vdash c : \forall j. \kappa_2$$

By induction:

$$\Delta \vdash c[c'/\alpha] : \forall j. \kappa_2$$

By the kind application rule and the definition of substitution:

$$\Delta \vdash c[\kappa][c'/\alpha] : \kappa_2[\kappa/j]$$

13. $\Delta[\alpha:\kappa'] \vdash \Lambda j. c : \forall j. \kappa$. Then by assumption, $\Delta[\alpha:\kappa'], j \vdash c : \kappa$. By induction $\Delta, j \vdash c[c'/\alpha] : \kappa$. Hence by the kind abstraction rule and the definition of substitution $\Delta \vdash (\Lambda j. c)[c'/\alpha] : \kappa$.

■

A few structural properties of **LIL** contexts are important as well. I state the weakening property and give a sketch of the proof. I remain informal about re-ordering properties of typing contexts throughout, but in the absence of dependent types or kinds it is straightforward (though tedious) to formalize them. I will also occasionally informally use simultaneous weakening to combine whole contexts: e.g. Δ, Δ' where the intersection of the domains is empty. It should be clear that this form of weakening can also be formalized by an induction on the second context, appealing to the unary forms of the weakening lemmas to incrementally construct the new context.

Lemma 2 (LIL Weakening)

1. If $\Delta \vdash \kappa \ \mathbf{ok}$ and $j \notin \Delta$, then $\Delta, j \vdash \kappa \ \mathbf{ok}$.
2. If $\Delta \vdash \kappa \ \mathbf{ok}$ and $\Delta \vdash \kappa' \ \mathbf{ok}$ and $\alpha \notin \Delta$, then $\Delta, \alpha:\kappa' \vdash \kappa \ \mathbf{ok}$.
3. If $\Delta \vdash c : \kappa$ and $\Delta \vdash \kappa' \ \mathbf{ok}$ and $\alpha \notin \Delta$ then $\Delta, \alpha:\kappa' \vdash c : \kappa$.
4. If $\Delta \vdash c \equiv c' : \kappa$ and $\Delta \vdash \kappa' \ \mathbf{ok}$ and $\alpha \notin \Delta$ then $\Delta, \alpha:\kappa' \vdash c \equiv c' : \kappa$.

Proof (Sketch): Each of the proofs proceeds similarly by induction on the structure of typing derivations. For each typing rule, I inductively construct new sub-derivations for each premise and use the original rule to construct the new derivation. The side-condition on binding sites does not follow immediately: it is necessary to observe that it is always possible to use alpha-variance to choose an appropriate bound variable different from α . For premises of the form $\vdash \Delta \ \mathbf{ok}$, note that the derivation of $\vdash \Delta, \alpha:\kappa' \ \mathbf{ok}$ follows directly from the assumptions and the definition of context well-formedness.

■

I also state an inversion property of **LIL** operations.

Lemma 3 (LIL inversion)

1. If D is a derivation of $\Psi; \Delta; \Gamma \vdash opr : \tau \text{ opr}_{32}$, then D has a unique last rule for any choice of opr .
2. If D is a derivation of $\Psi; \Delta; \Gamma \vdash fopr : \phi \text{ opr}_{64}$, then D has a unique last rule for any choice of opr .

Proof: By inspection. No two operation rules apply to the same syntactic construct. ■

Chapter 4

Example: Floating point array flattening

4.1 An optimized array strategy

The use of the **LIL** type refinement operations discussed in the previous chapter is somewhat non-intuitive. Before attempting to give a full account of the translation of the **MIL**, it is useful to consider a small example demonstrating the type-analysis methodology used in the **LIL**. The optimization I will describe here, floating point array flattening, is one of several implemented by the **TILT** compiler. In this scheme, an array of boxed floating point numbers is implemented not as an array of pointers, but as an array of the actual 64-bit values. Polymorphic array operations dispatch on the type of the array contents to determine which primitives to use. In an informal λ_i^{ML} style notation, this corresponds to defining the following type:

$$\text{Array}_{\text{opt}}(\alpha) = \text{Typecase}(\alpha) \text{ of } \text{Boxed}(\phi) \Rightarrow \text{Array}_{64}(\phi) \\ | _ \Rightarrow \text{Array}_{32}(\alpha)$$

As an example, I will show how to write a term level array constructor `arrayopt` that can be used to construct such arrays.

4.2 LIL implementation of optimized arrays

Type analysis in the **LIL** is based on the idea of first encoding types as abstract syntax trees, and then writing functions that use the encodings to reconstruct the actual types or to choose different branches of code. In fact, it is not necessary to encode entire types: we may choose an encoding that captures only the information that is useful for our purposes.

More concretely, we first choose an encoding strategy that captures the features of interest about a type. For every type of interest (that is, every type which is to be analyzed) we construct two object level items: a constructor which serves as a *static encoding* of the type (**SE**), and a term which serves as a *dynamic encoding* of the type (**DE**). Static encodings of types are used during typechecking to reconstruct the encoded type and to connect the type to its dynamic encoding. Dynamic encodings are used to perform the actual runtime dispatch. For each of these two encodings, there must also be a corresponding classifier describing it: the static encoding of a type is

classified by the *static encoding kind* (**SEK**), and the dynamic encoding of a type is classified by its *dynamic encoding type* (**DET**).

4.2.1 Static encodings

At the kind level, sums and recursive kinds are used to build a type's static encoding. For the purposes of this example optimization, the encoding can be very simple: we need only to be able to distinguish between boxed 64-bit types and other types. This is reflected in the static encoding kind for our example:

$$T_{\text{opt}} \stackrel{\text{def}}{=} T_{64} + T_{32} \quad (* \text{ Boxed of } (\phi) \mid \text{ Other of } (\tau)*)$$

The T_{opt} kind classifies static encodings. The comment is meant to suggest the intuitive correspondence between this definition and an ML-style representation of an abstract syntax tree.

Constructors of kind T_{opt} encode types. In order to take advantage of these encodings, we replace all constructors by their encodings. For example, polymorphic functions must expect encoded types as their type arguments. In order to be able to typecheck applications of these functions, we must be able to reconstruct the original type from its encoding. This is done by writing an object level function to *interpret* static encodings into the types they encode. For the example at hand, this function simply boxes up the 64 bit types and leaves the 32 bit types unchanged. Hence:

$$\text{interp} : T_{\text{opt}} \rightarrow T_{32} \stackrel{\text{def}}{=} \lambda(\alpha : T_{\text{opt}}). \text{case } \alpha \text{ of } \begin{array}{l} \text{inj}_1 \beta \Rightarrow \text{Boxed}(\beta) \\ \mid \text{inj}_2 \beta \Rightarrow \beta \end{array}$$

The use of this interpretation function can be seen by considering the static encodings of boxed 64-bit floats and 32-bit ints:

$$\begin{array}{ll} C_{BF} : T_{\text{opt}} & \stackrel{\text{def}}{=} \text{inj}_1^{T_{\text{opt}}} \text{Float} \\ C_I : T_{\text{opt}} & \stackrel{\text{def}}{=} \text{inj}_2^{T_{\text{opt}}} \text{Int} \end{array}$$

Notice that the interpretation function maps each of these back to the original type. This allows us to reconstruct the original type from the encoding. An important point to note here is that this scheme does not guarantee that all boxed floating point types will be represented as a first injection: it is perfectly possible to stuff $\text{Boxed}(\text{Float})$ into the second arm of the sum. This is not a problem, since it merely causes arrays of such types to be represented in un-flattened form. In fact, one could look at this as a feature, since the optimization does not force you to choose one representation or the other.

Of course, simply writing the interpretation is fairly uninteresting by itself since it merely gets us back to where we were before we encoded types. The real benefit of encoding types in this fashion comes from other useful functions that we can write using encodings of types. In particular, the actual type of specialized arrays that was given informally above can now be written down directly.

$$\text{Array}_{\text{opt}} : T_{\text{opt}} \rightarrow T_{32} \stackrel{\text{def}}{=} \lambda(\alpha : T_{\text{opt}}). \text{case } \alpha \text{ of } \begin{array}{l} \text{inj}_1 \beta \Rightarrow \text{Array}_{64}(\beta) \\ \mid \text{inj}_2 \beta \Rightarrow \text{Array}_{32}(\beta) \end{array}$$

The $\text{Array}_{\text{opt}}$ constructor is a function that maps encoded types to optimized array types. It takes advantage of the additional information present in the encoded types to represent arrays of boxed

floats more efficiently. When the encodings are statically known, the optimized array type can be reduced directly to an underlying primitive array type.

$$\begin{aligned}\text{Array}_{\text{opt}}(C_{BF}) &\equiv \text{Array}_{64}(\text{Float}) \\ \text{Array}_{\text{opt}}(C_I) &\equiv \text{Array}_{32}(\text{Int})\end{aligned}$$

When the actual encoding is abstract (for example in a polymorphic function) the composite array constructor can not be reduced definitively to a primitive array type, since it is not statically known which sort of array can be expected.

4.2.2 Dynamic encodings

Static encodings account for type analysis at the type level by representing types as constructor level abstract syntax trees and by writing functions which analyze the structure of encodings to produce optimized results. Type analysis at the term level is accounted for in an analogous fashion using dynamic encodings. In addition to encoding every type at the constructor level, we also give a dynamic encoding of every type at the term level. The **DE** of a type can then be used to dispatch at runtime. A key point of this methodology is that the dynamic and static encodings of a type are not independent: they are related via the dynamic encoding type. This allows information gained via tests on the dynamic encoding to be reflect back onto the static encoding.

In fact, the **DET** of a type is simply another object level type function operating on the **SE** of the type in the same fashion as the interpretation function and the optimized array functions. For the example at hand, the type of the dynamic encoding of a type whose static encoding is c is given by $R(c)$, where R is defined as follows:

$$R : T_{\text{opt}} \rightarrow T_{32} \stackrel{\text{def}}{=} \lambda(\alpha : T_{\text{opt}}). (\text{case } \alpha \text{ of } \text{inj}_1 \beta \Rightarrow \text{Unit} \mid \text{inj}_2 \beta \Rightarrow \text{Void}) \\ + (\text{case } \alpha \text{ of } \text{inj}_1 \beta \Rightarrow \text{Void} \mid \text{inj}_2 \beta \Rightarrow \text{Unit})$$

This definition is a bit subtle, and is worth examining in detail. The principle high-level point is that dynamic encodings of types will be values of sum type, which can be dispatched on using the term level case construct. The subtlety arises in the types given for the two arms of the sum type here. Metaphorically speaking, each arm of the sum here can be thought of as serving as a “proof” of the identity of α . The reasoning behind this is to observe that there are no closed values of type **Void**. This tells us that any closed value of type **case** c of $\text{inj}_1 \beta \Rightarrow \text{Unit} \mid \text{inj}_2 \beta \Rightarrow \text{Void}$ must in fact be unit. But this in turn tells us (informally) that c can only be of the form $\text{inj}_1 \tau$.

So the informal reasoning of the previous paragraph tells us that after dispatching on the dynamic encoding of a type, we can informally infer information about its static encoding. In particular, if the dynamic encoding is a left injection, then the static encoding must similarly be a left injection; and similarly for right injections. In fact, as we shall see, the **vcase** constructs gives us a formal method for using this information.

Dynamic encodings of types are terms whose types are given by the operation of R on their static encoding. So for example, the dynamic encodings of boxed 64-bit floats and of 32-bit ints are given as follows:

$$\begin{aligned}V_{BF} : R(C_{BF}) &\stackrel{\text{def}}{=} \text{inj}_1 * \\ V_I : R(C_I) &\stackrel{\text{def}}{=} \text{inj}_2 *\end{aligned}$$

4.2.3 Type analyzing terms

To re-cap, static encodings of types (such as C_{BF}) provide type level representations of the abstract syntax of the analyzable types. Dynamic encodings of types (such as V_{BF}) provide term level representations of the abstract syntax of the analyzable types, and are classified by types which depend on the static encoding of the type.

In this framework, functions which analyze types (given in the **MIL** by special primitives) simply become functions which take as arguments the static and dynamic encodings of the analyzable types and dispatch on them appropriately. The type of the function to construct flattened arrays reflects this:

$$\text{array}_{\text{opt}} : \forall(\alpha:T_{\text{opt}}).R(\alpha) \rightarrow \text{Int} \rightarrow \text{interp}(\alpha) \rightarrow \text{Array}_{\text{opt}}(\alpha)$$

The type argument and the first term argument correspond to the static and dynamic encodings of the type of the contents of the array. The type of the initial value for the array is obtained by interpreting the encoded type using the interp function, and the return type is given by interpreting the encoded type using the $\text{Array}_{\text{opt}}$ function.

The actual implementation of the array creation function simply case analyzes the dynamic encoding to determine the appropriate primitive array operation to use. The code itself is straightforward: the subtlety lies in understanding why the code is well-typed.

$$\begin{aligned} \text{array}_{\text{opt}} &\stackrel{\text{def}}{=} \Lambda(\alpha:T_{\text{opt}}).\lambda(x_\alpha:R(\alpha)).\lambda(i:\text{Int}).\lambda(y:\text{interp}(\alpha)). \\ &\quad \text{case } x_\alpha \\ &\quad \quad \text{of } \text{inj}_1 r \Rightarrow \text{vcase } \alpha \text{ of } \text{inj}_1 \beta_1 \Rightarrow \text{array}_{64}[\beta_1](i, \text{unbox}(y)) \\ &\quad \quad \quad | \text{inj}_2 \beta_2 \Rightarrow \text{dead } r \\ &\quad \quad | \text{inj}_2 r \Rightarrow \text{vcase } \alpha \text{ of } \text{inj}_1 \beta_1 \Rightarrow \text{dead } r \\ &\quad \quad \quad | \text{inj}_2 \beta_2 \Rightarrow \text{array}_{32}[\beta_2](i, y) \end{aligned}$$

As expected, the type argument α has the kind of encoded types, T_{opt} . Similarly, the term argument x_α has the type of dynamic encodings of α , and will be instantiated with the dynamic encoding of α . The actual type of the contents of the array can only be referred to via the interp function, since it is unknown at compile time. Therefore, the y argument with which the array is to be initialized is given type $\text{interp}(\alpha)$. Similarly, the return type $\text{Array}_{\text{opt}}(\alpha)$ gives the type of the returned array as a function of what α turns out to be.

The key to understanding how this all works out is to observe the effect that the vcase constructs have on the types of the variables in the function. Consider just the first branch of the case analysis on x_α , where the body of the arm does a virtual case analysis of α . According to the typing rule for vcase , the second arm of the vcase will be type-checked with $\text{inj}_2 \beta_2$ substituted everywhere for α , *including in the context*. This means that whereas outside the arm the variable r has type

$$\text{case } \alpha \text{ of } \text{inj}_1 \beta \Rightarrow \text{Unit} \mid \text{inj}_2 \beta \Rightarrow \text{Void}$$

within the arm it has type

$$\text{case } \text{inj}_2 \beta_2 \text{ of } \text{inj}_1 \beta \Rightarrow \text{Unit} \mid \text{inj}_2 \beta \Rightarrow \text{Void}$$

which is equivalent to simply Void . This satisfies the typing rule for vcase , which requires the dead branch to exhibit a value of type Void as proof that the branch is in fact dead.

To illustrate this further, the following table shows the types for the variables r and y inside and outside of the first **vcase**, at the occurrences indicated in bold in the definition above.

Outside of vcase	Inside of vcase
y : case α of $\text{inj}_1 \beta \Rightarrow \text{Boxed}(\beta)$ $\text{inj}_2 \beta \Rightarrow \beta$	y : $\text{Boxed}(\beta_1)$
r : case α of $\text{inj}_1 \beta \Rightarrow \text{Unit}$ $\text{inj}_2 \beta \Rightarrow \text{Void}$	r : Void

As an exercise, the reader may verify that when called with the appropriate arguments, the optimized array function defined above does in fact reduce to the appropriate 32 or 64-bit array primitive based on the representation of the type chosen.

$$\begin{aligned}
\text{array}_{\text{opt}}[C_{BF}](V_{BF})(10)(\text{box}(0.0)) &\mapsto^* \text{array}_{64}[\text{Float}](10, 0.0) \\
\text{array}_{\text{opt}}[C_I](V_I)(10)(0) &\mapsto^* \text{array}_{32}[\text{Int}](10, 0)
\end{aligned}$$

Chapter 5

The MIL to LIL translation

The translation from the **MIL** into the **LIL** language is primarily interesting in that it makes the uses of type analysis in the **MIL** primitives explicit and adds representations for types so that this analysis can be done at the term level. This implements in a typed fashion what is done currently in an untyped setting. The type analysis methodology for this is essentially the same as that described in the previous chapter, but extended to handle additional type analyzing operations.

The first section of this chapter gives a high-level overview of the translation, introducing the relations and stating some of the major typing properties. These translations are developed in detail in successive sections, along with proofs of their soundness.

5.1 Translation overview

The translation of **MIL** programs in to **LIL** programs is defined by several inductively defined relations between elements of the **MIL** syntactic classes and their correspondents in the **LIL**. These relations may be broadly grouped into four classes: those concerned with the static encoding of constructors, those concerned with the dynamic encoding of constructors, those concerned with proper **MIL** types, and those concerned with the expression translation itself.

The first group consists of the kind and constructor translations, and the corresponding translation on constructor typing contexts.

Static encoding translations		
SEK	$ \kappa $	$\bullet \vdash \kappa \text{ ok}$
SE context	$ \Delta $	$\vdash \Delta \text{ ok}$
SE	$ c $	$ \Delta \vdash c : \kappa $

The kind translation replaces **MIL** kinds (closed by definition) with closed **LIL** kinds classifying the translation of **MIL** constructors of the original kind. The static encoding translation for contexts ($|\Delta|$) simply applies the kind translation across the range of a **MIL** typing context. This allows the statement of the desired typing property of the constructor translation: that if $\Delta \vdash c : \kappa$, then $|\Delta| \vdash |c| : |\kappa|$ (theorem 2). (There are of course a number of other auxiliary typing properties to be shown: these are covered in more detail in subsequent sections.)

The second group, primarily concerned with the dynamic encoding of constructors, consists of additional translations on constructors and typing contexts and a relation on constructors and

kinds.

Dynamic encoding translations		
DET	$\lfloor c:\kappa \rfloor$	$ \Delta \vdash \lfloor c:\kappa \rfloor : T_{32}$
DE context	$\lfloor \Delta \rfloor$	$ \Delta \vdash \lfloor \Delta \rfloor \text{ ok}$
DE	$\lfloor c \rfloor$	$\bullet; \Delta ; \lfloor \Delta \rfloor \vdash \lfloor c \rfloor : \lfloor c:\kappa \rfloor \text{ exp}$

The first translation, $\lfloor c:\kappa \rfloor$, is a relation between **MIL** constructor/kind pairs and **LIL** types. It defines the dynamic encoding type for a constructor c of kind κ . The inhabitants of this type are the dynamic encodings of constructors, and are produced by a relation between **MIL** constructors and **LIL** expressions: $\lfloor c \rfloor$. The dynamic encoding translation for contexts maps **LIL** constructor contexts to **MIL** term contexts. The domain of the new context is constructed using an injection from type variables to term variables, while the range is constructed using the static encoding type translation. It should be the case that if $\Delta \vdash c:\kappa$, then $\Psi; |\Delta|; \lfloor \Delta \rfloor \vdash \lfloor c \rfloor : \lfloor c:\kappa \rfloor \text{ exp}$ (theorem 3)

The third group is concerned with the translation of proper **MIL** types (as opposed to constructors).

Type translations		
Types	$ \tau $	$ \Delta \vdash \tau : T_{32}$
Term context	$ \Gamma $	$ \Delta \vdash \Gamma \text{ ok}$

The type translation is a relation between proper **MIL** types and **LIL** constructors of kind T_{32} , and the term context translation simply maps this translation across the range of **MIL** term contexts. It should be the case that if $\Delta \vdash \tau \text{ ok}$, then $|\Delta| \vdash |\tau| : T_{32}$.

The final group of translations relates **MIL** terms of various syntactic classes to their corresponding **LIL** terms.

Term translations		
Small values	$\Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv'$	$\Psi; \Delta ; \lfloor \Delta \rfloor, \Gamma \vdash sv' : \tau $
Float values	$\Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv'$	$\Psi; \Delta ; \lfloor \Delta \rfloor, \Gamma \vdash fv' : \text{Float}$
Operations	$\Delta; \Gamma \vdash opr : \tau \rightsquigarrow opr'$	$\Psi; \Delta ; \lfloor \Delta \rfloor, \Gamma \vdash opr' : \tau \text{ opr}_{32}$
64 bit Operations	$\Delta; \Gamma \vdash fopr : \text{Float} \rightsquigarrow fopr'$	$\Psi; \Delta ; \lfloor \Delta \rfloor, \Gamma \vdash fopr' : \text{Float opr}_{64}$
Expression	$\Delta; \Gamma \vdash e : \tau \rightsquigarrow e'$	$\Psi; \Delta ; \lfloor \Delta \rfloor, \Gamma \vdash e' : \tau \text{ exp}$

It is convenient to phrase these as typed translations since the additional type information is sometimes required in order to construct the appropriate **LIL** syntax. The intended typing properties of these translations should be clear. For example, it should be the case that if $\Delta; \Gamma \vdash e : \tau$ and $\Delta; \Gamma \vdash e : \tau \rightsquigarrow e'$ then $\Psi; |\Delta|; \lfloor \Delta \rfloor, |\Gamma| \vdash e' : |\tau| \text{ exp}$ (for appropriate heap contexts Ψ , theorem 8).

5.2 Static encodings of constructors

A static encoding is a **LIL** constructor which encodes information about the **MIL** constructor which it represents. The static encoding can be analyzed at the type level to determine what type it represents. It can also be translated by an object level interpretation function to determine the actual **LIL** type corresponding to the **MIL** type that it represents. This section develops these mechanisms, beginning with the translation of kinds.

5.2.1 The kind translation

The choice of what information to capture in the encoding depends entirely on the kind of type analysis optimization that is desired. In the **MIL**, there are two major uses of type analysis: an array flattening optimization and the *vararg* optimization. The array optimization is almost exactly as described in chapter 4, and requires the encoding strategy to distinguish between boxed floats and other types. The *vararg* optimization is described in chapter 2, and requires the ability to distinguish between records of various widths and other types so that the **vararg** and **onearg** operations can choose an appropriate calling convention. For the purposes of this translation then, it is sufficient for the encoding to capture three classes of types: records (with their widths), boxed floating point numbers, and all other types. In fact, the actual implementation sub-divides the last category into pointer types and non-pointer types in order to implement a further optimization on sums that is not discussed here, as it adds nothing substantial to the discussion.

This division of types into categories is apparent in the static encoding kind that we choose for the translation:

$$T_{\text{mil}} \stackrel{\text{def}}{=} T_{32\text{list}} + 1 + T_{32}$$

Intuitively, T_{mil} corresponds to an ML datatype as such:

$$\text{datatype } T_{\text{mil}} = \mathbf{Record} \text{ of } T_{32\text{list}} \mid \mathbf{BFloat} \mid \mathbf{Other} \text{ of } T_{32}$$

For presentational purposes, I will often use ML pattern matching style notation using this intuitive correspondence.

Recall that the **SEK** is the kind that classifies static encodings. What this definition tells us is that encoded constructors are either a list of types corresponding to the types of the fields of a record, a boxed floating point number, or some other unknown type. Note that since the only 64 bit type in the **MIL** is **Float**, it is not necessary to include any information in the second arm of the sum. If we wished to allow arbitrary 64 bit types to be flattened into arrays, we could replace the 1 in the second arm of the sum with a T_{64} .

The previous definition tells us what the kind of encodings of proper **MIL** types looks like. Arbitrary **MIL** kinds are translated into **LIL** kinds simply by replacing all occurrences of T_{32} with T_{mil} , and leaving the rest of the structure intact.

$$\begin{aligned} |T_{32}| &\stackrel{\text{def}}{=} T_{\text{mil}} \\ |\kappa_1 \rightarrow \kappa_2| &\stackrel{\text{def}}{=} |\kappa_1| \rightarrow |\kappa_2| \\ |\kappa_1 \times \kappa_2| &\stackrel{\text{def}}{=} |\kappa_1| \times |\kappa_2| \end{aligned}$$

Following the methodology described in chapter 4, the first important object level type function we must define to assist with the encoding is an interpretation function which captures the meaning of the encoding in terms of the underlying types. This is done by defining an interpretation function interp of kind $T_{\text{mil}} \rightarrow T_{32}$.

$$\begin{aligned} \text{interp} : T_{\text{mil}} \rightarrow T_{32} &\stackrel{\text{def}}{=} \lambda(\alpha : T_{\text{mil}}). \\ &\text{case } \alpha \\ &\quad \mathbf{Record } l \Rightarrow \times(l) \\ &\quad \mid \mathbf{BFloat} \Rightarrow \text{Boxed}(\text{Float}) \\ &\quad \mid \mathbf{Other } t \Rightarrow t \end{aligned}$$

For clarity, I use an ML style pattern matching notation based on the informal datatype definition given above. It should be completely clear how to translate this to formal syntax simply by replacing uses of **Record** l , for example, with $\text{inj}_1 l$.

Notice that this function is an object level function as opposed to a meta-level translation. This is necessary, since type abstraction means that the interpretation cannot always be statically computed.

5.2.2 The constructor translation

The actual constructor translation $|c|$ translates **MIL** constructors of kind κ to **LIL** constructors of kind $|\kappa|$. This means that constructors of kind T_{32} will be mapped to **LIL** constructors of kind T_{mil} . For clarity in the translation, I begin by defining some **LIL** functions that serve to construct static encodings of types. These serve as “constructors” for the ML style notation suggested by the datatype given above.

$$\begin{array}{lll} \mathbf{Record} & : T_{32} \text{list} \rightarrow T_{\text{mil}} & \stackrel{\text{def}}{=} \lambda(\alpha : T_{32} \text{list}). \text{inj}_1^{T_{\text{mil}}} \alpha \\ \mathbf{BFloat} & : T_{\text{mil}} & \stackrel{\text{def}}{=} \text{inj}_2^{T_{\text{mil}}} (\text{inj}_1 *) \\ \mathbf{Other} & : T_{32} \rightarrow T_{\text{mil}} & \stackrel{\text{def}}{=} \lambda(\alpha : T_{32}). \text{inj}_2^{T_{\text{mil}}} (\text{inj}_2 \alpha) \end{array}$$

Note that for syntactic clarity I frequently leave off the kind decoration on sum injections where it is obvious from context. Frequently, throughout this dissertation I will use boldface for names of defined forms such as these to distinguish them from primitive syntax.

The actual static encoding translation proceeds for the most part compositionally over the structure of constructors. Constructors of higher kind are translated directly by recursively translating their sub-components. All other constructors are encoded into the T_{mil} kind.

It is convenient to define object level functions for use in the translation and elsewhere that correspond exactly to the type analyzing primitives from the **MIL**: $\mathbf{Array}_c$ and $\mathbf{Vararg}_{c_1 \rightarrow c_2}$. This should again be familiar from the methodology developed in chapter 4.

Because of the special treatment of arrays of boxed floating point numbers, the **Array** type needs to analyze the type of values the array contains. This is implemented by translating the **Array** type into a sum switch on the type. If the type being encoded turns out at run time to be a boxed floating point number, then the case statement will return the 64 bit array type: otherwise, it returns the 32 bit version. Note that in the record case, it must explicitly reconstruct the record using the constituent field types.

$$\begin{array}{ll} \mathbf{Array} & : T_{\text{mil}} \rightarrow T_{32} \stackrel{\text{def}}{=} \lambda(\alpha : T_{\text{mil}}). \text{case } \alpha \\ & \quad \mathbf{Record } \beta \Rightarrow \mathbf{Array}_{32}(\times(\beta)) \\ & \quad \mathbf{BFloat} \Rightarrow \mathbf{Array}_{64}(\mathbf{Float}) \\ & \quad \mathbf{Other } t \Rightarrow \mathbf{Array}_{32} t \end{array}$$

Vararg types similarly become dispatches over the argument types to select the appropriate function type: either a flattened function type for small records; otherwise a standard record. Throughout this translation, we assume that only records with fewer than 3 fields get flattened into

$ \alpha $	$\stackrel{\text{def}}{=} \alpha$
$ \lambda(\alpha:\kappa).c $	$\stackrel{\text{def}}{=} \lambda(\alpha: \kappa). c $
$ c_1 c_2 $	$\stackrel{\text{def}}{=} c_1 c_2 $
$ \pi_1 c $	$\stackrel{\text{def}}{=} \pi_1 c $
$ \pi_2 c $	$\stackrel{\text{def}}{=} \pi_2 c $
$ \langle c_1, c_2 \rangle $	$\stackrel{\text{def}}{=} \langle c_1 , c_2 \rangle$
$ \text{Int} $	$\stackrel{\text{def}}{=} \mathbf{Other}(\text{Int})$
$ \text{Sum}_i(c_i, \dots, c_n) $	$\stackrel{\text{def}}{=} \mathbf{Other}(\bigvee [\text{Tag}(0), \dots, \text{Tag}(i-1), \times [\text{Tag}(i), \text{interp } c_i], \dots, \times [\text{Tag}(n), \text{interp } c_n]])$
$ \text{Sum}_i^j(c_i, \dots, c_n) $	$\stackrel{\text{def}}{=} \mathbf{Other}(\text{Tag}(j)) \quad j < i$
$ \text{Sum}_i^j(c_i, \dots, c_j, \dots, c_n) $	$\stackrel{\text{def}}{=} \mathbf{Other}(\times [\text{Tag}(j), \text{interp } c_j]) \quad i \leq j \leq n$
$ \text{Exn} $	$\stackrel{\text{def}}{=} \mathbf{Other}(\exists(\alpha::T_{32}).(\alpha \times (\text{Dyntag } \alpha)))$
$ \text{Dyntag}_c $	$\stackrel{\text{def}}{=} \mathbf{Other}(\text{Dyntag}(\text{interp } c))$
$ \text{Farray} $	$\stackrel{\text{def}}{=} \mathbf{Other}(\text{Array}_{64}\text{Float})$
$ \text{Boxf} $	$\stackrel{\text{def}}{=} \mathbf{BFloat}$
$ \text{Unit} $	$\stackrel{\text{def}}{=} \mathbf{Record}[]$
$ c^1 $	$\stackrel{\text{def}}{=} \mathbf{Record}[\text{interp } c]$
$ c_1 \times c_2 $	$\stackrel{\text{def}}{=} \mathbf{Record}[\text{interp } c_1 , \text{interp } c_2]$
$ c_1 \times \dots \times c_n $	$\stackrel{\text{def}}{=} \mathbf{Record}([\text{interp } c_1 , \dots, \text{interp } c_n])$
$ (c_1, \dots, c_n) \rightarrow c $	$\stackrel{\text{def}}{=} \mathbf{Other}((\text{interp } c_1 , \dots, \text{interp } c_n) \rightarrow \text{interp } c)$
$ \text{Array}_c $	$\stackrel{\text{def}}{=} \mathbf{Other}(\text{Array}(c))$
$ \text{Vararg}_{c_1 \rightarrow c_2} $	$\stackrel{\text{def}}{=} \mathbf{Other}(\text{Vararg}(c_1)(\text{interp } c_2))$
$ \mu(\alpha, \beta).(c_1, c_2) $	$\stackrel{\text{def}}{=} \langle \mathbf{Other}(\text{Rec}[1+1](f)(\text{inj}_1 *)), \mathbf{Other}(\text{Rec}[1+1](f)(\text{inj}_2 *)) \rangle$ where $f = \lambda(\rho:1+1 \rightarrow T_{32}).\lambda(\omega.1+1).$ case ω $\text{inj}_1 - \Rightarrow \text{interp}(c_1 [\mathbf{Other}(\rho(\text{inj}_1 *))/\alpha, \mathbf{Other}(\rho(\text{inj}_2 *))/\beta])$ $\text{inj}_2 - \Rightarrow \text{interp}(c_2 [\mathbf{Other}(\rho(\text{inj}_1 *))/\alpha, \mathbf{Other}(\rho(\text{inj}_2 *))/\beta])$

Figure 5.1: The constructor translation (static encoding)

registers, but in general the number of fields is a machine dependent parameter of the translation.

$$\begin{aligned} \mathbf{Vararg} : T_{\text{mil}} \rightarrow T_{32} \rightarrow T_{32} &\stackrel{\text{def}}{=} \lambda(\alpha:T_{\text{mil}}).\lambda(\beta:T_{32}).\text{case } \alpha \\ &\quad \mathbf{Record } l \Rightarrow \\ &\quad (\text{case } l \\ &\quad \quad [] \Rightarrow () \rightarrow \beta \\ &\quad \quad |[t] \Rightarrow (t) \rightarrow \beta \\ &\quad \quad |[t_1, t_2] \Rightarrow (t_1, t_2) \rightarrow \beta \\ &\quad \quad |_- \Rightarrow (\times(l)) \rightarrow \beta \\ &\quad \quad | \mathbf{BFloat} \Rightarrow (\mathbf{BoxedFloat}) \rightarrow \beta \\ &\quad \quad | \mathbf{Other } t \Rightarrow (t) \rightarrow \beta \end{aligned}$$

Note that the **Vararg** function as defined above is asymmetric in its two arguments: it expects the static encoding of the argument type, but for the return type expects the type itself. This reflects the fact that the result depends only on the argument type and not the result type. An alternative (and equally valid) definition for the **Vararg** primitive can be given that expects static encodings for both arguments (since after all, the argument passed to the **Vararg** function will almost always be the interpretation of a static encoding).

$$\begin{aligned} \mathbf{Vararg} : T_{\text{mil}} \rightarrow T_{\text{mil}} \rightarrow T_{32} &\stackrel{\text{def}}{=} \lambda(\alpha:T_{\text{mil}}).\lambda(\beta:T_{\text{mil}}).\text{case } \alpha \\ &\quad \mathbf{Record } l \Rightarrow \\ &\quad (\text{case } l \\ &\quad \quad [] \Rightarrow () \rightarrow \text{interp}(\beta) \\ &\quad \quad |[t] \Rightarrow (t) \rightarrow \text{interp}(\beta) \\ &\quad \quad |[t_1, t_2] \Rightarrow (t_1, t_2) \rightarrow \text{interp}(\beta) \\ &\quad \quad |_- \Rightarrow (\times(l)) \rightarrow \text{interp}(\beta) \\ &\quad \quad | \mathbf{BFloat} \Rightarrow (\mathbf{BoxedFloat}) \rightarrow \text{interp}(\beta) \\ &\quad \quad | \mathbf{Other } t \Rightarrow (t) \rightarrow \text{interp}(\beta) \end{aligned}$$

From a semantic standpoint, the two definitions are equivalent. In practice however, using the previous definition is likely to provide noticeable performance benefits in type-checking since it provides a definition site for the interpretation of the result type. In the absence of this, a graph reduction implementation is required to avoid repeated reductions. While this is certainly possible (and may be desired in any case), careful design of the defined forms used in the translation can yield substantial performance benefits for very small cost in effort. While I will for the most part defer discussion of implementation issues until part 2 of this thesis, it is nonetheless worth pointing out briefly here that a careful theoretical design can greatly impact the ease and efficiency of the implementation.

Using these definitions, the constructor translation itself (figure 5.1) is almost completely straightforward. For the most part, the translation proceeds compositionally over the constructors in the obvious manner, applying the static encoding constructors at the leaves of the syntax tree. The case for sum types is more interesting, since it makes the representation of sums as tagged unions explicit. Note also that the known sum type from the **MIL** disappears. Known sums are no longer necessary because the tagging has been made explicit, and hence the de-structuring of sum values can be done via record selection, instead of via the **proj** primitive.

The most syntactically complex rule of the translation defines the compilation of **MIL** style recursive types into **LIL** style parametric recursive types. The essential idea behind this translation

is to replace uses of multiple mutually recursive types with the fix point of a single parameterized constructor. The first parameter of the constructor (as usual) is the recursive variable. The second parameter serves as a selector parameter which indicates which field of the original mutually recursive type is desired. The translation therefore proceeds by translating the bodies of the recursive types, replacing uses of the multiple mutually recursive variables with partial applications of the new recursive variable to appropriate sum injections indicating the selection of the appropriate arm. The body of the function of which a fixpoint is to be taken uses its parameter to select the appropriate arm of the recursive type to return. The final tuple of constructors is created by instantiating the fixpoint once for each of the different arms of the original constructor.

5.2.3 Translation of typing contexts

The kind translation extends in a natural way to define a translation on **MIL** constructor level typing contexts.

$$\begin{aligned} |\bullet| &\stackrel{\text{def}}{=} \bullet \\ |\Delta, \alpha:\kappa| &\stackrel{\text{def}}{=} |\Delta|, \alpha:|\kappa| \end{aligned}$$

5.2.4 Proofs of soundness for the constructor and kind translation

Lemma 4 (Well-formedness of T_{mil})

For all well formed Δ such that the bound variables T_{mil} are not in Δ ,

$$\Delta \vdash T_{\text{mil}} \text{ ok}$$

Proof: By construction. Note that for any given Δ , the definitions may be alpha-varied such that the variable condition is satisfied. ■

Lemma 5 (Soundness of kind translation (1))

$$\bullet \vdash |\kappa| \text{ ok}$$

Proof: By induction over the structure of κ , and lemma 4. ■

Corollary 1

If $\vdash \Delta \text{ ok}$ then $\Delta \vdash |\kappa| \text{ ok}$.

Proof: By lemma 5 and weakening. ■

Lemma 6

If $\alpha \notin \Delta'$ then $\alpha \notin |\Delta|$.

Proof: Note that Δ and $|\Delta|$ have the same domain by construction. ■

Lemma 7 (Soundness of the context translation)

If $\vdash \Delta \text{ ok}$ then $\vdash |\Delta| \text{ ok}$.

Proof: By induction over the structure of Δ .

1. If $\Delta = \bullet$ then $|\Delta| = \bullet$ and $\vdash \bullet \text{ ok}$

2. If $\Delta = \Delta', \alpha : \kappa$ then:

By assumption:

$\vdash \Delta' \text{ ok}$, $\alpha \notin \Delta'$, and $\vdash \kappa \text{ ok}$

By induction:

$\vdash |\Delta'| \text{ ok}$

By corollary 1:

$|\Delta'| \vdash |\kappa| \text{ ok}$

By lemma 6:

$\alpha \notin |\Delta'|$

By construction:

$\vdash |\Delta'|, \alpha : |\kappa| \text{ ok}$

By definition:

$\vdash |\Delta| \text{ ok}$

■

The previous result is used to prove a slightly stronger result about the soundness of the kind translation.

Theorem 1 (Soundness of kind translation (2))

If $\vdash \Delta \text{ ok}$ then $|\Delta| \vdash |\kappa| \text{ ok}$.

Proof: By lemma 7, $\vdash |\Delta| \text{ ok}$, so by corollary 1, $|\Delta| \vdash |\kappa| \text{ ok}$.

■

To assist in the proof of soundness of the constructor translation, I formalize the typing properties of the definitions in the following lemma:

Lemma 8 (Well-formedness of definitions)

For all well formed Δ such that the bound variables of the respective defined forms are not in Δ ,

$$\begin{aligned} \Delta &\vdash \text{interp} : T_{\text{mil}} \rightarrow T_{32} \\ \Delta &\vdash \text{Other} : T_{32} \rightarrow T_{\text{mil}} \\ \Delta &\vdash \text{BFloat} : T_{\text{mil}} \\ \Delta &\vdash \text{Record} : T_{32\text{list}} \rightarrow T_{\text{mil}} \\ \Delta &\vdash \text{Array} : T_{\text{mil}} \rightarrow T_{32} \\ \Delta &\vdash \text{Vararg} : T_{\text{mil}} \rightarrow T_{32} \rightarrow T_{32} \end{aligned}$$

Proof: By construction. Note that for any given Δ , the definitions may be alpha-varied such that the variable condition is satisfied.

■

Using this lemma, the soundness of the constructor translation follows straightforwardly.

Theorem 2 (Soundness of the constructor translation)

If $\Delta \vdash c : \kappa$ then $|\Delta| \vdash |c| : |\kappa|$.

Proof: By induction on the structure of typing derivations.

1. Suppose $\Delta, \alpha:\kappa \vdash \alpha:\kappa$. Then by induction $\vdash |\Delta, \alpha:\kappa|$ **ok**. By definition $|\Delta, \alpha:\kappa| = |\Delta|, \alpha:\kappa$, and $|\alpha| = \alpha$, and so by construction $|\Delta, \alpha:\kappa| \vdash \alpha:\kappa$.
2. Suppose $\Delta \vdash \lambda(\alpha::\kappa_1).c:\kappa_1 \rightarrow \kappa_2$. By theorem 1, $\Delta \vdash |\kappa_1|$ **ok**, and by induction, $|\Delta, \alpha:\kappa_1| \vdash |c|:\kappa_2$. But $|\Delta, \alpha:\kappa_1| = |\Delta|, \alpha:\kappa_1|$, and by lemma 6, $\alpha \notin |\Delta|$. So by construction, $|\Delta| \vdash |\lambda(\alpha::\kappa_1).c|:\kappa_1 \rightarrow \kappa_2$.
3. Suppose $\Delta \vdash c_1 c_2:\kappa_2$. By induction, $|\Delta| \vdash |c_1|:\kappa_1 \rightarrow \kappa_2$ and $|\Delta| \vdash |c_2|:\kappa_1$. Therefore by construction $|\Delta| \vdash |c_1 c_2|:\kappa_2$.
4. Suppose $\Delta \vdash \pi_1 c:\kappa_1$. By induction, $|\Delta| \vdash |c|:\kappa_1 \times \kappa_2$. So by construction, $|\Delta| \vdash |\pi_1 c|:\kappa_1$.
5. Suppose $\Delta \vdash \pi_2 c:\kappa_2$. By a similar argument to the previous case, $|\Delta| \vdash |\pi_2 c|:\kappa_2$.
6. Suppose $\Delta \vdash \langle c_1, c_2 \rangle:\kappa_1 \times \kappa_2$. By induction, $|\Delta| \vdash |c_1|:\kappa_1$ and $|\Delta| \vdash |c_2|:\kappa_2$. So by construction, $|\Delta| \vdash |\langle c_1, c_2 \rangle|:\kappa_1 \times \kappa_2$.
7. Suppose $\Delta \vdash \mathbf{Int}:\mathbf{T}_{32}$. Note that $|\mathbf{T}_{32}| = \mathbf{T}_{\text{mil}}$, so it suffices to show that $|\Delta| \vdash |\mathbf{Int}|:\mathbf{T}_{\text{mil}}$. By lemma 8, $|\Delta| \vdash \mathbf{Other}:\mathbf{T}_{32} \rightarrow \mathbf{T}_{\text{mil}}$, and by the **Int** axiom, $|\Delta| \vdash \mathbf{Int}:\mathbf{T}_{32}$. So by the application rule, $|\Delta| \vdash \mathbf{Other}(\mathbf{Int}):\mathbf{T}_{\text{mil}}$.
8. The other primitive type constructors follow similarly as with **Int**.
9. Suppose $\Delta \vdash \mu(\alpha, \beta).(c_1, c_2):\mathbf{T}_{32} \times \mathbf{T}_{32}$. It suffices to show that

$$\begin{aligned}
& |\Delta| \vdash \langle \mathbf{Other}(\mathbf{rec}_{1+1}\langle f, \mathbf{inj}_1 * \rangle), \mathbf{Other}(\mathbf{rec}_{1+1}\langle f, \mathbf{inj}_2 * \rangle) \rangle:\mathbf{T}_{\text{mil}} \times \mathbf{T}_{\text{mil}} \\
& \text{where } f = \lambda(\rho:1+1 \rightarrow \mathbf{T}_{32}).\lambda(\omega.1+1). \\
& \quad \text{case } \omega \\
& \quad \mathbf{inj}_1 - \Rightarrow \text{interp}(|c_1|[\mathbf{Other}(\rho(\mathbf{inj}_1 *))/\alpha, \mathbf{Other}(\rho(\mathbf{inj}_2 *))/\beta]) \\
& \quad \mathbf{inj}_2 - \Rightarrow \text{interp}(|c_2|[\mathbf{Other}(\rho(\mathbf{inj}_1 *))/\alpha, \mathbf{Other}(\rho(\mathbf{inj}_2 *))/\beta])
\end{aligned}$$

But note that it suffices to show that $|\Delta| \vdash f:(1+1 \rightarrow \mathbf{T}_{32}) \rightarrow (1+1) \rightarrow \mathbf{T}_{32}$, since the desired derivation can then be produced by applying the **rec** rule to show the well-formedness of the new recursive types; the application rule and lemma 8 to show the well-formedness of the applications of **Other**, and the pairing rule to show the well-formedness of the pair.

By induction, $|\Delta, \alpha:\mathbf{T}_{32}, \beta:\mathbf{T}_{32}| \vdash |c_1|:\mathbf{T}_{\text{mil}}$ and by weakening, the freshness assumption, and the definition of the context translation $|\Delta|, \alpha:\mathbf{T}_{\text{mil}}, \beta:\mathbf{T}_{\text{mil}}, \rho:(1+1 \rightarrow \mathbf{T}_{32}) \vdash |c_1|:\mathbf{T}_{\text{mil}}$. By the substitution lemma for the **LIL** (lemma 1):

$$|\Delta|, \rho:(1+1 \rightarrow \mathbf{T}_{32}) \vdash |c_1|[\mathbf{Other}(\rho(\mathbf{inj}_1 *))/\alpha, \mathbf{Other}(\rho(\mathbf{inj}_2 *))/\beta]:\mathbf{T}_{\text{mil}}$$

By the application typing rule therefore:

$$|\Delta|, \rho:(1+1 \rightarrow \mathbf{T}_{32}) \vdash \text{interp}(|c_1|[\mathbf{Other}(\rho(\mathbf{inj}_1 *))/\alpha, \mathbf{Other}(\rho(\mathbf{inj}_2 *))/\beta]):\mathbf{T}_{32}$$

A similar argument holds for c_2 . Therefore, by applying the typing rule for case and for repeated lambda abstraction (again using freshness and weakening for ω), we construct a well-formedness derivation for f .

■

5.2.5 Additional properties of the constructor translation

Commutation with substitution

It is an additional property of the constructor translation that it commutes with substitution: i.e. that it is compositional. This property is necessary for the proofs of a number of subsequent theorems, and so I give a proof of it here.

Lemma 9 (The constructor translation commutes with substitution)

$$|c|[[c'/\alpha]] = |c[c'/\alpha]].$$

Proof: By induction on c .

1. If c is a constant, then $\alpha \notin fvc, fv|c|$, so $|c|[[c'/\alpha]] = |c| = |c[c'/\alpha]]$.
2. If c is a variable α' :
 - (a) If $\alpha \neq \alpha'$ then as in the previous case.
 - (b) If $\alpha = \alpha'$ then $|\alpha|[[c'/\alpha]] = \alpha[[c'/\alpha]] = |c'| = |\alpha[c'/\alpha]]$.
3. If $c = \lambda(\beta:\kappa).c_2$, then by definition

$$|\lambda(\beta:\kappa).c_2|[[c'/\alpha]] = (\lambda(\beta:|\kappa|).|c_2|)[[c'/\alpha]] = \lambda(\beta:|\kappa|).(|c_2|[[c'/\alpha]])$$

By induction and the definition of substitution:

$$\lambda(\beta:|\kappa|).(|c_2|[[c'/\alpha]]) = \lambda(\beta:|\kappa|).(|c_2[c'/\alpha]|) = |\lambda(\beta:\kappa).(c_2[c'/\alpha])|$$

Note, we rely on alpha-variance to ensure non-capture.

4. If $c = c_1 c_2, \pi_i c, \langle c_1, c_2 \rangle$, the proof proceeds similarly.

■

The SE translation respects equivalence

Another important property of the **SE** translation is that equivalent **MIL** constructors translate to equivalent **LIL** constructors. In addition to being important for subsequent proofs, this lemma is an important “sanity-check” on the translation.

Lemma 10 (The constructor translation respects equivalence)

If $\Delta \vdash c_1 \equiv c_2 : \kappa$, then $|\Delta| \vdash |c_1| \equiv |c_2| : |\kappa|$.

Proof: (By induction on equivalence derivations) All of the structural and type constructor equivalence rules are unchanged from the **MIL** to the **LIL**, and the proof follows straightforwardly by induction. I give the reflexivity rule, the pair rule and the beta rule as examples. All of the primitive types follow directly by reflexivity or by the structural application rule.

1. Suppose $\Delta \vdash c \equiv c : \kappa$ by reflexivity.

By assumption:

$$\Delta \vdash c : \kappa$$

By theorem 2:

$$|\Delta| \vdash |c| : |\kappa|$$

By reflexivity:

$$|\Delta| \vdash |c| \equiv |c| : |\kappa|$$

2. Suppose $\Delta \vdash (\lambda(\alpha:\kappa_1).c_1)(c_2) \equiv c_1[c_2/\alpha] : \kappa_2$.

By assumption:

$$\vdash \kappa_1 \text{ ok}$$

$$\Delta, \alpha:\kappa_1 \vdash c_1 : \kappa_2$$

$$\Delta \vdash c_2 : \kappa_1$$

By theorems 1, and 2:

$$|\Delta| \vdash |\kappa_1| \text{ ok}$$

$$|\Delta, \alpha:\kappa_1| \vdash |c_1| : |\kappa_2|$$

$$|\Delta| \vdash |c_2| : |\kappa_1|$$

By the λ beta rule and the definition of the translations:

$$|\Delta| \vdash |(\lambda(\alpha:\kappa_1).c_1)(c_2)| \equiv |c_1| [|c_2|/\alpha] : |\kappa_2|$$

Finally, by lemma 9:

$$|\Delta| \vdash |(\lambda(\alpha:\kappa_1).c_1)(c_2)| \equiv |c_1[c_2/\alpha]| : |\kappa_2|$$

■

5.3 Dynamic encoding of constructors

The two translations given above, $|\kappa|$ and $|c|$, define the kind of the static encoding of a constructor and the encoding itself, respectively. The second element of the mapping from **MIL** to **LIL** is giving dynamic encodings for constructors in order to permit the move from a type-passing interpretation to a type erasure interpretation.

A constructor's dynamic encoding is a **LIL** term which also encodes the same information about the original **MIL** constructor, but at the term level instead of at the constructor level. Dynamic encodings are used to implement type dispatch at runtime. This section defines the dynamic encoding translation.

5.3.1 Dynamic encoding types

The dynamic encoding translation of a constructor c is notated as $\lfloor c \rfloor$. (The notation is intended to be suggestive of the fact that the translation moves “down” a level, from constructors to terms.) The type of the dynamic encoding of a constructor is driven by its kind: constructors of pair kind get represented by terms of pair type, etc. However, the type of the dynamic encoding of a constructor also depends on its static encoding: this is what captures the connection between static and dynamic encodings. For a **MIL** kind κ classifying a constructor c , I notate the dynamic encoding type for the constructor as $\lfloor c:\kappa \rfloor$.

$$\begin{aligned} \lfloor c:T_{32} \rfloor & \stackrel{\text{def}}{=} R(|c|) \\ \lfloor c:\kappa_1 \rightarrow \kappa_2 \rfloor & \stackrel{\text{def}}{=} \forall[\alpha:\kappa_1] (\lfloor \alpha:\kappa_1 \rfloor \rightarrow \lfloor c\alpha:\kappa_2 \rfloor) \\ \lfloor c:\kappa_1 \times \kappa_2 \rfloor & \stackrel{\text{def}}{=} \lfloor \pi_1 c:\kappa_1 \rfloor \times \lfloor \pi_2 c:\kappa_2 \rfloor \end{aligned}$$

The **DET** translation is defined in terms of a type level function R of type $T_{\text{mil}} \rightarrow T_{32}$. This function is defined shortly, but for the sake of understanding the translation it is easier to view this abstractly as the type of representations of constructors of kind T_{32} .

For constructors of pair kind, the **DET** is straightforward: a pair type with fields constructed by projecting out the two halves of the pair and applying the **DET** translation to each of them at the sub-component kind.

Constructor functions are slightly more interesting. The essential idea is that constructor functions (functions from types to types) will become term functions (functions from terms to terms). Accordingly, the **DET** of a constructor of arrow kind is a term level function type. However, note that this requires us to construct the **DET** of the argument type and the result type of the function, which in turn requires us to have the static encoding of the argument to the type function (which of course is not yet available). The solution to this dilemma is to observe that the **DE** function in question must be *polymorphic* over all possible static encodings: the **DET** of the argument and result can then be given in terms of the eventual encoded type passed to the function. This point is essential to the methodological goal of maintaining the connection between static encodings and dynamic encodings.

So far, I have avoided discussing the particulars of the actual encodings of types (that is, of constructors of kind T_{32}). Here what is required is something of the nature of the representation types of λ_r [CWM98]: the type of representation of a specific type. Unlike in λ_r , these types are not primitive here. Instead, they are programmed directly in the **LIL** in the following fashion.

$$\begin{aligned}
 R : T_{\text{mil}} \rightarrow T_{32} &\stackrel{\text{def}}{=} \lambda(\alpha : T_{\text{mil}}). \\
 &\text{case } \alpha \\
 &\quad (\mathbf{Record} \ \beta \Rightarrow \\
 &\quad \quad (\text{case } \beta \ ([] \Rightarrow \mathbf{Unit} \mid _ \Rightarrow \mathbf{Void}) \\
 &\quad \quad + \text{case } \beta \ ([_,] \Rightarrow \mathbf{Unit} \mid _ \Rightarrow \mathbf{Void}) \\
 &\quad \quad + \text{case } \beta \ ([_, _] \Rightarrow \mathbf{Unit} \mid _ \Rightarrow \mathbf{Void}) \\
 &\quad \quad + \text{case } \beta \ (_ \Rightarrow \mathbf{Unit} \mid _ \Rightarrow \mathbf{Void}) \\
 &\quad \quad \mid _ \Rightarrow \mathbf{Void}) \\
 &\quad + \text{case } \alpha \ (\mathbf{BFloat} \Rightarrow \mathbf{Unit} \mid _ \Rightarrow \mathbf{Void}) \\
 &\quad + \text{case } \alpha \ (\mathbf{Other} \ _ \Rightarrow \mathbf{Unit} \mid _ \Rightarrow \mathbf{Void})
 \end{aligned}$$

This definition is somewhat subtle. At the top level, the representations of types are always sums, indicating whether the type being represented is a record, boxed float, or an undistinguished type. The value carried in each branch of the sum serves as a witness to the identity of the original type. So for example in the boxed float case the carried value will have type $\text{case } \alpha \ (\mathbf{BFloat} \ _ \Rightarrow \mathbf{Unit} \mid _ \Rightarrow \mathbf{Void})$. Since the **Void** type is uninhabited, having a value of this type means that α can only be the static representation of **Boxed(Float)**. This information can be reflected back into the type system via the special **vcase** construct, whereby code can use this witness to refine α .

5.3.2 Notational issues

Unfortunately, the named-form syntactic restrictions imposed on **LIL** terms add a level of inessential complexity to the dynamic encoding translations. Constructors in the **MIL** are not in named-form: in fact, in the absence of singleton kinds or some other definitional mechanism there is in general no equivalent named-form for an arbitrary constructor. This means that translating **MIL** constructors into **LIL** terms requires the dynamic encoding essential to the translation to occur simultaneously

Pairs of expressions $\langle e_1, e_2 \rangle$	
$\langle sv_1, sv_2 \rangle$	$\stackrel{\text{def}}{=} \text{let } x = \langle sv_1, sv_2 \rangle \text{ in } x$
$\langle sv_1, \text{let } x = i \text{ in } e_2 \rangle$	$\stackrel{\text{def}}{=} \text{let } x = i \text{ in } \langle sv_1, e_2 \rangle \quad (x \notin \text{fv}(sv_1))$
$\langle \text{let } x = i \text{ in } e_1, e_2 \rangle$	$\stackrel{\text{def}}{=} \text{let } x = i \text{ in } \langle e_1, e_2 \rangle \quad (x \notin \text{fv}(e_2))$
Projection from expressions $\text{select}^i e$	
$\text{select}^i sv$	$\stackrel{\text{def}}{=} \text{let } x = \text{select}^i sv \text{ in } x$
$\text{select}^i(\text{let } x = i \text{ in } e)$	$\stackrel{\text{def}}{=} \text{let } x = i \text{ in } (\text{select}^i e)$
Sum injection of expressions $\text{inj}_i^c e$	
$\text{inj}_i^c sv$	$\stackrel{\text{def}}{=} \text{let } y = \langle \text{tag}_i, sv \rangle \text{ in let } x = \text{inj}_i^c y \text{ in } x$
$\text{inj}_i^c(\text{let } x = i \text{ in } e)$	$\stackrel{\text{def}}{=} \text{let } x = i \text{ in } (\text{inj}_i^c e)$
Application of expressions $e_1[c]e_2$	
$sv_1[c]sv_2$	$\stackrel{\text{def}}{=} \text{let } x = sv_1[c]sv_2 \text{ in } x$
$sv_1[c](\text{let } x = i \text{ in } e_2)$	$\stackrel{\text{def}}{=} \text{let } x = i \text{ in } (sv_1[c]e_2) \quad (x \notin \text{fv}(sv_1))$
$(\text{let } x = i \text{ in } e_1)[c]e_2$	$\stackrel{\text{def}}{=} \text{let } x = i \text{ in } (e_1[c]e_2) \quad (x \notin \text{fv}(e_2))$

Figure 5.2: Derived non-named-form expressions

$\frac{\Psi; \Delta; \Gamma \vdash e_1 : \tau_1 \text{ exp} \quad \Psi; \Delta; \Gamma \vdash e_2 : \tau_2 \text{ exp}}{\Psi; \Delta; \Gamma \vdash \langle e_1, e_2 \rangle : \tau_1 \times \tau_2 \text{ dexp}}$		$\Psi; \Delta; \Gamma \vdash e : \tau_0 \times \tau_1 \text{ exp} \quad (i \in 0, 1)$
$\frac{\Delta \vdash c \equiv +[\tau_1, \dots, \times[\text{Tag}(i), \tau_i], \dots, \tau_n] : T_{32} \quad \Psi; \Delta; \Gamma \vdash e : \tau_i \text{ exp} \quad (i \in 1 \dots n)}{\Psi; \Delta; \Gamma \vdash \text{inj}_i^c e : c \text{ dexp}}$		$\frac{\Delta \vdash c : \kappa \quad \Psi; \Delta; \Gamma \vdash e_1 : \forall(\alpha : \kappa). \tau_1 \rightarrow \tau_2 \text{ exp} \quad \Psi; \Delta; \Gamma \vdash e_2 : \tau_1[c/\alpha] \text{ exp}}{\Psi; \Delta; \Gamma \vdash e_1[c]e_2 : \tau_2[c/\alpha] \text{ dexp}}$

Figure 5.3: Typing rules for derived (non-named-form) expressions

$\downarrow \alpha \downarrow$	$\stackrel{\text{def}}{=} x_\alpha$
$\downarrow \lambda(\alpha:\kappa).c \downarrow$	$\stackrel{\text{def}}{=} \lambda[\alpha:\kappa](x_\alpha:\downarrow \alpha:\kappa \downarrow). \downarrow c \downarrow$
$\downarrow c_1 c_2 \downarrow$	$\stackrel{\text{def}}{=} \downarrow c_1 \downarrow \llbracket \downarrow c_2 \downarrow \rrbracket (\downarrow c_2 \downarrow)$
$\downarrow \pi_i c \downarrow$	$\stackrel{\text{def}}{=} \text{select}_i \downarrow c \downarrow$
$\downarrow \langle c_1, c_2 \rangle \downarrow$	$\stackrel{\text{def}}{=} \langle \downarrow c_1 \downarrow, \downarrow c_2 \downarrow \rangle$
$\downarrow \text{Unit} \downarrow$	$\stackrel{\text{def}}{=} \text{inj}_1^{R(\text{Unit})}(\text{inj}_1 \langle \rangle)$
$\downarrow \times [c] \downarrow$	$\stackrel{\text{def}}{=} \text{inj}_1^{R(\times [c])}(\text{inj}_2(\text{inj}_1 \langle \rangle))$
$\downarrow c_1 \times c_2 \downarrow$	$\stackrel{\text{def}}{=} \text{inj}_1^{R(c_1 \times c_2)}(\text{inj}_2(\text{inj}_2(\text{inj}_1 \langle \rangle)))$
$\downarrow \times [c_1, c_2, \dots, c_n] \downarrow$	$\stackrel{\text{def}}{=} \text{inj}_1^{R(\times [c_1, c_2, \dots, c_n])}(\text{inj}_2(\text{inj}_2(\text{inj}_2 \langle \rangle)))$
$\downarrow \text{Boxf} \downarrow$	$\stackrel{\text{def}}{=} \text{inj}_2^{R(\text{Boxf})} \langle \rangle$
$\downarrow c \text{ as } \mu(\alpha, \beta).(c_1, c_2) \downarrow$	$\stackrel{\text{def}}{=} \langle \text{inj}_3^{R(\pi_1 c)} \langle \rangle, \text{inj}_3^{R(\pi_2 c)} \langle \rangle \rangle$
$\downarrow c \text{ otherwise} \downarrow$	$\stackrel{\text{def}}{=} \text{inj}_3^{R(c)} \langle \rangle$

Figure 5.4: The dynamic encoding translation

with a mostly un-interesting “lifting” process which pulls bindings out of expressions and returns them to named form.

In order to isolate the core issues of the translation from these syntactic issues, it is useful to define derived syntactic forms (figure 5.2) that allow the translation to make use of apparently arbitrary expressions without concerning itself with named form. Derived typing rules (figure 5.3) for the extended syntax with an accompanying proof of soundness (section 5.3.5) permit the free use of the extended syntax in the translation.

Only a small number of derived expression forms are needed for the purposes of the dynamic encoding translation: additional definitions are possible. All variables are assumed to be fresh in the translation: this guarantees that variables will not conflict as bindings are “lifted”. This could also be dealt with by explicitly alpha-varying terms as necessary.

5.3.3 Dynamic encodings

The term-level encoding translation of constructors is given in figure 5.4 and is for the most part fairly intuitive. Note that the translation takes advantage of the defined forms from section 5.3.2 for presentational clarity: without this, it is necessary for the translation to handle explicitly the lifting of bindings. In point of fact, this sort of issue arises in many parts of the implementation of the translation from **MIL** to **LIL** and is handled using a monadic structure similar to the defined syntax used here.

For the purposes of the translation I assume an unspecified anti-symmetric injection from type variables into a set of term variables disjoint from those produced elsewhere in the translation. I refer to the term variables produced by this injection informally as being “indexed” by the type variable. Note that the translation does not rely on the ability to recover the original type variable from the indexed term variable.

Type functions become term functions which take both the static and dynamic representations

of the argument and return the dynamic representation of the result. Applications are modified in the corresponding fashion. Type level pairs and projections map to term level pairs and projections. For constructors of kind T_{32} , it is necessary to construct the type representations by injecting the appropriate witnesses into the sum type described above.

5.3.4 Dynamic encoding of type contexts

Since the **DE** translation is defined on open terms, the question naturally arises of what the appropriate notion of typing context for dynamic encodings should be. Just as the **DE** translation maps constructors to types, the corresponding context translation maps constructor contexts to term contexts using the **DET** translation defined above. Again, I use the down arrow syntax ($\downarrow \cdot$) to suggest the movement “down” a level in the syntactic hierarchy.

$$\begin{aligned} \downarrow \bullet &\stackrel{\text{def}}{=} \bullet \\ \downarrow \Delta, \alpha : \kappa &\stackrel{\text{def}}{=} \downarrow \Delta, x_\alpha : \downarrow \alpha : \kappa \end{aligned}$$

Note that the types produced by the **DET** translation will contain references to the type variables from the original context: that is, the context produced by the translation will have free type variables. These variables are described by the context produced by the static encoding of the original context: that is, $\downarrow \Delta \vdash \downarrow \Delta \text{ ok}$.

5.3.5 Proofs for the dynamic encoding translations

Proof of soundness of the DET translation

I begin by showing the well-formedness of the definition of representation types.

Lemma 11 (Well-formedness of R)

For all well formed Δ such that the bound variables R are not in Δ , $\Delta \vdash R : T_{\text{mil}} \rightarrow T_{32}$.

Proof: By construction. Note that for any given Δ , the definition may be alpha-varied such that the variable condition is satisfied. ■

Using this and the soundness theorems for the **SE** and **SEK** translations, I show the soundness of the **DET** translation.

Lemma 12 (Soundness of the dynamic encoding type translation)

If $\Delta \vdash c : \kappa$ then $\downarrow \Delta \vdash \downarrow c : \kappa : T_{32}$.

Proof: By induction on kinds κ .

1. Suppose $\kappa = T_{32}$. Then by definition, $\downarrow c : \kappa = R(\downarrow c)$. By theorem 2, $\downarrow \Delta \vdash \downarrow c : T_{\text{mil}}$, and so by lemma 11 and the application typing rule, $\downarrow \Delta \vdash R(\downarrow c) : T_{32}$.
2. Suppose $\kappa = \kappa_1 \rightarrow \kappa_2$. Then by definition, $\downarrow c : \kappa = \forall [\alpha : \kappa_1] (\downarrow \alpha : \kappa_1 \vdash \downarrow c \alpha : \kappa_2)$. It suffices to show that this is well-formed via the formation rule for universals. By theorem 1, $\downarrow \Delta \vdash \downarrow \kappa_1 \text{ ok}$. By the variable rule, $\Delta, \alpha : \kappa_1 \vdash \alpha : \kappa_1$, and so by induction, $\downarrow \Delta, \alpha : \kappa_1 \vdash \downarrow \alpha : \kappa_1 : T_{32}$. Again using the variable rule along with the assumption, we obtain that $\Delta, \alpha : \kappa_1 \vdash c \alpha : \kappa_2$ by the application rule, and so by induction, $\downarrow \Delta, \alpha : \kappa_1 \vdash \downarrow c \alpha : \kappa_2 : T_{32}$.

We therefore may invoke the formation rule for universals to produce a derivation of the desired result.

3. Suppose $\kappa = \kappa_1 \rightarrow \kappa_2$. Then by definition, $\downarrow c:\kappa = \downarrow \pi_1 c:\kappa_1 \times \downarrow \pi_2 c:\kappa_2$. It suffices to show that this is well-formed via the formation rule for pair types by showing that the two component types are well-formed. This follows immediately by induction. ■

As with the **SE** translation, the fact that the **DET** translation commutes with substitution is useful for subsequent proofs.

Lemma 13 (The DET translation commutes with substitution)

$\downarrow c:\kappa \downarrow \llbracket c'/\alpha \rrbracket = \downarrow c[c'/\alpha]:\kappa$.

Proof: By induction on kinds κ .

1. If $\kappa = T_{32}$, then $\downarrow c:\kappa = R(|c|)$. Since R is closed, $R(|c|)\llbracket c'/\alpha \rrbracket = R(|c|\llbracket c'/\alpha \rrbracket)$. By lemma 9, $|c|\llbracket c'/\alpha \rrbracket = |c[c'/\alpha]|$. Finally, by definition, $R(|c[c'/\alpha]|) = \downarrow c[c'/\alpha]:T_{32}$.
2. If $\kappa = \kappa_1 \rightarrow \kappa_2$ then by definition

$$\begin{aligned} \downarrow c:\kappa_1 \rightarrow \kappa_2 \downarrow \llbracket c'/\alpha \rrbracket &\stackrel{\text{def}}{=} (\forall [\beta:\kappa_1] (\downarrow \beta:\kappa_1 \downarrow) \rightarrow (\downarrow c\beta:\kappa_2 \downarrow)) \llbracket c'/\alpha \rrbracket \\ &\stackrel{\text{def}}{=} \forall [\beta:\kappa_1] (\downarrow \beta:\kappa_1 \downarrow) \rightarrow (\downarrow c\beta:\kappa_2 \downarrow \llbracket c'/\alpha \rrbracket) \end{aligned}$$

Note that $\alpha \neq \beta$ by the assumption of freshness in the translation, and hence

$$\downarrow \beta:\kappa_1 \downarrow \llbracket c'/\alpha \rrbracket = \downarrow \beta:\kappa_1 \downarrow$$

By induction, $\downarrow c\beta:\kappa_2 \downarrow \llbracket c'/\alpha \rrbracket = \downarrow (c[c'/\alpha])\beta:\kappa_2 \downarrow$, again using the fact that $\alpha \neq \beta$. Finally, we observe that

$$\forall [\beta:\kappa_1] (\downarrow \beta:\kappa_1 \downarrow) \rightarrow (\downarrow (c[c'/\alpha])\beta:\kappa_2 \downarrow) = \downarrow c[c'/\alpha]:\kappa_1 \rightarrow \kappa_2 \downarrow$$

3. If $\kappa = \kappa_1 \times \kappa_2$ then by definition,

$$\begin{aligned} \downarrow c:\kappa_1 \times \kappa_2 \downarrow \llbracket c'/\alpha \rrbracket &\stackrel{\text{def}}{=} (\downarrow \pi_1 c:\kappa_1 \downarrow \times \downarrow \pi_2 c:\kappa_2 \downarrow) \llbracket c'/\alpha \rrbracket \\ &\stackrel{\text{def}}{=} (\downarrow \pi_1 c:\kappa_1 \downarrow \llbracket c'/\alpha \rrbracket) \times (\downarrow \pi_2 c:\kappa_2 \downarrow \llbracket c'/\alpha \rrbracket) \end{aligned}$$

By induction, $(\downarrow \pi_i c:\kappa_i \downarrow \llbracket c'/\alpha \rrbracket) = \downarrow \pi_i c[c'/\alpha]:\kappa_i \downarrow$, so

$$\begin{aligned} (\downarrow \pi_1 c:\kappa_1 \downarrow \llbracket c'/\alpha \rrbracket) \times (\downarrow \pi_2 c:\kappa_2 \downarrow \llbracket c'/\alpha \rrbracket) &\stackrel{\text{def}}{=} \downarrow \pi_1 c[c'/\alpha]:\kappa_1 \downarrow \times \downarrow \pi_2 c[c'/\alpha]:\kappa_2 \downarrow \\ &\stackrel{\text{def}}{=} \downarrow c[c'/\alpha]:\kappa_1 \times \kappa_2 \downarrow \end{aligned}$$

■

Proof of soundness of the DE typing context translation

Lemma 14

If $\alpha \notin \Delta$ then $x_\alpha \notin \lfloor \Delta \rfloor$.

Proof: By assumption, the injection from type variables to term variables is anti-symmetric with respect to variable equality. Therefore, by anti-symmetry, if $\beta \neq \alpha$ then $x_\beta \neq x_\alpha$. Since the set of variables in the domain of Δ does not contain α , the image of the set of under the injection does not contain x_α . ■

Lemma 15 (Soundness of the context translation)

If $\vdash \Delta$ **ok** then $\lfloor \Delta \rfloor \vdash \lfloor \Delta \rfloor$ **ok**.

Proof: By induction over the structure of Δ .

1. If $\Delta = \bullet$ then $\lfloor \Delta \rfloor = \bullet$ and $\bullet \vdash \bullet$ **ok**
2. If $\Delta = \Delta', \alpha : \kappa$ then by assumption, $\vdash \Delta'$ **ok**, $\alpha \notin \Delta'$, and $\vdash \kappa$ **ok**. By the variable rule, $\Delta', \alpha : \kappa \vdash \alpha : \kappa$, and so by lemma 12, $\lfloor \Delta \rfloor, \alpha : \lfloor \kappa \rfloor \vdash \lfloor \alpha : \kappa \rfloor : T_{32}$. By induction, $\lfloor \Delta' \rfloor \vdash \lfloor \Delta' \rfloor$ **ok**, and by weakening and lemma 6 $\lfloor \Delta' \rfloor, \alpha : \lfloor \kappa \rfloor \vdash \lfloor \Delta' \rfloor$ **ok**.

Finally by the formation rule for contexts $\lfloor \Delta' \rfloor, \alpha : \lfloor \kappa \rfloor \vdash \lfloor \Delta' \rfloor, x_\alpha : \lfloor \alpha : \kappa \rfloor$ **ok**. Note that the side condition follows from lemma 14. ■

Proof of soundness of the DE translation

Lemma 16 (Inversion of derived expression derivations)

For any derivation D of $\Psi; \Delta; \Gamma \vdash e : \tau$ **dexp** the last rule of D is uniquely determined by e .

Proof: By inspection. Note that for any choice of e , exactly one rule applies. ■

Since the **DE** translation relies on the extended syntax defined in section 5.3.2, it is first necessary to show that well-typedness as a derived form corresponds with well-typedness of expressions.

Lemma 17 (Derived expression rules)

If $\Psi; \Delta; \Gamma \vdash e : \tau$ **dexp** then $\Psi; \Delta; \Gamma \vdash e : \tau$ **exp**.

Proof: By induction on the definition of the derived forms. The proof is straightforward: I give here one example case in detail.

1. $\Psi; \Delta; \Gamma \vdash \langle e_1, e_2 \rangle : \tau_1 \times \tau_2$ **dexp** By inversion (lemma 16, $\Psi; \Delta; \Gamma \vdash e_1 : \tau_1$ **exp** and $\Psi; \Delta; \Gamma \vdash e_2 : \tau_2$ **exp**). There are three possible cases in the definition of the derived form.
 - (a) $e_1 = sv_1, e_2 = sv_2$. Then $\langle sv_1, sv_2 \rangle \stackrel{\text{def}}{=} \text{let } x = \langle sv_1, sv_2 \rangle \text{ in } x$, which is well-formed by construction using the assumptions and the freshness of x .
 - (b) $e_1 = sv_1, e_2 = \text{let } x = i \text{ in } e'_2$. Then $\langle sv_1, e_2 \rangle \stackrel{\text{def}}{=} \text{let } x = i \text{ in } \langle sv_1, e'_2 \rangle$. By inverting the second assumption (3), $\Psi; \Delta; \Gamma \vdash i : \tau_x$ **opr**₃₂ and $\Psi; \Delta; \Gamma, x : \tau_x \vdash e'_2 : \tau_2$ **exp**. By weakening the first assumption (lemma 2), $\Psi; \Delta; \Gamma, x : \tau_x \vdash sv_1 : \tau_1$ **exp**. By the derived typing rule for pairs $\Psi; \Delta; \Gamma, x : \tau_x \vdash \langle sv_1, e'_2 \rangle : \tau_1 \times \tau_2$ **dexp**, and hence by induction $\Psi; \Delta; \Gamma, x : \tau_x \vdash \langle sv_1, e'_2 \rangle : \tau_1 \times \tau_2$ **exp**. Finally, by construction using the operation rule, $\Psi; \Delta; \Gamma \vdash \text{let } x = i \text{ in } \langle sv_1, e'_2 \rangle : \tau_1 \times \tau_2$ **exp**.

(c) $e_1 = \text{let } x = i \text{ in } e'_1, e_2 = sv_2$. This case proceeds exactly as the previous case.

2. The other cases proceed similarly. ■

Finally, I show the soundness of the dynamic encoding translation, using the soundness theorems for the **SEK**, **SE**, and **DET** translations, along with the commutativity lemma for the **DET** translation and the soundness lemmas for contexts.

Theorem 3 (Soundness of the dynamic encoding translation)

If $\Delta \vdash c : \kappa$ and $\vdash \Psi \text{ ok}$ then $\Psi; |\Delta|; |\Delta| \vdash |c| : |\kappa| \text{ exp}$.

Proof: By induction on derivations of $\Delta \vdash c : \kappa$.

1. $\Delta[\alpha:\kappa] \vdash \alpha : \kappa$.

By assumption:

$\vdash \Delta[\alpha:\kappa] \text{ ok}$

$\vdash \Psi \text{ ok}$

By lemmas 7 and 15:

$\vdash |\Delta[\alpha:\kappa]| \text{ ok}$

$|\Delta[\alpha:\kappa]| \vdash |\Delta[\alpha:\kappa]| \text{ ok}$

By definition:

$|\alpha| \stackrel{\text{def}}{=} x_\alpha$

$|\Delta[\alpha:\kappa]| \stackrel{\text{def}}{=} |\Delta| [x_\alpha : |\alpha:\kappa|]$

Hence by the variable rule:

$\Psi; |\Delta[\alpha:\kappa]|; |\Delta| [x_\alpha : |\alpha:\kappa|] \vdash x_\alpha : |\alpha:\kappa| \text{ exp}$

2. $\Delta \vdash \lambda(\alpha:\kappa_1).c : \kappa_1 \rightarrow \kappa_2$.

By assumption:

$\vdash \kappa_1 \text{ ok}, \alpha \notin \text{fv}(\Delta)$

$\Delta, \alpha:\kappa_1 \vdash c : \kappa_2$

By theorem 1 and lemma 12:

$|\Delta, \alpha:\kappa_1| \vdash |\kappa| \text{ ok}$

$|\Delta, \alpha:\kappa_1| \vdash |\alpha:\kappa_1| : T_{32}$

By induction:

$\Psi; |\Delta, \alpha:\kappa_1|; |\Delta, \alpha:\kappa_1| \vdash |c| : |\kappa_2| \text{ exp}$

So by the function introduction rule:

$\Psi; |\Delta|; |\Delta| \vdash |\lambda\alpha:\kappa_1.c| : |\lambda\alpha:\kappa_1.c:\kappa_1 \rightarrow \kappa_2| \text{ exp}$

3. $\Delta \vdash c_1 c_2 : \kappa_2$.

By assumption:

$\Delta \vdash c_1 : \kappa_1 \rightarrow \kappa_2$

$\Delta \vdash c_2 : \kappa_1$

By induction:

$$\Psi; |\Delta|; |\Delta| \vdash |c_1| : |c_1 : \kappa_1 \rightarrow \kappa_2| \quad \mathbf{exp}$$

$$\Psi; |\Delta|; |\Delta| \vdash |c_2| : |c_2 : \kappa_1| \quad \mathbf{exp}$$

By theorem 2:

$$|\Delta| \vdash |c_2| : |\kappa_1|$$

Note that $|c_1 : \kappa_1 \rightarrow \kappa_2| \stackrel{\text{def}}{=} \forall[\alpha : |\kappa_1|] (|\alpha : \kappa_1| \rightarrow |c_1 \alpha : \kappa_2|)$.

By lemma 13:

$$|\alpha : \kappa_1| [|c_2| / \alpha] = |c_2 : \kappa_1|$$

$$|c_1 \alpha : \kappa_2| [|c_2| / \alpha] = |c_1 c_2 : \kappa_2|$$

(Note $\alpha \notin fv(c_1)$ by freshness).

Therefore, by the derived rule for application of expressions:

$$\Psi; |\Delta|; |\Delta| \vdash c_1 c_2 : |c_1 c_2 : \kappa_2| \quad \mathbf{dexp}$$

Finally, observe that by lemma 17:

$$\Psi; |\Delta|; |\Delta| \vdash c_1 c_2 : |c_1 c_2 : \kappa_2| \quad \mathbf{exp}$$

4. $\Delta \vdash \pi_1 c : \kappa_1$.

By assumption:

$$\Delta \vdash c : \kappa_1 \times \kappa_2$$

By induction

$$\Psi; |\Delta|; |\Delta| \vdash |c| : |c : \kappa_1 \times \kappa_2| \quad \mathbf{exp}$$

By definition:

$$|c : \kappa_1 \times \kappa_2| = |\pi_1 c : \kappa_1| \times |\pi_2 c : \kappa_2|$$

By the typing rule for the $\mathbf{select}^i e$ derived form:

$$\Psi; |\Delta|; |\Delta| \vdash \mathbf{select}^1 |c| : |\pi_1 c : \kappa_1| \quad \mathbf{dexp}$$

By lemma 17:

$$\Psi; |\Delta|; |\Delta| \vdash \mathbf{select}^1 |c| : |\pi_1 c : \kappa_1| \quad \mathbf{exp}$$

5. $\Delta \vdash \pi_2 c : \kappa_2$. The proof proceeds as in the previous case.

6. $\Delta \vdash \langle c_1, c_2 \rangle : \kappa_1 \times \kappa_2$

By assumption:

$$\Delta \vdash c_1 : \kappa_1$$

$$\Delta \vdash c_2 : \kappa_2$$

By induction

$$\Psi; |\Delta|; |\Delta| \vdash |c_1| : |c_1 : \kappa_1| \quad \mathbf{exp}$$

$$\Psi; |\Delta|; |\Delta| \vdash |c_2| : |c_2 : \kappa_2| \quad \mathbf{exp}$$

By the derived pair introduction rule:

$$\Psi; |\Delta|; |\Delta| \vdash \langle |c_1|, |c_2| \rangle : |c_1 : \kappa_1| \times |c_2 : \kappa_2| \quad \mathbf{dexp}$$

By theorem 17:

$$\Psi; |\Delta|; |\Delta| \vdash \langle |c_1|, |c_2| \rangle : |c_1 : \kappa_1| \times |c_2 : \kappa_2| \quad \mathbf{exp}$$

By the definitions of the translations and the beta rule for pairs:

$$\Psi; |\Delta|; |\Delta| \vdash |\langle c_1, c_2 \rangle| : |\langle c_1, c_2 \rangle : \kappa_1 \times \kappa_2| \quad \mathbf{exp}$$

7. The rest of the cases follow directly by construction, with appeals to previous lemmas where necessary. ■

5.4 Type translations

5.4.1 Proper MIL types

The **DE** and **SE** translations complete the apparatus needed to translate proper **MIL** constructors. Constructors of kind T_{32} map to constructors of kind T_{mil} , which permits analysis of the form of the original constructor. However, since proper types are not analyzable in the **MIL** (only constructors of kind T_{32} can be analyzed) the translation of proper **MIL** types does not map into the T_{mil} kind, but instead goes directly to kind T_{32} .

$$\begin{aligned}
|\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n](\tau_1, \dots, \tau_m)(k) \rightarrow \tau| &\stackrel{\text{def}}{=} \\
&\forall[\alpha_1:|\kappa_1|, \dots, \alpha_n:|\kappa_n|] \\
&(|\alpha_1:\kappa_1|, \dots, |\alpha_n:\kappa_n|, |\tau_1|, \dots, |\tau_m|)(\text{Float}_0, \dots, \text{Float}_{k-1}) \rightarrow |\tau| \\
|T(c)| &\stackrel{\text{def}}{=} \text{interp } |c| \\
|\tau_1 \times \dots \times \tau_n| &\stackrel{\text{def}}{=} \times[|\tau_1|, \dots, |\tau_n|]
\end{aligned}$$

The most interesting piece of this translation is the treatment of polymorphic functions. Additional arguments corresponding to the dynamic representations of the type arguments are added to the parameter list of the function. Other types are translated compositionally. Notice though that the translation of a constructor used as a type is the *interpretation* of the translation of the constructor. This reflects the fact that proper **MIL** types are not available for analysis - they simply serve as classifiers.

5.4.2 The term typing context translation

The type and kind translations can be used to give a definition of the translation of a **MIL** typing context in the obvious manner.

$$\begin{aligned}
|\bullet| &\stackrel{\text{def}}{=} \bullet \\
|\Gamma, x:\tau| &\stackrel{\text{def}}{=} |\Gamma|, x:|\tau| \\
|\Gamma, x_{64}| &\stackrel{\text{def}}{=} |\Gamma|, x_{64}:\text{Float}
\end{aligned}$$

5.4.3 Proofs of soundness for the proper type translations

Theorem 4 (Soundness of the type translation)

If $\Delta \vdash \tau \text{ ok}$ then $|\Delta| \vdash |\tau| : T_{32}$

Proof: By induction on τ . We proceed by cases:

- $\Delta \vdash T(c) \text{ ok}$

By assumption:

$$\Delta \vdash c : T_{32}$$

By theorem 2:

$$|\Delta| \vdash |c| : T_{32}$$

Note that $|T_{32}| = T_{\text{mil}}$

By lemma 8:

$$|\Delta| \vdash \text{interp} : T_{\text{mil}} \rightarrow T_{32}$$

By the application rule:

$$|\Delta| \vdash \text{interp}(|c|) : T_{32}$$

- $\Delta \vdash \tau_1 \times \dots \times \tau_n \text{ ok}$

By assumption:

$$\Delta \vdash \tau_1 \text{ ok} \dots \Delta \vdash \tau_n \text{ ok}$$

By induction:

$$|\Delta| \vdash |\tau_1| : T_{32} \dots |\Delta| \vdash |\tau_n| : T_{32}$$

By construction:

$$|\Delta| \vdash \times[|\tau_1|, \dots, |\tau_n|] \text{ ok} T_{32}$$

- $\Delta \vdash \forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n](\tau_1, \dots, \tau_m)(k) \rightarrow \tau \text{ ok}$

By assumption:

$$\Delta \vdash \kappa_i \text{ ok} \quad i \in 1 \dots n$$

$$\Delta[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n] \vdash \tau_i \text{ ok} \quad i \in 1 \dots m$$

$$\Delta[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n] \vdash \tau \text{ ok}$$

By theorem 1 and lemma 12:

$$|\Delta| \vdash \kappa_i \text{ ok} \quad i \in 1 \dots n$$

$$|\Delta[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]| \vdash |\alpha_i:\kappa_i| : T_{32} \quad i \in 1 \dots n$$

By induction:

$$|\Delta[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]| \vdash |\tau_i| : T_{32} \quad i \in 1 \dots m$$

$$|\Delta[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]| \vdash |\tau| : T_{32}$$

By the function type introduction rule:

$$|\Delta| \vdash \left\{ \begin{array}{l} \forall[\alpha_1:|\kappa_1|, \dots, \alpha_n:|\kappa_n|] \\ (\downarrow\alpha_1:\kappa_1\downarrow, \dots, \downarrow\alpha_n:\kappa_n\downarrow, |\tau_1|, \dots, |\tau_m|)(\text{Float}_1 \dots \text{Float}_k) \rightarrow |\tau| \end{array} \right\} : T_{32}$$

■

Lemma 18

1. If $x \notin \Gamma$ then $x \notin |\Gamma|$.

2. If $x_{64} \notin \Gamma$ then $x_{64} \notin |\Gamma|$.

Proof: Note that Γ and $|\Gamma|$ have the same domain by construction.

■

Lemma 19 (Term context translation)

If $\Delta \vdash \Gamma \text{ ok}$ then $|\Delta| \vdash |\Gamma| \text{ ok}$.

Proof: By induction on derivations, using the domain lemma (lemma 18), the soundness lemma for the constructor translation (lemma 7) and the soundness theorem for the type translation (theorem 4). ■

Corollary 2

If $\Delta \vdash \Gamma$ **ok** then $|\Delta| \vdash |\Delta|, |\Gamma|$ **ok**.

Proof: By induction on Γ , using lemmas 18, 19 and 15. Note that we assume that indexed variables x_α are drawn from a “fresh” set, and hence there is no interference between the domains of $|\Delta|$ and $|\Gamma|$. ■

5.4.4 The type translation respects equivalence

Lemma 20 (The type translation respects equivalence)

If $\Delta \vdash \tau_1 \equiv \tau_2 : T_{32}$, then $|\Delta| \vdash |\tau_1| \equiv |\tau_2| : T_{32}$.

Proof: By induction on equivalence derivations, appealing to lemma 10 and the structural equivalence rule for application in the base case $(T(c))$. ■

5.5 The term translation

The translation from **MIL** terms to **LIL** terms is not much more complicated than the constructor translation, but is syntactically somewhat more cumbersome. In order to ensure that the appropriate type information is available for constructing **LIL** terms, it is convenient to phrase the translations on terms as typed translations of the form $\Delta; \Gamma \vdash e : \tau \rightsquigarrow e'$ indicating that in the typing context $\Delta; \Gamma$, a **MIL** term e of type τ translates to a **LIL** term e' .

5.5.1 Definitions for the term translation

In order to more concisely state the translation itself, I first define a number of **LIL** functions implementing type analysis primitives. In a practical implementation, it may be desirable to inline some or all of these functions: for the most part however this is purely a policy decision trading off code size against function calls. Note though that in general it is not always possible to eliminate the lambda abstraction since the type analysis mechanism requires that certain types be a variable. In the case that the type to be analyzed is an application of a variable, the lambda abstraction cannot be eliminated.

The first defined form is the optimized array constructor, essentially as described in the example from chapter 4. I again use ML style pattern matching at the term level for clarity of presentation.

Note that the definitions are closed and hence can be written directly as code.

```

array  $\stackrel{\text{def}}{=}$ 
  code $[\alpha:T_{\text{mil}}](\lambda(r_\alpha: R(\alpha), i:\text{Int}, v:\text{interp}(\alpha))():\text{Array}(\alpha).$ 
    case  $r_\alpha$ 
      Record  $x \Rightarrow \text{vcase } \alpha (\text{Record } _ \Rightarrow \text{array}_{\text{interp}(\alpha)}(i, v) \mid _ \Rightarrow \text{dead } x)$ 
      BFloat  $x \Rightarrow \text{vcase } \alpha (\text{BFloat } _ \Rightarrow \text{array}_{\text{Float}}(i, \text{unbox } v) \mid _ \Rightarrow \text{dead } x)$ 
      Other  $x \Rightarrow \text{vcase } \alpha (\text{Other } _ \Rightarrow \text{array}_{\text{interp}(\alpha)}(i, v) \mid _ \Rightarrow \text{dead } x)$ 

```

Here the **DE** argument, r_α is the dynamic encoding of the constructor represented by α . As can be seen from the definition of R , the type of r_α is therefore a term-level sum describing the top level structure of the encoded type. The appropriate branch is thereby chosen via an ordinary case statement and evaluated. The **vcase** serves to refine the type argument to reflect the identity of the type back into the type system.

In addition to the array creation function, I define specialized update and subscript functions for the optimized arrays.

```

upd  $\stackrel{\text{def}}{=}$ 
  code $[\alpha:T_{\text{mil}}](r_\alpha: R(\alpha), a: \text{Array}(\alpha), i:\text{Int}, v:\text{interp}(\alpha))():\text{Unit}.$ 
    case  $r_\alpha$ 
      Record  $x \Rightarrow \text{vcase } \alpha (\text{Record } _ \Rightarrow \text{upd}_{\text{interp}(\alpha)}(a, i, v) \mid _ \Rightarrow \text{dead } x)$ 
      BFloat  $x \Rightarrow \text{vcase } \alpha (\text{BFloat } _ \Rightarrow \text{upd}_{\text{Float}}(a, i, \text{unbox } v) \mid _ \Rightarrow \text{dead } x)$ 
      Other  $x \Rightarrow \text{vcase } \alpha (\text{Other } _ \Rightarrow \text{upd}_{\text{interp}(\alpha)}(a, i, v) \mid _ \Rightarrow \text{dead } x)$ 

```

```

sub  $\stackrel{\text{def}}{=}$ 
  code $[\alpha:T_{\text{mil}}](r_\alpha: R(\alpha), a: \text{Array}(\alpha), i:\text{Int})():\text{interp}(\alpha).$ 
    case  $r_\alpha$ 
      Record  $x \Rightarrow \text{vcase } \alpha (\text{Record } _ \Rightarrow \text{sub}_{\text{interp}(\alpha)}(a, i) \mid _ \Rightarrow \text{dead } x)$ 
      BFloat  $x \Rightarrow \text{vcase } \alpha (\text{BFloat } _ \Rightarrow \text{box}(\text{sub}_{\text{Float}}(a, i)) \mid _ \Rightarrow \text{dead } x)$ 
      Other  $x \Rightarrow \text{vcase } \alpha (\text{Other } _ \Rightarrow \text{sub}_{\text{interp}(\alpha)}(a, i) \mid _ \Rightarrow \text{dead } x)$ 

```

Notice that all of the array primitives, in the case that the array is a flattened float array, will perform boxing or unboxing. If the type is statically known, both the branching and the boxing can be optimized away, but in general they cannot be avoided. This makes it very important to recognize statically reducible uses of type dispatch.

The **vararg** and **onearg** primitives from the **MIL** correspond in a similar way to the defined **vararg** and **onearg** functions in the **LIL**. Notice that as with the **Vararg** type function, I choose to make **vararg** polymorphic over the encoding of the argument type, but the *interpretation* of the encoded return type (that is, the return type itself).

$$\begin{aligned}
\mathbf{vararg} &\stackrel{\text{def}}{=} \\
&\text{code}[\alpha:\mathbf{T}_{\text{mil}}, \rho:\mathbf{T}_{32}](r_\alpha: R(\alpha), f:(\text{interp}(\alpha) \rightarrow \rho))(): \mathbf{Vararg}(\alpha)(\rho). \\
&\text{case } r_\alpha(\\
&\quad \mathbf{Record } y \Rightarrow \mathbf{vcase } \alpha (\\
&\quad \quad \mathbf{Record } \alpha' \Rightarrow \text{case } y \\
&\quad \quad | \text{inj}_0 x \Rightarrow \mathbf{vcase } \alpha' \\
&\quad \quad \quad ([\] \Rightarrow \lambda().(f\langle \rangle)) \\
&\quad \quad | _ \Rightarrow \mathbf{dead } x) \\
&\quad \quad | \text{inj}_1 x \Rightarrow \mathbf{vcase } \alpha' \\
&\quad \quad \quad ([\beta] \Rightarrow \lambda(y:\beta).(f\langle y \rangle)) \\
&\quad \quad | _ \Rightarrow \mathbf{dead } x) \\
&\quad \quad | \text{inj}_2 x \Rightarrow \mathbf{vcase } \alpha' \\
&\quad \quad \quad ([\beta_1, \beta_2] \Rightarrow \lambda(y_1:\beta_1, y_2:\beta_2).(f\langle y_1, y_2 \rangle)) \\
&\quad \quad | _ \Rightarrow \mathbf{dead } x) \\
&\quad \quad | \text{inj}_3 _ \Rightarrow f \\
&\quad | _ \Rightarrow \mathbf{dead } y) \\
&| _ \Rightarrow f)
\end{aligned}$$

The implementations of **vararg** and **onearg** are more complicated than those of the array primitives. In addition to dispatching on the top level form of the type of the argument, they must also distinguish between record types with different numbers of fields. In the case that the argument type is a record type, both **vararg** and **onearg** must consider the number of fields of the record to determine whether or not to emit a wrapper function, and what the argument types should be. For **vararg**, this wrapper function packages its arguments into a record and passes it on to the original function. The **onearg** construct simply reverses this.

$$\begin{aligned}
\mathbf{onearg} &\stackrel{\text{def}}{=} \\
&\text{code}[\alpha:\mathbf{T}_{\text{mil}}, \rho:\mathbf{T}_{32}]\lambda(r_\alpha: R(\alpha), f: \mathbf{Vararg}(\alpha)(\rho))(): (\text{interp}(\alpha) \rightarrow \rho). \\
&\text{case } r_\alpha(\\
&\quad \mathbf{Record } y \Rightarrow \mathbf{vcase } \alpha (\\
&\quad \quad \mathbf{Record } \alpha' \Rightarrow \text{case } y \\
&\quad \quad | \text{inj}_0 x \Rightarrow \mathbf{vcase } \alpha' \\
&\quad \quad \quad ([\] \Rightarrow \lambda(y:\mathbf{unit}).(f())) \\
&\quad \quad | _ \Rightarrow \mathbf{dead } x) \\
&\quad \quad | \text{inj}_1 x \Rightarrow \mathbf{vcase } \alpha' \\
&\quad \quad \quad ([\beta] \Rightarrow \lambda(\langle y \rangle: \times [\beta]).(f(y))) \\
&\quad \quad | _ \Rightarrow \mathbf{dead } x) \\
&\quad \quad | \text{inj}_2 x \Rightarrow \mathbf{vcase } \alpha' \\
&\quad \quad \quad ([\beta_1, \beta_2] \Rightarrow \lambda(\langle y_1, y_2 \rangle: \times [\beta_1, \beta_2]).(f(y_1, y_2))) \\
&\quad \quad | _ \Rightarrow \mathbf{dead } x) \\
&\quad \quad | \text{inj}_3 _ \Rightarrow f \\
&\quad | _ \Rightarrow \mathbf{dead } y) \\
&| _ \Rightarrow f)
\end{aligned}$$

I assign a **LIL** heap label to each of the named definitions: for example, **vararg** is the label given to the **vararg** definition.

5.5.2 The term level translation

The complete **MIL** to **LIL** term translation is given in section 5.6, but a selection of example rules from the translation are discussed below to illustrate the important points of the translation.

Small values

The small value translation judgement is of the form $\Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv'$, and is for the most part uncomplicated.

Injectons into non-value carrying arms of sum types are translated into uses of union types, as described previously. Note that injections into value-carrying arms are not values in the **MIL** since they require allocation. Their translation is therefore described in the operation translation.

$$\frac{}{\Delta; \Gamma \vdash \text{inj_tag}_{\text{Sum}_i(\vec{c})}^j : \text{Sum}_i^j(\vec{c}) \rightsquigarrow \text{inj_union}_{|\text{Sum}_i(\vec{c})|} \text{tag}_j}$$

Float values

The float value translation judgement is of the form $\Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv'$. Float values in the **MIL** and the **LIL** correspond exactly: hence the translation leaves the terms unchanged.

64 bit operations

The 64 bit operation translation judgement is of the form $\Delta \vdash i : \text{Float} \rightsquigarrow i'$, and is also relatively uncomplicated. For example, the floating point subscript operation changes in only minor syntactic ways.

$$\frac{\Delta; \Gamma \vdash sv_1 : \text{Farray} \rightsquigarrow sv'_1 \quad \Delta; \Gamma \vdash sv_2 : \text{Int} \rightsquigarrow sv'_2}{\Delta; \Gamma \vdash \text{fsub}(sv_1, sv_2) : \text{Float} \rightsquigarrow \text{sub}_{\text{Float}}(sv'_1, sv'_2)}$$

32 bit operations

The 32 bit operation translation judgement is of the form $\Delta; \Gamma \vdash i : \tau \rightsquigarrow i'$. Much of the interesting work of the translation takes place in the operations.

Although the most significant part of the **MIL** to **LIL** translation is the representation of type analysis within the term language, the definitions given in section 5.5.1 isolate all of the uses of type analysis. Consequently, the translation of type analyzing primitives such as **vararg** is very straightforward: they simply become calls to the code definitions via the appropriate heap labels.

$$\frac{\Delta; \Gamma \vdash sv : c_1 \rightarrow c_2 \rightsquigarrow sv'}{\Delta; \Gamma \vdash \text{vararg}_{c_1 \rightarrow c_2} sv : \text{Vararg}_{c_1 \rightarrow c_2} \rightsquigarrow \text{call vararg}[|c_1|, \text{interp } |c_2|](|c_1|, sv')}$$

Notice that as previously discussed, I pass the interpretation of the encoded return type to the **vararg** code, instead of the encoded type itself. The dynamic encoding of the argument type is passed as an additional argument to the function to allow the dispatch on types to take place.

Sum types become union types, with the tagging of records made explicit.

$$\frac{\Delta; \Gamma \vdash sv : c_j \rightsquigarrow sv'}{\Delta; \Gamma \vdash \text{inj}_{\text{Sum}_i(\vec{c})}^j sv : \text{Sum}_i^j(\vec{c}) \rightsquigarrow \text{inj_union}_{|\text{Sum}_i(\vec{c})|} \langle \text{tag}_j, sv' \rangle}$$

As discussed in the constructor translation, the known sum type from the **MIL** disappears in the **LIL**. Uses of the projection from this type become second projections out of a tag-value pair.

$$\frac{\Delta; \Gamma \vdash sv : \text{Sum}_i^j(\vec{c}) \rightsquigarrow sv'}{\Delta; \Gamma \vdash \text{proj}_{\text{Sum}_i(\vec{c})}^j sv : c_j \rightsquigarrow \text{select}^2 sv'}$$

In the **LIL** the exception type is no longer primitive, instead being replaced by uses of existentials and pairs.

$$\frac{\Delta; \Gamma \vdash sv_1 : \text{Dyntag}_c \rightsquigarrow sv'_1 \quad \Delta; \Gamma \vdash sv_2 : c \rightsquigarrow sv'_2}{\Delta; \Gamma \vdash \text{inj_dyn}_c(sv_1, sv_2) : \text{Exn} \rightsquigarrow \text{pack} \langle sv'_1, sv'_2 \rangle \text{ as } \exists \alpha : \text{T}_{32}. (\text{Dyntag}(\alpha) \times \alpha) \text{ hiding } |c|}$$

Since the exception type is closed, I will frequently refer to it via the following definition for brevity:

$$\mathbf{Dyn} \stackrel{\text{def}}{=} \exists \alpha : \text{T}_{32}. (\text{Dyntag}(\alpha) \times \alpha)$$

One of the most illustrative rules is the rule for the translation of applications, since it makes clear the “plumbing” work necessary to make type analysis explicit.

$$\frac{\begin{array}{l} \Delta; \Gamma \vdash sv : \forall [\alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n] (\tau_1, \dots, \tau_m)(k) \rightarrow \tau \rightsquigarrow sv' \\ \Delta; \Gamma \vdash sv_i : \tau_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \rightsquigarrow sv'_i \quad \Delta; \Gamma \vdash fv_i : \text{Float} \rightsquigarrow fv'_i \end{array}}{\Delta; \Gamma \vdash sv[c_1, \dots, c_n](sv_1, \dots, sv_m)(fv_1, \dots, fv_k) : \tau[c_1/\alpha_1, \dots, c_n/\alpha_n] \rightsquigarrow sv'[\lfloor c_1 \rfloor, \dots, \lfloor c_n \rfloor](\lfloor c_1 \rfloor, \dots, \lfloor c_n \rfloor, sv'_1, \dots, sv'_m)(fv'_1, \dots, fv'_k)}$$

Notice that each type argument is passed to the function once as a type argument in its static encoding, and once as a term argument in its dynamic encoding. The rest of the arguments are translated using the small value and float value translations and are passed as usual. The application translation will of course be paralleled exactly by the translation of functions in the expression translation below.

The term translation judgements

The term level translation judgement is of the form $\Delta; \Gamma \vdash e : \tau \rightsquigarrow e'$. For the most part, the expression translation simply appeals to the various other judgements. The only interesting case is that of polymorphic functions.

$$\frac{\begin{array}{l} \Delta; \Gamma[f : \forall [\alpha :: \kappa] (\vec{\tau})(|x_f|) \rightarrow \tau_r, \alpha :: \kappa, \vec{x} : \vec{\tau}, \vec{x}_f] \vdash e_f : \tau_r \rightsquigarrow e'_f \\ \Delta; \Gamma[f : \forall [\alpha :: \kappa] (\vec{\tau})(|x_f|) \rightarrow \tau_r] \vdash e : \tau \rightsquigarrow e' \end{array}}{\Delta; \Gamma \vdash \text{let}_{\tau} \text{rec}_{\tau_r} f[\alpha :: \kappa](\vec{x} : \vec{\tau})(\vec{x}_f). e_f \text{ in } e : \forall [\alpha :: \kappa] (\vec{\tau})(|x_f|) \rightarrow \tau \rightsquigarrow \text{let}_{|\tau|} \text{rec}_{|\tau_r|} f[\alpha :: \kappa](\vec{x} : \vec{\tau})(\vec{x}_f : \text{Float}). e'_f \text{ in } e'}$$

As expected from the application case, each type argument is both lifted to the static encoding kind and used to construct a dynamic encoding type for a new term level argument serving as its dynamic encoding.

Programs

The final **LIL** program is produced from a closed **MIL** expression by wrapping the result of the expression translation with a heap binding the defined forms used in the translation.

$$\begin{array}{llll}
 d_i \stackrel{\text{def}}{=} & \bullet, & & \\
 \text{sub} & : \forall[\alpha:T_{\text{mil}}] \text{Code}[R(\alpha), \mathbf{Array}(\alpha), \mathbf{Int}][](\text{interp}(\alpha)) & \mapsto & \mathbf{sub}, \\
 \text{upd} & : \forall[\alpha:T_{\text{mil}}] \text{Code}[R(\alpha), \mathbf{Array}(\alpha), \mathbf{Int}, \text{interp}(\alpha)][](\mathbf{unit}) & \mapsto & \mathbf{upd}, \\
 \text{array} & : \forall[\alpha:T_{\text{mil}}] \text{Code}[R(\alpha), \mathbf{Int}, \text{interp}(\alpha)][](\mathbf{Array}(\alpha)) & \mapsto & \mathbf{array}, \\
 \text{vararg} & : \forall[\alpha:T_{\text{mil}}, \rho:T_{32}] \text{Code}[R(\alpha), (\text{interp}(\alpha) \rightarrow \rho)][](\mathbf{Vararg}(\alpha)(\rho)) & \mapsto & \mathbf{vararg}, \\
 \text{onearg} & : \forall[\alpha:T_{\text{mil}}, \rho:T_{32}] \text{Code}[R(\alpha), \mathbf{Vararg}(\alpha)(\rho)][](\text{interp}(\alpha) \rightarrow \rho) & \mapsto & \mathbf{onearg}
 \end{array}$$

$$\frac{\bullet; \bullet \vdash e : \tau \rightsquigarrow e'}{\vdash e : \tau \rightsquigarrow \text{letrec } d_i \text{ in } e' \text{ prog}}$$

5.5.3 Proof that the term level translation preserves typing

Small Values

Theorem 5 (Small value translation)

If $\Delta; \Gamma \vdash sv : \tau, \vdash \Psi \text{ ok}$ and $\Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv'$ then $\Psi; |\Delta|; \downarrow\Delta|, |\Gamma| \vdash sv' : |\tau|$

Proof: By induction on terms.

1. If $sv = x$ then

By assumption:

$$\begin{array}{l}
 \Delta; \Gamma, x:\tau \vdash x : \tau \\
 \vdash \Psi \text{ ok}
 \end{array}$$

By inverting the assumption:

$$\Delta \vdash \Gamma, x:\tau \text{ ok}$$

By lemma 19:

$$|\Delta| \vdash \downarrow\Delta|, |\Gamma, x:\tau| \text{ ok}$$

By the variable rule and the definition of $|\Gamma|$:

$$\Psi; |\Delta|; \downarrow\Delta|, |\Gamma|, x:|\tau| \vdash x : |\tau|$$

2. If $sv = i$ then

By assumption:

$$\begin{array}{l}
 \Delta; \Gamma \vdash i : \mathbf{Int} \\
 \vdash \Psi \text{ ok}
 \end{array}$$

By inverting the assumption:

$$\Delta \vdash \Gamma \text{ ok}$$

By lemma 19:

$$|\Delta| \vdash \downarrow \Delta|, |\Gamma| \text{ ok}$$

By the int rule:

$$\Psi; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash i : \text{Int}$$

3. If $sv = \text{unroll}_c sv$ then

By assumption:

$$\begin{aligned} \Delta; \Gamma \vdash sv : \pi_i \mu(\alpha, \beta)(c_1, c_2) \\ \Delta \vdash c \equiv \pi_i \mu(\alpha, \beta)(c_1, c_2) : T_{32} \\ \Delta \vdash sv : \pi_i \mu(\alpha, \beta)(c_1, c_2) \rightsquigarrow sv' \end{aligned}$$

By induction:

$$\Psi; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash sv' : |\pi_i \mu(\alpha, \beta)(c_1, c_2)|$$

By lemma 10:

$$|\Delta| \vdash |c| \equiv |\pi_i \mu(\alpha, \beta)(c_1, c_2)| : T_{\text{mil}}$$

By definition:

$$\begin{aligned} |\pi_i \mu(\alpha, \beta)(c_1, c_2)| &= \text{Rec}[1 + 1](f)(\text{inj}_i *) \\ \text{where } f &= \lambda(\rho:1 + 1 \rightarrow T_{32}).\lambda(\omega.1 + 1). \end{aligned}$$

case ω

$$\begin{aligned} \text{inj}_1 - &\Rightarrow \text{interp}(|c_1|[\mathbf{Other}(\rho(\text{inj}_1 *))/\alpha, \mathbf{Other}(\rho(\text{inj}_2 *))/\beta]) \\ \text{inj}_2 - &\Rightarrow \text{interp}(|c_2|[\mathbf{Other}(\rho(\text{inj}_1 *))/\alpha, \mathbf{Other}(\rho(\text{inj}_2 *))/\beta]) \end{aligned}$$

Therefore, by the unroll rule:

$$\Psi; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash \text{unroll}_{|c|} sv' : f(\text{Rec}[1 + 1](f))(\text{inj}_i *)$$

It suffices to show that:

$$f(\text{Rec}[1 + 1](f))(\text{inj}_i *) = \text{interp}(|c_i|[\pi_1 \mu(\alpha, \beta)(c_1, c_2)/\alpha, \pi_2 \mu(\alpha, \beta)(c_1, c_2)/\beta])$$

By lemma 9:

$$\begin{aligned} |c_i|[\pi_1 \mu(\alpha, \beta)(c_1, c_2)/\alpha, \pi_2 \mu(\alpha, \beta)(c_1, c_2)/\beta] &= \\ |c_i|[\pi_1 \mu(\alpha, \beta)(c_1, c_2)/\alpha, \pi_2 \mu(\alpha, \beta)(c_1, c_2)/\beta] & \end{aligned}$$

By definition:

$$\begin{aligned} |\pi_1 \mu(\alpha, \beta)(c_1, c_2)| &= \mathbf{Other}(\text{Rec}[1 + 1](f)(\text{inj}_1 *)) \\ |\pi_2 \mu(\alpha, \beta)(c_1, c_2)| &= \mathbf{Other}(\text{Rec}[1 + 1](f)(\text{inj}_2 *)) \end{aligned}$$

So it suffices to show that:

$$\begin{aligned} f(\text{Rec}[1 + 1](f))(\text{inj}_i *) &= \\ \text{interp}(|c_i|[\mathbf{Other}(\text{Rec}[1 + 1](f)(\text{inj}_1 *))/\alpha, \mathbf{Other}(\text{Rec}[1 + 1](f)(\text{inj}_2 *))/\beta]) & \end{aligned}$$

By definition of f :

$$\begin{aligned} f(\text{Rec}[1 + 1](f))(\text{inj}_i *) &= \\ \text{case } (\text{inj}_i *) & \\ \text{inj}_1 - &\Rightarrow \\ \text{interp}(|c_1|[\mathbf{Other}(\text{Rec}[1 + 1](f)(\text{inj}_1 *))/\alpha, \mathbf{Other}(\text{Rec}[1 + 1](f)(\text{inj}_2 *))/\beta]) & \\ \text{inj}_2 - &\Rightarrow \\ \text{interp}(|c_2|[\mathbf{Other}(\text{Rec}[1 + 1](f)(\text{inj}_1 *))/\alpha, \mathbf{Other}(\text{Rec}[1 + 1](f)(\text{inj}_2 *))/\beta]) & \end{aligned}$$

If $i = 1$, then by the beta rule:

$$f(\text{Rec}[1+1](f))(\text{inj}_1 *) = \text{interp}(|c_1|[\mathbf{Other}(\text{Rec}[1+1](f)(\text{inj}_1 *))/\alpha, \mathbf{Other}(\text{Rec}[1+1](f)(\text{inj}_2 *))/\beta])$$

If $i = 2$, then by the beta rule:

$$f(\text{Rec}[1+1](f))(\text{inj}_2 *) = \text{interp}(|c_2|[\mathbf{Other}(\text{Rec}[1+1](f)(\text{inj}_1 *))/\alpha, \mathbf{Other}(\text{Rec}[1+1](f)(\text{inj}_2 *))/\beta])$$

4. The **roll** case follows similarly.

5. If $sv = \text{inj_tag}_{\text{Sum}_i(\vec{c})}^j$ then:

By inversion of the first assumption:

$$\Delta \vdash \Gamma \text{ ok}$$

By lemma 19:

$$|\Delta| \vdash |\Delta|, |\Gamma| \text{ ok}$$

Therefore by the tag rule:

$$\Psi; |\Delta|; |\Delta|, |\Gamma| \vdash \text{tag}_j : \text{Tag}(j)$$

Note that

$$|\text{Sum}_i(c_i, \dots, c_n)| = \bigvee [\text{Tag}(0), \dots, \text{Tag}(i-1), \dots, \times [\text{Tag}(i), \text{interp } |c_i|], \dots, \times [\text{Tag}(n), \text{interp } |c_n|]]$$

So by the union formation rule :

$$\Psi; |\Delta|; |\Delta|, |\Gamma| \vdash \text{inj_union}_{|\text{Sum}_i(\vec{c})|} \text{tag}_j : |\text{Sum}_i(\vec{c})|$$

■

Float values

Theorem 6 (Float value translation)

If $\Delta; \Gamma \vdash fv : \text{Float}, \vdash \Psi \text{ ok}$ and $\Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv'$ then $|\Delta|; |\Delta|, |\Gamma| \vdash fv' : \text{Float}$

Proof:

1. If $fv = x_f$ then

By assumption:

$$\begin{aligned} \Delta; \Gamma, x_f \vdash x : \text{Float} \\ \vdash \Psi \text{ ok} \end{aligned}$$

By inverting the assumption:

$$\Delta \vdash \Gamma, x_f \text{ ok}$$

By lemma 19:

$$|\Delta| \vdash |\Delta|, |\Gamma, x_f| \text{ ok}$$

By the variable rule and the definition of $|\Gamma|$:

$$\Psi; |\Delta|; |\Delta|, |\Gamma|, x_f \vdash x_f : \text{Float}$$

2. $fv = r$.

By assumption:

$$\begin{aligned} \Delta; \Gamma \vdash r : \text{Float} \\ \vdash \Psi \text{ ok} \end{aligned}$$

By inverting the assumption:

$$\Delta \vdash \Gamma \text{ ok}$$

By lemma 19:

$$|\Delta| \vdash |\Delta|, |\Gamma| \text{ ok}$$

By the float rule:

$$\Psi; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash r : \text{Float}$$

■

64 bit operations

Theorem 7 (Float operation translation)

If $\Delta; \Gamma \vdash i : \text{Float}$, $\vdash \Psi \text{ ok}$, and $\Delta; \Gamma \vdash i : \text{Float} \rightsquigarrow i'$ then $|\Delta|; \downarrow \Delta|, |\Gamma| \vdash i' : \text{Float opr}_{64}$

Proof:

1. If $i = fv$ then:

By inversion of assumptions:

$$\Delta; \Gamma \vdash fv : \text{Float}$$

$$\Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv'$$

By theorem 6:

$$\Psi; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash fv' : \text{Float}$$

By the float value inclusion rule:

$$\Psi; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash fv' : \text{Float opr}_{64}$$

2. If $i = \text{unboxf } sv$ then:

By inversion of assumptions:

$$\Delta; \Gamma \vdash sv : \text{Boxf}$$

$$\Delta; \Gamma \vdash sv : \text{Boxf} \rightsquigarrow sv'$$

By theorem 5:

$$\Psi; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash sv' : \text{BoxedFloat}$$

By the unbox rule:

$$\Psi; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash \text{unbox } sv' : \text{Float opr}_{64}$$

3. If $i = \text{fsub}(sv_1, sv_2)$ then:

By inversion of assumptions:

$$\Delta; \Gamma \vdash sv_1 : \text{Farray}$$

$$\Delta; \Gamma \vdash sv_2 : \text{Int}$$

$$\Delta; \Gamma \vdash sv_1 : \text{Farray} \rightsquigarrow sv'_1$$

$$\Delta; \Gamma \vdash sv_2 : \text{Int} \rightsquigarrow sv'_2$$

By theorem 5:

$$\Psi; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash sv'_1 : \text{Array}_{64}\text{Float}$$

$$\Psi; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash sv'_2 : \text{Int}$$

Hence by the 64 bit subscript rule :

$$\Psi; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash \text{sub}_{\text{Float}}(sv'_1, sv'_2) : \text{Float opr}_{64}$$

■

Operations and expressions

I begin with a lemma to capture the appropriate typing conditions for the definitions used by the translation of the type analyzing primitives.

Lemma 21 (Definitions)

1. If $\vdash \Psi$ **ok** then

$$\Psi \vdash \mathbf{sub} : \forall[\alpha:T_{\text{mil}}] \text{Code}[R(\alpha), \mathbf{Array}(\alpha), \mathbf{Int}][](\text{interp}(\alpha)) \text{ hval}$$

2. If $\vdash \Psi$ **ok** then

$$\Psi \vdash \mathbf{upd} : \forall[\alpha:T_{\text{mil}}] \text{Code}[R(\alpha), \mathbf{Array}(\alpha), \mathbf{Int}, \text{interp}(\alpha)][](\mathbf{unit}) \text{ hval}$$

3. If $\vdash \Psi$ **ok** then

$$\Psi \vdash \mathbf{array} : \forall[\alpha:T_{\text{mil}}] \text{Code}[R(\alpha), \mathbf{Int}, \text{interp}(\alpha)][](\mathbf{Array}(\alpha)) \text{ hval}$$

4. If $\vdash \Psi$ **ok** then

$$\Psi \vdash \mathbf{vararg} : \forall[\alpha:T_{\text{mil}}, \beta:T_{32}] \text{Code}[R(\alpha), \text{interp}(\alpha) \rightarrow \beta][][](\mathbf{Vararg}(\alpha)(\beta)) \text{ hval}$$

5. If $\vdash \Psi$ **ok** then

$$\Psi \vdash \mathbf{onearg} : \forall[\alpha:T_{\text{mil}}, \beta:T_{32}] \text{Code}[R(\alpha), \mathbf{Vararg}(\alpha)(\beta)][](\text{interp}(\alpha) \rightarrow \beta) \text{ hval}$$

Proof: By construction. ■

I define a heap context Ψ_i that gives types for the above definitions as follows:

$$\begin{aligned} \Psi_i &\stackrel{\text{def}}{=} \bullet, \\ \mathbf{sub} &: \forall[\alpha:T_{\text{mil}}] \text{Code}[R(\alpha), \mathbf{Array}(\alpha), \mathbf{Int}][](\text{interp}(\alpha)), \\ \mathbf{upd} &: \forall[\alpha:T_{\text{mil}}] \text{Code}[R(\alpha), \mathbf{Array}(\alpha), \mathbf{Int}, \text{interp}(\alpha)][](\mathbf{unit}), \\ \mathbf{array} &: \forall[\alpha:T_{\text{mil}}] \text{Code}[R(\alpha), \mathbf{Int}, \text{interp}(\alpha)][](\mathbf{Array}(\alpha)), \\ \mathbf{vararg} &: \forall[\alpha:T_{\text{mil}}, \rho:T_{32}] \text{Code}[R(\alpha), (\text{interp}(\alpha) \rightarrow \rho)][][](\mathbf{Vararg}(\alpha)(\rho)), \\ \mathbf{onearg} &: \forall[\alpha:T_{\text{mil}}, \rho:T_{32}] \text{Code}[R(\alpha), \mathbf{Vararg}(\alpha)(\rho)][](\text{interp}(\alpha) \rightarrow \rho) \end{aligned}$$

Lemma 22 (Initial context)

$\vdash \Psi_i$ **ok**

Proof: By construction. ■

Operations and expressions are mutually recursive: hence their soundness theorems are stated most naturally together. Since the translation of operations generates calls to the code functions implementing type analysis, the soundness theorem for the translation is defined relative to the initial context defined above.

Theorem 8 (Operation and expression translation)

1. If $\Delta; \Gamma \vdash i : \tau$ and $\Delta; \Gamma \vdash i : \tau \rightsquigarrow i' : \tau$ then $\Psi_i; |\Delta|; \downarrow \Delta \vdash i' : |\tau| \text{ opr}_{32}$

2. If $\Delta; \Gamma \vdash e : \tau$ and $\Delta; \Gamma \vdash e : \tau \rightsquigarrow e' : \tau$ then $\Psi_i; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash e' : |\tau|$ **exp**

Proof:

The proof proceeds by simultaneous induction on operations and expressions. We begin with the operations.

1. If $i = sv$ then:

By inversion of assumptions:

$$\Delta; \Gamma \vdash sv : \tau$$

$$\Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv'$$

By theorem 5:

$$\Psi_i; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash sv' : |\tau|$$

By the small value expression inclusion rule:

$$\Psi_i; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash sv' : |\tau| \text{ **opr**}_{32}$$

2. If $i = \text{vararg}_{c_1 \rightarrow c_2} sv$

By inversion of assumptions:

$$\Delta \vdash c_1 : T_{32}$$

$$\Delta \vdash c_2 : T_{32}$$

$$\Delta; \Gamma \vdash sv : c_1 \rightarrow c_2$$

$$\Delta; \Gamma \vdash sv : c_1 \rightarrow c_2 \rightsquigarrow sv'$$

By theorem 2:

$$|\Delta| \vdash |c_1| : T_{\text{mil}}$$

$$|\Delta| \vdash |c_2| : T_{\text{mil}}$$

By lemma 8:

$$|\Delta| \vdash \text{interp } |c_2| : T_{32}$$

By theorem 3 and weakening:

$$\Psi_i; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash |c_1| : R(c_1) \quad (\text{Note } \downarrow c_1 : T_{32} \models R(c_1)).$$

By theorem 5:

$$\Psi_i; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash sv' : |c_1 \rightarrow c_2|$$

By the initial context lemma (22):

$$\Psi_i; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash \text{onearg} : \forall [\alpha : T_{\text{mil}}, \rho : T_{32}] \text{Code}[R(\alpha), \text{interp}(\alpha) \rightarrow \beta][\text{Vararg}(\alpha)(\rho)]$$

By the code call rule:

$$\Psi_i; |\Delta|; \downarrow \Delta|, |\Gamma| \vdash \text{call vararg}[|c_1|, \text{interp } |c_2|](|c_1|, sv') : \text{Vararg}(|c_1|)(\text{interp } |c_2|) \text{ **opr**}_{32}$$

3. If $i = \text{onearg}_{c_1 \rightarrow c_2} sv$

By inversion of assumptions:

$$\Delta \vdash c_1 : T_{32}$$

$$\Delta \vdash c_2 : T_{32}$$

$$\Delta; \Gamma \vdash sv : \text{Vararg}_{c_1 \rightarrow c_2}$$

$$\Delta; \Gamma \vdash sv : \text{Vararg}_{c_1 \rightarrow c_2} \rightsquigarrow sv'$$

By theorem 2:

$$|\Delta| \vdash |c_1| : T_{\text{mil}}$$

$$|\Delta| \vdash |c_2| : T_{\text{mil}}$$

By lemma 8:

$$|\Delta| \vdash \text{interp } |c_2| : T_{32}$$

By theorem 3 and weakening:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash |c_1| : R(|c_1|) \quad (\text{Note } \downarrow c_1 : T_{32} \models R(|c_1|)).$$

By theorem 5:

$$\begin{aligned} \Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv' : \mathbf{Vararg}(|c_1|)(\text{interp } |c_2|) \\ (\text{Note } |\mathbf{Vararg}_{c_1 \rightarrow c_2}| = \mathbf{Vararg}(|c_1|)(\text{interp } |c_2|)) \end{aligned}$$

By the initial context lemma (22):

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \text{vararg} : \forall[\alpha : T_{\text{mil}}, \rho : T_{32}] \text{Code}[R(\alpha), \mathbf{Vararg}(\alpha)(\rho)][](\text{interp}(\alpha) \rightarrow \beta)$$

By the code call rule:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \text{call onearg}[|c_1|, \text{interp } |c_2|](\downarrow c_1, sv') : \text{interp } |c_1| \rightarrow \text{interp } |c_2| \quad \mathbf{opr}_{32}$$

4. If $i = \text{boxf } fv$ then

By inversion of assumptions:

$$\begin{aligned} \Delta; \Gamma \vdash fv : \text{Float} \\ \Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv' \end{aligned}$$

By theorem 6:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash fv' : \text{Float}$$

By the box rule:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \text{box } fv' : \text{BoxedFloat}$$

5. If $i = \text{proj}_{\text{Sum}_i(\vec{c})}^j sv$ then

By inversion of assumptions:

$$\begin{aligned} \Delta \vdash \text{Sum}_i^j(\vec{c}) : T_{32} \\ \Delta; \Gamma \vdash sv : \text{Sum}_i^j(c_i, \dots, c_j, \dots, c_n) \\ \Delta; \Gamma \vdash sv : \text{Sum}_i^j(c_i, \dots, c_j, \dots, c_n) \rightsquigarrow sv' \end{aligned}$$

By theorem 5:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv' : |\text{Sum}_i^j(c_i, \dots, c_j, \dots, c_n)|$$

By definition:

$$|\text{Sum}_i^j(c_1, \dots, c_j, \dots, c_n)| = \times[\text{Tag}(j), \text{interp } |c_j|]$$

Hence by the selection rule:

$$\begin{aligned} \Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \text{select}^2 sv' : \text{interp } |c_j| \\ (\text{Note the implicit type inclusion, hence } \text{interp } |c_j| = |T(c_j)|) \end{aligned}$$

6. If $i = \text{inj}_{\text{Sum}_i(c_i, \dots, c_j, \dots, c_n)}^j sv$ then

By inverting the assumptions:

$$\begin{aligned} \Delta \vdash \text{Sum}_i(c_i, \dots, c_j, \dots, c_n) : T_{32} \\ \Delta; \Gamma \vdash sv : c_j \\ \Delta; \Gamma \vdash sv : c_j \rightsquigarrow sv' \end{aligned}$$

By theorem 5:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv' : |c_j|$$

By the pair and tag rules:

$$\Psi_i; |\Delta|; \downarrow \Delta \vdash \langle \mathbf{tag}_j, sv' \rangle : \times [\mathbf{Tag}(j), |c_j|]$$

By definition:

$$|\mathbf{Sum}_i(c_1, \dots, c_n)| = \bigvee [\mathbf{Tag}(0), \dots, \mathbf{Tag}(i-1), \dots, \times [\mathbf{Tag}(i), |c_i|], \dots, \times [\mathbf{Tag}(n), |c_n|]]$$

By the union injection rule:

$$\Psi_i; |\Delta|; \downarrow \Delta \vdash \mathbf{inj_union}_{|\mathbf{Sum}_i(\vec{c})|} \langle \mathbf{tag}_j, sv' \rangle : |\mathbf{Sum}_i(\vec{c})|$$

7. If $i = \langle sv'_1, \dots, sv'_n \rangle$ then:

By inverting assumptions:

$$\begin{aligned} \Delta; \Gamma \vdash sv_i : \tau_i \quad i \in 1 \dots n \\ \Delta; \Gamma \vdash sv_i : \tau_i \rightsquigarrow sv'_i \quad i \in 1 \dots n \end{aligned}$$

By theorem 5:

$$\Psi_i; |\Delta|; \downarrow \Delta \vdash sv'_i : |\tau_i| \quad i \in 1 \dots n$$

By the tuple rule:

$$\Psi_i; |\Delta|; \downarrow \Delta \vdash \langle sv'_1, \dots, sv'_n \rangle : \times [|\tau_1|, \dots, |\tau_n|]$$

8. If $i = \mathbf{select}^i sv$ then:

By inverting assumptions:

$$\begin{aligned} \Delta; \Gamma \vdash sv : \tau_1 \times \dots \times \tau_n \\ \Delta; \Gamma \vdash sv : \tau_1 \times \dots \times \tau_n \rightsquigarrow sv' \end{aligned}$$

By theorem 5:

$$\Psi_i; |\Delta|; \downarrow \Delta \vdash sv' : |\tau_1 \times \dots \times \tau_n|$$

By the select rule:

$$\Psi_i; |\Delta|; \downarrow \Delta \vdash \mathbf{select}^i sv' : |\tau_i|$$

9. If $i = \mathbf{case}_\tau(sv) (x_1.e_1, \dots, x_n.e_n)$ then:

By inverting assumptions:

$$\begin{aligned} \Delta \vdash \tau : \mathbf{T}_{32} \\ \Delta; \Gamma \vdash sv : \mathbf{Sum}_i(\vec{c}) \\ \Delta; \Gamma[x_j : \mathbf{Sum}_i^j(\vec{c})] \vdash e_j : \tau \\ \Delta; \Gamma \vdash sv : \mathbf{Sum}_i(\vec{c}) \rightsquigarrow sv' \\ \Delta; \Gamma[x_j : \mathbf{Sum}_i^j(\vec{c})] \vdash e_j : \tau \rightsquigarrow e'_j \end{aligned}$$

By theorem 4:

$$|\Delta| \vdash |\tau| : \mathbf{T}_{32}$$

By theorem 5:

$$\Psi_i; |\Delta|; \downarrow \Delta \vdash sv' : |\mathbf{Sum}_i(\vec{c})|$$

By induction:

$$\Psi_i; |\Delta|; \downarrow \Delta \vdash \Gamma[x_j : \mathbf{Sum}_i^j(\vec{c})] \vdash e'_j : |\tau| \mathbf{exp}$$

Note that:

$$\begin{aligned} |\Gamma[x_j : \mathbf{Sum}_i^j(\vec{c})]| &= |\Gamma[x_j : |\mathbf{Sum}_i^j(\vec{c})|]| \\ |\mathbf{Sum}_i(c_1, \dots, c_n)| &= \bigvee [\mathbf{Tag}(0), \dots, \mathbf{Tag}(i-1), \dots, \times [\mathbf{Tag}(i), |c_i|], \dots, \times [\mathbf{Tag}(n), |c_n|]] \\ |\mathbf{Sum}_i^j(c_1, \dots, c_n)| &= \mathbf{Tag}(j) \quad j < i \\ |\mathbf{Sum}_i^j(c_1, \dots, c_n)| &= \times [\mathbf{Tag}(j), |c_j|] \quad i \leq j \leq n \end{aligned}$$

Hence by the case introduction rule:

$$\Psi_i; |\Delta|; \downarrow \Delta \vdash \text{case}_{|\tau|}(sv')(x_1.e'_1, \dots, x_n.e'_n) : |\tau| \text{ exp}$$

10. If $i = \text{handle}_\tau(e_1, x.e_2)$ then:

By inverting assumptions:

$$\begin{aligned} \Delta \vdash \tau : T_{32} \\ \Delta; \Gamma \vdash e_1 : \tau \\ \Delta; \Gamma[x : \text{Exn}] \vdash e_2 : \tau \\ \Delta; \Gamma \vdash e_1 : \tau \rightsquigarrow e'_1 \\ \Delta; \Gamma[x : \text{Exn}] \vdash e_2 : \tau \rightsquigarrow e'_2 \end{aligned}$$

By theorem 4:

$$|\Delta| \vdash |\tau| : T_{32}$$

By induction:

$$\begin{aligned} \Psi_i; |\Delta|; \downarrow \Delta \vdash e'_1 : |\tau| \text{ exp} \\ \Psi_i; |\Delta|; \downarrow \Delta \vdash \Gamma[x : \text{Exn}] \vdash e'_2 : |\tau| \text{ exp} \end{aligned}$$

By definition:

$$|\Gamma[x : \text{Exn}]| = |\Gamma[x : \text{Dyn}]|$$

So by the handle rule :

$$\Psi_i; |\Delta|; \downarrow \Delta \vdash \text{handle}_{|\tau|}(e'_1, x.e'_2) : |\tau| \text{ opr}_{32}$$

11. If $i = sv[c_1, \dots, c_n](sv_1, \dots, sv_m)(fv_1, \dots, fv_k)$ then:

By inverting assumptions:

$$\begin{aligned} \Delta; \Gamma \vdash sv : \forall[\alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n](\tau_1, \dots, \tau_m)(k) \rightarrow \tau \\ \Delta; \Gamma \vdash sv_i : \tau_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \\ \Delta; \Gamma \vdash fv_i : \text{Float} \\ \Delta; \Gamma \vdash sv : \forall[\alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n](\tau_1, \dots, \tau_m)(k) \rightarrow \tau \rightsquigarrow sv' \\ \Delta; \Gamma \vdash sv_i : \tau_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \rightsquigarrow sv'_i \\ \Delta; \Gamma \vdash fv_i : \text{Float} \rightsquigarrow fv'_i \end{aligned}$$

By theorems 5 and 6:

$$\begin{aligned} \Psi_i; |\Delta|; \downarrow \Delta \vdash sv' : \forall[\alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n](\tau_1, \dots, \tau_m)(k) \rightarrow \tau \\ \Psi_i; |\Delta|; \downarrow \Delta \vdash sv'_i : \tau_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \\ \Psi_i; |\Delta|; \downarrow \Delta \vdash fv'_i : \text{Float} \end{aligned}$$

By definition:

$$\begin{aligned} |\forall[\alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n](\tau_1, \dots, \tau_m)(k) \rightarrow \tau| = \\ \forall[\alpha_1 :: |\kappa_1|, \dots, \alpha_n :: |\kappa_n|](\downarrow \alpha_1 :: \kappa_1 \downarrow, \dots, \downarrow \alpha_n :: \kappa_n \downarrow, |\tau_1|, \dots, |\tau_m|)(\text{Float}_0 \dots \text{Float}_k) \rightarrow |\tau| \end{aligned}$$

By theorem 2:

$$|\Delta| \vdash |c_i| : |\kappa| \quad i \in 1 \dots n$$

Therefore by the type instantiation rule :

$$\begin{aligned} \Psi_i; |\Delta|; \downarrow \Delta \vdash sv'[|c_1|, \dots, |c_n|] : \\ ((\downarrow \alpha_1 :: \kappa_1 \downarrow, \dots, \downarrow \alpha_n :: \kappa_n \downarrow, |\tau_1|, \dots, |\tau_m|)(\text{Float}_0 \dots \text{Float}_k) \rightarrow |\tau|)[|c_1|/\alpha_1, \dots, |c_n|/\alpha_n] \end{aligned}$$

To apply the function application rule, it suffices to show that:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \downarrow c_i \downarrow : \downarrow \alpha_i : \kappa_i \downarrow \llbracket |c_1|/\alpha_1, \dots, |c_n|/\alpha_n \rrbracket \quad i \in 1 \dots n$$

and

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv'_i : |\tau_i| \llbracket |c_1|/\alpha_1, \dots, |c_n|/\alpha_n \rrbracket \quad i \in 1 \dots m$$

By theorem 3:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \downarrow c_i \downarrow : \downarrow c_i : \kappa_i \downarrow \quad i \in 1 \dots n$$

By lemma 13:

$$\downarrow \alpha_i : \kappa_i \downarrow \llbracket |c_1|/\alpha_1, \dots, |c_n|/\alpha_n \rrbracket = \downarrow \alpha_i [c_1/\alpha_1, \dots, c_n/\alpha_n] : \kappa_i \downarrow = \downarrow c_i : \kappa_i \downarrow$$

We know from above that:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv'_i : |\tau_i [c_1/\alpha_1, \dots, c_n/\alpha_n]| \quad i \in 1 \dots m$$

But by lemma 9:

$$|\tau_i| \llbracket |c_1|/\alpha_1, \dots, |c_n|/\alpha_n \rrbracket = |\tau_i [c_1/\alpha_1, \dots, c_n/\alpha_n]|$$

Therefore, by the application rule :

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv' \llbracket |c_1|, \dots, |c_n| \rrbracket (\downarrow c_1 \downarrow, \dots, \downarrow c_n \downarrow, sv'_1, \dots, sv'_m) (fv'_1, \dots, fv'_k) : |\tau| \llbracket |c_1|/\alpha_1, \dots, |c_n|/\alpha_n \rrbracket$$

Finally, note that by lemma 9:

$$|\tau| \llbracket |c_1|/\alpha_1, \dots, |c_n|/\alpha_n \rrbracket = |\tau [c_1/\alpha_1, \dots, c_n/\alpha_n]|$$

12. If $i = \text{raise}_\tau sv$ then:

By inverting assumptions:

$$\Delta \vdash \tau : T_{32}$$

$$\Delta; \Gamma \vdash sv : \text{Exn}$$

$$\Delta; \Gamma \vdash sv : \text{Exn} \rightsquigarrow sv'$$

By theorem 4:

$$|\Delta| \vdash |\tau| : T_{32}$$

By theorem 5:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv' : |\text{Exn}|$$

Hence by the raise introduction rule:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \text{raise}_{|\tau|} sv' : |\tau|$$

13. If $i = \text{mkexntag}_c$ then:

By inverting assumptions:

$$\Delta \vdash c : T_{32}$$

$$\Delta \vdash \Gamma \text{ ok}$$

By theorem 4:

$$|\Delta| \vdash |c| : T_{32}$$

Hence by the dyntag intro rule:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \text{dyntag}_{|c|} : \text{Dyntag}(|c|)$$

14. If $i = \text{exncase}_\tau(sv) (sv_1 \Rightarrow x_1.e_1, - \Rightarrow e_2)$ then:

By inverting assumptions:

$$\begin{aligned}
& \Delta; \Gamma \vdash sv : \mathbf{Exn} \\
& \Delta \vdash \tau : T_{32} \\
& \Delta; \Gamma \vdash sv_1 : \mathbf{Dyntag}_{c_1} \\
& \Delta; \Gamma[x_1:c_1] \vdash e_1 : \tau \\
& \Delta; \Gamma \vdash e_2 : \tau \\
& \Delta; \Gamma \vdash sv : \mathbf{Exn} \rightsquigarrow sv' \\
& \Delta; \Gamma \vdash sv_1 : \mathbf{Dyntag}_{c_1} \rightsquigarrow sv'_1 \\
& \Delta; \Gamma[x_1:c_1] \vdash e_1 : \tau \rightsquigarrow e'_1 \\
& \Delta; \Gamma \vdash e_2 : \tau \rightsquigarrow e'_2
\end{aligned}$$

By theorem 4:

$$|\Delta| \vdash |\tau| : T_{32}$$

By theorem 5:

$$\begin{aligned}
& \Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash sv' : |\mathbf{Exn}| \\
& \Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash sv'_1 : |\mathbf{Dyntag}_{c_1}|
\end{aligned}$$

By induction:

$$\begin{aligned}
& \Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma[x_1:c_1]| \vdash e'_1 : |\tau| \mathbf{exp} \\
& \Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash e'_2 : |\tau| \mathbf{exp}
\end{aligned}$$

Hence by the dyncase introduction rule:

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash \mathbf{dyncase}_\tau(sv) (sv'_1 \Rightarrow x_1.e'_1, - \Rightarrow e'_2) : |\tau| \mathbf{opr}_{32}$$

15. If $i = \mathbf{sub}_c(sv_1, sv_2)$ then:

By inverting assumptions:

$$\begin{aligned}
& \Delta \vdash c : T_{32} \\
& \Delta; \Gamma \vdash sv_1 : \mathbf{Array}_c \\
& \Delta; \Gamma \vdash sv_2 : \mathbf{Int} \\
& \Delta; \Gamma \vdash sv_1 : \mathbf{Array}_c \rightsquigarrow sv'_1 \\
& \Delta; \Gamma \vdash sv_2 : \mathbf{Int} \rightsquigarrow sv'_2
\end{aligned}$$

By theorem 4:

$$|\Delta| \vdash |c| : T_{32}$$

By theorem 5:

$$\begin{aligned}
& \Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash sv'_1 : |\mathbf{Array}_c| \\
& \Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash sv'_2 : |\mathbf{Int}|
\end{aligned}$$

By theorem 3:

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash \downarrow c \downarrow : R(|c|)$$

By lemma 22:

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash \mathbf{sub} : \forall [\alpha : T_{\text{mil}}] \mathbf{Code}[R(\alpha), \mathbf{Array}(\alpha), \mathbf{Int}][](\text{interp}(\alpha))$$

Hence by the code call rule :

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash \mathbf{call} \mathbf{sub}[|c|](\downarrow c \downarrow, sv'_1, sv'_2) : \text{interp } |c| \mathbf{opr}_{32}$$

16. If $i = \mathbf{upd}_c(sv_1, sv_2, sv_3)$ then:

By inverting assumptions:

$$\begin{aligned}
&\Delta \vdash c : T_{32} \\
&\Delta; \Gamma \vdash sv_1 : \mathbf{Array}_c \\
&\Delta; \Gamma \vdash sv_2 : \mathbf{Int} \\
&\Delta; \Gamma \vdash sv_3 : c \\
&\Delta; \Gamma \vdash sv_1 : \mathbf{Array}_c \rightsquigarrow sv'_1 \\
&\Delta; \Gamma \vdash sv_2 : \mathbf{Int} \rightsquigarrow sv'_2 \\
&\Delta; \Gamma \vdash sv_3 : c \rightsquigarrow sv'_3
\end{aligned}$$

By theorem 4:

$$|\Delta| \vdash |c| : T_{32}$$

By theorem 5:

$$\begin{aligned}
&\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv'_1 : |\mathbf{Array}_c| \\
&\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv'_2 : |\mathbf{Int}| \\
&\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv'_3 : |c|
\end{aligned}$$

By theorem 3:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash |c| : R(|c|) \mathbf{exp}$$

By lemma 22:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \mathbf{upd} : \forall[\alpha : T_{\text{mil}}] \mathbf{Code}[R(\alpha), \mathbf{Array}(\alpha), \mathbf{Int}, \text{interp}(\alpha)][](\mathbf{unit})$$

Hence by the code call rule :

$$|\Delta|; |\Delta|, |\Gamma| \vdash \mathbf{call\ upd}[|c|](|c|, sv'_1, sv'_2, sv'_3) : \mathbf{Unit\ opr}_{32}$$

17. If $i = \mathbf{array}_c(sv_1, sv_2)$ then:

By inverting assumptions:

$$\begin{aligned}
&\Delta \vdash c : T_{32} \\
&\Delta; \Gamma \vdash sv_1 : \mathbf{Int} \\
&\Delta; \Gamma \vdash sv_2 : c \\
&\Delta; \Gamma \vdash sv_1 : \mathbf{Int} \rightsquigarrow sv'_1 \\
&\Delta; \Gamma \vdash sv_2 : c \rightsquigarrow sv'_2
\end{aligned}$$

By theorem 4:

$$|\Delta| \vdash |c| : T_{32}$$

By theorem 5:

$$\begin{aligned}
&\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv'_1 : |\mathbf{Int}| \\
&\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv'_2 : |c|
\end{aligned}$$

By theorem 3:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash |c| : R(|c|)$$

By lemma 22:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \mathbf{array} : \forall[\alpha : T_{\text{mil}}] \mathbf{Code}[R(\alpha), \mathbf{Int}, \text{interp}(\alpha)][](\mathbf{Array}(\alpha))$$

Hence by the code call rule :

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \mathbf{call\ array}[|c|](|c|, sv'_1, sv'_2) : \mathbf{Array}(\text{interp } |c|) \mathbf{opr}_{32}$$

18. If $i = \mathbf{fupd}(sv_1, sv_2, fv)$ then:

By inverting assumptions:

$$\begin{aligned}
&\Delta; \Gamma \vdash sv_1 : \mathbf{Farray} \\
&\Delta; \Gamma \vdash sv_2 : \mathbf{Int} \\
&\Delta; \Gamma \vdash fv : \\
&\Delta; \Gamma \vdash sv_1 : \mathbf{Farray} \rightsquigarrow sv'_1 \\
&\Delta; \Gamma \vdash sv_2 : \mathbf{Int} \rightsquigarrow sv'_2 \\
&\Delta; \Gamma \vdash fv : \mathbf{Float} \rightsquigarrow fv'
\end{aligned}$$

By theorems 5 and 6:

$$\begin{aligned}
&\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv'_1 : |\mathbf{Farray}| \\
&\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv'_2 : |\mathbf{Int}| \\
&\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash fv' : |\mathbf{Float}|
\end{aligned}$$

By the 64 bit array update rule:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash \mathbf{upd}_{\mathbf{Float}}(sv'_1, sv'_2, fv') : \mathbf{Unit} \mathbf{opr}_{32}$$

19. If $i = \mathbf{array}(sv, fv)$ then:

By inverting assumptions:

$$\begin{aligned}
&\Delta; \Gamma \vdash sv : \mathbf{Int} \\
&\Delta; \Gamma \vdash fv : \\
&\Delta; \Gamma \vdash sv : \mathbf{Int} \rightsquigarrow sv' \\
&\Delta; \Gamma \vdash fv : \mathbf{Float} \rightsquigarrow fv'
\end{aligned}$$

By theorems 5 and 6:

$$\begin{aligned}
&\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv' : |\mathbf{Int}| \\
&\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash fv' : |\mathbf{Float}|
\end{aligned}$$

By the 64 bit array creation rule:

$$|\Delta|; |\Delta|, |\Gamma| \vdash \mathbf{upd}_{\mathbf{Float}}(sv', fv') : \mathbf{Farray}(\mathbf{Float}) \mathbf{opr}_{32}$$

For expressions

1. If $e = sv$ then:

By inversion of assumptions:

$$\begin{aligned}
&\Delta; \Gamma \vdash sv : \tau \\
&\Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv'
\end{aligned}$$

By theorem 5:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv' : |\tau|$$

By the small value expression inclusion rule:

$$\Psi_i; |\Delta|; |\Delta|, |\Gamma| \vdash sv' : |\tau| \mathbf{exp}$$

2. If $e = \mathbf{let}_\tau x = i \mathbf{in} e$ then:

By inversion of assumptions:

$$\begin{aligned}
&\Delta; \Gamma \vdash i : \tau_i \\
&\Delta; \Gamma, x:\tau_i \vdash e : \tau \\
&\Delta; \Gamma \vdash i : \tau_i \rightsquigarrow i' \\
&\Delta; \Gamma, x:\tau_i \vdash e : \tau \rightsquigarrow e'
\end{aligned}$$

By induction:

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash i' : |\tau_i| \text{ opr}_{32}$$

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma|, x : |\tau_i| \vdash e' : |\tau| \text{ exp} \quad (\text{Note that } |\Gamma, x : \tau_i| = |\Gamma|, x : |\tau_i|)$$

By the operation rule:

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash \text{let}_{|\tau|} x = i' \text{ in } e' : |\tau| \text{ exp}$$

3. If $e = \text{let}_{\tau} x_f = i \text{ in } e$ then:

By inversion of assumptions:

$$\Delta; \Gamma \vdash i : \text{Float}$$

$$\Delta; \Gamma, x_f \vdash e : \tau$$

$$\Delta; \Gamma \vdash i : \text{Float} \rightsquigarrow i'$$

$$\Delta; \Gamma, x_f \vdash e : \tau \rightsquigarrow e'$$

By induction and theorem 7:

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash i' : \text{Float} \text{ opr}_{64}$$

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma|, x_f : \text{Float} \vdash e' : |\tau| \text{ exp} \quad (\text{Note that } |\Gamma, x_f| = |\Gamma|, x_f : \text{Float})$$

By the 64 bit operation rule:

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash \text{let}_{|\tau|} x_f = i' \text{ in } e' : |\tau| \text{ exp}$$

4. if $e = \text{let}_{\tau} \text{rec}_{\tau_r} f[\alpha::\kappa](x::\tau)(\vec{x}_f).e_f \text{ in } e$ then:

By inversion of assumptions:

$$\Delta, \alpha::\kappa; \Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r, x::\tau, \vec{x}_f] \vdash e_f : \tau_r$$

$$\Delta; \Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r] \vdash e : \tau$$

$$\Delta, \alpha::\kappa; \Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r, x::\tau, \vec{x}_f] \vdash e_f : \tau_r \rightsquigarrow e'_f$$

$$\Delta; \Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r] \vdash e : \tau \rightsquigarrow e'$$

By induction:

$$\Psi_i; |\Delta, \alpha::\kappa|; \downarrow \Delta, \alpha::\kappa \downarrow, |\Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r, x::\tau, \vec{x}_f]| \vdash e'_f : |\tau_r| \text{ exp}$$

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r]| \vdash e' : |\tau| \text{ exp}$$

By definition:

$$\downarrow \Delta, \alpha::\kappa \downarrow = \downarrow \Delta \downarrow, x_{\alpha} : \vec{\alpha}::\kappa \downarrow$$

And

By definition:

$$|\Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r, x::\tau, \vec{x}_f]| =$$

$$|\Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r, x : |\tau|, x_f : \vec{\text{Float}}]| =$$

$$|\Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(\downarrow \alpha::\kappa \downarrow, |\vec{\tau}|)(\vec{\text{Float}}) \rightarrow |\tau_r|, x : |\tau|, \overrightarrow{x_f : \text{Float}}]|$$

Hence by the rec rule:

$$\Psi_i; |\Delta|; \downarrow \Delta \downarrow, |\Gamma| \vdash$$

$$\text{let}_{|\tau|} \text{rec}_{|\tau_r|} f[\alpha::\kappa](\overrightarrow{x_{\alpha} : \downarrow \alpha::\kappa \downarrow}, \overrightarrow{x::\tau})(\overrightarrow{x_f : \text{Float}}).e'_f \text{ in } e' : |\tau|$$

■

Programs

The correctness of the program translation follows almost immediately.

Theorem 9 (Programs)

If $\bullet; \bullet \vdash e : \tau$ and $\vdash e : \tau \rightsquigarrow \text{let } d_i \text{ in } e' \text{ prog}$ then $\bullet \vdash \text{let } d_i \text{ in } e : |\tau|$.

Proof: (By construction)

By lemma 21 and the heap well-formedness rule:

$$\Psi_i \vdash d_i \text{ ok}$$

By theorem 8

$$\Psi_i; \bullet; \bullet \vdash e' : |\tau| \text{ exp}$$

So by the program well-formedness rule

$$\bullet \vdash \text{let } d_i \text{ in } e : |\tau|$$

■

5.6 The complete term level translation

Small Values

$$\boxed{\Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv'}$$

$$\frac{}{\Delta; \Gamma[x:\tau] \vdash x : \tau \rightsquigarrow x}$$

$$\frac{}{\Delta; \Gamma \vdash i : \text{Int} \rightsquigarrow i}$$

$$\frac{\begin{array}{l} \Delta \vdash c \equiv \pi_i \mu(\alpha, \beta)(c_1, c_2) : T_{32} \\ \Delta \vdash sv : \pi_i \mu(\alpha, \beta)(c_1, c_2) \rightsquigarrow sv' \end{array}}{\Delta; \Gamma \vdash \text{unroll}_c sv : c_i[\pi_1 \mu(\alpha, \beta)(c_1, c_2)/\alpha, \pi_2 \mu(\alpha, \beta)(c_1, c_2)/\beta] \rightsquigarrow \text{unroll}_{\text{interp}|c|} sv'}$$

$$\frac{\begin{array}{l} \Delta \vdash c \equiv \pi_i \mu(\alpha, \beta)(c_1, c_2) : T_{32} \\ \Delta; \Gamma \vdash sv : c_i[\pi_1 \mu(\alpha, \beta)(c_1, c_2)/\alpha, \pi_2 \mu(\alpha, \beta)(c_1, c_2)/\beta] \rightsquigarrow sv' \end{array}}{\Delta; \Gamma \vdash \text{roll}_c sv : \pi_i \mu(\alpha, \beta)(c_1, c_2) \rightsquigarrow \text{roll}_{|c|} sv'}$$

$$\frac{}{\Delta; \Gamma \vdash \text{inj_tag}_{\text{Sum}_i(\vec{c})}^j : \text{Sum}_i^j(\vec{c}) \rightsquigarrow \text{inj_union}_{|\text{Sum}_i(\vec{c})|} \text{tag}_j}$$

Float values

$$\boxed{\Delta \vdash fv : \text{Float} \rightsquigarrow fv'}$$

$$\frac{}{\Delta; \Gamma \vdash r : \text{Float} \rightsquigarrow r}$$

$$\frac{}{\Delta; \Gamma[x_f] \vdash x_f : \text{Float} \rightsquigarrow x_f}$$

$$\frac{\Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv'}{\Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv'}$$

$$\frac{\Delta; \Gamma \vdash sv : \text{Boxf} \rightsquigarrow sv'}{\Delta; \Gamma \vdash \text{unboxf } sv : \text{Float} \rightsquigarrow \text{unbox } sv'}$$

$$\frac{\Delta; \Gamma \vdash sv_1 : \text{Farray} \rightsquigarrow sv'_1 \quad \Delta; \Gamma \vdash sv_2 : \text{Int} \rightsquigarrow sv'_2}{\Delta; \Gamma \vdash \text{fsub}(sv_1, sv_2) : \text{Float} \rightsquigarrow \text{sub}_{\text{Float}}(sv'_1, sv'_2)}$$

Operations

$\Delta; \Gamma \vdash i : \tau \rightsquigarrow i'$

$$\frac{\Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv'}{\Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv'}$$

$$\frac{\Delta; \Gamma \vdash sv : c_1 \rightarrow c_2 \rightsquigarrow sv'}{\Delta; \Gamma \vdash \text{vararg}_{c_1 \rightarrow c_2} sv : \text{Vararg}_{c_1 \rightarrow c_2} \rightsquigarrow \text{call vararg}[|c_1|, \text{interp } |c_2|](\lfloor c_1 \rfloor, sv')}$$

$$\frac{\Delta; \Gamma \vdash sv : \text{Vararg}_{c_1 \rightarrow c_2} \rightsquigarrow sv'}{\Delta; \Gamma \vdash \text{onearg}_{c_1 \rightarrow c_2} sv : c_1 \rightarrow c_2 \rightsquigarrow \text{call onearg}[|c_1|, \text{interp } |c_2|](\lfloor c_1 \rfloor, sv')}$$

$$\frac{\Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv'}{\Delta; \Gamma \vdash \text{boxf } fv : \text{Boxf} \rightsquigarrow \text{box } fv'}$$

$$\frac{\Delta; \Gamma \vdash sv : \text{Sum}_i^j(\vec{c}) \rightsquigarrow sv'}{\Delta; \Gamma \vdash \text{proj}_{\text{Sum}_i(\vec{c})}^j sv : c_j \rightsquigarrow \text{select}^2 sv'}$$

$$\frac{\Delta; \Gamma \vdash sv : c_j \rightsquigarrow sv'}{\Delta; \Gamma \vdash \text{inj}_{\text{Sum}_i(\vec{c})}^j sv : \text{Sum}_i^j(\vec{c}) \rightsquigarrow \text{inj_union}_{|\text{Sum}_i(\vec{c})|} \langle \text{tag}_j, sv' \rangle}$$

$$\begin{array}{c}
\frac{\Delta; \Gamma \vdash sv_i : \tau_i \rightsquigarrow sv'_i}{\Delta; \Gamma \vdash \langle sv_1, \dots, sv_n \rangle : \tau_1 \times \dots \times \tau_n \rightsquigarrow \langle sv'_1, \dots, sv'_n \rangle} \\
\\
\frac{\Delta; \Gamma \vdash sv : \tau_1 \times \dots \times \tau_n \rightsquigarrow sv'}{\Delta; \Gamma \vdash \mathbf{select}^i sv : \tau_i \rightsquigarrow \mathbf{select}^i sv'} \\
\\
\frac{\Delta; \Gamma \vdash sv : \mathbf{Sum}_i(\vec{c}) \rightsquigarrow sv' \quad \Delta; \Gamma[x_j : \mathbf{Sum}_i^j(\vec{c})] \vdash e_j : \tau \rightsquigarrow e'_j}{\Delta; \Gamma \vdash \mathbf{case}_\tau(sv)(x_1.e_1, \dots, x_n.e_n) : \tau \rightsquigarrow \mathbf{case}_{|\tau|}(sv')(x_1.e'_1, \dots, x_n.e'_n)} \\
\\
\frac{\Delta; \Gamma \vdash e_1 : \tau \rightsquigarrow e'_1 \quad \Delta; \Gamma[x : \mathbf{Exn}] \vdash e_2 : \tau \rightsquigarrow e'_2}{\Delta; \Gamma \vdash \mathbf{handle}_\tau(e_1, x.e_2) : \tau \rightsquigarrow \mathbf{handle}_{|\tau|}(e'_1, x.e'_2)} \\
\\
\frac{\Delta; \Gamma \vdash sv : \forall[\alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n](\tau_1, \dots, \tau_m)(k) \rightarrow \tau \rightsquigarrow sv' \quad \Delta; \Gamma \vdash sv_i : \tau_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \rightsquigarrow sv'_i \quad \Delta; \Gamma \vdash fv_i : \mathbf{Float} \rightsquigarrow fv'_i}{\Delta; \Gamma \vdash sv[c_1, \dots, c_n](sv_1, \dots, sv_m)(fv_1, \dots, fv_k) : \tau[c_1/\alpha_1, \dots, c_n/\alpha_n] \rightsquigarrow sv'[[c_1], \dots, [c_n]](\lfloor c_1 \rfloor, \dots, \lfloor c_n \rfloor, sv'_1, \dots, sv'_m)(fv'_1, \dots, fv'_k)} \\
\\
\frac{\Delta; \Gamma \vdash sv : \mathbf{Exn} \rightsquigarrow sv'}{\Delta; \Gamma \vdash \mathbf{raise}_\tau sv : \tau \rightsquigarrow \mathbf{raise}_{|\tau|} sv'} \\
\\
\frac{\Delta; \Gamma \vdash sv_1 : \mathbf{Dyntag}_c \rightsquigarrow sv'_1 \quad \Delta; \Gamma \vdash sv_2 : c \rightsquigarrow sv'_2}{\Delta; \Gamma \vdash \mathbf{inj_dyn}_c(sv_1, sv_2) : \mathbf{Exn} \rightsquigarrow \mathbf{pack}\langle sv'_1, sv'_2 \rangle \mathbf{as} \exists \alpha : \mathbf{T}_{32}. (\mathbf{Dyntag}(\alpha) \times \alpha) \mathbf{hiding} |c|} \\
\\
\frac{}{\Delta; \Gamma \vdash \mathbf{mkexntag}_c : \mathbf{Dyntag}_c \rightsquigarrow \mathbf{dyntag}_{|c|}} \\
\\
\frac{\Delta; \Gamma \vdash sv_1 : \mathbf{Dyntag}_{c_1} \rightsquigarrow sv'_1 \quad \Delta; \Gamma[x_1 : c_1] \vdash e_1 : \tau \rightsquigarrow e'_1 \quad \Delta; \Gamma \vdash e_2 : \tau \rightsquigarrow e'_2}{\Delta; \Gamma \vdash \mathbf{exncase}_\tau(sv)(sv_1 \Rightarrow x_1.e_1, - \Rightarrow e_2) : \tau \rightsquigarrow \mathbf{dyncase}_\tau(sv)(sv'_1 \Rightarrow x_1.e'_1, - \Rightarrow e'_2)} \\
\\
\frac{\Delta; \Gamma \vdash sv_1 : \mathbf{Array}_c \rightsquigarrow sv'_1 \quad \Delta; \Gamma \vdash sv_2 : \mathbf{Int} \rightsquigarrow sv'_2}{\Delta; \Gamma \vdash \mathbf{sub}_c(sv_1, sv_2) : c \rightsquigarrow \mathbf{call sub}[|c|](\lfloor c \rfloor, sv'_1, sv'_2)}
\end{array}$$

$$\begin{array}{c}
\Delta; \Gamma \vdash sv_1 : \text{Array}_c \rightsquigarrow sv'_1 \quad \Delta; \Gamma \vdash sv_2 : \text{Int} \rightsquigarrow sv'_2 \\
\Delta; \Gamma \vdash sv_3 : c \rightsquigarrow sv'_3 \\
\hline
\Delta; \Gamma \vdash \text{upd}_c(sv_1, sv_2, sv_3) : \text{Unit} \rightsquigarrow \text{call upd}[|c|](|c|, sv'_1, sv'_2, sv'_3)
\end{array}$$

$$\begin{array}{c}
\Delta; \Gamma \vdash sv_1 : \text{Farray} \rightsquigarrow sv'_1 \quad \Delta; \Gamma \vdash sv_2 : \text{Int} \rightsquigarrow sv'_2 \\
\Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv' \\
\hline
\Delta; \Gamma \vdash \text{fupd}(sv_1, sv_2, fv) : \text{Unit} \rightsquigarrow \text{upd}_{\text{Float}}(sv'_1, sv'_2, fv')
\end{array}$$

$$\begin{array}{c}
\Delta; \Gamma \vdash sv_1 : \text{Int} \rightsquigarrow sv'_1 \quad \Delta; \Gamma \vdash sv_2 : c \rightsquigarrow sv'_2 \\
\hline
\Delta; \Gamma \vdash \text{array}_c(sv_1, sv_2) : \text{Array}_c \rightsquigarrow \text{call array}[|c|](|c|, sv'_1, sv'_2)
\end{array}$$

$$\begin{array}{c}
\Delta; \Gamma \vdash sv : \text{Int} \rightsquigarrow sv' \quad \Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv' \\
\hline
\Delta; \Gamma \vdash \text{farray}(sv, fv) : \text{Farray} \rightsquigarrow \text{array}_{\text{Float}}(sv', fv')
\end{array}$$

Expressions

$\Delta; \Gamma \vdash e : \tau \rightsquigarrow e'$
--

$$\begin{array}{c}
\Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv' \\
\hline
\Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv'
\end{array}$$

$$\begin{array}{c}
\Delta; \Gamma \vdash i : \tau' \rightsquigarrow i' \quad \Delta; \Gamma[x : \tau'] \vdash e : \tau \rightsquigarrow e' \\
\hline
\Delta; \Gamma \vdash \text{let}_\tau x = i \text{ in } e : \tau \rightsquigarrow \text{let}_{|\tau|} x = i' \text{ in } e'
\end{array}$$

$$\begin{array}{c}
\Delta; \Gamma \vdash i : \text{Float} \rightsquigarrow i' \quad \Delta; \Gamma[x_f] \vdash e : \tau \rightsquigarrow e' \\
\hline
\Delta; \Gamma \vdash \text{let}_\tau x_f = i \text{ in } e : \tau \rightsquigarrow \text{let}_{|\tau|} x_f = i' \text{ in } e'
\end{array}$$

$$\begin{array}{c}
\Delta; \Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r, \alpha::\kappa, x::\vec{\tau}, \vec{x}_f] \vdash e_f : \tau_r \rightsquigarrow e'_f \\
\Delta; \Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r] \vdash e : \tau \rightsquigarrow e' \\
\hline
\Delta; \Gamma \vdash \text{let}_\tau \text{rec}_{\tau_r} f[\alpha::\kappa](x::\vec{\tau})(\vec{x}_f).e_f \text{ in } e : \tau \rightsquigarrow \\
\text{let}_{|\tau|} \text{rec}_{|\tau_r|} f[\alpha::\kappa](\overrightarrow{x_\alpha} : \overrightarrow{|\alpha::\kappa|}, \overrightarrow{x::\vec{\tau}})(\overrightarrow{x_f : \text{Float}}).e'_f \text{ in } e'
\end{array}$$

Chapter 6

Closure converted LIL

6.1 Introduction

Closure conversion is a common strategy for implementing languages with higher order functions. The essential idea is to replace terms that evaluate via substitution of values for free variables with terms that can be evaluated by referring to an environment providing definitions for free variables. Since the latter evaluation model matches more closely that of actual machines, closure conversion is probably the most standard technique used for compiling higher-order functions.

In the **LIL**, closure conversion is the process by which function definitions are replaced with existentially abstracted pairs of environments and code pointers; and by which function applications are replaced with projections and code calls. Since the **LIL** is a typed language, the closure conversion translation must maintain well-typedness of terms. This means, among other things, that in addition to turning free *term* variables into additional function arguments, the closure conversion translation must also turn free *constructor* variables into additional function constructor arguments. However, since the **LIL** language admits a type erasure interpretation, the type environments do not need to be represented at runtime, which simplifies the problem of typed closure conversion.

The fact that **LIL** types can be erased at runtime means that in the underlying assembly language semantics, polymorphic instantiation can be a value, whereas in the **MIL** a polymorphic instantiation corresponds to passing actual runtime data and hence can not be thought of as a value in the sense of requiring no computation (as opposed to being *valuable* in the sense of being freely duplicable). This fact is important since it allows closure conversion to instantiate polymorphic functions directly when closures are built, which simplifies the theoretical framework for typed closure conversion.

There is extensive literature on typed closure conversion [CWM98, MWCG97, MMH96, MH98], and the **LIL** closure conversion algorithm adds nothing essential to this previous work. In this chapter I will briefly describe the important translation steps and prove a soundness theorem without going into any great detail.

6.2 The closure conversion translation

There are a number of different strategies possible for implementing closure conversion for recursive functions. Morrisett and Harper [MH98] give a fairly comprehensive overview of the different

approaches, and in section 9.3 I discuss some of the trade-offs between the different approaches. The translation as described here uses the “recursive code” translation, which implements closures via simple existentials and pairs, but requires a primitive notion of self-calling code. Fortunately (by design), the **LIL** language supports all of these notions.

In the next sections, I give a high-level overview of each part of the translation, and list a few key translation rules. In general, the closure conversion translation rewrites terms by first rewriting their sub-terms, and then reconstructing the term intact. Only in the specific cases of functions, function types, and function applications is a change made to the immediate structure of a term or type.

6.2.1 Constructors

Closure conversion does not change kinds in any way, so the first translation I discuss is the translation of constructors. I notate the translation of a **LIL** constructor c under closure-conversion as $|c|$. Almost every case of the constructor translation is uninteresting except for the case for a polymorphic function type.

$$\begin{aligned} & |\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \rightarrow (c_1)(c_2)(c_3)| \stackrel{\text{def}}{=} \\ & \exists[\alpha_e:\text{T}_{32}]. \times [\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \text{Code}(\alpha_e :: |c_1|)(|c_2|)(|c_3|), \alpha_e] \end{aligned}$$

This case is actually quite illuminating since it clearly demonstrates the closure conversion process, and is therefore worth exploring more closely.

Recall the type of the arrow constructor in the **LIL**: $\rightarrow:\text{T}_{32}\text{list} \rightarrow \text{T}_{64}\text{list} \rightarrow \text{T}_{32} \rightarrow \text{T}_{32}$. The types c_1 , c_2 , and c_3 are therefore the 32 bit argument type list, the 64 bit argument type list, and the return type, respectively. The translation of the arrow type simply translates these types compositionally, and then adds an additional type to the head of the 32 bit argument type list. This type, α_e , is the type of the environment containing the free variables of the function that are to be passed to it as an additional argument. I use the informal $::$ notation to mean addition to the head of a list.

The type $\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n] \text{Code}(\alpha_e :: |c_1|)(|c_2|)(|c_3|)$ therefore describes a polymorphic code function which in addition to the 32 bit arguments described by $|c_1|$, expects a first argument described by α_e . Referring to this type as τ_c , the type $\times[\tau_c, \alpha_e]$ classifies a function of this type paired with an actual environment of the same type: α_e .

What then is α_e ? The important point of the closure conversion translation in a typed setting is that the types of the free variables of a function are not apparent from its type. The environment type is *abstract* outside of the local scope where the function is defined. All that need be apparent from the type is that the type of the environment expected by the function and the type of the environment provided in the closure coincide.

In order to capture this abstraction, the standard technique is to use existential quantification. This then is the last part of the translation process. The final type is produced by existentially quantifying over the environment type, α_e .

6.2.2 Typing contexts

The translation on typing contexts Ψ and Γ is defined in the obvious way by mapping the type translation across the ranges of the contexts.

6.2.3 Terms

Term translations		
Small values	$\Psi; \Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv'$	$ \Psi ; \Delta; \Gamma \vdash sv' : \tau $
Float values	$\Psi; \Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv'$	$ \Psi ; \Delta; \Gamma \vdash fv' : \text{Float}$
Operations	$\Psi; \Delta; \Gamma \vdash opr : \tau \rightsquigarrow opr'$	$ \Psi ; \Delta; \Gamma \vdash opr' : \tau \text{ opr}_{32}$
64 bit Operations	$\Psi; \Delta; \Gamma \vdash fopr : \text{Float} \rightsquigarrow fopr'$	$ \Psi ; \Delta; \Gamma \vdash fopr' : \text{Float opr}_{64}$
Expression	$\Psi; \Delta; \Gamma \vdash e : \tau \rightsquigarrow d \text{ in } e'$	$ \Psi , \Psi(d); \bullet; \bullet \vdash d \text{ ok}$ $ \Psi , \Psi(d); \Delta; \Gamma \vdash e' : \tau \text{ exp}$
Programs	$\Psi \vdash p : \tau \rightsquigarrow p'$	$\Psi \vdash p' : \tau $

The high-level shape of the term level translation should be apparent from the discussion of the type translation. The essential task of the translation is to replace functions with existentially quantified code/environment pairs, and to replace function calls with closure-projections and code calls.

The term level translation is a typed translation, since the types of free variables is needed to construct the appropriate environment types. Additionally, it is very important in the **LIL** that the substitutions associated with type refinement operations be carried out as part of the translation, since these change the set of free type variables of expressions (as well as affecting equalities between them).

The general form of the closure conversion translation therefore exactly parallels the typing rules for the **LIL**. In addition to checking the type of terms however, the closure conversion relation also “produces” a new closure converted term. Moreover, in the case of the expression translation, a new set of heap values is also produced containing the new code replacing the functions from the original term.

The interesting cases from the translation arise from application and function abstraction. In order to simplify the syntactic structure of the translation, it is convenient to catch applications in the expression translation rather than in the more obvious operation translation. Consequently, the operation translation has no case for application.

Function application

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash g : \forall[\alpha_1::\kappa_1, \dots, \alpha_n::\kappa_n](\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n) \rightarrow \tau \rightsquigarrow sv' \\
\Psi; \Delta; \Gamma \vdash sv_i : \tau_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \rightsquigarrow sv'_i \quad \Psi; \Delta; \Gamma \vdash fv_i : \phi_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \rightsquigarrow fv'_i \\
\Psi; \Delta; \Gamma, x : \tau[c_1/\alpha_1, \dots, c_n/\alpha_n] \vdash e : \tau' \rightsquigarrow d \text{ in } e' \\
\hline
\Psi; \Delta; \Gamma \vdash \text{let } x = g[c_1, \dots, c_n](sv_1, \dots, sv_m)(fv_1, \dots, fv_k) \text{ in } e : \tau' \rightsquigarrow \\
d \text{ in} \\
\text{let } [\alpha_e, x_c] = \text{unpack } sv' \text{ in} \\
\text{let } f = \text{select}^1 x_c \text{ in} \\
\text{let } x_e = \text{select}^2 x_c \text{ in} \\
\text{let } x = \text{call } f[|c_1|, \dots, |c_n|](x_e, sv'_1, \dots, sv'_m)(fv'_1, \dots, fv'_k) \text{ in } e'
\end{array}$$

Since function definitions are replaced with closure definitions by the translation, the variable g will be bound to a closure after the translation. The new term replacing the application first unpacks the closure to get at the underlying code/environment pair, and then selects out the components. The code pointer thus extracted is then passed the environment as an argument along with all of the original arguments.

Function abstraction

The translation of function definitions is by far the most complicated part of closure conversion. The complete translation rule is given in figure 6.1. Conceptually, it can be separated into two steps: producing code to construct the environment and the closure, and wrapping the function body in code to extract the components of the environment.

Creating the closure

The environment must contain all of the free variables of the function: that is, all of the free variables of the function body except the function arguments and the function itself. Since the **LIL** as defined lacks heterogeneous tuples, the 64 bit free variables must be boxed before they can be placed in the environment. Consequently, the closure creation code begins by boxing up all of the free 64 bit variables, then creates a record consisting of all the 32 bit variables followed by the boxed 64 bit variables. This record is the function’s environment.

In addition to closing up the function over the free term variables via the environment, it is also necessary to abstract away all of the free type variables. These are “packaged” up into the closure by immediately instantiating the label for the new code function with the free type variables from the original function. Since the type parameters do not actually exist at runtime, there is no need to store them separately in the closure via a kind existential.

To create the final closure value, this partially instantiated label is written into a two field record with the environment which is then packed into the existential type, hiding the type of the environment. The original function name is bound to the resulting package.

Restoring free variables from the environment

The code from the previous section, executed at the site of the original function definition, serves to create the closure which is then passed around in place of the function. A second code sequence must also be added as a prelude to the function body to restore the free variables of the function from the environment. This code can be seen in the **fcode** definition from the function translation rule.

The initial segment of the un-packaging code simply projects out the 32 bit variables from the environment. The second segment projects out and un-boxes the boxed 64 bit variables from the environment. The final segment creates a closure to bind the function name to for internal uses. This consists of allocating a tuple to hold the code pointer and the environment, and packing it into the existential. The necessity of recreating the closure tuple inside of recursive functions is the principal disadvantage of the “recursive-code” approach to closure conversion. Note though that this is only necessary for *escaping* occurrences of the recursion variable: non-escaping occurrences can be replaced by uses of the code pointer and the original environment. This will be discussed in more detail in section 9.3.

6.2.4 Heap Values and Programs

Code functions are already closed, and hence are not directly closure-converted. However, in the **LIL** as defined, they may contain normal functions as sub-terms. Therefore, the bodies of code functions are translated just as any other expression.

$$\begin{array}{l}
\Delta \vdash \kappa_i \text{ ok} \quad (i \in 1 \dots k) \\
\Delta, \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k \vdash \tau_i : T_{32} \quad (i \in 1 \dots m) \\
\Delta, \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k \vdash \phi_i : T_{64} \quad (i \in 1 \dots n) \\
\Delta, \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k \vdash \tau : T_{32} \\
\Delta \vdash \tau_r \equiv \forall[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k](\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n) \rightarrow \tau : T_{32} \\
\Psi; \Delta, \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k; \Gamma, f:\tau, x_1:\tau_1, \dots, x_m:\tau_m, z_1:\phi_1, \dots, z_n:\phi_n \vdash e_f : \tau \rightsquigarrow d \text{ in } e'_f \\
\Psi; \Delta; \Gamma[f : \tau_r] \vdash e : \tau \rightsquigarrow d' \text{ in } e' \\
fv_{32}(e_f) \setminus \{x_1, \dots, x_m, f\} = \{y_1, \dots, y_q\} \quad fv_{64}(e_f) \setminus \{z_1, \dots, z_n\} = \{r_1, \dots, r_p\} \\
fv_t(e_f) \setminus \{\alpha_1, \dots, \alpha_k\} = \{\beta_1, \dots, \beta_n\} \\
\tau_e \stackrel{\text{def}}{=} \times[|\Gamma(y_1)|, \dots, |\Gamma(y_q)|, \text{Boxed } |\Gamma(r_1)|, \dots, \text{Boxed } |\Gamma(r_p)|] \\
\mathbf{fcode} \stackrel{\text{def}}{=} \mathbf{code}[\beta_1:\Delta(\beta_1), \dots, \beta_l:\Delta(\beta_l), \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k] \\
(x_e:\tau_e, x_1:|\tau_1|, \dots, x_m:|\tau_m|) \\
(z_1:|\phi_1|, \dots, z_n:|\phi_n|). \\
\text{let } y_1 = \text{select}^1 x_e \text{ in} \\
\dots \\
\text{let } y_q = \text{select}^q x_e \text{ in} \\
\text{let } b_1 = \text{select}^{q+1} x_e \text{ in} \\
\text{let } r_1 = \text{unbox } b_1 \text{ in} \\
\dots \\
\text{let } b_p = \text{select}^{q+p+1} x_e \text{ in} \\
\text{let } r_p = \text{unbox } b_p \text{ in} \\
\text{let } f_c = \langle \mathbf{fcode}[\beta_1, \dots, \beta_l], x_e \rangle \text{ in} \\
\text{let } f = \text{pack } f_c \text{ as } |\tau| \text{ hiding } \tau_e \text{ in} \\
e'_f
\end{array}$$

$$\begin{array}{l}
\Psi; \Delta; \Gamma \vdash \text{let}_\tau \text{rec}_{\tau_r} f[\alpha::\kappa](x::\tau)(\vec{x}_f).e_f \text{ in } e : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau \rightsquigarrow \\
d, \mathbf{fcode}:|\tau_r| \mapsto \mathbf{fcode}, d' \\
\text{let } b_1 = \text{box } r_1 \text{ in} \\
\dots \\
\text{let } b_p = \text{box } r_p \text{ in} \\
\text{let } x_e = \langle y_1, \dots, y_q, b_1, \dots, b_p \rangle \text{ in} \\
\text{let } f_c = \langle \mathbf{fcode}[\beta_1, \dots, \beta_l], x_e \rangle \text{ in} \\
\text{let } f = \text{pack } f_c \text{ as } |\tau| \text{ hiding } \tau_e \text{ in} \\
e'
\end{array}$$

Figure 6.1: Closure converting recursive function definitions

$$\frac{\Psi; \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k; x_1:\tau_1, \dots, x_m:\tau_m, z_1:\phi_1, \dots, z_n:\phi_n \vdash e:\tau \rightsquigarrow d \text{ in } e'}{\Psi \vdash \text{code}[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k](x_1:\tau_1, \dots, x_m:\tau_m)(z_1:\phi_1, \dots, z_n:\phi_n).e:} \\
\forall[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k] \text{Code}(\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n)(\tau) \rightsquigarrow \\
d \text{ in code}_{|\tau|}[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k](x_1:|\tau_1|, \dots, x_m:|\tau_m|)(z_1:|\phi_1|, \dots, z_n:|\phi_n|).e'$$

Heaps are translated by translating the individual heap values, lifting out any new heap values to the top level.

$$\frac{\Psi \vdash \epsilon \rightsquigarrow \epsilon \text{ in } \epsilon}{\Psi \vdash \epsilon \rightsquigarrow \epsilon \text{ in } \epsilon} \\
\frac{\Psi[\ell:\tau] \vdash hval:\tau \rightsquigarrow d_\ell \text{ in } hval' \quad \Psi[\ell:\tau] \vdash d \rightsquigarrow d_h \text{ in } d'}{\Psi[\ell:\tau] \vdash d, \ell:\tau \mapsto hval \rightsquigarrow d_\ell, d_h \text{ in } d', \ell:|\tau| \mapsto hval'}$$

The translation on programs rewrites the heap and the expression body to form a new program.

$$\frac{\Psi, \Psi(d) \vdash d \rightsquigarrow d_h \text{ in } d' \quad \Psi, \Psi(d); \bullet; \bullet \vdash e:\tau \rightsquigarrow d_e \text{ in } e'}{\Psi \vdash \text{letrec } d \text{ in } e:\tau \rightsquigarrow \text{letrec } d', d_h, d_e \text{ in } e'}$$

6.3 Soundness of closure conversion

The proof of type preservation for the closure conversion translation is fairly straightforward, as only a few sub-terms are actually changed. In general, most of the cases follow directly by induction, reconstructing the original term or derivation from the inductively re-written sub-terms or sub-derivations. Only in the few cases where the immediate term is re-written does the proof involve anything more substantial.

6.3.1 Constructors

As with the **MIL** to **LIL** translation, I begin by proving the soundness of the constructor translation, and then proceed with auxiliary lemmas showing that the constructor translation commutes with substitution, and respects equivalence. The latter two lemmas are useful for the term level proofs.

Theorem 10 (Soundness of the constructor translation)

If $\Delta \vdash c:\kappa$ then $\Delta \vdash |c|:\kappa$.

Proof: By induction on the structure of types. For all types and constructors except the type of polymorphic functions, the translation leaves the structure of the type the same, and so the proof follows straightforwardly by induction, producing the new derivation from the inductively obtained sub-derivations using the same rule as in the original derivation. The only interesting rule is the rule for the type of polymorphic functions.

Suppose the type being translated is of the form $\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \rightarrow (c_1)(c_2)(c_3)$.

By inversion of the assumed derivation:

$$\begin{aligned} \Delta \vdash \kappa_i \text{ ok} \quad i \in 1 \dots n \\ \Delta, \alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n \vdash c_1 : \mathsf{T}_{32}\mathsf{list} \\ \Delta, \alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n \vdash c_2 : \mathsf{T}_{32}\mathsf{list} \\ \Delta, \alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n \vdash c_3 : \mathsf{T}_{32} \end{aligned}$$

By induction:

$$\begin{aligned} \Delta, \alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n \vdash |c_1| : \mathsf{T}_{32}\mathsf{list} \\ \Delta, \alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n \vdash |c_2| : \mathsf{T}_{32}\mathsf{list} \\ \Delta, \alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n \vdash |c_3| : \mathsf{T}_{32} \end{aligned}$$

By weakening:

$$\begin{aligned} \Delta, \alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n, \alpha_e:\mathsf{T}_{32} \vdash |c_1| : \mathsf{T}_{32}\mathsf{list} \\ \Delta, \alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n, \alpha_e:\mathsf{T}_{32} \vdash |c_2| : \mathsf{T}_{32}\mathsf{list} \\ \Delta, \alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n, \alpha_e:\mathsf{T}_{32} \vdash |c_3| : \mathsf{T}_{32} \end{aligned}$$

(Note that α_e can always be chosen appropriately)

By construction (cons):

$$\Delta, \alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n, \alpha_e:\mathsf{T}_{32} \vdash (\alpha_e :: |c_1|) : \mathsf{T}_{32}\mathsf{list}$$

By construction (application and universal introduction):

$$\Delta, \alpha_e:\mathsf{T}_{32} \vdash \forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \mathsf{Code}(\alpha_e :: |c_1|)(|c_2|)(|c_3|) : \mathsf{T}_{32}\mathsf{list}$$

By construction (pairing and existential introduction):

$$\Delta \vdash \exists[\alpha_e:\mathsf{T}_{32}]. \times [\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \mathsf{Code}(\alpha_e :: |c_1|)(|c_2|)(|c_3|), \alpha_e] : \mathsf{T}_{32}$$

■

Lemma 23 (The type translation commutes with substitution)

$$|c[c'/\alpha]| = |c|[[c'/\alpha]]$$

Proof: By induction on c .

1. If c is a constant, then $\alpha \notin \mathit{fvc}, \mathit{fv}|c|$, so $|c|[[c'/\alpha]] = |c| = |c[c'/\alpha]|$.
2. If c is a variable α' :
 - (a) If $\alpha \neq \alpha'$ then as in the previous case.
 - (b) If $\alpha = \alpha'$ then $|\alpha|[[c'/\alpha]] = \alpha[[c'/\alpha]] = |c'| = |\alpha[c'/\alpha]|$.
3. If $c = \lambda(\beta:\kappa).c_2$, then by definition

$$|\lambda(\beta:\kappa).c_2|[[c'/\alpha]] = (\lambda(\beta:|\kappa|).|c_2|)[[c'/\alpha]] = \lambda(\beta:|\kappa|).(|c_2|[[c'/\alpha]])$$

By induction and the definition of substitution:

$$\lambda(\beta:|\kappa|).(|c_2|[[c'/\alpha]]) = \lambda(\beta:|\kappa|).(|c_2[c'/\alpha]|) = |\lambda(\beta:\kappa).(c_2[c'/\alpha])|$$

Note, I rely on alpha-variance to ensure non-capture.

4. If $c = \forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \rightarrow (c_1)(c_2)(c_3)$ then by the definition of the translation and of substitution

$$\begin{aligned} & |\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \rightarrow (c_1)(c_2)(c_3)|[[c'/\alpha]] \\ & \stackrel{\text{def}}{=} (\exists[\alpha_e:\mathsf{T}_{32}]. \times [\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \mathsf{Code}(\alpha_e :: |c_1|)(|c_2|)(|c_3|), \alpha_e])[[c'/\alpha]] \\ & \stackrel{\text{def}}{=} \exists[\alpha_e:\mathsf{T}_{32}]. \times [\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \mathsf{Code}(\alpha_e :: |c_1|[[c'/\alpha]])(|c_2|[[c'/\alpha]])(|c_3|[[c'/\alpha]]), \alpha_e] \end{aligned}$$

Note that I rely on alpha-variance to avoid capture.

Finally, by induction, and by the definition of substitution:

$$\begin{aligned} & \exists[\alpha_e : T_{32}]. \times [\forall[\alpha_1 : \kappa_1, \dots, \alpha_n : \kappa_n]. \mathbf{Code}(\alpha_e :: |c_1|[[c'/\alpha]])(|c_2|[[c'/\alpha]])(|c_3|[[c'/\alpha]]), \alpha_e] \\ & \stackrel{\text{def}}{=} \exists[\alpha_e : T_{32}]. \times [\forall[\alpha_1 : \kappa_1, \dots, \alpha_n : \kappa_n]. \mathbf{Code}(\alpha_e :: |c_1[c'/\alpha]])(|c_2[c'/\alpha]])(|c_3[c'/\alpha]]), \alpha_e] \\ & \stackrel{\text{def}}{=} |\forall[\alpha_1 : \kappa_1, \dots, \alpha_n : \kappa_n]. \rightarrow (c_1)(c_2)(c_3)[c'/\alpha]| \end{aligned}$$

5. If $c = c_1 c_2, \pi_i c, \langle c_1, c_2 \rangle$, the proof proceeds similarly. ■

Lemma 24 (The type translation respects equivalence)

If $\Delta \vdash c_1 \equiv c_2 : \kappa$ then $\Delta \vdash |c_1| \equiv |c_2| : \kappa$.

Proof: (By induction on equivalence derivations) By induction on equivalence derivations. The proof is straightforward with almost all cases identical to the proof of lemma 10 from chapter 5. The only cases where the derivation changes is in the case that the last rule used is the structural rule for constructor application $c_1 c_2$. Inductively, equivalence derivations exist for the sub-terms. If the applications match the form of a polymorphic function (that is, $\forall[\kappa_1](\dots(\forall[\kappa_n](\rightarrow(c_x)(c_z)(c_r))))$), then the application, reflexivity, pair, and existential equivalence rules must be applied to construct the equivalence derivation for the translation. If the applications do not match such a form, then the application rule suffices. All of the primitive types follow directly by reflexivity or by the structural application rule. I give some example cases here.

1. Suppose $\Delta \vdash c \equiv c : \kappa$ by reflexivity.

By assumption:

$$\Delta \vdash c : \kappa$$

By theorem 10:

$$\Delta \vdash |c| : \kappa$$

By reflexivity:

$$\Delta \vdash |c| \equiv |c| : \kappa$$

2. Suppose $\Delta \vdash (\lambda(\alpha : \kappa_1).c_1)(c_2) \equiv c_1[c_2/\alpha] : \kappa_2$.

By assumption:

$$\Delta \vdash \kappa_1 \mathbf{ok}$$

$$\Delta, \alpha : \kappa_1 \vdash c_1 : \kappa_2$$

$$\Delta \vdash c_2 : \kappa_1$$

By assumption, and by theorem 2:

$$\Delta \vdash \kappa_1 \mathbf{ok}$$

$$\Delta, \alpha : \kappa_1 \vdash |c_1| : \kappa_2$$

$$\Delta \vdash |c_2| : \kappa_1$$

By the λ beta rule and the definition of the translations:

$$\Delta \vdash |(\lambda(\alpha : \kappa_1).c_1)(c_2)| \equiv |c_1|[[c_2]/\alpha] : \kappa_2$$

Finally, by lemma 9:

$$\Delta \vdash |(\lambda(\alpha : \kappa_1).c_1)(c_2)| \equiv |c_1[c_2/\alpha]| : \kappa_2$$

3. Suppose $\Delta \vdash \forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \rightarrow (c_1)(c_2)(c_3) \equiv \forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \rightarrow (c'_1)(c'_2)(c'_3) : T_{32}$.

By induction and weakening

$$\Delta, \alpha_e : T_{32}, \alpha_1 : \kappa_1 \dots \alpha_n : \kappa_n \vdash c_1 \equiv c'_1 : T_{32} \text{list}$$

$$\Delta, \alpha_e : T_{32}, \alpha_1 : \kappa_1 \dots \alpha_n : \kappa_n \vdash c_2 \equiv c'_2 : T_{32} \text{list}$$

$$\Delta, \alpha_e : T_{32}, \alpha_1 : \kappa_1 \dots \alpha_n : \kappa_n \vdash c_3 \equiv c'_3 : T_{32}$$

By reflexivity

$$\Delta, \alpha_e : T_{32}, \alpha_1 : \kappa_1 \dots \alpha_n : \kappa_n \vdash \alpha_e \equiv \alpha_e : T_{32}$$

By the pair and existential rules

$$\begin{aligned} \Delta \vdash \exists[\alpha_e : T_{32}]. \times [\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \text{Code}(\alpha_e :: |c_1|)(|c_2|)(|c_3|), \alpha_e] \\ \equiv \exists[\alpha_e : T_{32}]. \times [\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \text{Code}(\alpha_e :: |c'_1|)(|c'_2|)(|c'_3|), \alpha_e] : T_{32} \end{aligned}$$

■

6.3.2 Typing contexts

Theorem 11 (Soundness of the context translations)

1. If $\Delta \vdash \Gamma \text{ ok}$ then $\Delta \vdash |\Gamma| \text{ ok}$.
2. If $\vdash \Psi \text{ ok}$ then $\vdash |\Psi| \text{ ok}$.

Proof: By induction on typing contexts. The proof proceeds by cases on contexts.

1. Suppose $\Delta \vdash \Gamma \text{ ok}$.

- If $\Gamma = \bullet$, then its translation is empty, and hence is trivially well-formed.
- If $\Gamma = \Gamma', x:\tau$, then by induction, $\Delta \vdash |\Gamma'| \text{ ok}$. By theorem 10, $\Delta \vdash |\tau| \text{ ok}$, and since the domain is unchanged by the translation, $x \notin \Gamma'$. Therefore, by construction $\Delta \vdash |\Gamma', x:\tau| \text{ ok}$.

2. The proof for heap contexts Ψ proceeds identically.

■

6.3.3 Terms

For terms, I begin by stating the soundness theorem for values and 64 bit operations. For these syntactic classes, the statement of the theorem is straightforward.

Theorem 12 (Values and 64 bit operations)

1. If $\Psi; \Delta; \Gamma \vdash sv : \tau$ and $\Psi; \Delta; \Gamma \vdash sv : \tau \rightsquigarrow sv'$ then $|\Psi|; \Delta; |\Gamma| \vdash sv' : |\tau|$.
2. If $\Psi; \Delta; \Gamma \vdash fv : \text{Float}$ and $\Psi; \Delta; \Gamma \vdash fv : \text{Float} \rightsquigarrow fv'$ then $|\Psi|; \Delta; |\Gamma| \vdash fv' : \text{Float}$.
3. If $\Psi; \Delta; \Gamma \vdash fopr : \text{Float opr}_{64}$ and $\Psi; \Delta; \Gamma \vdash fopr : \text{Float} \rightsquigarrow fopr'$ then

$$|\Psi|; \Delta; |\Gamma| \vdash fopr' : \text{Float opr}_{64}$$

Proof (Sketch): By induction on terms. The closure conversion translation for each of these syntactic classes leaves the top level structure of the term intact, reconstructing it from the inductively re-written sub-terms. Therefore, the proof constructs a new derivation in each case by using the same inference rule from the original derivation on the inductively re-written sub-terms. Kinds are unchanged by the translation, and for constructors I appeal to theorem 10. ■

The theorem for 32 bit operations and expressions is more complicated, since the expression translation potentially yields new heap bindings (code functions) to be lifted out. The soundness theorem for expressions states that both the bindings and the new expression are well-formed in the heap context obtained by extending the translation of the original heap context with the bindings for the additional heap values produced by the translation.

Theorem 13 (32 bit operations and expressions)

1. *If*

$$\Psi; \Delta; \Gamma \vdash opr : \tau \text{ opr}_{32} \text{ and } \Psi; \Delta; \Gamma \vdash opr : \tau \rightsquigarrow opr'$$
then

$$|\Psi|; \Delta; |\Gamma| \vdash opr' : |\tau| \text{ opr}_{32}$$
2. *If*

$$\Psi; \Delta; \Gamma \vdash e : \tau \text{ exp and } \Psi; \Delta; \Gamma \vdash e : \tau \rightsquigarrow d \text{ in } e'$$
then

$$|\Psi|, \Psi(d) \vdash d \text{ ok and } |\Psi|, \Psi(d); \Delta; |\Gamma| \vdash \text{in } e' : |\tau| \text{ exp}$$

Proof: By induction on terms.

1. The proof for operations proceeds exactly as in the previous theorem, appealing to the induction hypothesis to obtain well-typedness derivations for sub-terms, and constructing a new derivation with the same rule as in the original derivation. Note that the case for applications is trivially true, since the operation translation is not defined on applications.
2. The proof for expressions also proceeds similarly, except in the cases binding applications and functions where the translation does incremental work.

Suppose $\text{let } x = g[c_1, \dots, c_n](sv_1, \dots, sv_m)(fv_1, \dots, fv_k) \text{ in } e$

By assumption:

$$\begin{aligned} & \Psi; \Delta; \Gamma \vdash sv : \forall[\alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n](\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_k) \rightarrow \tau \\ & \Delta \vdash c_i : \kappa_i \quad i \in 1 \dots n \\ & \Psi; \Delta; \Gamma \vdash sv_i : \tau_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \quad i \in 1 \dots m \\ & \Psi; \Delta; \Gamma \vdash fv_i : \phi_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \quad i \in 1 \dots k \\ & \Psi; \Delta; \Gamma, x : \tau[c_1/\alpha_1, \dots, c_n/\alpha_n] \vdash e : \tau' \text{ exp} \\ & \Psi; \Delta; \Gamma \vdash g : \forall[\alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n](\tau_1, \dots, \tau_m)(k) \rightarrow \tau \rightsquigarrow sv' \\ & \Psi; \Delta; \Gamma \vdash sv_i : \tau_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \rightsquigarrow sv'_i \\ & \Psi; \Delta; \Gamma \vdash fv_i : \phi_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \rightsquigarrow fv'_i \\ & \Psi; \Delta; \Gamma, x : \tau[c_1/\alpha_1, \dots, c_n/\alpha_n] \vdash e : \tau' \rightsquigarrow d \text{ in } e' \end{aligned}$$

It suffices to show that:

$$\begin{aligned}
|\Psi|, \Psi(d); \Delta; |\Gamma| \vdash & \text{let}[\alpha_e, x_c] = \text{unpack } sv' \text{ in} & : |\tau'| \text{ exp} \\
& \text{let } f = \text{select}^1 x_c \text{ in} \\
& \text{let } x_e = \text{select}^2 x_c \text{ in} \\
& \text{let } x = \text{call } f[|c_1|, \dots, |c_n|](x_e, sv'_1, \dots, sv'_m)(fv'_1, \dots, fv'_k) \\
& \text{in } e'
\end{aligned}$$

Let

$$\tau_c = \times [\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \text{Code}(\alpha_e, |\tau_1|, \dots, |\tau_m|)(|\phi_1|, \dots, |\phi_k|)(|\tau|), \alpha_e]$$

and

$$\tau_f = \forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \text{Code}(\alpha_e, |\tau_1|, \dots, |\tau_m|)(|\phi_1|, \dots, |\phi_k|)(|\tau|)$$

.

So by the expression typing rules, it suffices to show:

$$\begin{aligned}
|\Psi|, \Psi(d); \Delta, \alpha_e:\text{T}_{32}; |\Gamma| \vdash & sv' : \exists[\alpha_e:\text{T}_{32}]. \tau_c \\
|\Psi|, \Psi(d); \Delta, \alpha_e:\text{T}_{32}; |\Gamma|, x_c:\tau_c \vdash & \text{select}^1 x_c : \tau_f \text{ opr}_{32} \\
|\Psi|, \Psi(d); \Delta, \alpha_e:\text{T}_{32}; |\Gamma|, x_c:\tau_c \vdash & \text{select}^2 x_c : \alpha_e \text{ opr}_{32} \\
|\Psi|, \Psi(d); \Delta, \alpha_e:\text{T}_{32}; |\Gamma|, x_c:\tau_c, f:\tau_f, x_e:\alpha_e \vdash & \text{call } f[\vec{c_i}](x_e, \vec{sv'_j})(\vec{fv'_l}) : |\tau|[\vec{c_i}/\alpha_i] \text{ opr}_{32} \\
& i \in 1 \dots n, j \in 1 \dots m, l \in 1 \dots k \\
|\Psi|, \Psi(d); \Delta, \alpha_e:\text{T}_{32}; |\Gamma|, x_c:\tau_c, f:\tau_f, x_e:\alpha_e, x:\tau \vdash & e' : |\tau'| \text{ exp}
\end{aligned}$$

By theorem 12:

$$|\Psi|, \Psi(d); \Delta; |\Gamma| \vdash sv : |\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n](\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_k) \rightarrow \tau|$$

And by definition:

$$\begin{aligned}
|\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n](\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_k) \rightarrow \tau| & \stackrel{\text{def}}{=} \\
\exists[\alpha_e:\text{T}_{32}]. \times [\forall[\alpha_1:\kappa_1, \dots, \alpha_n:\kappa_n]. \text{Code}(\alpha_e, |\tau_1|, \dots, |\tau_m|)(|\phi_1|, \dots, |\phi_k|)(|\tau|), \alpha_e]
\end{aligned}$$

So by the unpack rule:

$$|\Psi|, \Psi(d); \Delta, \alpha_e:\text{T}_{32}; |\Gamma| \vdash sv' : \exists[\alpha_e:\text{T}_{32}]. \tau_c \quad \text{For some } \tau_c$$

By the variable rule:

$$|\Psi|, \Psi(d); \Delta, \alpha_e:\text{T}_{32}; |\Gamma|, x_c:\tau_c \vdash x_c : \times [\tau_f, \alpha_e]$$

So by the select rule:

$$\begin{aligned}
|\Psi|, \Psi(d); \Delta, \alpha_e:\text{T}_{32}; |\Gamma|, x_c:\tau_c \vdash & \text{select}^1 x_c : \tau_f \text{ opr}_{32} \\
|\Psi|, \Psi(d); \Delta, \alpha_e:\text{T}_{32}; |\Gamma|, x_c:\tau_c \vdash & \text{select}^2 x_c : \alpha_e \text{ opr}_{32}
\end{aligned}$$

By the variable rule:

$$|\Psi|, \Psi(d); \Delta, \alpha_e:\text{T}_{32}; |\Gamma|, x_c:\tau_c, f:\tau_f, x_e:\alpha_e \vdash f : \tau_f$$

By theorems 10 and 12:

$$\begin{aligned}
\Delta \vdash |c_i| : \kappa_i \quad & i \in 1 \dots n \\
|\Psi|, \Psi(d); \Delta; |\Gamma| \vdash sv'_i : |\tau_i[c_1/\alpha_1, \dots, c_n/\alpha_n]| \quad & i \in 1 \dots m \\
|\Psi|, \Psi(d); \Delta; |\Gamma| \vdash fv'_i : |\phi_i[c_1/\alpha_1, \dots, c_n/\alpha_n]| \quad & i \in 1 \dots k
\end{aligned}$$

By the commutation lemma (lemma 23):

$$\begin{aligned}
|\Psi|, \Psi(d); \Delta; |\Gamma| \vdash sv'_i : |\tau_i[|c_1|/\alpha_1, \dots, |c_n|/\alpha_n]| \quad & i \in 1 \dots m \\
|\Psi|, \Psi(d); \Delta; |\Gamma| \vdash fv'_i : |\phi_i[|c_1|/\alpha_1, \dots, |c_n|/\alpha_n]| \quad & i \in 1 \dots k
\end{aligned}$$

Therefore by the **call** rule :

$$|\Psi|, \Psi(d); \Delta, \alpha_e : T_{32}; |\Gamma|, x_c : \tau_c, f : \tau_f, x_e : \alpha_e \vdash \mathbf{call} f[|\vec{c}_i|](x_e, \vec{sv}'_j)(\vec{fv}'_l) : |\tau|[\overline{|\vec{c}_i|/\alpha_i}] \mathbf{opr}_{32}$$

$$i \in 1 \dots n, j \in 1 \dots m, l \in 1 \dots k$$

Finally, by induction:

$$|\Psi|, \Psi(d); \Delta, \alpha_e : T_{32}; |\Gamma|, x_c : \tau_c, f : \tau_f, x_e : \alpha_e, x : |\tau|[\overline{|\vec{c}_1|/\alpha_1}, \dots, |\vec{c}_n|/\alpha_n] \vdash e' : |\tau'| \mathbf{exp}$$

3. The case for function abstraction proceeds in a similar fashion, building up well-typedness derivations for the **fcode** definitions and for the closure creation code.

■

An additional lemma relating translations of heap contexts to translations of heaps is needed for subsequent theorems.

Lemma 25 ($\Psi(d)$)

If $\Psi \vdash d \rightsquigarrow d_h \text{ in } d'$ then $|\Psi(d)| = \Psi(d')$.

Proof: By induction on d.

1. If $d = \epsilon$ then $d' = \epsilon$ and $|\Psi(d)| = \epsilon = \Psi(d')$.
2. If $d = d_r, \ell : \tau \mapsto hval$ then:

By assumption:

$$\Psi[\ell : \tau] \vdash hval : \tau \rightsquigarrow d_\ell \text{ in } hval'$$

$$\Psi[\ell : \tau] \vdash d_r \rightsquigarrow d_h \text{ in } d'_r$$

$$\Psi[\ell : \tau] \vdash d, \ell : \tau \mapsto hval \rightsquigarrow d_\ell, d_h \text{ in } d'_r, \ell : |\tau| \mapsto hval'$$

By definition:

$$|\Psi(d)|$$

$$= |\Psi(d_r, \ell : \tau \mapsto hval)|$$

$$= |\Psi(d_r), \ell : \tau|$$

$$= |\Psi(d_r)|, \ell : |\tau|$$

By induction:

$$|\Psi(d_r)| = \Psi(d'_r)$$

so

$$|\Psi(d_r)|, \ell : |\tau|$$

$$= \Psi(d'_r), \ell : |\tau|$$

$$= \Psi(d'_r, \ell : |\tau| \mapsto hval')$$

■

As with expressions, translating heap values may yield additional heap bindings. The soundness theorem for the heap value translation states that the both the new heap value and the new bindings are well-formed in the translation of the original context extended with bindings for the new heap entries.

Lemma 26 (Soundness of the heap value translation)

If

$$\Psi \vdash hval : \tau \mathbf{hval} \text{ and } \Psi \vdash hval : \tau \rightsquigarrow d \text{ in } hval'$$

then

$$|\Psi|, \Psi(d) \vdash d \mathbf{ok} \text{ and } |\Psi|, \Psi(d) \vdash hval' : |\tau| \mathbf{hval}$$

Proof: By construction.

By assumption:

$$\Psi \vdash \text{code}_\tau[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k](x_1:\tau_1, \dots, x_m:\tau_m)(z_1:\phi_1, \dots, z_n:\phi_n).e : \forall[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k] \text{Code}(\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n)(\tau) \text{ hval}$$

and

$$\Psi \vdash \text{code}_\tau[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k](x_1:\tau_1, \dots, x_m:\tau_m)(z_1:\phi_1, \dots, z_n:\phi_n).e : \forall[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k] \text{Code}(\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n)(\tau) \rightsquigarrow d \text{ in } \text{code}_{|\tau|}[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k](x_1:|\tau_1|, \dots, x_m:|\tau_m|)(z_1:|\phi_1|, \dots, z_n:|\phi_n|).e'$$

By inversion:

$$\begin{aligned} & \bullet \vdash \kappa_i \text{ ok} \quad (i \in 1 \dots k) \\ & \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k \vdash \tau_i : T_{32} \quad (i \in 1 \dots m) \\ & \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k \vdash \phi_i : T_{64} \quad (i \in 1 \dots n) \\ & \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k \vdash \tau : T_{32} \\ & \Psi; \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k; x_1:\tau_1, \dots, x_m:\tau_m, z_1:\phi_1, \dots, z_n:\phi_n \vdash e : \tau \text{ exp} \\ & \Psi; \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k; x_1:\tau_1, \dots, x_m:\tau_m, z_1:\phi_1, \dots, z_n:\phi_n \vdash e : \tau \rightsquigarrow d \text{ in } e' \end{aligned}$$

By theorem 10:

$$\begin{aligned} & \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k \vdash |\tau_i| : T_{32} \quad (i \in 1 \dots m) \\ & \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k \vdash |\phi_i| : T_{64} \quad (i \in 1 \dots n) \\ & \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k \vdash |\tau| : T_{32} \end{aligned}$$

By theorem 13:

$$\begin{aligned} & |\Psi|, \Psi(d) \vdash d \text{ ok} \quad \text{and} \\ & |\Psi|, \Psi(d); \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k; x_1:|\tau_1|, \dots, x_m:|\tau_m|, z_1:|\phi_1|, \dots, z_n:|\phi_n| \vdash e' : |\tau| \text{ exp} \end{aligned}$$

Note that:

$$\begin{aligned} & |\forall[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k] \text{Code}(\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n)(\tau)| \\ & = \forall[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k] \text{Code}(|\tau_1|, \dots, |\tau_m|)(|\phi_1|, \dots, |\phi_n|)(|\tau|) \end{aligned}$$

So by the heap value rule :

$$\begin{aligned} & |\Psi|, \Psi(d) \vdash \text{code}_{|\tau|}[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k](x_1:|\tau_1|, \dots, x_m:|\tau_m|)(z_1:|\phi_1|, \dots, z_n:|\phi_n|).e' : \\ & \quad \forall[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k] \text{Code}(\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n)(\tau) \text{ hval} \end{aligned}$$

■

Lemma 27 (Soundness of the heap translation)

If $\Psi \vdash d \text{ ok}$ and $\Psi \vdash d \rightsquigarrow d_h \text{ in } d'$ then $|\Psi|, \Psi(d_h) \vdash d_h, d' \text{ ok}$.

Proof: By induction on d . We proceed by cases on d .

1. ϵ

By assumption:

$$\vdash \Psi \text{ ok}$$

By theorem 11:

$$\vdash |\Psi| \text{ ok}$$

By the empty heap rule:

$$|\Psi| \vdash \epsilon \text{ ok}$$

(Note that $\Psi(\epsilon) = \bullet$)

2. $d, \ell : \tau \mapsto \text{hval}$

By assumption:

$$\begin{aligned}\Psi[\ell:\tau] \vdash d, \ell:\tau \mapsto hval \text{ ok} \\ \Psi[\ell:\tau] \vdash d, \ell:\tau \mapsto hval \rightsquigarrow d_\ell, d_h \text{ in } d', \ell:|\tau| \mapsto hval'\end{aligned}$$

By inverting assumptions:

$$\begin{aligned}\Psi[\ell:\tau] \vdash hval:\tau \text{ hval} \\ \Psi[\ell:\tau] \vdash hval:\tau \rightsquigarrow d_\ell \text{ in } hval' \\ \Psi[\ell:\tau] \vdash d \text{ ok} \\ \Psi[\ell:\tau] \vdash d \rightsquigarrow d_h \text{ in } d'\end{aligned}$$

By induction:

$$|\Psi[\ell:\tau]|, \Psi(d_h) \vdash d_h, d' \text{ ok}$$

By lemma 26:

$$\begin{aligned}|\Psi[\ell:\tau]|, \Psi(d_\ell) \vdash d_\ell \text{ ok and} \\ |\Psi[\ell:\tau]|, \Psi(d_\ell) \vdash hval':|\tau| \text{ hval}\end{aligned}$$

By lemma weakening:

$$\begin{aligned}|\Psi[\ell:\tau]|, \Psi(d_h), \Psi(d_\ell) \vdash d_h, d' \text{ ok} \\ |\Psi[\ell:\tau]|, \Psi(d_h), \Psi(d_\ell) \vdash d_\ell \text{ ok}\end{aligned}$$

By definition:

$$\begin{aligned}|\Psi[\ell:\tau]|, \Psi(d_h), \Psi(d_\ell) \vdash d_h, d', d_\ell \text{ ok} \\ \text{and } |\Psi[\ell:\tau]| = |\Psi|[\ell:|\tau|]\end{aligned}$$

So by the non-empty heap rule:

$$|\Psi[\ell:\tau]|, \Psi(d_\ell, d_h) \vdash d_\ell, d_h, d', \ell:|\tau| \mapsto hval' \text{ ok}$$

■

Theorem 14 (Programs)

If $\Psi \vdash p:\tau$ and $\Psi \vdash p:\tau \rightsquigarrow p':\tau$ then $|\Psi| \vdash p':|\tau|$

Proof:

By definition, p is of the form **letrec** d **in** e .

By assumption:

$$\begin{aligned}\Psi, \Psi(d) \vdash d \text{ ok} \\ \Psi, \Psi(d) \vdash d \rightsquigarrow d_h \text{ in } d' \\ \Psi, \Psi(d); \bullet; \bullet \vdash e:\tau \text{ exp} \\ \Psi, \Psi(d); \bullet; \bullet \vdash e:\tau \rightsquigarrow d_e \text{ in } e'\end{aligned}$$

By lemma 25:

$$|\Psi(d)| = \Psi(d')$$

So by definition:

$$\begin{aligned}|\Psi, \Psi(d)| \\ = |\Psi|, |\Psi(d)| \\ = |\Psi|, \Psi(d')\end{aligned}$$

By lemma 27:

$$|\Psi, \Psi(d)|, \Psi(d_h) \vdash d_h, d' \text{ ok}$$

And by lemma 13:

$$|\Psi, \Psi(d)|, \Psi(d_e); \bullet; \bullet \vdash e':|\tau| \text{ exp}$$

So by the above argument:

$|\Psi|, \Psi(d'), \Psi(d_h) \vdash d_h, d' \text{ ok}$
and $|\Psi|, \Psi(d'), \Psi(d_e); \bullet; \bullet \vdash e' : |\tau| \text{ exp}$

By weakening:

$|\Psi|, \Psi(d'), \Psi(d_h), \Psi(d_e) \vdash d_h, d' \text{ ok}$
and $|\Psi|, \Psi(d'), \Psi(d_h), \Psi(d_e); \bullet; \bullet \vdash e' : |\tau| \text{ exp}$

So by the program rule :

$|\Psi| \vdash \text{letrec } d', d_h, d_e \text{ in } e' : |\tau|$

■

Chapter 7

TILTAL

Closure converted **LIL** code is sufficiently low-level that it is practical to translate it directly into a typed assembly language. In this chapter I present such a language, and define a simple translation from **LIL** into it. I call this assembly language TILT Typed Assembly Language (**TILTAL**), in reference to the original typed assembly language (**TAL**) [MWCG98]. The overall structure of **TILTAL** is very similar to the stack version of **TAL** [MCGW02].

The presentation of the **TILTAL** target language is intended to suggest how the ideas used to implement type analysis in the **LIL** translate down to the assembly code level. The actual implementation targets the **TALx86** infrastructure [MCG⁺99] which is significantly different in many ways. I will discuss the **TALx86** type system in more detail in the implementation chapters.

7.1 TILTAL overview

I begin by introducing the **TILTAL** language in general terms and describing the typing judgements which define its static semantics. The syntax for the language is described in figure 7.1. The kind and type levels are almost entirely the same as that of the **LIL** and are mostly elided. Constructs for describing the type of stacks and a kind ST classifying them are added. Code function types are eliminated and replaced with a code type which describes the stack and register format expected by a code segment. “Nonsense” types ns_{32} and ns_{64} are introduced to describe uninitialized stacks slots, with corresponding constants ns_{32} and ns_{64} .

A table describing the typing judgements for **TILTAL** is given in figure 7.2. The complete list of inference rules for these judgements can be found in appendix C. Most **TILTAL** terms are judged well-typed with respect to a heap context Ψ which binds labels at closed types; a constructor/kind context Δ which binds kind variables and binds constructor variables at kinds; and a register file type Γ which describes the state of the abstract machine by mapping register names to types.

In the the next several sections, I will introduce the important syntactic classes in **TILTAL** and briefly describe their use. For the most part, readers familiar with the existing literature on typed assembly language will find nothing substantially different.

7.1.1 Stacks

The treatment of stacks and stack types in **TILTAL** is much the same as in previous work [MCGW02]. I give a brief introduction to the syntax and the concepts here.

k	$::= \dots \mid ST$
c, τ, ϕ, σ	$::= \dots \mid \boxed{\text{delete} \rightarrow}$ $\mid \epsilon \mid \tau_1 \triangleright_{32} c \mid \phi_1 \triangleright_{64} c \mid \sigma_1 \circ \sigma_2 \mid \text{sptr } \sigma$ $\mid ns_{64} \mid ns_{32} \mid \Gamma \rightarrow 0$
Γ	$::= \{\mathbf{r}_1:\tau_1, \mathbf{r}_2:\tau_2, \mathbf{r}_e:\tau_e, \mathbf{r}_t:\tau_t, \mathbf{f}_1:\phi_1, \mathbf{f}_2:\phi_2, \mathbf{sp}:\sigma\}$
$\mathbf{r}$	$::= \mathbf{r}_1 \mid \mathbf{r}_2 \mid \mathbf{r}_e \mid \mathbf{r}_t$
$\mathbf{f}$	$::= \mathbf{f}_1 \mid \mathbf{f}_2$
q	$::= \text{roll}_c \mid \text{unroll}_c \mid \text{inj_union}_{(i,c)} \mid \text{forgetunion}$ $\mid \text{pack}[\tau]c \mid \text{inj_dyn}_\tau$
s	$::= \epsilon \mid w \triangleright_{32} s \mid l \triangleright_{64} s$
l	$::= \mathbf{r} \mid ns_{64}$
w	$::= \ell \mid i \mid ns_{32} \mid \text{tag}_i \mid w[\vec{c}] \mid \text{sptr } i \mid q \ w$
fv	$::= l \mid \mathbf{f} \mid \mathbf{sp}(i)$
sv	$::= \mathbf{r} \mid \mathbf{sp}(i) \mid w \mid sv[\vec{c}] \mid q \ sv$
i	$::= \text{mov } \mathbf{r}, sv \mid \text{loadr } \mathbf{r}, \mathbf{r}(i) \mid \text{store } \mathbf{r}(i), sv$ $\mid \text{malloc } \mathbf{r}[\tau_1, \dots, \tau_n] \langle sv_1, \dots, sv_n \rangle \mid \text{malloc}_\tau \mathbf{r}(sv_1, sv_2)$ $\mid \text{malloc}_\phi \mathbf{r}, fv \mid \text{malloc}_\phi \mathbf{r}(sv, fv)$ $\mid \text{call } sv \mid \text{brtag}_i \mathbf{r}, sv \mid \text{brtgd}_i \mathbf{r}, sv$ $\mid \text{brdyn } \mathbf{r}, sv_1, sv_2 \mid \text{dyntag}_c \mathbf{r}$ $\mid \text{swrites } \mathbf{sp}(i), sv \mid \text{salloc } n \mid \text{sfree } n$ $\mid \text{mov } \mathbf{r}, \mathbf{sp} \mid \text{mov } \mathbf{sp}, sv$ $\mid \text{unpack}[\alpha, \mathbf{r}], sv$ $\mid \text{sub}_\tau \mathbf{r}, sv_1, sv_2 \mid \text{upd}_\tau sv_1, sv_2, sv_3$ $\mid \text{fmov } \mathbf{f}, fv \mid \text{floadr } \mathbf{f}, \mathbf{r} \mid \text{fstore } \mathbf{r}, fv$ $\mid \text{fswrites } \mathbf{sp}(i), fv$ $\mid \text{sub}_\phi \mathbf{f}, sv_1, sv_2 \mid \text{upd}_\phi sv_1, sv_2, fv$
ti	$::= \text{vcase}[\alpha_1. \text{dead } sv, \alpha_2] c$ $\mid \text{vcase}[\alpha_1, \alpha_2. \text{dead } sv] c$ $\mid \text{refine}[\langle \beta, \gamma \rangle] c \mid \text{refine}[\text{fold } \beta] c$
I	$::= \text{ret} \mid \text{jmp } sv \mid \text{halt}_\tau \mid i; I \mid ti; I$
$hval$	$::= \langle \vec{w} \rangle \mid [\vec{w}] \mid [\vec{l}] \mid l \mid \text{dtag}$ $\mid \text{code}[\alpha_1::\kappa_1, \dots, \alpha_n::\kappa_n].\Gamma.I$
H	$::= \{\ell_1:\tau_1 \mapsto hval_1, \dots, \ell_n:\tau_n \mapsto hval_n\}$
R	$::= \{\mathbf{r}_1 \mapsto w_1, \mathbf{r}_2 \mapsto w_2, \mathbf{r}_e \mapsto w_e, \mathbf{r}_t \mapsto w_t,$ $\mathbf{f}_1 \mapsto l_1, \mathbf{f}_2 \mapsto l_2: \mathbf{sp} \mapsto s\}$
P	$::= (H, R, I)$
Ψ	$::= \bullet \mid \Psi, \ell:\tau$
Δ	$::= \bullet \mid \Delta, j \mid \Delta, \alpha:\kappa$
Γ	$::= \{\mathbf{r}_1:\tau_1, \mathbf{r}_2:\tau_2, \mathbf{r}_e:\tau_e, \mathbf{r}_t:\tau_t, \mathbf{f}_1:\phi_1, \mathbf{f}_2:\phi_2, \mathbf{sp}:\sigma\}$

Figure 7.1: TILTAL syntax

TILTAL typing judgements	
Kinds	$\Delta \vdash \kappa \text{ ok}$
Constructors	$\Delta \vdash c : \kappa$
Equivalence	$\Delta \vdash c \equiv c' : \kappa$
Coercions	$\Delta \vdash q : \tau \Rightarrow \tau'$
Stacks	$\Psi; \Delta \vdash s : \sigma$
64-bit values	$\Psi; \Delta \vdash l : \phi$
32-bit values	$\Psi; \Delta \vdash w : \tau$
64-bit operands	$\Psi; \Delta; \Gamma \vdash fv : \phi$
32-bit operands	$\Psi; \Delta; \Gamma \vdash sv : \tau$
Instructions	$\Psi; \Delta; \Gamma \vdash i \Rightarrow \Gamma'$
Instruction Sequences	$\Psi; \Delta; \Gamma \vdash I : \tau$
Heap values	$\Psi \Delta; \Gamma \vdash hval : \tau \text{ hval}$
Heaps	$\vdash H : \Psi$
Well-formed Register File	$\Psi \vdash R : \Gamma$
Well-formed Program	$\vdash (H, R, I) \text{ ok}$

Figure 7.2: Judgements defining well-typedness of **TILTAL** programs

Stack types and values in **TILTAL** are used to represent the actual control stack of a program. Stack slots store return addresses of functions, as well as variables which cannot be spilled into registers. At the value level a stack is either empty (ϵ), or a value pushed onto a stack ($w \triangleright_{32} s$ or $l \triangleright_{64} s$). Stacks in **TILTAL** are permitted to contain either 32 or 64-bit values. Pointers into the stack are permitted in a limited fashion. The meta-variable s is used to stand for stack values, and the meta-variable σ is used for the type of stacks.

Stack types are built up from empty stacks using stack addition and composition. The type ϵ describes an empty stack, the type $\tau \triangleright_{32} \sigma$ describes a stack described by σ with a 32 bit value described by τ pushed on top, and similarly for $\phi \triangleright_{64} \sigma$ when ϕ describes a 64 bit value. In this way, stack types closely parallel the structure of actual stacks. In addition to these constructs, stacks types may also be constructed by concatenating two stack types. Such a compound stack, written $\sigma_1 \circ \sigma_2$, is equivalent to the stack obtained by prepending all of σ_1 to σ_2 . Finally, the type $\text{sptr}(\sigma)$ describes a pointer to a stack described by σ . This facility is used to compile exception handlers, which must maintain pointers into the stack.

For syntactic convenience, I define several operations on stack types for use in the static semantics. I define stack type lookup operations $\sigma[i]_{32}$ and $\sigma[i]_{64}$ which find the appropriately size type in the stack σ located i words from the top of the stack.

Definition 1 (Stack type 32 bit subscript)

$$\begin{aligned}
(\tau \triangleright_{32} \sigma)[0]_{32} &\stackrel{\text{def}}{=} \tau \\
(\tau \triangleright_{32} \sigma)[n+1]_{32} &\stackrel{\text{def}}{=} \sigma[n]_{32} \\
(\phi \triangleright_{64} \sigma)[n+2]_{32} &\stackrel{\text{def}}{=} \sigma[n]_{32}
\end{aligned}$$

Definition 2 (Stack type 64 bit subscript)

$$\begin{aligned}
(\phi \triangleright_{64} \sigma)[0]_{64} &\stackrel{\text{def}}{=} \phi \\
(\tau \triangleright_{32} \sigma)[n+1]_{64} &\stackrel{\text{def}}{=} \sigma[n]_{64} \\
(\phi \triangleright_{64} \sigma)[n+2]_{64} &\stackrel{\text{def}}{=} \sigma[n]_{64}
\end{aligned}$$

In a similar manner, I define stack type updates $(\sigma)[i]_{32} \leftarrow \tau$ and $(\sigma)[i]_{64} \leftarrow \phi$ which replace the type at word offset i in σ with τ or ϕ respectively.

Definition 3 (Stack type 32 bit update)

$$\begin{aligned}
(\tau \triangleright_{32} \sigma)[0]_{32} \leftarrow \tau' &\stackrel{\text{def}}{=} \tau' \triangleright_{32} \sigma \\
(\phi \triangleright_{64} \sigma)[0]_{32} \leftarrow \tau' &\stackrel{\text{def}}{=} \tau' \triangleright_{32} ns_{32} \triangleright_{32} \sigma \\
(\phi \triangleright_{64} \sigma)[1]_{32} \leftarrow \tau' &\stackrel{\text{def}}{=} ns_{32} \triangleright_{32} \tau' \triangleright_{32} \sigma \\
(\tau \triangleright_{32} \sigma)[n+1]_{32} \leftarrow \tau' &\stackrel{\text{def}}{=} (\sigma)[n]_{32} \leftarrow \tau' \\
(\phi \triangleright_{64} \sigma)[n+2]_{32} \leftarrow \tau' &\stackrel{\text{def}}{=} (\sigma)[n]_{32} \leftarrow \tau'
\end{aligned}$$

Definition 4 (Stack type 64 bit update)

$$\begin{aligned}
(\tau_1 \triangleright_{32} \tau_2 \triangleright_{32} \sigma)[0]_{64} \leftarrow \phi' &\stackrel{\text{def}}{=} \phi' \triangleright_{64} \sigma \\
(\tau_1 \triangleright_{32} \phi \triangleright_{64} \sigma)[0]_{64} \leftarrow \phi' &\stackrel{\text{def}}{=} \phi' \triangleright_{64} ns_{32} \triangleright_{32} \sigma \\
(\phi \triangleright_{64} \sigma)[0]_{64} \leftarrow \phi' &\stackrel{\text{def}}{=} \phi' \triangleright_{64} \sigma \\
(\phi \triangleright_{64} \sigma)[1]_{64} \leftarrow \phi' &\stackrel{\text{def}}{=} ns_{32} \triangleright_{32} (ns_{32} \triangleright_{32} \sigma)[0]_{64} \leftarrow \phi' \\
(\tau \triangleright_{32} \sigma)[n+1]_{64} \leftarrow \phi' &\stackrel{\text{def}}{=} (\sigma)[n]_{64} \leftarrow \phi' \\
(\phi \triangleright_{64} \sigma)[n+2]_{64} \leftarrow \phi' &\stackrel{\text{def}}{=} (\sigma)[n]_{64} \leftarrow \phi'
\end{aligned}$$

Finally, I define a size operation on stack types in the obvious manner.

Definition 5 (Stack type size)

$$\begin{aligned}
|\epsilon| &\stackrel{\text{def}}{=} 0 \\
|\tau \triangleright_{32} \sigma| &\stackrel{\text{def}}{=} 1 + |\sigma| \\
|\phi \triangleright_{64} \sigma| &\stackrel{\text{def}}{=} 2 + |\sigma|
\end{aligned}$$

7.1.2 Values

Small values in **TILTAL** are syntactically divided into 32 and 64-bit sizes, written with the meta-variables w and l respectively. The types of 32-bit values have kind T_{32} and are written using the meta-variable τ , whereas the types of 64-bit values have kind T_{64} and are written with the meta-variable ϕ . The two forms of 64-bit values are IEEE floating point numbers or nonsense. 32-bit

values include integers, sum tags, pointers into the stack, nonsense values, and labels. Coercions applied to values (qw) change the type of the value but have no runtime effect. Labels serve as pointers to values allocated on the heap. As with the version of the **LIL** extended with allocation primitives, a heap type Ψ maps labels to the types of the values to which they point.

7.1.3 Operands: Registers and stack slots

Instruction operands are similarly divided into 32 and 64-bit versions. In addition to simple values as described above, operands can be registers or stack slots.

There are three sorts of registers in **TILTAL**: 32-bit registers, 64-bit registers, and the stack register. By convention, I reserve one of the 32-bit registers for the exception record (as described in section 8.1.2). Registers are used to hold the operands and intermediate values, and to pass arguments to functions. The registers given in the version of **TILTAL** described here are not specialized to the x86 platform, which has a very idiosyncratic register architecture.

The stack register is used to maintain the control stack. Reading and writing stack slots is done using offsets from the front of the stack given in 32 bit strides. Allocation of and de-allocation of stack space is handled by special instructions. Newly allocated stack space contains nonsense values which can subsequently be overwritten. Stack slots are used for the same purposes as the 32 and 64-bit registers, as well as for holding return addresses. The `call` instruction pushes a return address onto the stack, which can subsequently be used by the `ret` instruction.

7.1.4 Instructions and instruction sequences

The instruction set described here is very small and focuses mainly on the instructions that are interesting from a typing standpoint. Instructions are provided for moving values between the heap, the stack, and the register file as well as allocating and deallocating space on the stack and the heap. Primitive array instructions are provided to handle bounds checking internally. Instructions such as the array operations and the memory allocation operation are intended to be implemented as assembler macros, since they are not directly provided by the machine.

The typing rules for instructions are somewhat different from the usual sorts of typing rules in that instead of ascribing a type to an instruction, they describe the effect that the instruction has on the type of the state. The judgement $\Psi; \Delta; \Gamma \vdash i \Rightarrow \Gamma'$ indicates that the instruction i , when executed in a state described by a register file type Γ , will result in a state described by Γ' . Note that instructions do not change the type of values in the heap, nor the set of free types.

In addition to standard instructions, there are several type instructions which have no runtime effect. The `vcase` instructions correspond to the `vcase` constructs of the **LIL**, and the `refine` instructions correspond to the path refinement constructs. These instructions can be eliminated when types are removed. Note that I syntactically segregate the type instructions from the ordinary instructions, and give typing rules for them only as elements of code sequences. This simplifies things greatly, since typing the refinement instructions individually would presumably require the instruction typing judgement to return a substitution with all of the related complications.

Instruction sequences are a sequence of instructions and type instructions, terminated by either `ret`, `jmp sv`, or `halt $\tau$` . The return instruction expects the return address to be the top-most element of the stack, and pops it off before returning.

7.1.5 Heap values, heaps, and register files

Values that are not guaranteed to fit in 32 or 64-bit slots are always allocated on the heap. Aggregate heap values include tuples of 32-bit values, arrays of 32 or 64-bit values. Individual 64-bit values are also permitted to be allocated on the heap to allow them to be “boxed” and used in contexts expecting 32-bit values. The **dtag** value serves as a place-holder for exception tags, where the label serves as the unique identifier of the tag. Finally, code blocks are also allocated in the heap.

Heaps are mappings from labels to heap values. Values in the heap may contain references to other values in the heap, but must be closed with respect to types.

Register files are mappings from register names to values of the appropriate size and type for the given register. For example, register files map 32 bit register names to 32 bit word values. Register files are classified by register file types Γ .

7.1.6 Programs

A **TILTAL** program is a triple consisting of a heap H , a register file R , and an instruction sequence I . Execution of **TILTAL** program is defined by transitions between **TILTAL** programs. The complete dynamic semantics for **TILTAL** is given in appendix C.

7.1.7 Typing contexts

The form of heap contexts Ψ and constructor contexts Δ are essentially unchanged from the **LIL**. As before, heap contexts give types to labels, and constructor contexts bind kind variables, and bind constructor variables at their kinds.

The term level context from the **LIL** is replaced by a register file type which describes the type of the registers and stack of the **TILTAL** abstract machine. The type of the register file (and consequently the stack as well) is given by a register file type Γ which maps each register to the type of its contents. In addition to being used as part of the static semantics, register file types are included into the type level to describe the pre-conditions of code segments. The type $\Gamma \rightarrow 0$ describes a code segment which expects a register file described by Γ .

I use the notation $\Gamma(\mathbf{r})$ to indicate the type assigned by Γ to a register $\mathbf{r}$, and similarly for float registers $\mathbf{f}$ and the stack register **sp**. I use the notation $\Gamma\{\mathbf{r}:\tau\}$ to indicate the register file type obtained by replacing the type of $\mathbf{r}$ in Γ with τ , and analogously for float registers and stack registers.

7.2 TILTAL derived forms

7.2.1 Partial instruction sequences

In order to allow **TILTAL** expressions to be built up somewhat compositionally during the translation of **LIL** code, I define a derived notion of partial instruction sequences.

Definition 6 (Partial instruction sequence)

A partial instruction sequence S is a series of ordinary (non-refining) instructions not terminated by a jump, return or halt. That is,

$$S ::= \epsilon \mid i; S$$

I define a judgement defining well-typedness as a partial instruction sequence: $\Psi; \Delta; \Gamma \vdash S \Rightarrow \Gamma'$ with inference rules as follows:

$$\frac{}{\Psi; \Delta; \Gamma \vdash \epsilon \Rightarrow \Gamma}$$

$$\frac{\Psi; \Delta; \Gamma \vdash i \Rightarrow \Gamma' \quad \Psi; \Delta; \Gamma' \vdash S \Rightarrow \Gamma''}{\Psi; \Delta; \Gamma \vdash i; S \Rightarrow \Gamma''}$$

Definition 7 (Partial instruction sequence composition)

A partial instruction sequence S can be composed with a second partial instruction sequence S' to form a new instruction sequence $S; S'$ as follows:

$$\begin{aligned} \epsilon; S' &\stackrel{\text{def}}{=} S' \\ (i; S); S' &\stackrel{\text{def}}{=} i; (S; S') \end{aligned}$$

The result of composing well-formed partial instruction sequences is well-formed.

Lemma 28 (Soundness of partial instruction sequence composition)

If $\Psi; \Delta; \Gamma \vdash S \Rightarrow \Gamma'$ and $\Psi; \Delta; \Gamma' \vdash S' \Rightarrow \Gamma''$ then $\Psi; \Delta; \Gamma \vdash S; S' \Rightarrow \Gamma''$

Proof: By induction on S . ■

Definition 8 (Partial instruction sequence completion)

A partial instruction sequence S can be completed with an instruction sequence I to form a new instruction sequence $S; I$ as follows:

$$\begin{aligned} \epsilon; I &\stackrel{\text{def}}{=} I \\ (i; S); I &\stackrel{\text{def}}{=} i; (S; I) \end{aligned}$$

The result of completing a well-formed partial instruction sequence with a well-formed instruction sequence is itself well-formed: in this sense, partial instruction sequence completion is sound.

Lemma 29 (Soundness of partial instruction sequence completion)

If $\Psi; \Delta; \Gamma \vdash S \Rightarrow \Gamma'$ and $\Psi; \Delta; \Gamma' \vdash I \text{ ok}$ then $\Psi; \Delta; \Gamma \vdash S; I \text{ ok}$

Proof: By induction on S . ■

7.2.2 Heap fragments

For convenience in the **LIL** to **TILTAL** translation, I also define a notion of partial, or open heaps, called heap fragments. Heap fragments are simply incomplete heaps, with their own derived typing judgement. Translations of **LIL** terms generally produce heap fragments in addition to **TILTAL** terms, reflecting the fact that the translation may place items in the static data segment.

Definition 9 (Heap fragment)

A heap fragment F is a syntactic heap, which may have open labels.

I define a judgement defining well-typedness as a heap fragment of the form $\Psi \vdash F : \Psi'$.

$$\frac{\vdash \Psi, \Psi' \text{ ok} \quad \Psi, \Psi' \vdash hval_i : \Psi'(\ell_i) \text{ hval}}{\Psi \vdash \{\ell_1 : \tau_1 \mapsto hval_1, \dots, \ell_n : \tau_n \mapsto hval_n\} : \Psi'}$$

Two heap fragments F_1 and F_2 can be combined to form a new heap fragment $F_1; F_2$.

Definition 10 (Heap fragment combination)

$$\begin{aligned} & \{\ell_1 : \tau_1 \mapsto hval_1, \dots, \ell_n : \tau_n \mapsto hval_n\}; \{\ell_{n+1} : \tau_{n+1} \mapsto hval_{n+1}, \dots, \ell_m : \tau_m \mapsto hval_m\} \\ & \stackrel{\text{def}}{=} \{\ell_1 : \tau_1 \mapsto hval_1, \dots, \ell_n : \tau_n \mapsto hval_n, \ell_{n+1} : \tau_{n+1} \mapsto hval_{n+1}, \dots, \ell_m : \tau_m \mapsto hval_m\} \end{aligned}$$

The result of combining two well-formed heap fragments with disjoint labels is a well-formed heap-fragment.

Lemma 30 (Soundness of heap fragment combination)

If $\Psi \vdash F_1 : \Psi_1$ and $\Psi \vdash F_2 : \Psi_2$ and the domains of F_1 and F_2 are disjoint, then $\Psi \vdash F_1; F_2 : \Psi_1, \Psi_2$

Proof: By induction on F_1 . ■

A heap fragment F can be added to a heap H to produce a new heap $F + H$.

Definition 11 (Heap fragment incorporation)

$$\begin{aligned} & \{\ell_1 : \tau_1 \mapsto hval_1, \dots, \ell_n : \tau_n \mapsto hval_n\} + \{\ell_{n+1} : \tau_{n+1} \mapsto hval_{n+1}, \dots, \ell_m : \tau_m \mapsto hval_m\} \\ & \stackrel{\text{def}}{=} \{\ell_1 : \tau_1 \mapsto hval_1, \dots, \ell_n : \tau_n \mapsto hval_n, \ell_{n+1} : \tau_{n+1} \mapsto hval_{n+1}, \dots, \ell_m : \tau_m \mapsto hval_m\} \end{aligned}$$

The result of incorporating a well-formed heap fragment with a well-formed heap with disjoint labels is well-formed.

Lemma 31 (Soundness of heap fragment incorporation)

If $\vdash H : \Psi$ and $\Psi \vdash F : \Psi'$ and the domains of H and F are disjoint, then $\vdash H + F : \Psi, \Psi'$

Proof: By induction on F . ■

Chapter 8

The LIL to TILTAL translation

The final stage of compilation translates **LIL** programs to **TILTAL** code. This process makes the control structure of the program explicit via call, return, and jump instructions, and replaces variables with stack slots and registers. Since **TILT** uses a stack to allocate activation records, it is also necessary to account for this in the translation. Stacks have been dealt with before in the context of typed assembly language, and the language design approach taken in **TILTAL** is not substantially different in nature from the that previously mapped out by Morrisett, et al. [MWCG98, MCGW02] (though the translation methodology is quite different). The actual code generation is not unusual except in that it preserves type information.

This chapter gives a detailed presentation of the translation of **LIL** programs into **TILTAL** programs, and proves the soundness of this translation with respect to an abstract model of register allocation.

8.1 The constructor level

The type translation from **LIL** to **TILTAL** is relatively simple compared to the translation from the **MIL**. For the most part, the **TILTAL** type system is an extension of the **LIL** system. A few notable changes at the type level are discussed in more detail below.

8.1.1 Kinds

Every **LIL** kind is a **TILTAL** kind, and since no change is made in the constructor translation that changes kinds, the translation on kinds is the identity. In general, I treat every **LIL** kind as its own translation.

8.1.2 The constructor translation

The complete constructor translation is given in figure 8.1. Note that the $\rightarrow$ type constructor has no translation, since the translation is only defined on closure-converted programs. This could be made explicit by giving a syntactic variant of the **LIL** for closure converted programs and defining the closure conversion translation from chapter 6 as a translation between programs in the original **LIL** and programs in the new syntactic variant. However, it is much simpler to consider each variant as a refinement of a broader syntactic class. In general, throughout the translation I will

$ \alpha $	$\stackrel{\text{def}}{=} \alpha$
$ \lambda(\alpha:\kappa).c $	$\stackrel{\text{def}}{=} \lambda(\alpha:\kappa). c $
$ c_1 c_2 $	$\stackrel{\text{def}}{=} c_1 c_2 $
$ \pi_1 c $	$\stackrel{\text{def}}{=} \pi_1 c $
$ \pi_2 c $	$\stackrel{\text{def}}{=} \pi_2 c $
$ \text{inj}_i^k c $	$\stackrel{\text{def}}{=} \text{inj}_i^k c $
$ \text{case}(c, [\alpha_1.c_1, \dots, \alpha_n.c_n]) $	$\stackrel{\text{def}}{=} \text{case}(c , [\alpha_1. c_1 , \dots, \alpha_n. c_n])$
$ \langle c_1, c_2 \rangle $	$\stackrel{\text{def}}{=} \langle c_1 , c_2 \rangle$
$ \text{fold}_{\mu j.k} c $	$\stackrel{\text{def}}{=} \text{fold}_{\mu j.k} c $
$ \text{pr}(j, \alpha:\kappa, \rho:(j \rightarrow \kappa'), \iota: \text{in } c) $	$\stackrel{\text{def}}{=} \text{pr}(j, \alpha:\kappa, \rho:(j \rightarrow \kappa'), \iota: \text{in } c)$
$ \text{Float} $	$\stackrel{\text{def}}{=} \text{Float}$
$ \text{Int} $	$\stackrel{\text{def}}{=} \text{Int}$
$ \text{Boxed} $	$\stackrel{\text{def}}{=} \text{Boxed}$
$ \text{Void} $	$\stackrel{\text{def}}{=} \text{Void}$
$ \times $	$\stackrel{\text{def}}{=} \times$
$ \rightarrow $	$\stackrel{\text{def}}{=} \text{UNDEFINED}$
$ \text{Code} $	$\stackrel{\text{def}}{=}$
$\lambda(\alpha_{32}:\text{T}_{32}\text{list}, \alpha_{64}:\text{T}_{64}\text{list}, \alpha_r:\text{T}_{32}).\forall[\rho_1:ST, \rho_2:ST].$ $\{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:\mathbf{Exnptr}(\rho_2), \mathbf{r}_t:ns_{32},$ $\mathbf{sp}:\mathbf{cont}(\alpha_{32} @_{32} \alpha_{64} @_{64} \epsilon)(\rho_1)(\rho_2)(\alpha_r) \triangleright_{32} (\alpha_{32} @_{32} (\alpha_{64} @_{64} \rho_1 \circ \rho_2))\} \rightarrow 0$	
$ c[\kappa] $	$\stackrel{\text{def}}{=} c [\kappa]$
$ \Lambda j.c $	$\stackrel{\text{def}}{=} \Lambda j. c $
$ \exists $	$\stackrel{\text{def}}{=} \exists$
$ \forall $	$\stackrel{\text{def}}{=} \forall$
$ \text{Rec} $	$\stackrel{\text{def}}{=} \text{Rec}$
$ \bigvee $	$\stackrel{\text{def}}{=} \bigvee$
$ \text{Array}_{32} $	$\stackrel{\text{def}}{=} \text{Array}_{32}$
$ \text{Array}_{64} $	$\stackrel{\text{def}}{=} \text{Array}_{64}$
$ \text{Dyntag} $	$\stackrel{\text{def}}{=} \text{Dyntag}$
$ \text{Dyn} $	$\stackrel{\text{def}}{=} \text{Dyn}$
$ \text{Tag} $	$\stackrel{\text{def}}{=} \text{Tag}$

Figure 8.1: The constructor translation

assume that the **LIL** programs under consideration are of the closure-converted variant and will not remark upon this further.

The only interesting work of the constructor translation is to translate the type of code functions into code sequence types of the form $\Gamma \rightarrow 0$. This can be thought of as replacing all uses of the **Code** constant by a closed defined form of the same kind. This is especially convenient, since the soundness of the translation follows almost directly after showing the well-formedness of this definition, since any well-formedness derivation of a **LIL** constructor can be turned into a well-formedness derivation of its translation by replacing all uses of the **Code** axiom with uses of the appropriately weakened well-formedness derivation for the definition.

The translation of the **Code** type

In order to give the appropriate definition for the **Code** type, it is first necessary to discuss the calling conventions used by the translation, since these must be encoded in the types of code sequences. For simplicity in this translation, I use a so-called “caller pops” calling convention in which a called function is expected to return the stack with the space allocated for its in-arguments intact. This convention is not especially convenient for doing tail-call elimination: consequently the actual implementation uses a “callee-pops” calling convention in which a called function is expected to de-allocate the space allocated for its in-arguments before returning. The formal translation uses the caller-pops convention for simplicity.

The treatment of exception handlers in the translation is slightly different from that used by Morrisett et al. [MCGW02] in which two registers are dedicated to the exception handling mechanism. To correspond more closely with the current **TILT** implementation, I reserve a single register to hold the current exception frame. Exception frames contain the address of the current handler, and a pointer to the handler’s stack. Old exception frames are stored on the top of the handler stack.

Definition 12 (**Exnptr** and **Exnhndler**)

$$\begin{aligned} \mathbf{Exnhndler} &\stackrel{\text{def}}{=} \lambda(\rho:ST). \\ &\quad \forall[].\{\mathbf{r}_1:\mathbf{Dyn}, \mathbf{r}_2:ns_{32}, \mathbf{r}_t:ns_{32}, \mathbf{r}_e:ns_{32}, \mathbf{sp}:\rho, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}\} \rightarrow 0 \\ \mathbf{Exnptr} &\stackrel{\text{def}}{=} \lambda(\rho:ST). \times [\mathbf{Exnhndler}(\rho), \rho] \end{aligned}$$

The type of the exception handler itself is parameterized by the type of the handler stack. Notice that the handler expects an exception packet in $\mathbf{r}_1$.

Code function types in **TILTAL** must also encode the type of a function’s continuation in the form of a return address type. I define a type $\mathbf{cont}:ST \rightarrow ST \rightarrow T_{32}$ which describes the machine state expected by the calling code upon return. Continuation types are parameterized over the number of words of in-argument space on the stack, the return type, the remainder of the stack above the next exception handler, and the remainder of the stack below the next handler. The return value is passed in $\mathbf{r}_t$ by convention. The return continuation cannot assume anything about the argument slots, which must be coerced to a nonsense value before returning.

Definition 13 (cont)

$$\begin{aligned} \mathbf{cont} &\stackrel{\text{def}}{=} \lambda(\alpha_i:\mathbf{nat}, \rho_1:ST, \rho_2:ST, \alpha_{\mathbf{ret}}:T_{32}). \\ &\quad \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:\mathbf{Exnptr}(\rho_2), \mathbf{r}_t:\alpha_{\mathbf{ret}}, \\ &\quad \mathbf{sp}:ns_{32}^{\alpha_i} \circ \rho_1 \circ \rho_2\} \rightarrow 0 \end{aligned}$$

The translation of the **Code** type using these auxiliary definitions is given in definition 14 below. Code function types take three arguments corresponding to a list of 32 bit arguments, a list of 64 bit arguments, and a return type. In addition the new code sequence is polymorphic over two stack variables: ρ_1 and ρ_2 . In the translation of a function type, the second stack variable, ρ_2 , classifies the stack tail expected by the current exception handler, while the first stack variable, ρ_1 , classifies the stack between the arguments and the last handler. Functions expect all of their arguments on the stack, and the exception frame in register $\mathbf{r}_e$. The top most item on the stack upon entry to the function is the address of the code to which the function should return (that is, the function's continuation). Function return values are returned in register $\mathbf{r}_t$.

Definition 14

$$\begin{aligned} |\mathbf{Code}| &\stackrel{\text{def}}{=} \lambda(\alpha_{32}:T_{32}\mathbf{list}, \alpha_{64}:T_{64}\mathbf{list}, \alpha_r:T_{32}). \forall[\rho_1:ST, \rho_2:ST]. \\ &\quad \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:\mathbf{Exnptr}(\rho_2), \mathbf{r}_t:ns_{32}, \\ &\quad \mathbf{sp}:\mathbf{cont}(|\alpha_{32} @_{32} \alpha_{64} @_{64} \epsilon|)(\rho_1)(\rho_2)(\alpha_r) \triangleright_{32} (\alpha_{32} @_{32} (\alpha_{64} @_{64} \rho_1 \circ \rho_2))\} \rightarrow 0 \end{aligned}$$

The construction of the stack type in definition 14 uses defined operators $@_{32}$ and $@_{64}$ which prepend lists of 32 bit and 64 bit types (respectively) to stack types. It is straightforward to define these as object level recursors within the type system.

Definition 15 ($c @_{32} \sigma$)

$c @_{32} \sigma = f(c)\sigma$ where

$$\begin{aligned} \mathbf{f}:T_{32}\mathbf{list} \rightarrow ST \rightarrow ST &\stackrel{\text{def}}{=} \lambda(\alpha:T_{32}\mathbf{list}). \lambda(\rho:ST). \mathbf{pr}(j, \alpha:1 + T_{32} \times j, \rho:(j \rightarrow ST), \mathbf{in} \\ &\quad \mathbf{case}(\alpha) \\ &\quad \mathbf{inj}_1 \star \Rightarrow \sigma \\ &\quad \mathbf{inj}_2 \beta \Rightarrow ns_{32} \triangleright_{32} (\rho\beta)) \end{aligned}$$

Definition 16 ($c @_{64} \sigma$)

$c @_{64} \sigma = f(c)\sigma$ where

$$\begin{aligned} \mathbf{f}:T_{64}\mathbf{list} \rightarrow ST \rightarrow ST &\stackrel{\text{def}}{=} \lambda(\alpha:T_{64}\mathbf{list}). \lambda(\rho:ST). \mathbf{pr}(j, \alpha:1 + T_{32} \times j, \rho:(j \rightarrow ST), \mathbf{in} \\ &\quad \mathbf{case}(\alpha) \\ &\quad \mathbf{inj}_1 \star \Rightarrow \sigma \\ &\quad \mathbf{inj}_2 \beta \Rightarrow ns_{32} \triangleright_{64} (\rho\beta)) \end{aligned}$$

The definition also refers to the length function for stacks $|\cdot|:ST \rightarrow \mathbf{nat}$. It is straightforward to define this using primitive recursion in the same fashion as the iterated stack type and the list append functions.

8.1.3 Soundness of the type translation

Lemma 32 (Well-formedness of the Code definition)

If $\vdash \Delta \text{ ok}$ then $\Delta \vdash \mathbf{Code} : T_{32}\text{list} \rightarrow T_{64}\text{list} \rightarrow T_{32} \rightarrow T_{32}$.

Proof: By construction, $\vdash \mathbf{Code} : T_{32}\text{list} \rightarrow T_{64}\text{list} \rightarrow T_{32} \rightarrow T_{32}$. The result then follows by repeated weakening. ■

Theorem 15 (Soundness of the constructor translation)

For a **LIL** constructor c , if $\Delta \vdash c : \kappa$ then $\Delta \vdash |c| : \kappa$.

Proof: By construction. Every derivation D of $\Delta \vdash c : \kappa$ can be turned into a derivation of $\Delta \vdash |c| : \kappa$ by replacing every use of the **Code** axiom with a derivation for **Code** obtained via lemma 32. ■

Theorem 16 (The translation respects substitution)

$|c| [|c'|/\alpha] = |c[c'/\alpha]|$

Proof: The code definition (**Code**) and the **Code** primitive are both closed, so

$$|\mathbf{Code}| [|c'|/\alpha] = \mathbf{Code} [|c'|/\alpha] = \mathbf{Code} = |\mathbf{Code}| = |\mathbf{Code}[c'/\alpha]|$$

All other constructors remain unchanged by the translation, so the proof follows trivially. ■

Theorem 17 (The translation respects equivalence)

If $\Delta \vdash c \equiv c' : \kappa$ then $\Delta \vdash |c| \equiv |c'| : \kappa$

Proof: By construction. Every derivation of $\Delta \vdash c \equiv c' : \kappa$ can be turned into a derivation of $\Delta \vdash |c| \equiv |c'| : \kappa$ by replacing every use of the reflexivity axiom on **Code** with a use of the reflexivity axiom on **Code**. ■

8.2 The term level: preliminaries

Most of the work of the translation from **LIL** to **TILTAL** takes place at the level of terms. The next few sections will introduce some of the key concepts used in the translation.

8.2.1 Register and stack slot allocation

One of the key issues in translating from a variable binding/substitution based language to a register transfer style language is how to efficiently manage a fixed set of registers. This is the problem of register allocation. For the purposes of the formal translation, I choose to leave the specifics of register and stack slot allocation abstract. The problem of register allocation has been studied extensively, and this dissertation does not add anything substantially new to the discussion. In practice, standard techniques can be applied to determine a suitable mapping and the translation uses this information abstractly. This has the advantage of essentially making the translation parametric over the choice of allocation methods. Nonetheless, I wish to be able to show that the

translation as I define it is sound. Consequently, I will place certain typing requirements on the allocation method.

The term level translation is defined with respect to a notion of an abstract allocator, which I generally write as $\mathcal{A}$. An allocator is an object that is responsible for mapping program variables (both 32 and 64 bit) to registers and stack slots of the appropriate size. In order to avoid relying on any particular register allocation technology, I generally remain agnostic as to the internal data structures of the allocator, defining only a set of operations which any allocator must support.

Definition 17

An allocator is an object $\mathcal{A}$ with the following associated operations:

- For every 32 bit variable x , $\mathcal{A}(x) = \mathbf{r}$ or $\mathcal{A}(x) = \mathbf{sp}(i)$.
- For every 64 bit variable x_f , $\mathcal{A}(x_f) = \mathbf{f}$ or $\mathcal{A}(x_f) = \mathbf{sp}(i)$.
- $\text{frmsz}(\mathcal{A})$ is a natural number.
- For every **LIL** typing context Γ and stack type σ , $|\Gamma|_{\mathcal{A}}^{\sigma} = \Gamma_A$ for some register file type Γ_A .

The most basic such operation is to give a location for a variable. For a 32 bit variable x , I write the location associated with x by an allocator $\mathcal{A}$ as $\mathcal{A}(x)$. Where appropriate, I sometimes write $\mathcal{A}[x \rightarrow \mathbf{r}]$ when $\mathcal{A}(x) = \mathbf{r}$ or $\mathcal{A}[x \rightarrow \mathbf{sp}(i)]$ when $\mathcal{A}(x) = \mathbf{sp}(i)$. Similarly, for a 64 bit variable x_{64} , I write the location associated with x_{64} by the allocator $\mathcal{A}$ as $\mathcal{A}(x_{64})$. Where appropriate, I sometimes write $\mathcal{A}[x_{64} \rightarrow \mathbf{f}]$ when $\mathcal{A}(x_{64}) = \mathbf{f}$ or $\mathcal{A}[x_{64} \rightarrow \mathbf{sp}(i)]$ when $\mathcal{A}(x_{64}) = \mathbf{sp}(i)$. It is perfectly reasonable (and likely) that an allocator will map multiple variables to the same registers and stack slots ¹

An allocator defines and manages the stack frame for a function. Consequently, I require that an allocator answer queries about the size (in 32 bit words) of the current frame: written $\text{frmsz}(\mathcal{A})$.

An allocator is responsible for managing the stack and register resources used to implement a **LIL** program in the **TILTAL** abstract machine. At various points in the program, it is necessary to write out the type of the abstract machine. Consequently, the second operation required of an allocator is to be able to generate a register file type from a given typing context Γ . Note though that an allocator is only responsible for managing the current *frame* and the registers - it should be parametric with respect to the rest of the stack. With this in mind, I define the translation of a **LIL** typing context Γ under an allocator $\mathcal{A}$ with respect to a stack tail σ to be the register file $|\Gamma|_{\mathcal{A}}^{\sigma}$.

This completely defines an allocator: an object $\mathcal{A}$ which assigns locations to variables, has a defined frame size, and maps every context/stack type pair to a register file type. However, not every allocator is sufficient for the purposes of the translation, since there are many incoherent allocators which satisfy this definition. In order to state the soundness of the **LIL** to **TILTAL** translation, I define the notion of a *good allocator* that satisfies certain conditions.

Definition 18 (Good allocator for a context)

Let Γ_A be $|\Gamma|_{\mathcal{A}}^{\sigma}$, where Γ is a context and σ is a stack type, and let $\mathbf{r}_t$, $\mathbf{f}_t$, and $\mathbf{r}_e$ be designated machine registers. Then I say that an allocator $\mathcal{A}$ is a good allocator for Γ if:

¹While for semantic correctness it is to be hoped that this will only occur for variables whose live ranges do not overlap, this is *not* required by the translation so long as the typing requirements discussed below are fulfilled.

1. For all x , $\mathcal{A}(x) \neq \mathbf{r}_e$ and $\mathcal{A}(x) \neq \mathbf{r}_t$ and for all x_{64} , $\mathcal{A}(x) \neq \mathbf{f}_t$.
2. $\Gamma_A(\mathbf{sp}) = \sigma_f \circ \sigma$ and $\text{frmsz}(\mathcal{A}) = |\sigma_f|$.
3. $|\bullet|_{\mathcal{A}}^{\sigma} = \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:ns_{32}, \mathbf{r}_t:ns_{32}, \mathbf{sp}:ns_{32}^{\overline{n}} \circ \sigma\}$ where $n = \text{frmsz}(\mathcal{A})$.
4. If $\Gamma = \Gamma_1, x:\tau, \Gamma_2$ then:
 - (a) $\mathcal{A}$ is a good allocator for Γ_1, Γ_2 .
 - (b) If $\mathcal{A}(x) = \mathbf{r}$ then $|\Gamma|_{\mathcal{A}}^{\sigma} = |\Gamma_1, \Gamma_2|_{\mathcal{A}}^{\sigma} \{\mathbf{r}:|\tau|\}$
 - (c) If $\mathcal{A}(x) = \mathbf{sp}(i)$ then $|\Gamma|_{\mathcal{A}}^{\sigma} = |\Gamma_1, \Gamma_2|_{\mathcal{A}}^{\sigma} \{\mathbf{sp}:(\sigma')[i]_{32} \leftarrow |\tau|\}$ where $|\Gamma_1, \Gamma_2|_{\mathcal{A}}^{\sigma}(\mathbf{sp}) = \sigma'$.
5. If $\Gamma = \Gamma_1, x_{64}:\phi, \Gamma_2$ then:
 - (a) $\mathcal{A}$ is a good allocator for Γ_1, Γ_2 .
 - (b) If $\mathcal{A}(x_{64}) = \mathbf{f}$, $|\Gamma|_{\mathcal{A}}^{\sigma} = |\Gamma_1, \Gamma_2|_{\mathcal{A}}^{\sigma} \{\mathbf{f}:|\phi|\}$
 - (c) If $\mathcal{A}(x) = \mathbf{sp}(i)$ then $|\Gamma|_{\mathcal{A}}^{\sigma} = |\Gamma_1, \Gamma_2|_{\mathcal{A}}^{\sigma} \{\mathbf{sp}:(\sigma')[i]_{64} \leftarrow |\phi|\}$ where $|\Gamma_1, \Gamma_2|_{\mathcal{A}}^{\sigma}(\mathbf{sp}) = \sigma'$.

The first three requirements ensure that the allocator doesn't interfere with registers and stack slots outside of its control, and that all unused registers and stack slots are given the nonsense type by the translation.

The special registers $\mathbf{r}_t$, $\mathbf{f}_t$ $\mathbf{r}_e$ are not available for general allocation. The first two are reserved for temporary computation and returned values, and the last for exception frames.

The second requirement of a good allocator is that the stack type in the machine state returned from the translation of a context must be an extension of the stack provided, and that the size of the extension must match the result of the $\text{frmsz}()$ query.

The third requirement of a good allocator is that it not impose extraneous requirements on the state: that is, that the translation of an empty typing context is empty.

The important requirement to be a good allocator is that the locations returned from variable location queries must be coherent with the machine state obtained by translating a context containing the query variables. This constraint is expressed in the last two good allocator constraints.

There are two properties of interest that follow from these constraints. The first can be summarized informally as follows: if $\Gamma(x) = \tau$ then $|\Gamma|_{\mathcal{A}}^{\sigma}(\mathcal{A}(x)) = |\tau|$: that is, if the allocator assigns a variable from a typing context to a location, then the type assigned to that location in the translation of the context must be the translation of the variable type. More precisely, if $\Gamma_A = |\Gamma|_{\mathcal{A}}^{\sigma}$ then:

1. If for every x such that $\Gamma = \Gamma_1, x:\tau, \Gamma_2$:
 - (a) If $\mathcal{A}(x) = \mathbf{r}$, then $\Gamma_A(\mathbf{r}) = |\tau|$.
 - (b) If $\mathcal{A}(x) = \mathbf{sp}(i)$ and $\Gamma_A(\mathbf{sp}) = \sigma'$ then $\sigma'[i]_{32} = |\tau|$.
2. If for every x such that $\Gamma = \Gamma_1, x_{64}:\phi, \Gamma_2$:
 - (a) If $\mathcal{A}(x_{64}) = \mathbf{f}$, then $\Gamma_A(\mathbf{f}) = |\phi|$.
 - (b) If $\mathcal{A}(x) = \mathbf{sp}(i)$ and $\Gamma_A(\mathbf{sp}) = \sigma'$ then $\sigma'[i]_{64} = |\phi|$.

The second property that follows from the definition is that the extension of a typing context by a single variable translates to the same machine state as the translation of the original context with the type of the location for the new variable updated with the appropriate type. Informally, $|\Gamma, x:\tau|_{\mathcal{A}}^{\sigma} = |\Gamma|_{\mathcal{A}}^{\sigma}\{\mathcal{A}(x):|\tau|\}$. More precisely:

1. If $\Gamma = \Gamma', x:\tau$ then:

- (a) If $\mathcal{A}(x) = \mathbf{r}$ then $|\Gamma|_{\mathcal{A}}^{\sigma} = |\Gamma'|_{\mathcal{A}}^{\sigma}\{\mathbf{r}:|\tau|\}$
- (b) If $\mathcal{A}(x) = \mathbf{sp}(i)$ then $|\Gamma|_{\mathcal{A}}^{\sigma} = |\Gamma'|_{\mathcal{A}}^{\sigma}\{\mathbf{sp}:(\sigma')[i]_{32} \leftarrow |\tau|\}$ where $|\Gamma'|_{\mathcal{A}}^{\sigma}(\mathbf{sp}) = \sigma'$.

2. If $\Gamma = \Gamma', x_{64}$ then:

- (a) If $\mathcal{A}(x_{64}) = \mathbf{f}$, $|\Gamma|_{\mathcal{A}}^{\sigma} = |\Gamma'|_{\mathcal{A}}^{\sigma}\{\mathbf{f}:|\phi|\}$
- (b) If $\mathcal{A}(x) = \mathbf{sp}(i)$ then $|\Gamma|_{\mathcal{A}}^{\sigma} = |\Gamma'|_{\mathcal{A}}^{\sigma}\{\mathbf{sp}:(\sigma')[i]_{64} \leftarrow |\phi|\}$ where $|\Gamma'|_{\mathcal{A}}^{\sigma}(\mathbf{sp}) = \sigma'$.

A third (indirect) consequence of the good allocator definition is that the translation of contexts commutes with substitution.

Lemma 33 (Context translation substitution)

If $\mathcal{A}$ is a good allocator for Γ then

$$|\Gamma|_{\mathcal{A}}^{\sigma}[|c|/\alpha] = |\Gamma[c/\alpha]|_{\mathcal{A}}^{\sigma[|c|/\alpha]}$$

Proof: (By induction on Γ).

1. If $\Gamma = \bullet$ then

$$\begin{aligned} |\Gamma|_{\mathcal{A}}^{\sigma}[|c|/\alpha] &= \\ \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:ns_{32}, \mathbf{r}_t:ns_{32}, \mathbf{sp}:ns_{32}^{\bar{n}} \circ \sigma\}[|c|/\alpha] &= \\ \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:ns_{32}, \mathbf{r}_t:ns_{32}, \mathbf{sp}:ns_{32}^{\bar{n}} \circ \sigma[|c|/\alpha]\} &= \\ |\Gamma[c/\alpha]|_{\mathcal{A}}^{\sigma[|c|/\alpha]} \end{aligned}$$

2. If $\Gamma = \Gamma', x:\tau$ then by the good allocator assumption (using informal notation)

$$\begin{aligned} |\Gamma|_{\mathcal{A}}^{\sigma}[|c|/\alpha] &= (\text{by definition}) \\ (|\Gamma'|_{\mathcal{A}}^{\sigma}\{\mathcal{A}(x):|\tau|\})[|c|/\alpha] &= (\text{by the good allocator assumption}) \\ (|\Gamma'|_{\mathcal{A}}^{\sigma}[|c|/\alpha])\{\mathcal{A}(x):|\tau[|c|/\alpha]|\} &= (\text{by definition}) \\ (|\Gamma'|_{\mathcal{A}}^{\sigma}[|c|/\alpha])\{\mathcal{A}(x):|\tau[c/\alpha]|\} &= (\text{by lemma 16}) \\ |\Gamma'[c/\alpha]|_{\mathcal{A}}^{\sigma[|c|/\alpha]}\{\mathcal{A}(x):|\tau[c/\alpha]|\} &= (\text{by induction}) \\ |\Gamma[c/\alpha]|_{\mathcal{A}}^{\sigma[|c|/\alpha]} \end{aligned}$$

■

Being a good allocator is a very weak requirement on allocators. In particular, it does not guarantee that the choice of locations is semantically well-behaved. For example, the allocator which assigns every variable of a given type to the same location (regardless of liveness) is a perfectly good allocator by this definition: the fact that such an allocator is a poor choice practically is irrelevant from a type-soundness standpoint.

The notion of a good allocator for a context Γ extends naturally to a notion of a good allocator for an expression.

Definition 19 (Good allocator for an expression)

If D is a derivation of $\Psi; \Delta; \Gamma; \mathcal{A} \vdash e : \tau$ then I say that an allocator $\mathcal{A}$ is a good allocator for e if for all sub-derivations of D of the form $\Psi; \Delta'; \Gamma' \vdash e' : \tau'$ (for some $\Delta', \Gamma', e', \tau'$), $\mathcal{A}$ is a good allocator for Γ' .

Note that this implies that a good allocator for e is a good allocator for all sub-expressions of e . Also note that this implies that $\mathcal{A}$ is a good allocator for Γ .

Since expressions also occur as sub-derivations of operations (in case expressions and handlers) it is necessary to define an analogous property for operations.

Definition 20 (Good allocator for an operation)

If D is a derivation of $\Psi; \Delta; \Gamma; \mathcal{A} \vdash opr : \tau$ **opr**₃₂ then I say that an allocator $\mathcal{A}$ is a good allocator for opr if for all sub-derivations of D of the form $\Psi; \Delta'; \Gamma' \vdash e' : \tau'$ (for some $\Delta', \Gamma', e', \tau'$), $\mathcal{A}$ is a good allocator for Γ' .

Note that this implies that a good allocator for an expression e is a good allocator for all the operations in e , and that a good allocator for an operation opr is a good allocator for all expressions in opr .

8.2.2 Derived instructions

It is convenient in the course of the translation to make use of some derived instructions not provided in the core instruction set but which can be defined as sequence of core instructions. In a complete implementation these instructions might be added as primitive.

The first defined instructions coerce registers and stack slots to the nonsense type.

Definition 21 (Junk instructions)

$$\begin{aligned} \mathbf{junk} \ r & \stackrel{\text{def}}{=} \mathbf{mov} \ r, ns_{32} \\ \mathbf{fjunk} \ f & \stackrel{\text{def}}{=} \mathbf{fmov} \ f, ns_{64} \\ \mathbf{sjunk} \ \mathbf{sp}(i) & \stackrel{\text{def}}{=} \mathbf{swrite} \ \mathbf{sp}(i), ns_{32} \\ \mathbf{fsjunk} \ \mathbf{sp}(i) & \stackrel{\text{def}}{=} \mathbf{fswrite} \ \mathbf{sp}(i), ns_{64} \end{aligned}$$

For brevity, a single instruction is defined to coerce all registers to the nonsense type.

Definition 22 (Register coercion)

$$\begin{aligned} \mathbf{junkregs} & \stackrel{\text{def}}{=} \mathbf{junk} \ r_1; \\ & \quad \mathbf{junk} \ r_2; \\ & \quad \mathbf{junk} \ f_1; \\ & \quad \mathbf{junk} \ f_2; \\ & \quad \epsilon \end{aligned}$$

Lemma 34 (Derived instructions)

For well-formed contexts Ψ , Δ , and Γ , where $\Gamma(\mathbf{sp}) = \sigma$

- $\Psi; \Delta; \Gamma \vdash \mathbf{junk} \ r \Rightarrow \Gamma\{\mathbf{r}:ns_{32}\}$

- $\Psi; \Delta; \Gamma \vdash \mathbf{fjunk} \ f \Rightarrow \Gamma\{\mathbf{f}:ns_{64}\}$
- $\Psi; \Delta; \Gamma \vdash \mathbf{sjunk} \ \mathbf{sp}(i) \Rightarrow \Gamma\{\mathbf{sp}:(\sigma)[i]_{32} \leftarrow ns_{32}\}$
- $\Psi; \Delta; \Gamma \vdash \mathbf{fsjunk} \ \mathbf{sp}(i) \Rightarrow \Gamma\{\mathbf{sp}:(\sigma)[i]_{64} \leftarrow ns_{64}\}$
- $\Psi; \Delta; \Gamma \vdash \mathbf{junkregs} \ \mathbf{r} \Rightarrow \Gamma\{\mathbf{r}_1:ns_{32}\}\{\mathbf{r}_2:ns_{32}\}\{\mathbf{f}_1:ns_{64}\}\{\mathbf{f}_2:ns_{64}\}$

Proof: By construction. ■

Finally, I give an inductive definition of an instruction to coerce a range of stack slots to the nonsense type.

Definition 23 (Stack range coercion)

$$\begin{aligned} \mathbf{junkstack} \ n \dots n &\stackrel{\text{def}}{=} \mathbf{sjunk} \ \mathbf{sp}(n); \\ &\epsilon \\ \mathbf{junkstack} \ n \dots m \quad (n < m) &\stackrel{\text{def}}{=} \mathbf{sjunk} \ \mathbf{sp}(m); \\ &\mathbf{junkstack} \ n \dots (m-1) \end{aligned}$$

Lemma 35 (Stack range coercion)

For natural numbers n and m where $n < m$ and for well-formed contexts Ψ , Δ , and Γ , where $\Gamma(\mathbf{sp}) = \sigma$:

$$\Psi; \Delta; \Gamma \vdash \mathbf{junkstack} \ n \dots m \Rightarrow \Gamma\{\mathbf{sp}:\sigma'\}$$

where $\sigma' = (((\sigma)[n]_{32} \leftarrow ns_{32}) \dots)[m]_{32} \leftarrow ns_{32}$

Proof: By induction on m . If $n = m$, then the derived instruction sequence is $\mathbf{sjunk} \ \mathbf{sp}(m):\epsilon$, and the result follows directly by lemma 34. If $n < m$, then by induction, we get an instruction sequence which junks $\mathbf{sp}(n), \dots, \mathbf{sp}(m-1)$. After appending on the additional $\mathbf{sjunk} \ \mathbf{sp}(m)$ instruction, the result again follows by lemma 34. ■

In addition to these composite instructions, I also define a more general form of the move instruction which targets either registers or stack slots.

Definition 24 (Arbitrary mov and fmov)

$$\begin{aligned} \mathbf{srmov} \ \mathbf{r}, sv &\stackrel{\text{def}}{=} \mathbf{mov} \ \mathbf{r}, sv \\ \mathbf{srmov} \ \mathbf{sp}(i), sv &\stackrel{\text{def}}{=} \mathbf{swrite} \ \mathbf{sp}(i), sv \\ \mathbf{srfmov} \ \mathbf{f}, fv &\stackrel{\text{def}}{=} \mathbf{mov} \ \mathbf{f}, fv \\ \mathbf{srfmov} \ \mathbf{sp}(i), fv &\stackrel{\text{def}}{=} \mathbf{fswrite} \ \mathbf{sp}(i), fv \end{aligned}$$

Lemma 36 (Arbitrary mov)

For a location $dest_{32}$, and for an operand sv such that $\Psi; \Delta; \Gamma \vdash sv : \tau$

- If $dest_{32} = \mathbf{r}$ then $\Psi; \Delta; \Gamma \vdash \mathbf{srmov} \ \mathbf{r}, sv \Rightarrow \Gamma\{\mathbf{r}:\tau\}$
- If $dest_{32} = \mathbf{sp}(i)$ and $\Gamma(\mathbf{sp}) = \sigma$ then $\Psi; \Delta; \Gamma \vdash \mathbf{srmov} \ \mathbf{sp}(i), sv \Rightarrow \Gamma\{\mathbf{sp}:(\sigma)[i]_{32} \leftarrow \tau\}$

Proof: By construction. ■

Lemma 37 (Arbitrary fmov)

For a location $dest_{32}$, and for an operand fv such that $\Psi; \Delta; \Gamma \vdash fv : \phi$

- If $dest_{32} = \mathbf{f}$ then $\Psi; \Delta; \Gamma \vdash \mathbf{fsmov} \ \mathbf{f}, fv \Rightarrow \Gamma\{\mathbf{f}:\phi\}$
- If $dest_{32} = \mathbf{sp}(i)$ and $\Gamma(\mathbf{sp}) = \sigma$ then $\Psi; \Delta; \Gamma \vdash \mathbf{fsmov} \ \mathbf{sp}(i), fv \Rightarrow \Gamma\{\mathbf{sp}:(\sigma)[i]_{64} \leftarrow \phi\}$

Proof: By construction. ■

To simplify the translation (as well as eliminate un-necessary moves), I define register to register **mov** instructions which coerce the source register to ns_{32} after the move whenever the registers are not aliases.

Definition 25 (Move and junk)

$$\begin{aligned} \mathbf{movj} \ \mathbf{r}, \mathbf{r} &\stackrel{\text{def}}{=} \epsilon \\ \mathbf{movj} \ \mathbf{r}, \mathbf{r}' \ (\mathbf{r} \neq \mathbf{r}') &\stackrel{\text{def}}{=} \mathbf{mov} \ \mathbf{r}, \mathbf{r}'; \\ &\quad \mathbf{junk} \ \mathbf{r}' \\ \mathbf{fmovj} \ \mathbf{f}, \mathbf{f} &\stackrel{\text{def}}{=} \epsilon \\ \mathbf{fmovj} \ \mathbf{f}, \mathbf{f}' \ (\mathbf{f} \neq \mathbf{f}') &\stackrel{\text{def}}{=} \mathbf{fmov} \ \mathbf{f}, \mathbf{f}'; \\ &\quad \mathbf{fjunk} \ \mathbf{f}' \end{aligned}$$

Lemma 38 (32 bit move and junk)

For registers $\mathbf{r}$ and $\mathbf{r}'$ $\Psi; \Delta; \Gamma \vdash \mathbf{movj} \ \mathbf{r}, \mathbf{r}' \Rightarrow \Gamma\{\mathbf{r}':ns_{32}\}\{\mathbf{r}:\Gamma(\mathbf{r}')\}$

Proof: By construction.

If the registers are aliases, then the instruction sequence is empty, and

$$\Gamma\{\mathbf{r}':ns_{32}\}\{\mathbf{r}:\Gamma(\mathbf{r}')\} = \Gamma\{\mathbf{r}:ns_{32}\}\{\mathbf{r}:\Gamma(\mathbf{r})\} = \Gamma$$

If the registers are different, then **mov** and **junk** instructions are emitted, which update the register file type accordingly. ■

Lemma 39 (64 bit move and junk)

For registers $\mathbf{f}$ and $\mathbf{f}'$ $\Psi; \Delta; \Gamma \vdash \mathbf{fmovj} \ \mathbf{f}, \mathbf{f}' \Rightarrow \Gamma\{\mathbf{f}':ns_{64}\}\{\mathbf{f}:\Gamma(\mathbf{f}')\}$

Proof: By construction. If the registers are aliases, then the instruction sequence is empty, and $\Gamma = \Gamma\{\mathbf{f}':ns_{64}\}\{\mathbf{f}:\Gamma(\mathbf{f}')\}$. If the registers are different, then **fmov** and **fjunk** instructions are emitted, which update the register file type accordingly. ■

For exception handlers, I define a code sequence for copying a stack segment. I begin by defining a code sequence **stackcopy** (σ_1, σ_2) which copies the portion of the stack described by σ_2 into pre-allocated space below the already copied σ_1 . Note that **stackcopy** (σ_1, σ_2) is only defined when σ_1 and σ_2 have well-defined sizes: that is, when they are composed solely of pushes and of compositions of stacks with well-defined sizes.

$$\begin{aligned}
\mathbf{stackcopy}(\sigma, \epsilon) &\stackrel{\text{def}}{=} \epsilon \\
\mathbf{stackcopy}(\sigma_1, \tau \triangleright_{32} \sigma_2) &\stackrel{\text{def}}{=} \mathbf{swrite} \mathbf{sp}(|\sigma_1|), \mathbf{sp}(|\sigma_1 \circ \tau \triangleright_{32} \sigma_2| + |\sigma_1|); \\
&\quad \mathbf{stackcopy}(\sigma_1 \circ (\tau \triangleright_{32} \epsilon), \sigma_2) \\
\mathbf{stackcopy}(\sigma_1, \phi \triangleright_{64} \sigma_2) &\stackrel{\text{def}}{=} \mathbf{fswrite} \mathbf{sp}(|\sigma_1|), \mathbf{sp}(|\sigma_1 \circ \phi \triangleright_{64} \sigma_2| + |\sigma_1|); \\
&\quad \mathbf{stackcopy}(\sigma_1 \circ (\phi \triangleright_{64} \epsilon), \sigma_2)
\end{aligned}$$

Lemma 40 (stackcopy)

For well-formed stack types σ_1, σ_2 and σ such that $|\sigma_1|$ and $|\sigma_2|$ are well-defined:

$$\Psi; \Delta; \Gamma\{\mathbf{sp}:(\sigma_1 \circ ns_{32}^{|\sigma_2|}) \circ (\sigma_1 \circ \sigma_2) \circ \sigma\} \vdash \mathbf{stackcopy}(\sigma_1, \sigma_2) \Rightarrow \Gamma\{\mathbf{sp}:(\sigma_1 \circ \sigma_2) \circ (\sigma_1 \circ \sigma_2) \circ \sigma\}$$

Proof: By induction on σ_2 .

- If σ_2 is empty, then $|\sigma_2| = 0$, so

$$(\sigma_1 \circ ns_{32}^{|\sigma_2|}) \circ (\sigma_1 \circ \sigma_2) \circ \sigma = \sigma_1 \circ \sigma_1 \circ \sigma$$

- If $\sigma_2 = \tau \triangleright_{32} \sigma'_2$ then:

$$\text{Let } \sigma'_1 = \sigma_1 \circ (\tau \triangleright_{32} \epsilon)$$

By induction:

$$\Psi; \Delta; \Gamma\{\mathbf{sp}:(\sigma'_1 \circ ns_{32}^{|\sigma'_2|}) \circ (\sigma'_1 \circ \sigma'_2) \circ \sigma\} \vdash \mathbf{stackcopy}(\sigma'_1, \sigma'_2) \Rightarrow \Gamma\{\mathbf{sp}:(\sigma'_1 \circ \sigma'_2) \circ (\sigma'_1 \circ \sigma'_2) \circ \sigma\}$$

$$\text{Note that } \sigma'_1 \circ \sigma'_2 = \sigma_1 \circ (\tau \triangleright_{32} \epsilon) \circ \sigma'_2 = \sigma_1 \circ \tau \triangleright_{32} \sigma'_2 = \sigma_1 \circ \sigma_2$$

So by the partial sequence instruction rule, it suffices to show that:

$$\Psi; \Delta; \Gamma\{\mathbf{sp}:(\sigma_1 \circ ns_{32}^{|\sigma_2|}) \circ (\sigma_1 \circ \sigma_2) \circ \sigma\} \vdash \mathbf{swrite} \mathbf{sp}(|\sigma_1|), \mathbf{sp}(|\sigma_1 \circ \tau \triangleright_{32} \sigma'_2| + |\sigma_1|) \Rightarrow \Gamma\{\mathbf{sp}:(\sigma_1 \circ (\tau \triangleright_{32} \epsilon) \circ ns_{32}^{|\sigma'_2|}) \circ (\sigma_1 \circ \sigma_2) \circ \sigma\}$$

But note that $|\sigma_1 \circ ns_{32}^{|\sigma_2|} \circ \sigma_1| = |\sigma_1 \circ \tau \triangleright_{32} \sigma'_2| + |\sigma_1|$, and hence

$$((\sigma_1 \circ ns_{32}^{|\sigma_2|}) \circ (\sigma_1 \circ \sigma_2) \circ \sigma)[|\sigma_1 \circ \tau \triangleright_{32} \sigma'_2| + |\sigma_1|]_{32} = \sigma_2[0]_{32} = \tau$$

$$\text{And since } ns_{32}^{|\sigma_2|} = ns_{32}^{1+|\sigma'_2|} = ns_{32} \triangleright_{32} ns_{32}^{|\sigma'_2|}$$

$$(((\sigma_1 \circ ns_{32}^{|\sigma_2|}) \circ (\sigma_1 \circ \sigma_2) \circ \sigma)[|\sigma_1|]_{32} \leftarrow \tau = (\sigma_1 \circ (\tau \triangleright_{32} \epsilon) \circ ns_{32}^{|\sigma'_2|}) \circ (\sigma_1 \circ \sigma_2) \circ \sigma$$

Which is what we wanted.

- If $\sigma_2 = \phi \triangleright_{64} \sigma'_2$ then the result follows by a similar argument. ■

The **stackcopy** definition is used to define a stack duplication code sequence **copyframe** σ which emits code to duplicate σ on the top of the stack.

$$\begin{aligned}
\mathbf{copyframe} \sigma &\stackrel{\text{def}}{=} \mathbf{salloc} |\sigma|; \\
&\quad \mathbf{stackcopy}(\epsilon, \sigma)
\end{aligned}$$

64-bit values	$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{64} fv : \phi \rightsquigarrow fv'$
32-bit values	$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau \rightsquigarrow sv'$
64 bit operations	$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash dest_{64} \leftarrow fopr : \phi \rightsquigarrow S$
32 bit operations	$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash dest_{32} \leftarrow opr : \tau \rightsquigarrow S$
Expressions	$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F$
Heap values	$\Psi \vdash_h hval : \tau \rightsquigarrow I; F$
Heaps	$\Psi \vdash d \rightsquigarrow H$
Programs	$\vdash p : \tau \rightsquigarrow P$

Figure 8.2: LIL to TILTAL translation judgements

Lemma 41 (copyframe)

For well-formed stack types σ and σ' , such that $|\sigma|$ is well-defined:

$$\Psi; \Delta; \Gamma\{\mathbf{sp}:\sigma \circ \sigma'\} \vdash \mathbf{copyframe} \sigma \Rightarrow \Gamma\{\mathbf{sp}:\sigma \circ \sigma \circ \sigma'\}$$

Proof: By construction.

By the salloc rule, it suffices to show that

$$\Psi; \Delta; \Gamma\{\mathbf{sp}:ns_{32}^{|\sigma|} \circ \sigma \circ \sigma'\} \vdash \mathbf{stackcopy}(\epsilon, \sigma) \Rightarrow \Gamma\{\mathbf{sp}:\sigma \circ \sigma \circ \sigma'\}$$

Which follows immediately by lemma 40. ■

8.3 The term translation

The judgement forms used for the translation of **LIL** programs to **TILTAL** programs are listed in figure 8.3. In addition to the usual **LIL** typing contexts, all of the expression level judgements are defined with respect to an allocator $\mathcal{A}$ and two stack types σ_1 and σ_2 . The allocator provides the locations of variables in registers and frame slots, and the two stack types keep track of the layout of the rest of the stack below the current frame: σ_1 describes the stack above the current handler and σ_2 describes the stack below it.

Within the bodies of functions, these stack types will generally refer to two additional free variables representing the two stack segments expected by the function (the segment above the enclosing handler, and the segment below it). By convention, I name these variable ρ_1 and ρ_2 , and the stack types are expected to be well-formed in a context including these variables in addition to free type variables from the original program. So for example, an invariant of the most of the translation judgements is that $|\Delta|, \rho_1:ST, \rho_2:ST \vdash \sigma_1:ST$, and similarly for σ_2 , where Δ is the current constructor context. For brevity, I will generally abbreviate this idiom as follows:

Definition 26

$$\Delta^{\rho_1, \rho_2} \stackrel{\text{def}}{=} |\Delta|, \rho_1:ST, \rho_2:ST$$

The next several sections will give detailed overviews of the individual translation judgements and discuss some of the more interesting translation rules, as well as stating and proving the relevant

soundness properties. The complete definition of the translation can be found at the end of the chapter.

8.3.1 Values

LIL 32 bit and 64 bit values translate into **TILTAL** 32 bit and 64 bit operands, respectively. The 32 bit value translation judgement $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau \rightsquigarrow sv'$ indicates that in heap context Ψ , constructor context Δ , term context Γ , allocator $\mathcal{A}$, and stack segments σ_1 and σ_2 ; a **LIL** small value sv of type τ translates to a **TILTAL** operand sv' . The 64 bit has an analogous interpretation.

32 bit values

For closed values, the translation does little. For example, the translation of an integer value:

$$\frac{}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} i : \text{Int} \rightsquigarrow i}$$

Similarly, for coerced values such as values injected into union types, the translation simply inductively translates the coerced value to an operand, and then applies the translated coercion.

The rules for translating variables are more interesting however, since they use the allocator to choose locations for the variable.

$$\frac{}{\Psi; \Delta; \Gamma[x:\tau]; \mathcal{A}[x \rightarrow \mathbf{r}], \sigma_1, \sigma_2 \vdash_{32} x : \tau \rightsquigarrow \mathbf{r}}$$

$$\frac{}{\Psi; \Delta; \Gamma[x:\tau]; \mathcal{A}[x \rightarrow \mathbf{sp}(i)], \sigma_2, \sigma_2 \vdash_{32} x : \tau \rightsquigarrow \mathbf{sp}(i)}$$

The translation of 32 bit **LIL** values to **TILTAL** operands is sound in the sense that given a well-behaved allocator (as defined in definition 18), a well-typed **LIL** value translates to a **TILTAL** operand that is well-typed in the translation of the term context under the allocator. More precisely:

Theorem 18 (Soundness of the small value translation)

If $\Psi; \Delta; \Gamma \vdash sv : \tau$
and $\mathcal{A}$ is a good allocator for Γ
and $\Delta^{\rho_1, \rho_2} \vdash \sigma_1 : ST$
and $\Delta^{\rho_1, \rho_2} \vdash \sigma_2 : ST$
and $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau \rightsquigarrow sv'$

then

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash sv' : |\tau|$$

Proof: (theorem 18) By induction on sv . The proof proceeds by cases.

1. Suppose $\Psi; \Delta; \Gamma[x:\tau] \vdash x : \tau$ and $\Psi; \Delta; \Gamma[x:\tau]; \mathcal{A}[x \rightarrow \mathbf{r}], \sigma_1, \sigma_2 \vdash_{32} x : \tau \rightsquigarrow \mathbf{r}$.

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \mathbf{r} : |\tau|$$

By assumption:

$$\Gamma = \Gamma_1, x:\tau, \Gamma_2$$

$$\mathcal{A}(x) = \mathbf{r}$$

So by the good allocator assumption:

$$|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}(\mathcal{A}(x)) = |\tau|$$

So by the register rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \mathbf{r} : |\tau|$$

2. Suppose $\Psi; \Delta; \Gamma[x:\tau] \vdash x : \tau$ and $\Psi; \Delta; \Gamma[x:\tau]; \mathcal{A}[x \rightarrow \mathbf{sp}(i)], \sigma_2, \sigma_2 \vdash_{32} x : \tau \rightsquigarrow \mathbf{sp}(i)$.

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \mathbf{sp}(i) : |\tau|$$

By assumption:

$$\Gamma = \Gamma_1, x:\tau, \Gamma_2$$

$$\mathcal{A}(x) = \mathbf{sp}(i)$$

So by the good allocator assumption:

$$|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}(\mathbf{sp}) = \sigma$$

$$\sigma[i]_{32} = |\tau|$$

So by the stack slot rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \mathbf{sp}(i) : |\tau|$$

3. Suppose $\Psi[\ell:\tau]; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} \ell : \tau \rightsquigarrow \ell$.

To show:

$$|\Psi[\ell:\tau]|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \ell : |\tau|$$

By definition, $|\Psi[\ell:\tau]| = |\Psi|, \ell:|\tau|$ so the result follows directly by the label rule.

4. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} i : \mathbf{Int} \rightsquigarrow i$.

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash i : \mathbf{Int}$$

This follows directly by the integer rule.

5. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} \mathbf{inj_union}_c sv : c \rightsquigarrow \mathbf{inj_union}_{(i,|c|)} sv'$.

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \mathbf{inj_union}_{(i,|c|)} sv' : |c|$$

By inversion:

$$\Delta \vdash c \equiv \bigvee[\dots, c_i, \dots] : \mathbf{T}_{32}$$

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : c_i \rightsquigarrow sv'$$

$$\Psi; \Delta; \Gamma \vdash sv : c_i$$

By induction:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash sv' : |c_i|$$

By theorem 17 and weakening:

$$\Delta^{\rho_1, \rho_1} \vdash |c| \equiv |\bigvee[\dots, c_i, \dots]| : \mathbf{T}_{32}$$

By construction (injunion):

$$\Delta^{\rho_1, \rho_2} \vdash \mathbf{inj_union}_{(i,|c|)} : |c_i| \Rightarrow |c|$$

By construction (coercion app):

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \mathbf{inj_union}_{(i,|c|)} sv' : |c|$$

6. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} \text{roll}_\tau sv : \tau \rightsquigarrow \text{roll}_{|\tau|} sv'$.

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \text{roll}_{|\tau|} sv' : |\tau|$$

By inversion:

$$\begin{aligned} \Delta \vdash \tau &\equiv \text{Rec}[\kappa](c)(c_p) : T_{32} \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv &: c(\text{Rec}[\kappa]c)c_p \rightsquigarrow sv' \\ \Psi; \Delta; \Gamma \vdash sv &: c(\text{Rec}[\kappa]c)c_p \end{aligned}$$

By induction:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash sv' : |c(\text{Rec}[\kappa]c)c_p|$$

By theorem 17 and weakening:

$$\Delta^{\rho_1, \rho_1} \vdash |\tau| \equiv |\text{Rec}[\kappa](c)(c_p)| : T_{32}$$

Note that $|\text{Rec}[\kappa](c)(c_p)| = \text{Rec}[\kappa](|c|)(|c_p|)$ and $|c(\text{Rec}[\kappa]c)c_p| = (|c|(\text{Rec}[\kappa]|c|)|c_p|)$.

By construction (roll):

$$\Delta^{\rho_1, \rho_2} \vdash \text{roll}_{|\tau|} : (|c|(\text{Rec}[\kappa]|c|)|c_p|) \Rightarrow |\tau|$$

By construction (coercion app):

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \text{roll}_{|\tau|} sv' : |\tau|$$

7. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} \text{unroll}_\tau sv : c(\text{Rec}[\kappa]c)c_p \rightsquigarrow \text{unroll}_{|\tau|} sv'$.

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \text{unroll}_{|\tau|} sv' : |c(\text{Rec}[\kappa]c)c_p|$$

By inversion:

$$\begin{aligned} \Delta \vdash \tau &\equiv \text{Rec}[\kappa](c)(c_p) : T_{32} \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv &: \tau \rightsquigarrow sv' \\ \Psi; \Delta; \Gamma \vdash sv &: \tau \end{aligned}$$

By induction:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash sv' : |\tau|$$

By theorem 17 and weakening:

$$\Delta^{\rho_1, \rho_1} \vdash |\tau| \equiv |\text{Rec}[\kappa](c)(c_p)| : T_{32}$$

Note that $|\text{Rec}[\kappa](c)(c_p)| = \text{Rec}[\kappa](|c|)(|c_p|)$ and $|c(\text{Rec}[\kappa]c)c_p| = (|c|(\text{Rec}[\kappa]|c|)|c_p|)$.

By construction (unroll):

$$\Delta^{\rho_1, \rho_2} \vdash \text{unroll}_{|\tau|} : |\tau| \Rightarrow (|c|(\text{Rec}[\kappa]|c|)|c_p|)$$

By construction (coercion app):

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \text{unroll}_{|\tau|} sv' : |c(\text{Rec}[\kappa]c)c_p|$$

8. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} \text{pack } sv \text{ as } \tau \text{ hiding } c : \tau \rightsquigarrow (\text{pack}[|\tau|]|c|)sv'$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash (\text{pack}[|\tau|]|c|)sv' : |\tau|$$

By inversion:

$$\begin{aligned} \Delta \vdash \tau &\equiv \exists[\kappa](c') : T_{32} \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv &: c'c \rightsquigarrow sv' \\ \Psi; \Delta; \Gamma \vdash sv &: c'c \end{aligned}$$

By induction:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash sv' : |c'c|$$

By theorem 17 and weakening:

$$\Delta^{\rho_1, \rho_1} \vdash |\tau| \equiv |\exists[\kappa](c')| : T_{32}$$

Note that $|\exists[\kappa]c'| = \exists[\kappa]|c'|$ and $|c c'| = |c| |c'|$.

By construction (pack):

$$\Delta^{\rho_1, \rho_2} \vdash \text{pack}[|\tau|]|c| : |c'| |c| \Rightarrow |\tau|$$

By construction (coercion app):

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \text{pack}[|\tau|]|c|sv' : |\tau|$$

9. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv[c] : c'c \rightsquigarrow sv'[|c|]$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash sv[c] : |c'c|$$

By inversion:

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : c'c \rightsquigarrow sv'$$

$$\Psi; \Delta; \Gamma \vdash sv : \forall[\kappa](c')$$

$$\Delta \vdash c : \kappa$$

By induction:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash sv' : |\forall[\kappa](c')|$$

By theorem15 and weakening:

$$\Delta^{\rho_1, \rho_2} \vdash |c| : \kappa$$

Note that $|\forall[\kappa]c'| = \forall[\kappa]|c'|$ and $|c c'| = |c| |c'|$.

By construction (forall instantiation):

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash sv'[|c|] : |c'c|$$

■

64 bit values

The translation rules for 64 bit values is exactly analogous to that of 32 bit values. For variables:

$$\frac{}{\Psi; \Delta; \Gamma[x_f:\phi]; \mathcal{A}[x_f \rightarrow \mathbf{f}], \sigma_1, \sigma_2 \vdash_{64} x_f : \phi \rightsquigarrow \mathbf{f}}$$

$$\frac{}{\Psi; \Delta; \Gamma[x_f:\phi]; \mathcal{A}[x_f \rightarrow \mathbf{sp}(i)], \sigma_2, \sigma_2 \vdash_{32} x_f : \phi \rightsquigarrow \mathbf{sp}(i)}$$

And for constants:

$$\frac{}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_2, \sigma_2 \vdash_{32} \mathbf{r} : \text{Float} \rightsquigarrow \mathbf{r}}$$

The soundness theorem is stated and proved in much the same fashion as for 32 bit values except that no induction is required.

Theorem 19 (Soundness of the 64 bit value translation)

If $\Psi; \Delta; \Gamma \vdash fv : \phi$
 and $\mathcal{A}$ is a good allocator for Γ
 and $\Delta^{\rho_1, \rho_2} \vdash \sigma_1 : ST$
 and $\Delta^{\rho_1, \rho_2} \vdash \sigma_2 : ST$
 and $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} fv : \phi \rightsquigarrow fv'$

then

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash fv' : |\phi|$$

Proof: By construction. The proof proceeds by cases on the last rule of the translation derivation.

1. Suppose $\Psi; \Delta; \Gamma[x_f : \phi]; \mathcal{A}[x_f \rightarrow \mathbf{f}], \sigma_1, \sigma_2 \vdash_{64} x_f : \phi \rightsquigarrow \mathbf{f}$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \mathbf{f} : |\phi|$$

By assumption:

$$\Gamma = \Gamma_1, x_f : \phi, \Gamma_2$$

$$\mathcal{A}(x_f) = \mathbf{f}$$

So by the good allocator assumption:

$$|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}(\mathcal{A}(x_f)) = |\phi|$$

So by the float register rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \mathbf{f} : |\phi|$$

2. Suppose $\Psi; \Delta; \Gamma[x_f : \phi]; \mathcal{A}[x_f \rightarrow \mathbf{sp}(i)], \sigma_2, \sigma_2 \vdash_{32} x_f : \phi \rightsquigarrow \mathbf{sp}(i)$.

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \mathbf{sp}(i) : |\phi|$$

By assumption:

$$\Gamma = \Gamma_1, x : \tau, \Gamma_2$$

$$\mathcal{A}(x) = \mathbf{sp}(i)$$

So by the good allocator assumption:

$$|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}(\mathbf{sp}) = \sigma$$

$$\sigma[i]_{64} = |\phi|$$

So by the stack slot rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \mathbf{sp}(i) : |\phi|$$

3. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_2, \sigma_2 \vdash_{32} \mathbf{r} : \mathbf{Float} \rightsquigarrow \mathbf{r}$.

To show:

$$|\Psi[\ell : \tau]|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash \mathbf{r} : \mathbf{Float}$$

This follows immediately by construction.

■

8.3.2 Operations

The translation judgement for **LIL** operations takes the form $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash dest_{32} \leftarrow opr : \tau \rightsquigarrow S$. Whereas the value translation relates **LIL** values and **TILTAL** operands, the operation translation relates a **LIL** operation opr to both a **TILTAL** location $dest_{32}$ and a partial instruction sequence S (as defined in section 7.2.1). The idea behind this relation is that given the machine state assumptions encoded in the allocator ($\mathcal{A}$), the partial instruction sequence S implements the operation opr , leaving the result in the destination $dest_{32}$.

Structuring the translation in this manner is intended to improve the quality of the generated code by (among other things) eliminating un-necessary **mov** instructions. The **TILT** certifying implementation discussed in subsequent chapters uses this idea in a very similar form to that given here in the formal translation. A closely related approach to code generation is also taken by Dybvig et al. in their paper “Destination-Driven Code Generation” [DHB90].

Not all **LIL** operations are given translations by the operation translation. It simplifies the structure of the translation noticeably to translate only certain instructions (raise, handle, and the various case operations) as part of an expression. Consequently, these operations are treated as part of the expression translation.

In general, slightly different code must be produced when the destination is a stack slot versus a register. In order to keep the translation simple, I handle all stack slot destinations with a single rule.

$$\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r}_t \leftarrow opr : \tau \rightsquigarrow S}{\begin{array}{l} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{sp}(i) \leftarrow opr : \tau \rightsquigarrow S; \\ \quad \mathbf{swrite} \ \mathbf{sp}(i), \mathbf{r}_t; \\ \quad \mathbf{junk} \ \mathbf{r}_t \end{array}}$$

This rule simply translates the operation using the temporary register as the destination, then writes the value of the temporary register to the appropriate stack slot and clears the temporary register. In some cases, better code could be produced by adding additional rules for specific operation/destination pairs in which the intermediate usage of the temporary register could be eliminated.

For the inclusion of small values into the operation level, the translation produces an operand from the small value and moves it into the destination register.

$$\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau \rightsquigarrow sv'}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow sv : \tau \rightsquigarrow \mathbf{mov} \ \mathbf{r}, sv'}$$

Note that in many cases, a good allocator will have assigned $\mathbf{r}$ as the location for sv (for example, when sv is a variable which is not live after the operation). In such cases, un-necessary move instructions will be generated. As above, additional translation rules can be added to produce better code for such cases.

A more interesting rule to consider is the translation rule for the **call** operation, which invokes a code function on a list of 64 and 32 bit arguments. Code calls are implemented by writing the code arguments onto the stack and then calling the translated code sequence. Upon return, it is necessary to move the result from the temporary register $\mathbf{r}_t$ to the destination register (since the calling convention specifies that function results are returned in $\mathbf{r}_t$). Note the use of the **movj** pseudo-instruction to ensure correctness in the case that the destination register is $\mathbf{r}_t$.

$$\begin{array}{c}
|\Gamma|_{\mathcal{A}}^{\epsilon} = \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{r}_e:ns_{32}, \mathbf{r}_t:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{sp} : \sigma_f\} \\
\Delta^{\rho_1, \rho_2} \vdash \sigma_f \equiv \underbrace{ns_{32} \triangleright_{32} \cdots \triangleright_{32} ns_{32}}_m \triangleright_{32} \underbrace{ns_{64} \triangleright_{64} \cdots ns_{64}}_k \triangleright_{64} \sigma' : ST \\
\\
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \text{Code}(\tau_0, \dots, \tau_{m-1})(\phi_0, \dots, \phi_{k-1}) \rightarrow \tau \rightsquigarrow sv' \\
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_i : \tau_i \rightsquigarrow sv'_i \quad i \in 0 \dots m-1 \\
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} fv_i : \phi_i \rightsquigarrow fv'_i \quad i \in 0 \dots k-1 \\
\\
\hline
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \text{call } sv(sv_0, \dots, sv_{m-1})(fv_0, \dots, fv_{k-1}) : \tau \rightsquigarrow \\
\quad \text{fswrite } \mathbf{sp}(m+2*(k-1)), fv'_{k-1} \\
\quad \vdots \\
\quad \text{fswrite } \mathbf{sp}(m), fv'_0 \\
\quad \text{swrite } \mathbf{sp}(m-1), sv'_{m-1} \\
\quad \vdots \\
\quad \text{swrite } \mathbf{sp}(0), sv'_0 \\
\quad \mathbf{junkregs} \\
\quad \text{call } sv'[\sigma' \circ \sigma_1, \sigma_2] \\
\quad \mathbf{movj } \mathbf{r}, \mathbf{r}_t
\end{array}$$

This rule illustrates several important ideas used in the translation. Firstly, note that the translation requires that the translation of the typing context under the allocator be a register file in which all of the registers contain only junk. This reflects the fact that all registers in this calling convention are caller-save, and hence may be overwritten during the function call. A translation using a callee-save convention could allow some or all of the registers to be occupied. The implicit effect of this rule is to force any valid translation to spill registers using the spill rule before applying a call rule (or to never map registers at all). This demonstrates a key advantage of defining the translation parametrically with respect to allocation: there is a clean separation of concerns between the expectations of the translator and the mechanism by which the register allocator satisfies those expectations.

Secondly, note that the translation expects the allocator to include space for the out-arguments in the frame. This is not required by the allocator methodology, but is convenient for the purposes of avoiding a frame pointer, as well as permitting the allocator to potentially re-use temporary slots as out-arguments. The translation expresses this requirement on the allocator by premising the translation of a `call` operation on the availability of sufficient unused slots on the top of the frame:

$$\Delta^{\rho_1, \rho_2} \vdash \sigma_f \equiv \underbrace{ns_{32} \triangleright_{32} \cdots \triangleright_{32} ns_{32}}_m \triangleright_{32} \underbrace{ns_{64} \triangleright_{64} \cdots ns_{64}}_k \triangleright_{64} \sigma' : ST$$

A translation only exists if the allocator provides sufficient space in the frame.

This rule also illustrates the use of the derived `movj` partial instruction sequence from section 7.2.1. The technique used to handle stack slot destinations above implies that the translation must be prepared for the destination register to be an alias for the temporary register. This is handled here by the `movj` instruction, which performs a move on its argument registers and junks the source register only if the registers are different. Consequently, if $\mathbf{r} = \mathbf{r}_t$, then the result is left in $\mathbf{r}_t$; and if $\mathbf{r} \neq \mathbf{r}_t$, the result is moved to $\mathbf{r}$ and the $\mathbf{r}_t$ register is junked.

The rest of the operation translation rules proceed in a similar fashion. The complete translation rules can be found in section 8.4 below.

8.3.3 Soundness of the operation translations

The soundness theorem for the operation translation states that for a well-typed operation and a good allocator; whenever the operation is related to a destination and a partial code sequence by the translation, the partial code sequence is well-formed and leaves its result in the destination. Note that the translation assumes that the $\mathbf{r}_e$ register which is left unmapped by the allocator will contain an exception frame.

Theorem 20 (Soundness of the operation translation.)

If $\Psi; \Delta; \Gamma; \mathcal{A} \vdash \text{opr} : \tau \text{ opr}_{32}$ and $\mathcal{A}$ is a good allocator for Γ and $\Delta^{\rho_1, \rho_2} \vdash \sigma_1 : ST$ and $\Delta^{\rho_1, \rho_2} \vdash \sigma_2 : ST$ then:

1. If $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \text{opr} : \tau \rightsquigarrow S$
then
 $|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S \Rightarrow \Gamma_A\{\mathbf{r} : |\tau|\}$
where $|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}\{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_1)\} = \Gamma_A$
2. If $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{sp}(i) \leftarrow \text{opr} : \tau \rightsquigarrow S$
then
 $|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S \Rightarrow \Gamma_A\{\mathbf{sp} : \sigma'\}$
where $|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}\{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_1)\} = \Gamma_A$ and $\Gamma_A(\mathbf{sp}) = \sigma$ and $(\sigma)[i]_{32} \leftarrow |\tau| = \sigma'$.

Proof: By induction on derivations. The proof proceeds by cases on the last rule used. Note that only one rule applies when the destination is a stack slot, and that when the destination is a register, at most one rule applies for each instruction form. Also note that the good allocator assumption guarantees that $|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}(\mathbf{r}_t) = ns_{32}$, and hence for cases that modify $\mathbf{r}_t$, the output typing condition requires that $\mathbf{r}_t$ be coerced to ns_{32} .

Throughout the proof, I use Γ_A to refer to $|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}\{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_1)\}$

1. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1 \sigma_2 \vdash \mathbf{sp}(i) \leftarrow \text{opr} : \tau \rightsquigarrow S$. The stack slot rule is the only rule that applies, so

By inversion:

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r}_t \leftarrow \text{opr} : \tau \rightsquigarrow S$$

Let $\sigma = \Gamma_A(\mathbf{sp})$.

By induction:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S \Rightarrow \Gamma_A\{\mathbf{r}_t : |\tau|\}$$

By the stack write rule and lemma 34:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_t : |\tau|\} \vdash \mathbf{swrite} \mathbf{sp}(i), \mathbf{r}_t; \Rightarrow \Gamma_A\{\mathbf{sp} : (\sigma)[i]_{32} \leftarrow |\tau|\}\{\mathbf{r}_t : ns_{32}\}$$

$$\mathbf{junk} \mathbf{r}_t$$

So by lemma 28 (composition):

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S; \mathbf{swrite} \mathbf{sp}(i), \mathbf{r}_t; \mathbf{junk} \mathbf{r}_t \Rightarrow \Gamma_A\{\mathbf{sp} : (\sigma)[i]_{32} \leftarrow |\tau|\}\{\mathbf{r}_t : ns_{32}\}$$

2. Suppose $opr = sv$.

By assumption:

$$\begin{aligned} \Psi; \Delta; \Gamma; \mathcal{A} \vdash sv : \tau \text{ \textbf{opr}}_{32} \\ \Delta, \Gamma, \mathcal{A} \vdash \mathbf{r} \leftarrow sv : \tau \rightsquigarrow \text{mov } \mathbf{r}, sv' \end{aligned}$$

By inversion:

$$\begin{aligned} \Psi; \Delta; \Gamma; \mathcal{A} \vdash sv : \tau \\ \Delta, \Gamma, \mathcal{A} \vdash_{32} sv : \tau \rightsquigarrow sv' \end{aligned}$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{mov } \mathbf{r}, sv' \Rightarrow \Gamma_A\{\mathbf{r}:|\tau|\}$$

By theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv' : |\tau|$$

By the mov typing rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{mov } \mathbf{r}, sv' \Rightarrow \Gamma_A\{\mathbf{r}:|\tau|\}$$

3. Suppose $opr = \text{select } sv$

By assumption:

$$\begin{aligned} \Psi; \Delta; \Gamma; \mathcal{A} \vdash \text{select } sv : \tau_i \text{ \textbf{opr}}_{32} \\ \Delta, \Gamma, \mathcal{A} \vdash \mathbf{r} \leftarrow \text{select } sv : \tau_i \rightsquigarrow \text{mov } \mathbf{r}, sv' \\ \text{loadr } \mathbf{r}, \mathbf{r}(i) \end{aligned}$$

By inversion:

$$\begin{aligned} \Psi; \Delta; \Gamma \vdash sv : \times (\tau_0, \dots, \tau_i, \dots, \tau_n) \\ \Delta, \Gamma, \mathcal{A} \vdash_{32} sv : \times (\tau_0, \dots, \tau_i, \dots, \tau_n) \rightsquigarrow sv' \end{aligned}$$

To show:

$$\begin{aligned} |\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{mov } \mathbf{r}, sv' &\Rightarrow \Gamma_A\{\mathbf{r}:|\tau_i|\} \\ \text{loadr } \mathbf{r}, \mathbf{r}(i) & \end{aligned}$$

By theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv' : |\times (\tau_0, \dots, \tau_i, \dots, \tau_n)|$$

By definition:

$$|\times (\tau_0, \dots, \tau_i, \dots, \tau_n)| = \times (|\tau_0|, \dots, |\tau_i|, \dots, |\tau_n|)$$

By the mov and load typing rules:

$$\begin{aligned} |\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{mov } \mathbf{r}, sv' &\Rightarrow \Gamma_A\{\mathbf{r}:|\tau_i|\} \\ \text{loadr } \mathbf{r}, \mathbf{r}(i) & \end{aligned}$$

4. Suppose $opr = \text{dyntag}_c$.

By the new tag instruction rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{dyntag}_{|c|} \mathbf{r} \Rightarrow \Gamma_A\{\mathbf{r}:|\text{Dyntag}(c)|\}$$

5. Suppose $opr = \text{box } fv$.

By inversion:

$$\begin{aligned} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{64} fv : \text{Float} \rightsquigarrow fv' \\ \Psi; \Delta; \Gamma \vdash fv : \phi \end{aligned}$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{malloc}_{|\phi|} \mathbf{r}, fv' \Rightarrow \Gamma_A \{\mathbf{r}: |\text{Boxed}(\phi)|\}$$

By assumption $\mathcal{A}$ is a good allocator for Γ , so

By theorem 19:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash fv' : |\phi|$$

By the malloc_ϕ rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{malloc}_{|\phi|} \mathbf{r}, fv' \Rightarrow \Gamma_A \{\mathbf{r}: \text{Boxed}(|\phi|)\}$$

(Note, $|\text{Boxed}(\phi)| = \text{Boxed}(|\phi|)$)

6. Suppose $\text{opr} = \langle sv_1, \dots, sv_n \rangle$.

By inversion:

$$\begin{aligned} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \tau_i \rightsquigarrow sv'_i \\ \Psi; \Delta; \Gamma \vdash sv_i : \tau_i \end{aligned}$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{malloc} \mathbf{r}, [|\tau_1|, \dots, |\tau_n|] \langle sv'_1, \dots, sv'_n \rangle \Rightarrow \Gamma_A \{\mathbf{r}: |\tau_\times|\}$$

where $\tau_\times = \times[\tau_1, \dots, \tau_n]$

By theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv_i : |\tau_i|$$

So the result follows directly by the malloc rule.

7. Suppose $\text{opr} = \text{call } sv(sv_0, \dots, sv_{m-1})(fv_0, \dots, fv_{k-1})$.

By inversion:

$$\begin{aligned} |\Gamma|_A^\epsilon(\mathbf{sp}) = \{\mathbf{r}_1: ns_{32}, \mathbf{r}_2: ns_{32}, \mathbf{r}_e: ns_{32}, \mathbf{r}_t: ns_{32}, \mathbf{f}_1: ns_{64}, \mathbf{f}_2: ns_{64}, \mathbf{sp} : \sigma_f\} \\ \Delta^{\rho_1, \rho_2} \vdash \sigma_f \equiv \underbrace{ns_{32} \triangleright_{32} \dots \triangleright_{32} ns_{32}}_m \triangleright_{32} \underbrace{ns_{64} \triangleright_{64} \dots ns_{64}}_k \triangleright_{64} \sigma' : ST \end{aligned}$$

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \text{Code}(\tau_0, \dots, \tau_{m-1})(\phi_0, \dots, \phi_{k-1}) \rightarrow \tau \rightsquigarrow sv'$$

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_i : \tau_i \rightsquigarrow sv'_i \quad i \in 0 \dots m-1$$

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} fv_i : \phi_i \rightsquigarrow fv'_i \quad i \in 0 \dots k-1$$

$$\Psi; \Delta; \Gamma \vdash sv : \text{Code}[\tau_0, \dots, \tau_n][\phi_0, \dots, \phi_k](\tau)$$

$$\Psi; \Delta; \Gamma \vdash sv_i : \tau_i$$

$$\Psi; \Delta; \Gamma \vdash fv_i : \phi_i$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{fswrite } \mathbf{sp}(m + 2 * (k - 1)), fv'_{k-1} \Rightarrow \Gamma_A \{\mathbf{r}: |\tau|\}$$

$$\begin{aligned} & \vdots \\ & \text{fswrite } \mathbf{sp}(m), fv'_0 \\ & \text{swrite } \mathbf{sp}(m-1), sv'_{m-1} \\ & \vdots \\ & \text{swrite } \mathbf{sp}(0), sv'_0 \\ & \text{call } sv'[\sigma' \circ \sigma_1, \sigma_2] \\ & \text{movj } \mathbf{r}, \mathbf{r}_t \end{aligned}$$

The proof proceeds by stepping through the emitted instructions and applying the appropriate typing rules. For brevity, I will simply describe the register file type after each instruction

rather than restating the entire typing judgement. I also leave the inner inductions on m and k informal.

By theorems 19 and 18:

$$\begin{aligned} |\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash fv'_i : |\phi_i| \\ |\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_i : |\tau_i| \end{aligned}$$

By the good allocator assumption, $\Gamma_A(\mathbf{sp}) = \sigma_f \circ \sigma_1 \circ \sigma_2$.

By assumption:

$$\Delta^{\rho_1, \rho_2} \vdash \sigma_f \equiv \underbrace{ns_{32} \triangleright_{32} \cdots \triangleright_{32} ns_{32}}_m \triangleright_{32} \underbrace{ns_{64} \triangleright_{64} \cdots ns_{64}}_k \triangleright_{64} \sigma' : ST \quad \text{where } \sigma_f = |\Gamma|_A^\epsilon(\mathbf{sp})$$

Therefore, by the **swrite** and **fswrite** rules, it suffices to show that the remainder of the instruction sequence after the stack writes is well-typed assuming that the register file has the type:

$$\Gamma_A\{\mathbf{sp}:\sigma_c\} \quad \text{where } \sigma_c = \tau_0 \triangleright_{32} \cdots \triangleright_{32} \tau_{m-1} \triangleright_{32} \phi_0 \triangleright_{64} \cdots \phi_{k-1} \triangleright_{64} \sigma' \circ \sigma_1 \circ \sigma_2$$

By theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv : |\mathbf{Code}(\tau_0, \dots, \tau_{m-1})(\phi_0, \dots, \phi_{k-1}) \rightarrow \tau|$$

And by definition 14:

$$\begin{aligned} & |\mathbf{Code}(\tau_0, \dots, \tau_{m-1})(\phi_0, \dots, \phi_{k-1}) \rightarrow \tau| = \\ & \quad \forall[\rho_1:ST, \rho_2:ST]. \\ & \quad \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:\mathbf{Exnptr}(\rho_2), \mathbf{r}_t:ns_{32}, \\ & \quad \mathbf{sp}:\mathbf{cont}(|\sigma_{32} \circ \sigma_{64}|)(\rho_1)(\rho_2)(|\tau|) \triangleright_{32} (\sigma_{32} \circ (\sigma_{64} \circ \rho_1 \circ \rho_2))\} \rightarrow 0 \\ & \quad \text{where } \sigma_{32} = \tau_0 \triangleright_{32} \cdots \triangleright_{32} \tau_{m-1} \triangleright_{32} \epsilon \\ & \quad \text{and } \sigma_{64} = \phi_0 \triangleright_{64} \cdots \phi_{k-1} \triangleright_{64} \epsilon \end{aligned}$$

So by the instantiation rule, $sv[\sigma' \circ \sigma_1, \sigma_2]$ has type:

$$\begin{aligned} & \Gamma_c \rightarrow 0 \quad \text{where} \\ & \Gamma_c = \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:\mathbf{Exnptr}(\sigma_2), \mathbf{r}_t:ns_{32}, \\ & \quad \mathbf{sp}:\mathbf{cont}(|\sigma_{32} \circ \sigma_{64}|)(\sigma' \circ \sigma_1)(\sigma_2)(|\tau|) \triangleright_{32} (\sigma_{32} \circ (\sigma_{64} \circ \sigma' \circ \sigma_1 \circ \sigma_2))\} \end{aligned}$$

Let $\tau_{\mathbf{ret}} = \mathbf{cont}(|\sigma_{32} \circ \sigma_{64}|)(\sigma' \circ \sigma_1)(\sigma_2)(|\tau|)$

Note that $\Gamma_c = \Gamma_A\{\mathbf{sp}:\tau_{\mathbf{ret}} \triangleright_{32} \sigma_c\}$ and so the **call** rule applies. Therefore, it suffices to show that the remainder of the code sequence is well-typed under $\Gamma_{\mathbf{ret}}$, where $\tau_{\mathbf{ret}} = \Gamma_{\mathbf{ret}} \rightarrow 0$.

By lemma 38, the register file after applying the last move and junk instruction is

$$\Gamma_{\mathbf{ret}}\{\mathbf{r}_t:ns_{32}\}\{\mathbf{r}:|\tau|\}$$

But notice that by definition of **cont** (definition 13):

$$\begin{aligned} \Gamma_{\mathbf{ret}} = \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:\mathbf{Exnptr}(\sigma_2), \mathbf{r}_t:|\tau|, \\ \mathbf{sp}:ns_{32} \xrightarrow{m+2*k} \sigma' \circ \sigma_1 \circ \sigma_2\} \end{aligned}$$

So $\Gamma_{\mathbf{ret}} = \Gamma_A\{\mathbf{r}_t:|\tau|\}$.

Finally, recall that by the good allocator assumption:

$$\Gamma_A\{\mathbf{r}_t:ns_{32}\} = \Gamma_A$$

Hence:

$$\begin{aligned}
& \Gamma_{\text{ret}}\{\mathbf{r}_t:ns_{32}\}\{\mathbf{r}:|\tau|\} \\
&= \Gamma_A\{\mathbf{r}_t:|\tau|\}\{\mathbf{r}_t:ns_{32}\}\{\mathbf{r}:|\tau|\} \\
&= \Gamma_A\{\mathbf{r}_t:ns_{32}\}\{\mathbf{r}:|\tau|\} \\
&= \Gamma_A\{\mathbf{r}:|\tau|\}
\end{aligned}$$

8. Suppose $opr = \text{array}_\tau(sv_1, sv_2)$

By inversion:

$$\begin{aligned}
& \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \text{Int} \rightsquigarrow sv'_1 \\
& \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_2 : \tau \rightsquigarrow sv'_2 \\
& \Psi; \Delta; \Gamma \vdash sv_1 : \text{Int} \\
& \Psi; \Delta; \Gamma \vdash sv_2 : \tau
\end{aligned}$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{malloc}_{|\tau|} \mathbf{r}, sv'_1, sv'_2 \Rightarrow \Gamma_A\{\mathbf{r}:|\text{Array}_{32}(\tau)|\}$$

By theorem 18:

$$\begin{aligned}
& |\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_1 : |\text{Int}| \\
& |\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_2 : |\tau|
\end{aligned}$$

so by the malloc rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{malloc}_{|\tau|} \mathbf{r}, sv'_1, sv'_2 \Rightarrow \Gamma_A\{\mathbf{r}:|\text{Array}_{32}(\tau)|\}$$

9. Suppose $opr = \text{farray}_\tau(sv, fv)$

By inversion:

$$\begin{aligned}
& \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \text{Int} \rightsquigarrow sv' \\
& \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{64} fv : \phi \rightsquigarrow fv' \\
& \Psi; \Delta; \Gamma \vdash sv : \text{Int} \\
& \Psi; \Delta; \Gamma \vdash fv : \phi
\end{aligned}$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{fmalloc}_{|\phi|} \mathbf{r}, sv', fv' \Rightarrow \Gamma_A\{\mathbf{r}:|\text{Array}_{64}(\tau)|\}$$

By theorems 18 and 19:

$$\begin{aligned}
& |\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv' : |\text{Int}| \\
& |\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash fv' : |\tau|
\end{aligned}$$

so by the fmalloc rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{fmalloc}_{|\phi|} \mathbf{r}, sv', fv' \Rightarrow \Gamma_A\{\mathbf{r}:|\text{Array}_{64}(\tau)|\}$$

10. Suppose $opr = \text{sub}_\tau(sv_1, sv_2)$.

By inversion:

$$\begin{aligned}
& \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \text{Array}_{32}(\tau) \rightsquigarrow sv'_1 \\
& \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_2 : \text{Int} \rightsquigarrow sv'_2 \\
& \Psi; \Delta; \Gamma \vdash sv_1 : \text{Array}_{32}(\tau) \\
& \Psi; \Delta; \Gamma \vdash sv_2 : \text{Int}
\end{aligned}$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{sub}_{|\tau|} \mathbf{r}, sv'_1, sv'_2 \Rightarrow \Gamma_A\{\mathbf{r}:|\tau|\}$$

By theorem 18:

$$\begin{aligned} &|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_1 : |\mathbf{Array}_{32}(\tau)| \\ &|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_2 : \mathbf{Int} \end{aligned}$$

so by the sub rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{sub}_{|\tau|} \mathbf{r}, sv'_1, sv'_2 \Rightarrow \Gamma_A \{\mathbf{r}:|\tau|\}$$

11. Suppose $opr = \mathbf{upd}_\tau(sv_1, sv_2, sv_3)$.

By inversion:

$$\begin{aligned} &\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \mathbf{Array}_{32}(\tau) \rightsquigarrow sv'_1 \\ &\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_2 : \mathbf{Int} \rightsquigarrow sv'_2 \\ &\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_3 : \tau \rightsquigarrow sv'_3 \\ &\Psi; \Delta; \Gamma \vdash sv_1 : \mathbf{Array}_{32}(\tau) \\ &\Psi; \Delta; \Gamma \vdash sv_2 : \mathbf{Int} \\ &\Psi; \Delta; \Gamma \vdash sv_3 : \tau \end{aligned}$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{upd} sv'_1, sv'_2, sv'_3 \Rightarrow \Gamma_A \{\mathbf{r:Unit}\} \\ \mathbf{malloc} \mathbf{r}, [] \langle \rangle$$

By theorem 18:

$$\begin{aligned} &|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_1 : |\mathbf{Array}_{32}(\tau)| \\ &|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_2 : \mathbf{Int} \\ &|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_3 : |\tau| \end{aligned}$$

So by the upd rule:

$$\begin{aligned} &|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{upd} sv'_1, sv'_2, sv'_3 \\ &\Rightarrow \Gamma_A \end{aligned}$$

And by the malloc rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{malloc} \mathbf{r}, [] \langle \rangle \Rightarrow \Gamma_A \{\mathbf{r:Unit}\}$$

12. Suppose $opr = \mathbf{fupd}_\phi(sv_1, sv_2, fv)$.

By inversion:

$$\begin{aligned} &\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \mathbf{Array}_{32}(\tau) \rightsquigarrow sv'_1 \\ &\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_2 : \mathbf{Int} \rightsquigarrow sv'_2 \\ &\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{64} fv : \phi \rightsquigarrow fv' \\ &\Psi; \Delta; \Gamma \vdash sv_1 : \mathbf{Array}_{32}(\tau) \\ &\Psi; \Delta; \Gamma \vdash sv_2 : \mathbf{Int} \\ &\Psi; \Delta; \Gamma \vdash fv : \phi \end{aligned}$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{fupd} sv'_1, sv'_2, fv' \Rightarrow \Gamma_A \{\mathbf{r:Unit}\} \\ \mathbf{malloc} \mathbf{r}, [] \langle \rangle$$

By theorems 18 and 19:

$$\begin{aligned} &|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_1 : |\mathbf{Array}_{32}(\tau)| \\ &|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_2 : \mathbf{Int} \\ &|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash fv' : |\phi| \end{aligned}$$

So by the fupd rule:

$$\begin{aligned} & |\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{fupd } sv'_1, sv'_2, fv' \\ & \Rightarrow \Gamma_A \end{aligned}$$

And by the malloc rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{malloc } \mathbf{r}, \langle \rangle \Rightarrow \Gamma_A \{\mathbf{r}:\text{Unit}\}$$

■

A useful corollary states that state of the machine after executing the translation of an operation is the same as the machine state given by translating the extended context.

Corollary 3 (Operation context extension.)

If $\Psi; \Delta; \Gamma; \mathcal{A} \vdash \text{opr} : \tau \text{ opr}_{32}$
 and $\mathcal{A}$ is a good allocator for Γ and for $\Gamma, x:\tau$
 and $\Delta^{\rho_1, \rho_2} \vdash \sigma_1 : ST$
 and $\Delta^{\rho_1, \rho_2} \vdash \sigma_2 : ST$
 and $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathcal{A}(x) \leftarrow \text{opr} : \tau \rightsquigarrow S$

then

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S \Rightarrow \Gamma_B$$

where $\Gamma_A = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_1)\}$ and $\Gamma_B = |\Gamma, x:\tau|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_1)\}$

Proof: The proof follows trivially by applying theorem 20 and observing that the good allocator assumption implies that the resulting context is equal to the translation of the extended context. ■

The proof of soundness for the 64 bit operation translation follows exactly the same form as that of the 32 bit operation translation.

Theorem 21 (Soundness of the 64 bit operation translation.)

If $\Psi; \Delta; \Gamma; \mathcal{A} \vdash \text{fopr} : \phi \text{ opr}_{64}$ and $\mathcal{A}$ is a good allocator for Γ and $\Delta^{\rho_1, \rho_2} \vdash \sigma_1 : ST$ and $\Delta^{\rho_1, \rho_2} \vdash \sigma_2 : ST$ then:

1. If $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{f} \leftarrow \text{fopr} : \phi \rightsquigarrow S$
 then

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S \Rightarrow \Gamma_A \{\mathbf{f} : |\phi|\}$$

where $|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_1)\} = \Gamma_A$

2. If $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{sp}(i) \leftarrow \text{fopr} : \phi \rightsquigarrow S$
 then

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S \Rightarrow \Gamma_A \{\mathbf{sp} : \sigma'\}$$

where $|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_1)\} = \Gamma_A$ and $\Gamma_A(\mathbf{sp}) = \sigma$ and $(\sigma)[i]_{64} \leftarrow |\phi| = \sigma'$.

Proof: By induction on derivations. The proof proceeds by cases on the last rule used. Note that only one rule applies when the destination is a stack slot, and that when the destination is a register, at most one rule applies for each instruction form. Also note that the good allocator assumption guarantees that $|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}(\mathbf{f}_t) = ns_{32}$, and hence for cases that modify $\mathbf{f}_t$, the output typing condition requires that $\mathbf{f}_t$ be coerced to ns_{64} .

Throughout the proof, I use Γ_A to refer to $|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_1)\}$

1. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1 \sigma_2 \vdash \mathbf{sp}(i) \leftarrow fopr : \phi \rightsquigarrow S$. The stack slot rule is the only rule that applies, so

By inversion:

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{f}_t \leftarrow fopr : \phi \rightsquigarrow S$$

Let $\sigma = \Gamma_A(\mathbf{sp})$.

By induction:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S \Rightarrow \Gamma_A\{\mathbf{f}_t : |\phi|\}$$

By the stack write rule and lemma 34:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{f}_t : |\phi|\} \vdash \mathbf{fswrite} \ \mathbf{sp}(i), \mathbf{f}_t; \Rightarrow \Gamma_A\{\mathbf{sp}(\sigma)[i]_{64} \leftarrow |\phi|\}\{\mathbf{f}_t : ns_{64}\} \\ \mathbf{junk} \ \mathbf{f}_t$$

So by lemma 28 (composition):

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{f}_t : |\phi|\} \vdash S; \quad \Rightarrow \Gamma_A\{\mathbf{sp}(\sigma)[i]_{64} \leftarrow |\phi|\}\{\mathbf{f}_t : ns_{64}\} \\ \mathbf{fswrite} \ \mathbf{sp}(i), \mathbf{f}_t; \\ \mathbf{junk} \ \mathbf{f}_t$$

2. Suppose $fopr = fv$.

By assumption:

$$\Psi; \Delta; \Gamma; \mathcal{A} \vdash fv : \phi \ \mathbf{opr}_{64} \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{f} \leftarrow fv : \tau \rightsquigarrow \mathbf{fmov} \ \mathbf{f}, fv'$$

By inversion:

$$\Psi; \Delta; \Gamma; \mathcal{A} \vdash fv : \phi \\ \Psi; \Delta, \Gamma, \mathcal{A} \vdash_{64} fv : \phi \rightsquigarrow fv'$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{fmov} \ \mathbf{f}, fv' \Rightarrow \Gamma_A\{\mathbf{f} : |\phi|\}$$

By theorem 19:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash fv' : |\phi|$$

By the fmov typing rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{fmov} \ \mathbf{f}, sv' \Rightarrow \Gamma_A\{\mathbf{f} : |\phi|\}$$

3. Suppose $fopr = \mathbf{fsub}_\phi(sv_1, sv_2)$.

By inversion:

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \mathbf{Array}_{64}(\phi) \rightsquigarrow sv'_1 \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_2 : \mathbf{Int} \rightsquigarrow sv'_2 \\ \Psi; \Delta; \Gamma \vdash sv_1 : \mathbf{Array}_{64}(\phi) \\ \Psi; \Delta; \Gamma \vdash sv_2 : \mathbf{Int}$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{sub}_{|\phi|} \ \mathbf{f}, sv'_1, sv'_2 \Rightarrow \Gamma_A\{\mathbf{f} : |\phi|\}$$

By theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_1 : |\mathbf{Array}_{64}(\phi)| \\ |\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv'_2 : \mathbf{Int}$$

so by the fsub rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{sub}_{|\phi|} \ \mathbf{f}, sv'_1, sv'_2 \Rightarrow \Gamma_A\{\mathbf{f} : |\phi|\}$$

4. Suppose $fopr = \text{unbox } sv$

By inversion:

$$\begin{aligned} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \text{Boxed}(\phi) &\rightsquigarrow sv' \\ \Psi; \Delta; \Gamma \vdash sv : \text{Boxed } \phi \end{aligned}$$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{floadr } \mathbf{f}, sv' \Rightarrow \Gamma_A \{\mathbf{f}:|\phi|\}$$

By theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv' : |\text{Boxed}(\phi)|$$

so by the floadr rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{floadr } \mathbf{f}, sv' \Rightarrow \Gamma_A \{\mathbf{f}:|\phi|\}$$

■

As before, a useful corollary states that state of the machine after executing the translation of a 64 bit operation is the same as the machine state given by translating the extended context.

Corollary 4 (Operation context extension.)

If $\Psi; \Delta; \Gamma; \mathcal{A} \vdash fopr : \phi \text{ opr}_{32}$
and $\mathcal{A}$ is a good allocator for Γ and for $\Gamma, x_f : \tau$
and $\Delta^{\rho_1, \rho_2} \vdash \sigma_1 : ST$
and $\Delta^{\rho_1, \rho_2} \vdash \sigma_2 : ST$
and $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathcal{A}(x_f) \leftarrow fopr : \phi \rightsquigarrow S$
then

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S \Rightarrow \Gamma_B$$

where $\Gamma_A = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_1)\}$ and $\Gamma_B = |\Gamma, x_f : \phi|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_1)\}$

Proof: The proof follows trivially by applying theorem 21 and observing that the good allocator assumption implies that the resulting context is equal to the translation of the extended context. ■

8.3.4 Expressions

The general expression translation judgement takes the form $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F$. **LIL** expressions translate to **TILTAL** instruction sequences I , terminated either by a return, a jump, or a halt, depending on the context of occurrence of the expression. In addition, the translation of nested sub-expressions may produce a heap fragment F containing additional code blocks which must be allocated in the heap at the top level. All new labels generated by the translation are assumed to be fresh.

Occurrence parameters

As usual, the translation relies on an allocator and two stack types describing the stack layout below the current stack frame. Additionally, the translation is parameterized by a context of occurrence C indicating the context in which the translated term occurs. For the purposes of this translation, contexts range over return contexts, jump contexts, and halt contexts.

$$C :: = \text{ret} \mid \text{jmp } sv \mid \text{halt}$$

Return contexts indicate that the expression is a function body and should exit using the function return convention. Jump contexts indicate the expression occurs as a sub-block of a larger expression, and should exit by placing the return value in $\mathbf{r}_t$ and jumping to the provided address. Halt contexts indicate that the expression occurs as a program body and should exit by terminating the program.

Additional contexts are possible: for example the implementation uses a special context to efficiently translate boolean expressions whose only use is to decide branches. The implementation also uses this mechanism to implement proper tail-recursion: function calls occurring at the end of expressions translated in the return context are tail calls. The formal translation does not address this optimization since it adds nothing interesting to the development. As with the use of the destination parameter in the operation translation, this idea is closely related to the techniques used by Dybvig, et al [DHB90].

Most of the translation inference rules are parametric with respect to the context of occurrence. The exceptions are the rules for the small value base case of expressions, where the appropriate terminator for the instruction sequence is chosen.

Return values

For return contexts, the translation inference rule requires that the top element of the stack (below the frame) be an appropriate continuation type². The code sequence emitted moves the return value to $\mathbf{r}_t$, de-allocates the stack frame, coerces the registers and in-arguments to nonsense, and returns.

$$\frac{\begin{array}{l} \Delta^{\rho_1, \rho_2} \vdash \sigma_1 \equiv (\mathbf{cont}(\overline{m})(\sigma'_1)(\sigma_2)(|\tau_r|)) \triangleright_{32} \sigma'_1 : ST \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau_r \leadsto sv' \quad \text{frmsz}(\mathcal{A}) = n \end{array}}{\begin{array}{l} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\mathbf{ret}} sv : \tau_r \leadsto \mathbf{mov} \ \mathbf{r}_t, sv'; \\ \quad \mathbf{sfree} \ n; \\ \quad \mathbf{junkregs}; \\ \quad \mathbf{junkstack} \ 1 \dots m; \\ \quad \mathbf{ret}; \end{array}}$$

For expressions translated in a jump context, the context of occurrence carries with it the destination of the jump. The translation rule makes no explicit assumptions about the state of the stack, but insists that the type of the destination be appropriate for a jump. The translated code moves the result value to the temporary register and jumps to the destination. Note that the destination is assumed to have been previously instantiated with all of the necessary type variables: it is not sound to simply instantiate it with the current typing context Δ , since Δ may contain more type variables than the destination is expecting.

$$\frac{\begin{array}{l} |\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash sv_l : \Gamma\{\mathbf{r}_t : |\tau|\} \rightarrow 0 \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau \leadsto sv' \end{array}}{\begin{array}{l} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\mathbf{jmp} \ sv_l} sv : \tau \leadsto \mathbf{mov} \ \mathbf{r}_t, sv'; \\ \quad \mathbf{jmp} \ sv_l; \end{array}}$$

²It is an invariant of the translation as stated that this should always be true. Note that the soundness of the translation does not rely on this invariant, since the existence of a translation derivation is predicated on it being satisfied. However, a proof of completeness of the term translation would require that this assumption be shown valid.

Finally, expressions translated in a halt context move the result value to the temporary register, de-allocate the stack frame, and halt the program. Note that the translation insists that the stack segments below the current frame be empty.

$$\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \epsilon, \epsilon \vdash_{32} sv : \tau_r \rightsquigarrow sv' \quad \text{frmsz}(\mathcal{A}) = n}{\Psi; \Delta; \Gamma; \mathcal{A}, \epsilon, \epsilon \vdash_{\text{halt}} sv : \tau_r \rightsquigarrow \text{mov } \mathbf{r}_t, sv'; \\ \text{sfree } n; \\ \text{halt}_{\tau_r};}$$

Spilling

For the most part, all of the exception rules are syntax directed. However, I include one non-syntax directed rule in the translation to give the allocator more flexibility in its choices of locations. The non-deterministic *spill* rule allows a shift to a new good allocator so long as a code sequence can be found that re-arranges the frame appropriately.

$$\frac{\begin{array}{l} \Psi; \Delta; \Gamma; \mathcal{A}', \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I \\ |\Psi|; \Delta^{\rho_1, \rho_1}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash S \Rightarrow |\Gamma|_{\mathcal{A}'}^{\sigma_1 \circ \sigma_2} \\ \text{where } \mathcal{A}' \text{ is a good allocator for } e \end{array}}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow S; I}$$

This rule is intended to capture the fact that a register allocator may wish to insert code into the instruction stream at various points (such as spill and restore code). Since various instructions (such as function calls) insist that the register file be empty, this rule gives the allocator an opportunity to obey this constraint without forcing all variables live across the call to be permanently stack resident.

Operations

The translation of ordinary **LIL** operation binding demonstrates the interaction between the allocator and the destination based translation used for operations.

$$\frac{\begin{array}{l} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathcal{A}(x) \leftarrow opr : \tau_i \rightsquigarrow S \\ \Psi; \Delta; \Gamma, x : \tau_i; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I \end{array}}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let}_\tau x = opr \text{ in } e : \tau \rightsquigarrow S; \\ I}$$

The operation *opr* is translated with respect to the destination chosen by the allocator for *x* to produce a partial instruction sequence *S* that implements the operation and leaves the result in the chosen destination. The rest of the expression is translated to produce an instruction sequence *I*, to which *S* is prepended. Uses of *x* within *e* will be translated via the allocator to references to the chosen location for *x*.

Operations translated as part of the expression translation

Some operations are not given direct translations as part of the operation translation, but are instead translated directly by the expression translation. This is done to keep the operation trans-

$$\begin{array}{c}
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \text{Dyn} \rightsquigarrow sv' \quad \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \text{Dyntag}(\tau_1) \rightsquigarrow sv' \\
\Psi; \Delta; \Gamma, x_1 : \times [\text{Dyntag}(\tau_1), \tau_1]; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp } \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_1, \sigma_2]} e_1 : \tau_x \rightsquigarrow I_1 \\
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp } \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_1, \sigma_2]} e_2 : \tau_x \rightsquigarrow I_2 \\
\Psi; \Delta; \Gamma, x : \tau; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I \\
|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} = \Gamma_A \\
\hline
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let } x = \text{dyncase}(sv)(sv_1 \Rightarrow x_1.e_1, - \Rightarrow e_2) \text{ in } e : \tau \rightsquigarrow \\
\begin{array}{l}
\text{mov } \mathbf{r}_t, sv' \\
\text{brdyn } \mathbf{r}_t, sv'_1, \ell_1[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\
\text{junk } \mathbf{r}_t \\
I_2 \\
\ell_1 : \forall[|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A[\mathbf{r}_t : \times [\text{Dyntag}(|\tau_1|), |\tau_1|]] \\
\text{srmov } \mathcal{A}(x_1), \mathbf{r}_t \\
\text{junk } \mathbf{r}_t \\
I_1 \\
\ell_{\text{end}} : \forall[|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A[\mathbf{r}_t : |\tau_x|] \\
\text{srmov } \mathcal{A}(x), \mathbf{r}_t \\
\text{junk } \mathbf{r}_t \\
I
\end{array}
\end{array}$$

Figure 8.3: Exception case translation in **TILTAL**.

lation simpler, usually because the operations generate additional heap allocated code blocks. Sum and exception case analysis and handlers all fall into this category.

The translation of exception case analysis (figure 8.3) produces two new code blocks and a code sequence. The code sequence compares the tag of an exception in register $\mathbf{r}_t$ to a known tag, and if equal branches to the first additional code block: ℓ_1 . This code block moves the refined exception into the destination assigned to x_1 by the allocator, and then executes the translation of the body of the case statement e_1 . Notice though that the context parameter for the expression translation is changed to indicate the block should exit by jumping to the merge point ℓ_{end} . Since code blocks must be closed, all of the free type variables must be threaded through the blocks: hence the instantiation of the destination label $\ell_{\text{end}}[(\Pi_1 \Delta), \sigma_1, \sigma_2]$. (I use the notation $\Pi_1 \Delta$ to indicate the constructor variables bound in the domain of Δ .)

In the case that the branch is not taken, the other branch of the case is executed. Note that it too is translated in a context with the merge point as its jump destination. The remainder of the expression, e , is translated in the original context of occurrence, and then is emitted with the merge point label ℓ_{end} , with the free type variables suitably abstracted. Note that by convention the continuation expects its argument in $\mathbf{r}_t$.

The translation of sum cases follows essentially the same form, but with substantially more cases. The principle additional complication is that numerous branches must be generated to distinguish between the various tags and tagged record within the union type. No effort is made in the formal translation to emit these in an especially efficient fashion (for example by using general comparisons and intermediate branches instead of simple equality tests).

Exception handling is handled in a relatively inefficient manner for simplicity in the translation

(see figure 8.4). The translation insists that the allocator spill all registers to the stack at every handler site, and then pushes the old exception handler frame onto the stack and copies the entire frame over it. The handled expression is translated with a jump continuation that forwards control to a postlude which cleans up the stack and restores the old handler before jumping to the merge point. The handler code itself does a similar cleanup before executing the handler code, which is translated with the final merge point as its continuation.

Raising exceptions is also dealt with at the expression level, not because it requires additional code blocks to be emitted, but rather because exception raising code ends with an unconditional jump and hence does not fit into the syntactic category of partial instruction sequences. To raise an exception, the exception packet must be placed in $\mathbf{r}_1$, the handler stack must be restored from the exception frame before jumping to the handler, also obtained from the frame.

$$\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \text{Dyn} \rightsquigarrow sv'}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let } x = \text{raise}_\tau sv \text{ in } e : \tau \rightsquigarrow \begin{array}{l} \text{mov } \mathbf{r}_1, sv \\ \text{loadr } \mathbf{r}_t, \mathbf{r}_e(1) \\ \text{mov } \mathbf{sp}, \mathbf{r}_t \\ \text{loadr } \mathbf{r}_t, \mathbf{r}_e(0) \\ \text{jmp } \mathbf{r}_t \end{array}}$$

Type operations

The remainder of the expression translation inference rules deal with the type instructions, such as unpacking existentials and refining types using the LX based primitives. Since these constructs are essentially the same in **LIL** and **TILTAL**, the translation does little interesting work. For example, the refinement operation for pairs simply translates the rest of the expression and prepends a pair refinement instruction.

$$\frac{\begin{array}{l} \Delta \vdash c \equiv \alpha : \kappa_1 \times \kappa_2 \\ \left. \begin{array}{l} \Psi; \Delta, \beta : \kappa_1, \gamma : \kappa_2; \\ \Gamma[\langle \beta, \gamma \rangle / \alpha]; \\ \mathcal{A}, \sigma_1[\langle \beta, \gamma \rangle / \alpha], \sigma_2[\langle \beta, \gamma \rangle / \alpha] \end{array} \right\} \vdash_{C[\langle \beta, \gamma \rangle / \alpha]} e[\langle \beta, \gamma \rangle / \alpha] : \tau[\langle \beta, \gamma \rangle / \alpha] \rightsquigarrow I \end{array}}{\Psi; \Delta, \alpha : \kappa_1 \times \kappa_2; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let } \langle \beta, \gamma \rangle = c \text{ in } e : \tau \rightsquigarrow \text{refine } \langle \beta, \gamma \rangle = |c|; I}$$

The rest of the type instruction rules are similarly un-interesting and are not described here further. A complete listing of all of the expression translation is given in section 8.4.

8.3.5 Soundness of the expression translation

The soundness theorem for the expression translation says that if a well-formed **LIL** expression translates to an instruction sequence and a heap fragment, then both the instruction sequence and the heap fragment are well-formed. Note that the heap context Ψ is extended with the additional bindings introduced by the new heap when typechecking the instruction sequence. As before, I assume that the allocator is a good allocator for the expression, and that the stack segment parameters are well-formed.

$$\begin{array}{c}
\Psi; \Delta; \Gamma; \mathcal{A}, \epsilon, \sigma_3 \vdash_{\text{jmp } \ell_{\text{post}}[(\Pi_1 \Delta), \epsilon, \sigma_h]} e_1 : \tau_x \rightsquigarrow I_1; F_1 \\
\Psi; \Delta; \Gamma, x_2 : \text{Dyn}; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp } \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_f \circ \sigma_1, \sigma_2]} e_2 : \tau_x \rightsquigarrow I_2; F_2 \\
\Psi; \Delta; \Gamma, x : \tau_x; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F \\
\Gamma_A = \{\mathbf{r}_1 : ns_{32}, \mathbf{r}_2 : ns_{32}, \mathbf{r}_t : ns_{32}, \mathbf{r}_e : \mathbf{Exnptr}(\sigma_2), \mathbf{sp} : \sigma_f \circ \sigma_1 \circ \sigma_2, \mathbf{f}_1 : ns_{64}, \mathbf{f}_2 : ns_{64}\} \\
\text{where } \Gamma_A = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_2) \\
\text{and } \Gamma_B = |\Gamma|_{\mathcal{A}}^{\sigma_3} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_3) \\
\text{and } \sigma_f = |\Gamma|_{\mathcal{A}}^{\epsilon}(\mathbf{sp}) \\
\text{and } \sigma_3 = \mathbf{Exnptr}(\sigma_2) \triangleright_{32} \sigma_f \circ \sigma_1 \circ \sigma_2 \\
\text{and } \sigma_h = \sigma_f \circ \sigma_3 \\
\hline
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let } x = \text{handle}_{\tau_x}(e_1, x_2.e_2) \text{ in } e : \tau \rightsquigarrow \\
\text{push } \mathbf{r}_e \\
\text{stackcopy } (\mathbf{Exnptr}(\sigma_2)) \triangleright_{32} \sigma_f \\
\text{sfree } 1 \\
\text{malloc } \mathbf{r}_e[\mathbf{Exnhndler}(\sigma_h), \sigma_h] \langle \ell_{\text{handle}}[\Pi_1 \Delta, \rho_1, \rho_2], \mathbf{sp} \rangle \\
I_1 \\
\ell_{\text{handle}} : \forall [\Delta, \rho_1, \rho_2]. \{\mathbf{r}_1 : \text{Dyn}, \mathbf{r}_2 : ns_{32}, \mathbf{r}_t : ns_{32}, \mathbf{r}_e : ns_{32}, \mathbf{sp} : \sigma_h, \mathbf{f}_1 : ns_{64}, \mathbf{f}_2 : ns_{64}\} \\
\text{srmov } \mathcal{A}(x_2), \mathbf{r}_1 \\
\text{sfree}(\text{frmsz}(\mathcal{A})) \\
\text{pop } \mathbf{r}_e \\
\text{junk } \mathbf{r}_1 \\
I_2 \\
\ell_{\text{post}} : \forall [|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_B[\mathbf{r}_t : |\tau_x|] \\
\text{sfree}(\text{frmsz}(\mathcal{A})) \\
\text{pop } \mathbf{r}_e \\
\text{jmp } \ell_{\text{end}}[\Pi_1 \Delta, \sigma_1, \sigma_2] \\
\ell_{\text{end}} : \forall [|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A[\mathbf{r}_t : |\tau_x|] \\
\text{srmov } \mathcal{A}(x), \mathbf{r}_t \\
\text{junk } \mathbf{r}_t \\
I; \\
F; \\
F_1; \\
F_2
\end{array}$$

Figure 8.4: Exception handler implementation in **TILTAL**.

Theorem 22 (Soundness of the expression translation.)

If $\mathcal{A}$ is a good allocator for e
 and $\Psi; \Delta; \Gamma \vdash e : \tau \text{ exp}$
 and $\Delta^{\rho_1, \rho_2} \vdash \sigma_1 : ST$
 and $\Delta^{\rho_1, \rho_2} \vdash \sigma_2 : ST$
 and $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F$

then

$|\Psi| \vdash F : \Psi_F$
 and $|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash I \text{ ok}$
 where $\Gamma_A = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_2)\}$

Proof: By induction on derivations. The proof proceeds by cases on the last rule of the translation derivation. I generally identify the rule under consideration by listing the conclusion of the rule, except where ambiguous.

Throughout the proof, let $\Gamma_A = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_2)\}$.

1. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{ret}} sv : \tau_r \rightsquigarrow \text{mov } \mathbf{r}_t, sv';$
 $\text{sfree } n;$
 $\text{junkregs};$
 $\text{junkstack } 1 \dots m;$
 ret

To show:

$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{mov } \mathbf{r}_t, sv'; \quad \text{ok}$
 $\text{sfree } n;$
 $\text{junkregs};$
 $\text{junkstack } 1 \dots m;$
 ret

By assumption:

$\Psi; \Delta; \Gamma; \mathcal{A} \vdash sv : \tau_r$
 $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau_r \rightsquigarrow sv'$
 $\Delta^{\rho_1, \rho_2} \vdash \sigma_1 \equiv (\mathbf{cont}(\overline{m})(\sigma'_1)(\sigma_2)(|\tau_r|)) \triangleright_{32} \tau_1 \triangleright_{32} \dots \triangleright_{32} \tau_m \triangleright_{32} \sigma'_1 : ST$

So by theorem 18:

$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv' : |\tau_r|$

By the **mov** typing rule:

$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{mov } \mathbf{r}_t, sv' \Rightarrow \Gamma_A \{\mathbf{r}_t : |\tau_r|\}$

By the good allocator assumption, $\Gamma_A(\mathbf{sp}) = \sigma_f \circ (\sigma_1 \circ \sigma_2)$ and $\text{frmsz}((\cdot)\mathcal{A}) = |\sigma_f|$

So by the **sfree** typing rule:

$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \{\mathbf{r}_t : |\tau_r|\} \vdash \text{sfree } n \Rightarrow \Gamma_A \{\mathbf{r}_t : |\tau_r|\} \{\mathbf{sp} : (\sigma_1 \circ \sigma_2)\}$

By lemma 34:

$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \{\mathbf{r}_t : |\tau_r|\} \{\mathbf{sp} : (\sigma_1 \circ \sigma_2)\} \vdash \text{junkregs} \Rightarrow$
 $\Gamma_A \{\mathbf{r}_t : |\tau_r|\} \{\mathbf{sp} : (\sigma_1 \circ \sigma_2)\} \{\mathbf{r}_1 : ns_{32}\} \{\mathbf{r}_2 : ns_{32}\} \{\mathbf{f}_1 : ns_{64}\} \{\mathbf{f}_2 : ns_{64}\}$

Let $\Gamma'_A = \Gamma_A \{\mathbf{r}_t : |\tau_r|\} \{\mathbf{sp} : (\sigma_1 \circ \sigma_2)\} \{\mathbf{r}_1 : ns_{32}\} \{\mathbf{r}_2 : ns_{32}\} \{\mathbf{f}_1 : ns_{64}\} \{\mathbf{f}_2 : ns_{64}\}$

By assumption:

$\Delta^{\rho_1, \rho_2} \vdash \sigma_1 \equiv (\mathbf{cont}(\overline{m})(\sigma'_1)(\sigma_2)(|\tau_r|)) \triangleright_{32} \tau_1 \triangleright_{32} \dots \triangleright_{32} \tau_m \triangleright_{32} \sigma'_1 : ST$

So by lemma 35:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma'_A \vdash \mathbf{junkstack} \ 1 \dots m \Rightarrow \Gamma'_A \{ \mathbf{sp}: (ns_{32}^{\overline{m}}) \circ \sigma'_1 \circ \sigma_2 \}$$

And by definition:

$$\mathbf{cont}(\overline{m})(\sigma'_1)(\sigma_2)(|\tau_r|) = \Gamma_{\mathbf{ret}} \rightarrow 0$$

where

$$\Gamma_{\mathbf{ret}} = \{ \mathbf{r}_1: ns_{32}, \mathbf{r}_2: ns_{32}, \mathbf{f}_1: ns_{64}, \mathbf{f}_2: ns_{64}, \mathbf{r}_e: \mathbf{Exnptr}(\sigma_2), \mathbf{r}_t: |\tau_r|, \mathbf{sp}: ns_{32}^{\overline{m}} \circ \sigma'_1 \circ \sigma_2 \}$$

So by the return rule, it suffices to show that :

$$\Gamma_r \{ \mathbf{sp}: ns_{32}^{\overline{m}} \circ \sigma'_1 \circ \sigma_2 \} = \Gamma_{\mathbf{ret}}$$

By the good allocator assumption:

$$\Gamma_r(\mathbf{r}_e) = \Gamma_A(\mathbf{r}_e) = \mathbf{Exnptr}(\sigma_2)$$

So by definition:

$$\Gamma_r = \{ \mathbf{r}_1: ns_{32}, \mathbf{r}_2: ns_{32}, \mathbf{f}_1: ns_{64}, \mathbf{f}_2: ns_{64}, \mathbf{r}_e: \mathbf{Exnptr}(\sigma_2), \mathbf{r}_t: |\tau_r|, \mathbf{sp}: ns_{32}^{\overline{m}} \circ \sigma'_1 \circ \sigma_2 \} = \Gamma_{\mathbf{ret}}$$

2. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\mathbf{jmp} \ sv_l} sv : \tau \rightsquigarrow \mathbf{mov} \ \mathbf{r}_t, sv' \quad \mathbf{jmp} \ sv_l$

To show:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{mov} \ \mathbf{r}_t, sv' \ \mathbf{ok} \quad \mathbf{jmp} \ sv_l$$

By assumption:

$$\begin{aligned} & \Psi; \Delta; \Gamma; \mathcal{A} \vdash sv : \tau \\ & \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau \rightsquigarrow sv' \\ & |\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash sv_l : \Gamma \{ \mathbf{r}_t: |\tau| \} \rightarrow 0 \end{aligned}$$

So by theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv' : |\tau|$$

By the **mov** typing rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{mov} \ \mathbf{r}_t, sv' \Rightarrow \Gamma_A \{ \mathbf{r}_t: |\tau_r| \}$$

And by the **jmp** typing rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \{ \mathbf{r}_t: |\tau_r| \} \vdash \mathbf{jmp} \ sv'_l \ \mathbf{ok}$$

3. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \epsilon, \epsilon \vdash_{\mathbf{halt}} sv : \tau_r \rightsquigarrow \mathbf{mov} \ \mathbf{r}_t, sv' \quad \begin{array}{l} \mathbf{sfree} \ n \\ \mathbf{halt}_{\tau_r} \end{array}$

By inversion:

$$\begin{aligned} & \Psi; \Delta; \Gamma; \mathcal{A}, \epsilon, \epsilon \vdash_{32} sv : \tau_r \rightsquigarrow sv' \\ & \mathbf{frmsz}(\mathcal{A}) = n \\ & \Psi; \Delta; \Gamma; \mathcal{A} \vdash sv : \tau_r \end{aligned}$$

By theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv' : |\tau_r|$$

So by the **mov** typing rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{mov} \ \mathbf{r}_t, sv' \Rightarrow \Gamma_A \{ \mathbf{r}_t: |\tau_r| \}$$

By the good allocator assumption:

$$\Gamma_A(\mathbf{sp}) = |\Gamma|_{\mathcal{A}}^{\epsilon \circ \epsilon} = \sigma_f \circ \epsilon$$

where $|\sigma_f| = \text{frmsz}((\cdot)\mathcal{A}) = n$

So by the **sfree** typing rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_t:|\tau_r|\} \vdash \mathbf{sfree} \ n \Rightarrow \Gamma_A\{\mathbf{r}_t:|\tau_r|\}\{\mathbf{sp}:\epsilon\}$$

And by the **halt** typing rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_t:|\tau_r|\}\{\mathbf{sp}:\epsilon\} \vdash \mathbf{halt}_{|\tau_r|} \ \mathbf{ok}$$

4. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \mathbf{let}_\tau x = \mathit{opr} \ \mathbf{in} \ e : \tau \rightsquigarrow (S; I); F$.

By inversion:

$$\begin{aligned} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 &\vdash \mathcal{A}(x) \leftarrow \mathit{opr} : \tau_i \rightsquigarrow S \\ \Psi; \Delta; \Gamma, x:\tau_i; \mathcal{A}, \sigma_1, \sigma_2 &\vdash_C e : \tau \rightsquigarrow I\Psi; \Delta; \Gamma \vdash \mathit{opr} : \tau_i \ \mathbf{opr}_{32} \\ \Psi; \Delta; \Gamma, x:\tau_i &\vdash e : \tau \ \mathbf{exp} \end{aligned}$$

And by the good allocator assumption, $\mathcal{A}$ is a good allocator for Γ .

So by corollary 3

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S \Rightarrow \Gamma_B$$

Where $\Gamma_B = |\Gamma, x:\tau_i|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e: \mathbf{Exnptr}(\sigma_1)\}$

By the definition of a good allocator for an expression, $\mathcal{A}$ is a good allocator for the sub-expression e , so:

By induction:

$$\begin{aligned} |\Psi| &\vdash F : \Psi_F \\ |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_B &\vdash I \ \mathbf{ok} \end{aligned}$$

And by lemma 29 (partial instruction sequence completion) and heap context weakening:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash (S; I) \ \mathbf{ok}$$

5. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \mathbf{let}_\tau x_f = \mathit{fopr} \ \mathbf{in} \ e : \tau \rightsquigarrow (S; I); F$.

By inversion:

$$\begin{aligned} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 &\vdash \mathcal{A}(x_f) \leftarrow \mathit{fopr} : \phi \rightsquigarrow S \\ \Psi; \Delta; \Gamma, x_f:\phi; \mathcal{A}, \sigma_1, \sigma_2 &\vdash_C e : \tau \rightsquigarrow I\Psi; \Delta; \Gamma \vdash \mathit{fopr} : \phi \ \mathbf{opr}_{32} \\ \Psi; \Delta; \Gamma, x_f:\tau &\vdash e : \tau \ \mathbf{exp} \end{aligned}$$

And by the good allocator assumption, $\mathcal{A}$ is a good allocator for Γ .

So by corollary 4

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S \Rightarrow \Gamma_B$$

Where $\Gamma_B = |\Gamma, x_f:\phi|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e: \mathbf{Exnptr}(\sigma_1)\}$

By the definition of a good allocator for an expression, $\mathcal{A}$ is a good allocator for the sub-expression e , so:

By induction:

$$\begin{aligned} |\Psi| &\vdash F : \Psi_F \\ |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_B &\vdash I \ \mathbf{ok} \end{aligned}$$

And by lemma 29 (partial instruction sequence completion) and heap context weakening:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash (S; I) \ \mathbf{ok}$$

6. Suppose $\Psi; \Delta; \Gamma; \mathcal{A} \vdash_C \text{let}[\alpha, x] = \text{unpack } sv \text{ in } e : \tau \rightsquigarrow \text{unpack } [\alpha, \mathbf{r}_t] sv';$
 $\text{srmov } \mathcal{A}(x), \mathbf{r}_t;$
 $\text{junk } \mathbf{r}_t;$
 $I;$
 F

By inversion:

$$\begin{aligned} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \exists[\kappa][c] \rightsquigarrow sv' \\ \Delta, \alpha : \kappa; \Gamma, x : (c\alpha); \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F \\ \Psi; \Delta; \Gamma \vdash sv : \exists[\kappa](c) \\ \Psi; \Delta, \alpha : \kappa; \Gamma, x : (c\alpha) \vdash e : \tau' \text{ exp } \alpha \notin fv(\tau') \end{aligned}$$

By theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv' : |\exists[\kappa](c)|$$

Since $|\exists[\kappa](c)| = \exists[\kappa](|c|)$, by the unpack rule it suffices to show:

$$\begin{aligned} |\Psi|; \Delta, \alpha : \kappa^{\rho_1, \rho_2}; \Gamma_A \{ \mathbf{r}_t : c \alpha \} \vdash \text{srmov } \mathcal{A}(x), \mathbf{r}_t; \text{ok} \\ \text{junk } \mathbf{r}_t; \\ I \end{aligned}$$

By lemma 36, the machine state after the **srmov** instruction is (informally written):

$$\Gamma_A \{ \mathbf{r}_t : |c \alpha| \} \{ \mathcal{A}(x) : |c \alpha| \}$$

And by lemma 34, the machine state after the **junk** instruction is (also informally written):

$$\begin{aligned} \Gamma_A \{ \mathbf{r}_t : |c \alpha| \} \{ \mathcal{A}(x) : |c \alpha| \} \{ \mathbf{r}_t : ns_{32} \} = \Gamma_A \{ \mathcal{A}(x) : |c \alpha| \} \\ (\text{since } \Gamma_A(\mathbf{r}_t) = ns_{32}). \end{aligned}$$

But by the good allocator assumption:

$$\begin{aligned} \Gamma_A \{ \mathcal{A}(x) : |c \alpha| \} = \Gamma_B \\ \text{where } \Gamma_B = |\Gamma, x : c \alpha|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{ \mathbf{r}_e : \mathbf{Exnptr}(\sigma_1) \} \end{aligned}$$

Finally, by induction and heap context weakening:

$$\begin{aligned} |\Psi| \vdash F : \Psi_F \\ |\Psi|, \Psi_F; \Delta, \alpha : \kappa^{\rho_1, \rho_2}; \Gamma_B \vdash I \text{ ok} \end{aligned}$$

7. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let} \langle \beta, \gamma \rangle = c \text{ in } e : \tau \rightsquigarrow I; F$ and $\Delta \vdash c \equiv \langle c_1, c_2 \rangle : \kappa_1 \times \kappa_2$.

By inversion:

$$\begin{aligned} \Psi; \Delta, ; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e[c_1, c_2/\beta, \gamma] : \tau \rightsquigarrow I; F \\ \Psi; \Delta; \Gamma \vdash e[c_1, c_2/\beta, \gamma] : \tau \text{ exp} \end{aligned}$$

By induction:

$$\begin{aligned} |\Psi| \vdash F : \Psi_F \\ |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash I \text{ ok} \end{aligned}$$

8. Suppose

$$\begin{aligned} \Psi; \Delta, \alpha : \kappa_1 \times \kappa_2; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let} \langle \beta, \gamma \rangle = c \text{ in } e : \tau \rightsquigarrow \text{refine} \langle \beta, \gamma \rangle = |c|; \\ I; \\ F \end{aligned}$$

and $\Delta, \alpha : \kappa_1 \times \kappa_2 \vdash c \equiv \alpha : \kappa_1 \times \kappa_2$.

By inversion:

$$\left. \begin{array}{l} \Psi; \Delta, \beta:\kappa_1, \gamma:\kappa_2; \\ \Gamma[\langle\beta, \gamma\rangle/\alpha]; \\ \mathcal{A}, \sigma_1[\langle\beta, \gamma\rangle/\alpha], \sigma_2[\langle\beta, \gamma\rangle/\alpha] \end{array} \right\} \vdash_{C[\langle\beta, \gamma\rangle/\alpha]} e[\langle\beta, \gamma\rangle/\alpha] : \tau[\langle\beta, \gamma\rangle/\alpha] \rightsquigarrow I; F$$

$$\Psi; \Delta, \beta:\kappa_1, \gamma:\kappa_2, \Delta'; \Gamma[\langle\beta, \gamma\rangle/\alpha] \vdash e[\langle\beta, \gamma\rangle/\alpha] : \tau[\langle\beta, \gamma\rangle/\alpha] \text{ exp}$$

By assumption:

$$\Delta, \alpha:\kappa_1 \times \kappa_2^{\rho_1, \rho_2} \vdash \sigma_1 : ST$$

$$\Delta, \alpha:\kappa_1 \times \kappa_2^{\rho_1, \rho_2} \vdash \sigma_2 : ST$$

So by substitution :

$$\Delta, \beta:\kappa_1, \gamma:\kappa_2^{\rho_1, \rho_2} \vdash \sigma_1[\langle\beta, \gamma\rangle/\alpha] : ST$$

$$\Delta, \beta:\kappa_1, \gamma:\kappa_2^{\rho_1, \rho_2} \vdash \sigma_2[\langle\beta, \gamma\rangle/\alpha] : ST$$

So by induction :

$$|\Psi| \vdash F : \Psi_F$$

$$|\Psi|, \Psi_F; \Delta, \beta:\kappa_1, \gamma:\kappa_2^{\rho_1, \rho_2}; \Gamma_S \vdash I \text{ ok}$$

where $\Gamma_S = |\Gamma[\langle\beta, \gamma\rangle/\alpha]|_{\mathcal{A}}^{(\sigma_1 \circ \sigma_2)[\langle\beta, \gamma\rangle/\alpha]}$

By lemma 33:

$$\Gamma_S = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}[\langle\beta, \gamma\rangle/\alpha] = \Gamma_A[\langle\beta, \gamma\rangle/\alpha]$$

So by the pair refinement rule:

$$|\Psi|, \Psi_F; \Delta, \alpha:\kappa_1 \times \kappa_2^{\rho_1, \rho_2}; \Gamma_A \vdash \text{refine}\langle\beta, \gamma\rangle = |c|; \text{ok}$$

I

9. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let fold } \beta = c \text{ in } e : \tau \rightsquigarrow I; F$ and $\Delta \vdash c \equiv \text{fold}_{\mu j. \kappa} c' : \mu j. \kappa$.

By inversion:

$$\Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e[c'/\beta] : \tau \rightsquigarrow I; F$$

$$\Psi; \Delta; \Gamma \vdash e[c'/\beta] : \tau \text{ exp}$$

By induction:

$$|\Psi| \vdash F : \Psi_F$$

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash I \text{ ok}$$

10. Suppose

$$\Psi; \Delta, \alpha:\mu j. \kappa, \Delta'; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let fold } \beta = c \text{ in } e : \tau \rightsquigarrow \text{refine}[\text{fold } \beta]c;$$

$I;$
 F

and $\Delta, \alpha:\mu j. \kappa, \Delta' \vdash c \equiv \alpha : \mu j. \kappa$.

By inversion:

$$\left. \begin{array}{l} \Psi; \Delta, \beta:\kappa[\mu j. \kappa/j], \Delta'; \\ \Gamma[\text{fold}_{\mu j. \kappa} \beta/\alpha]; \\ \mathcal{A}, \\ \sigma_1[\text{fold}_{\mu j. \kappa} \beta/\alpha], \\ \sigma_2[\text{fold}_{\mu j. \kappa} \beta/\alpha] \end{array} \right\} \vdash_{C[\text{fold}_{\mu j. \kappa} \beta/\alpha]} e[\text{fold}_{\mu j. \kappa} \beta/\alpha] : \tau[\text{fold}_{\mu j. \kappa} \beta/\alpha] \rightsquigarrow I; F$$

$$\Psi; \Delta, \beta:\kappa[\mu j. \kappa/j], \Delta'; \Gamma[\text{fold}_{\mu j. \kappa} \beta/\alpha] \vdash e[\text{fold}_{\mu j. \kappa} \beta/\alpha] : \tau[\text{fold}_{\mu j. \kappa} \beta/\alpha] \text{ exp}$$

By assumption:

$$\begin{aligned}\Delta, \alpha: \mu j. \kappa, \Delta'^{\rho_1, \rho_2} &\vdash \sigma_1 : ST \\ \Delta, \alpha: \mu j. \kappa, \Delta'^{\rho_1, \rho_2} &\vdash \sigma_2 : ST\end{aligned}$$

So by substitution :

$$\begin{aligned}\Delta, \beta: \kappa[\mu j. \kappa / j], \Delta'^{\rho_1, \rho_2} &\vdash \sigma_1[|\mathbf{fold}_{\mu j. \kappa} \beta|/\alpha] : ST \\ \Delta, \beta: \kappa[\mu j. \kappa / j], \Delta'^{\rho_1, \rho_2} &\vdash \sigma_2[|\mathbf{fold}_{\mu j. \kappa} \beta|/\alpha] : ST\end{aligned}$$

So by induction :

$$\begin{aligned}|\Psi| &\vdash F : \Psi_F \\ |\Psi|, \Psi_F; \Delta, \beta: \kappa[\mu j. \kappa / j], \Delta'^{\rho_1, \rho_2}; \Gamma_S &\vdash I \mathbf{ok} \\ \text{where } \Gamma_S &= |\Gamma[\mathbf{fold}_{\mu j. \kappa} \beta/\alpha]|_{\mathcal{A}}^{(\sigma_1 \circ \sigma_2)}[|\mathbf{fold}_{\mu j. \kappa} \beta|/\alpha]\end{aligned}$$

By lemma 33:

$$\Gamma_S = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}[|\mathbf{fold}_{\mu j. \kappa} \beta|/\alpha] = \Gamma_A[|\mathbf{fold}_{\mu j. \kappa} \beta|/\alpha]$$

So by the fold refinement rule:

$$|\Psi|, \Psi_F; \Delta, \alpha: \mu j. \kappa, \Delta'^{\rho_1, \rho_2}; \Gamma_A \vdash \mathbf{refine}[\mathbf{fold} \beta]c; \mathbf{ok} \\ I$$

11. Suppose $\Psi; \Delta; \Gamma \vdash_C \mathbf{let}_{\tau} \mathbf{inj}_i \beta = (c, sv) \mathbf{in} e : \tau \rightsquigarrow I; F$ and $\Delta \vdash c \equiv \mathbf{inj}_i^{+[\kappa_1, \dots, \kappa_i, \dots, \kappa_n]} c' : +[\kappa_1, \dots, \kappa_i, \dots, \kappa_n]$.

By inversion:

$$\begin{aligned}\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 &\vdash_C e[c'/\beta] : \tau \rightsquigarrow I; F \\ \Psi; \Delta; \Gamma &\vdash e[c'/\beta] : \tau \mathbf{exp}\end{aligned}$$

By induction:

$$\begin{aligned}|\Psi| &\vdash F : \Psi_F \\ |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A &\vdash I \mathbf{ok}\end{aligned}$$

12. Suppose

$$\begin{aligned}\Psi; \Delta, \alpha: +[\kappa_1, \dots, \kappa_n], \Delta' &\vdash_C \mathbf{let}_{\tau} \mathbf{inj}_i \beta = (c, sv) \mathbf{in} e : \tau \rightsquigarrow \mathbf{refine}[\mathbf{inj}_i \beta]|c|, sv'; \\ &I; \\ &F\end{aligned}$$

and $\Delta, \alpha: \kappa_1 + \kappa_2, \Delta' \vdash c \equiv \alpha: \kappa_1 + \kappa_2$.

By inversion:

$$\left. \begin{aligned} &\Psi; \Delta, \beta: \kappa_i, \Delta'; \\ &\Gamma[\mathbf{inj}_i \beta/\alpha]; \\ &\mathcal{A}, \sigma_1[\mathbf{inj}_i \beta/\alpha], \sigma_2[\mathbf{inj}_i \beta/\alpha] \end{aligned} \right\} \vdash_C [\mathbf{inj}_i \beta/\alpha] e[\mathbf{inj}_i \beta/\alpha] : \tau[\mathbf{inj}_i \beta/\alpha] \rightsquigarrow I; F$$

$$\begin{aligned}\Psi; \Delta, \beta: \kappa_j, \Delta'; \Gamma[\mathbf{inj}_j \beta/\alpha]; \mathcal{A}, \sigma_1[\mathbf{inj}_j \beta/\alpha], \sigma_2[\mathbf{inj}_j \beta/\alpha] &\vdash_{32} sv[\mathbf{inj}_j \beta/\alpha] : \mathbf{Void} \rightsquigarrow sv' \\ j &\in 1 \dots i-1, i+1 \dots n \\ \Psi; \Delta, \beta: \kappa_i, \Delta'; \Gamma[\mathbf{inj}_i \beta/\alpha] &\vdash e[\mathbf{inj}_i \beta/\alpha] : \tau[\mathbf{inj}_i \beta/\alpha] \mathbf{exp} \\ \Psi; \Delta, \beta: \kappa_j, \Delta'; \Gamma[\mathbf{inj}_j \beta/\alpha] &\vdash sv[\mathbf{inj}_j \beta/\alpha] : \mathbf{Void} \\ j &\in 1 \dots i-1, i+1 \dots n\end{aligned}$$

By assumption:

$$\begin{aligned}\Delta, \alpha: +[\kappa_1, \dots, \kappa_n], \Delta'^{\rho_1, \rho_2} &\vdash \sigma_1 : ST \\ \Delta, \alpha: +[\kappa_1, \dots, \kappa_n], \Delta'^{\rho_1, \rho_2} &\vdash \sigma_2 : ST\end{aligned}$$

So by substitution :

$$\begin{aligned} \Delta, \beta:\kappa_i, \Delta'^{\rho_1, \rho_2} \vdash \sigma_1[|\text{inj}_i \beta|/\alpha] : ST \\ \Delta, \beta:\kappa_i, \Delta'^{\rho_1, \rho_2} \vdash \sigma_2[|\text{inj}_i \beta|/\alpha] : ST \\ j \in 1 \dots n \end{aligned}$$

So by induction :

$$\begin{aligned} |\Psi| \vdash F : \Psi_F \\ |\Psi|, \Psi_F; \Delta, \beta:\kappa_i, \Delta'^{\rho_1, \rho_2}; \Gamma_S \vdash I \text{ ok} \\ \text{where } \Gamma_S = |\Gamma[\text{inj}_i \beta/\alpha]|_{\mathcal{A}}^{(\sigma_1 \circ \sigma_2)[|\text{inj}_i \beta|/\alpha]} \end{aligned}$$

By lemma 33:

$$\Gamma_S = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}[|\text{inj}_i \beta|/\alpha] = \Gamma_A[|\text{inj}_i \beta|/\alpha]$$

By theorem 18 (small value translation):

$$\begin{aligned} |\Psi|; \Delta\beta:\kappa_i, \Delta'^{\rho_1, \rho_2}; \Gamma_V \vdash sv' : \text{Void} \\ j \in 1 \dots i-1, i+1 \dots n \\ \text{where } \Gamma_S = |\Gamma[\text{inj}_i \beta/\alpha]|_{\mathcal{A}}^{(\sigma_1 \circ \sigma_2)[|\text{inj}_i \beta|/\alpha]} \end{aligned}$$

So by the sum refinement rule:

$$|\Psi|, \Psi_F; \Delta, \alpha: + [\kappa_1, \dots, \kappa_n], \Delta'^{\rho_1, \rho_2}; \Gamma_A \vdash \text{refine}[\text{inj}_i \beta][c], sv'; \text{ok} \\ I$$

13. Suppose the last rule used was the **case** translation rule. The translation of a case statement produces a number of new blocks and an instruction sequence. By lemma 30 it suffices to show that each of these is well-formed when the heap-context is extended with the types of the new blocks. Note that I assume all labels are fresh. I leave informal the inner induction on the number of branches n .

Let

$$\begin{aligned} \Psi_F = \Psi(F; F_1; \dots; F_n), \\ \ell_0: \forall[|\Delta|, \rho_1: ST, \rho_2: ST]. \Gamma_A[\mathbf{r}_t: |\tau_0|], \\ \ell_1: \forall[|\Delta|, \rho_1: ST, \rho_2: ST]. \Gamma_A[\mathbf{r}_t: |\tau_1|], \\ \vdots \\ \ell_n: \forall[|\Delta|, \rho_1: ST, \rho_2: ST]. \Gamma_A[\mathbf{r}_t: |\tau_n|], \\ \ell_{\text{end}}: \forall[|\Delta|, \rho_1: ST, \rho_2: ST]. \Gamma_A[\mathbf{r}_t: |\tau|] \end{aligned}$$

(a) Heap fragments.

To show:

$$\begin{aligned} |\Psi| \vdash F; : \Psi(F; F_1; \dots; F_n) \\ F_1; \\ \vdots \\ F_n \end{aligned}$$

By inversion:

$$\begin{aligned} \Psi; \Delta; \Gamma, x_i: \tau_i; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp}} \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_1, \sigma_2] e_i : \tau \rightsquigarrow I_i; F_i \quad i \in 1 \dots n \\ \Psi; \Delta; \Gamma, x_i: \tau_i \vdash e_i : \tau \text{ exp} \\ \Psi; \Delta; \Gamma, x: \tau; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau_e \rightsquigarrow I; F \\ \Psi; \Delta; \Gamma, x: \tau \vdash e : \tau_e \text{ exp} \end{aligned}$$

By induction:

$$\begin{aligned} |\Psi| &\vdash F_i : \Psi_i \\ |\Psi| &\vdash F : \Psi' \end{aligned}$$

So by lemma 30:

$$|\Psi| \vdash F; F_1; \dots; F_n : \Psi', \Psi_1, \dots, \Psi_n$$

(b) The prelude code:

To show:

$$\begin{aligned} |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash & \text{mov } \mathbf{r}_t, sv' \quad \text{ok} \\ & \text{brtag}_0 \mathbf{r}_t, \ell_0[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\ & \vdots \\ & \text{brtag}_{k-1} \mathbf{r}_t, \ell_{k-1}[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\ & \text{brtgd}_k \mathbf{r}_t, \ell_k[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\ & \vdots \\ & \text{brtag}_{n-1} \mathbf{r}_t, \ell_{n-1}[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\ & \text{mov } \mathbf{r}_t, \text{forgetunion } \mathbf{r}_t \\ & \text{jmp } \ell_n[(\Pi_1 \Delta), \sigma_1, \sigma_2] \end{aligned}$$

By inversion:

$$\begin{aligned} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} & sv : \bigvee[\tau_1, \dots, \tau_n] \rightsquigarrow sv' \\ \Delta \vdash \tau_i \equiv \text{Tag}(i) : T_{32} & \quad i \in 1 \dots (j-1) \\ \Delta \vdash \tau_i \equiv \times[\text{Tag}(i), \tau'_i] : T_{32} & \quad i \in j \dots n \\ \Psi; \Delta; \Gamma \vdash sv : \bigvee[\tau_0, \dots, \tau_k] \end{aligned}$$

By theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv' : \bigvee[\tau_0, \dots, \tau_k]$$

So by the mov rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{mov } \mathbf{r}_t, sv' \Rightarrow \Gamma_A\{\mathbf{r}_t : \bigvee[\tau_0, \dots, \tau_k]\}$$

By theorem 17 and weakening:

$$\begin{aligned} \Delta^{\rho_1, \rho_2} \vdash |\tau_i| \equiv \text{Tag}(i) : T_{32} & \quad i \in 1 \dots (j-1) \\ \Delta^{\rho_1, \rho_2} \vdash |\tau_i| \equiv \times[\text{Tag}(i), \tau'_i] : T_{32} & \quad i \in j \dots n \end{aligned}$$

By assumption, $\ell_i : \forall[|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A[\mathbf{r}_t : |\tau_i|]$,

So the **brtag** and **brtgd** instructions are well-formed and have well-formed targets.

After $n-1$ tag comparisons, $\mathbf{r}_t : \bigvee[|\tau_n|]$. (There is an inner induction here about which I am being informal).

Therefore, by the **forgetunion** and the **mov** rules:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_t : \bigvee[|\tau_n|]\} \vdash \text{mov } \mathbf{r}_t, \text{forgetunion } \mathbf{r}_t \Rightarrow \Gamma_A\{\mathbf{r}_t : \bigvee[|\tau_n|]\}$$

So the final jump is well-formed as well.

(c) ℓ_i , for $i \in 0 \dots n$.

To show:

$$\begin{aligned} |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A[\mathbf{r}_t : |\tau_i|] \vdash & \text{srmov } \mathcal{A}(x_i), \mathbf{r}_t \text{ ok} \\ & \text{junk } \mathbf{r}_t \\ & I_i \end{aligned}$$

By inversion:

$$\begin{aligned} \Psi; \Delta; \Gamma, x_i : \tau_i; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp } \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_1, \sigma_2]} & e_i : \tau \rightsquigarrow I_i; F_i \quad i \in 1 \dots n \\ \Psi; \Delta; \Gamma, x_i : \tau_i \vdash e_i : \tau & \text{exp} \end{aligned}$$

By induction:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_i \vdash I_i \text{ ok} \quad i \in 1 \dots n$$

where $\Gamma_i = |\Gamma, x_i: \tau_i|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}$

By the good allocator assumption:

$$\Gamma_i(\mathbf{r}_t) = ns_{32} \text{ and}$$

$$\Gamma_i = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathcal{A}(x_i): |\tau_i|\} = \Gamma_A \{\mathcal{A}(x_i): |\tau_i|\}$$

So by lemmas 36 and 34:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \{\mathbf{r}_t: |\tau_i|\} \vdash \text{srmov } \mathcal{A}(x_i), \mathbf{r}_t \text{ ok} \quad i \in 1 \dots n$$

$$\text{junk } \mathbf{r}_t$$

$$I_i$$

(d) ℓ_{end}

To show:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A[\mathbf{r}_t: |\tau|] \vdash \text{srmov } \mathcal{A}(x), \mathbf{r}_t \text{ ok}$$

$$\text{junk } \mathbf{r}_t$$

$$I$$

By inversion:

$$\Psi; \Delta; \Gamma, x: \tau; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e: \tau_e \rightsquigarrow I; F$$

$$\Psi; \Delta; \Gamma, x: \tau \vdash e: \tau_e \text{ exp}$$

By induction:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_x \vdash$$

$$\text{okwhere } \Gamma_x = |\Gamma, x: \tau|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}$$

By the good allocator assumption:

$$\Gamma_x(\mathbf{r}_t) = ns_{32} \text{ and}$$

$$\Gamma_x = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathcal{A}(x): |\tau|\} = \Gamma_A \{\mathcal{A}(x): |\tau|\}$$

So by lemmas 36 and 34:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \{\mathbf{r}_t: |\tau|\} \vdash \text{srmov } \mathcal{A}(x), \mathbf{r}_t \text{ ok}$$

$$\text{junk } \mathbf{r}_t$$

$$I$$

14. Suppose the last rule used was the **dyncase** translation rule. The translation of an exception case statement produces a number of new blocks and an instruction sequence. By lemma 30 it suffices to show that each of these is well-formed when the heap-context is extended with the types of the new blocks. Note that I assume all labels are fresh.

Let

$$\Psi_F = \Psi(F; F_1; F_n),$$

$$\ell_1: \forall[|\Delta|, \rho_1: ST, \rho_2: ST]. \Gamma_A[\mathbf{r}_t: \times [\text{Dyntag}(|\tau_1|), |\tau_1|]]$$

$$\ell_{\text{end}}: \forall[|\Delta|, \rho_1: ST, \rho_2: ST]. \Gamma_A[\mathbf{r}_t: |\tau_x|]$$

(a) Heap fragments.

To show:

$$|\Psi| \vdash F; : \Psi(F; F_1; F_2)$$

$$F_1;$$

$$F_2$$

By inversion:

$$\begin{aligned}
& \Psi; \Delta; \Gamma, x_1 : \times [\text{Dyntag}(\tau_1), \tau_1]; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp } \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_1, \sigma_2]} e_1 : \tau_x \rightsquigarrow I_1; F_1 \\
& \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp } \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_1, \sigma_2]} e_2 : \tau_x \rightsquigarrow I_2; F_2 \\
& \Psi; \Delta; \Gamma, x : \tau_x; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F \\
& \Psi; \Delta; \Gamma, x_1 : \times [\text{Dyntag}(\tau_1), \tau_1] \vdash e_1 : \tau_x \text{ exp} \\
& \Psi; \Delta; \Gamma \vdash e_2 : \tau_x \text{ exp} \\
& \Psi; \Delta; \Gamma, x : \tau_x \vdash e : \tau \text{ exp}
\end{aligned}$$

By induction:

$$\begin{aligned}
& |\Psi| \vdash F_1 : \Psi_1 \\
& |\Psi| \vdash F_2 : \Psi_2 \\
& |\Psi| \vdash F : \Psi'
\end{aligned}$$

So by lemma 30:

$$|\Psi| \vdash F; F_1; F_2 : \Psi', \Psi_1, \Psi_2$$

(b) The “otherwise” case:

To show:

$$\begin{aligned}
& |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{mov } \mathbf{r}_t, sv' \quad \text{ok} \\
& \quad \text{brdyn } \mathbf{r}_t, sv'_1, \ell_1[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\
& \quad \text{junk } \mathbf{r}_t \\
& \quad I_2
\end{aligned}$$

By inversion:

$$\begin{aligned}
& \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \text{Dyn} \rightsquigarrow sv' \\
& \Psi; \Delta; \Gamma \vdash sv : \text{Dyn}
\end{aligned}$$

By theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv' : |\text{Dyn}|$$

So by the mov rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{mov } \mathbf{r}_t, sv' \Rightarrow \Gamma_A \{ \mathbf{r}_t : |\text{Dyn}| \}$$

By assumption:

$$\ell_1 : \forall [|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A \{ \mathbf{r}_t : \times [\text{Dyntag}(|\tau_1|), |\tau_1|] \}$$

So by the brdyn rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{brdyn } \mathbf{r}_t, sv'_1, \ell_1[(\Pi_1 \Delta), \sigma_1, \sigma_2] \Rightarrow \Gamma_A \{ \mathbf{r}_t : ||\text{Dyn}|| \}$$

And by lemma 34:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{junk } \mathbf{r}_t \Rightarrow \Gamma_A \{ \mathbf{r}_t : ||\text{Dyn}|| \} \{ \mathbf{r}_t : ns_{32} \}$$

Note that $\Gamma_A \{ \mathbf{r}_t : ||\text{Dyn}|| \} \{ \mathbf{r}_t : ns_{32} \} = \Gamma_A$.

By induction:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash I_2 \text{ ok}$$

So the result follows directly by the instruction sequencing rules.

(c) The “equal” case – ℓ_1 :

To show:

$$\begin{aligned}
& |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \{ \mathbf{r}_t : \times [\text{Dyntag}(|\tau_1|), |\tau_1|] \} \vdash \text{srmov } \mathcal{A}(x_1), \mathbf{r}_t \text{ ok} \\
& \quad \text{junk } \mathbf{r}_t \\
& \quad I_1
\end{aligned}$$

By inversion:

$$\begin{aligned}
& \Psi; \Delta; \Gamma, x_1 : \times [\text{Dyntag}(\tau_1), \tau_1]; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp } \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_1, \sigma_2]} e_1 : \tau_x \rightsquigarrow I_1; F_1 \\
& \Psi; \Delta; \Gamma, x_1 : \times [\text{Dyntag}(\tau_1), \tau_1] \vdash e_1 : \tau_x \text{ exp}
\end{aligned}$$

By induction:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_1 \vdash I_1 \text{ ok where } \Gamma_1 = |\Gamma, x_1: \tau_1|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}$$

By the good allocator assumption:

$$\Gamma_1(\mathbf{r}_t) = ns_{32} \text{ and}$$

$$\Gamma_1 = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{ \mathcal{A}(x_1): | \times [\text{Dyn tag}(\tau_1), \tau_1] | \} = \Gamma_A \{ \mathcal{A}(x_1): | \times [\text{Dyn tag}(|\tau_1|), |\tau_1|] | \}$$

So it suffices to show that:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \{ \mathbf{r}_t: | \times [\text{Dyn tag}(|\tau_1|), |\tau_1|] | \} \vdash \text{rmov } \mathcal{A}(x_i), \mathbf{r}_t \Rightarrow \Gamma_B$$

junk $\mathbf{r}_t$

$$\text{where } \Gamma_B = \Gamma_A \{ \mathcal{A}(x_1): | \times [\text{Dyn tag}(|\tau_1|), |\tau_1|] | \}$$

Which follows by lemmas 36 and 34.

(d) The continuation – ℓ_{end} :

To show:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A[\mathbf{r}_t: |\tau_x|] \vdash \text{rmov } \mathcal{A}(x), \mathbf{r}_t \text{ ok}$$

junk $\mathbf{r}_t$
 I

By inversion:

$$\Psi; \Delta; \Gamma, x: \tau_x; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e: \tau \rightsquigarrow I; F$$

$$\Psi; \Delta; \Gamma, x: \tau_x \vdash e: \tau \text{ exp}$$

By induction:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_x \vdash$$

$$\text{ok where } \Gamma_x = |\Gamma, x: \tau_x|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}$$

By the good allocator assumption:

$$\Gamma_x(\mathbf{r}_t) = ns_{32} \text{ and}$$

$$\Gamma_x = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{ \mathcal{A}(x): |\tau_x| \} = \Gamma_A \{ \mathcal{A}(x): |\tau_x| \}$$

So by lemmas 36 and 34:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \{ \mathbf{r}_t: |\tau_x| \} \vdash \text{rmov } \mathcal{A}(x), \mathbf{r}_t \text{ ok}$$

junk $\mathbf{r}_t$
 I

15. Suppose

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let } x = \text{raise } sv \text{ in } e: \tau \rightsquigarrow \text{mov } \mathbf{r}_1, sv$$

$\text{loadr } \mathbf{r}_t, \mathbf{r}_e(1)$
 $\text{mov } \mathbf{sp}, \mathbf{r}_t$
 $\text{loadr } \mathbf{r}_t, \mathbf{r}_e(0)$
 $\text{jmp } \mathbf{r}_t$

By inversion:

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv: \text{Dyn} \rightsquigarrow sv'$$

$$\Psi; \Delta; \Gamma \vdash sv: \text{Dyn}$$

By theorem 18:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash sv': | \text{Dyn} |$$

By assumption:

$$\Gamma_A(\mathbf{sp}) = \sigma_f \circ \sigma_1 \circ \sigma_2$$

$$\Gamma_A(\mathbf{r}_e) = \text{Exnptr}(\sigma_2)$$

By definition:

$$\begin{aligned}\mathbf{Exnptr}(\sigma_2) &= \times[\mathbf{Exnhandler}(\sigma_2), \sigma_2] \\ \mathbf{Exnhandler}(\sigma_2) &= \forall[].\{\mathbf{r}_1:\text{Dyn}, \mathbf{sp}:\sigma_2\} \rightarrow 0\end{aligned}$$

By the mov rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash \text{mov } \mathbf{r}_1, sv' \Rightarrow \Gamma_A\{\mathbf{r}_1:\text{Dyn}|\}$$

By the load rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_1:\text{Dyn}|\} \vdash \text{loadr } \mathbf{r}_t, \mathbf{r}_e(1) \Rightarrow \Gamma_A\{\mathbf{r}_1:\text{Dyn}|\}\{\mathbf{r}_t:\sigma_2\}$$

By the stack load rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_1:\text{Dyn}|\}\{\mathbf{r}_t:\sigma_2\} \vdash \text{mov } \mathbf{sp}, \mathbf{r}_t \Rightarrow \Gamma_A\{\mathbf{r}_1:\text{Dyn}|\}\{\mathbf{r}_t:\sigma_2\}\{\mathbf{sp}:\sigma_2\}$$

By the load rule:

$$\begin{aligned}|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_1:\text{Dyn}|\}\{\mathbf{r}_t:\sigma_2\}\{\mathbf{sp}:\sigma_2\} \vdash \text{loadr } \mathbf{r}_t, \mathbf{r}_e(0) \Rightarrow \\ \Gamma_A\{\mathbf{r}_1:\text{Dyn}|\}\{\mathbf{sp}:\sigma_2\}\{\mathbf{r}_t:\mathbf{Exnhandler}(\sigma_2)\}\end{aligned}$$

By the jump rule:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_1:\text{Dyn}|\}\{\mathbf{sp}:\sigma_2\}\{\mathbf{r}_t:\mathbf{Exnhandler}(\sigma_2)\} \vdash \text{jmp } \mathbf{r}_t \text{ ok}$$

16. Suppose the last rule used was the **handle** translation rule. The translation of an exception handler produces a number of new blocks and an instruction sequence. By lemma 30 it suffices to show that each of these is well-formed when the heap-context is extended with the types of the new blocks. Note that I assume all labels are fresh.

Let

$$\begin{aligned}\Gamma_B &= |\Gamma|_{\mathcal{A}}^{\sigma_3}\{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_3)\} \\ \sigma_f &= |\Gamma|_{\mathcal{A}}^e(\mathbf{sp}) \\ \sigma_3 &= \mathbf{Exnptr}(\sigma_2) \triangleright_{32} \sigma_f \circ \sigma_1 \circ \sigma_2 \\ \sigma_h &= \sigma_f \circ \sigma_3 \\ \Psi_F &= \Psi(F; F_1; F_n), \\ &\quad \ell_{\text{handle}}:\forall[\Delta, \rho_1, \rho_2].\{\mathbf{r}_1:\text{Dyn}, \mathbf{sp}:\sigma_h, \dots\} \\ &\quad \ell_{\text{post}}:\forall[|\Delta|, \rho_1:ST, \rho_2:ST].\Gamma_B[\mathbf{r}_t:|\tau_x|] \\ &\quad \ell_{\text{end}}:\forall[|\Delta|, \rho_1:ST, \rho_2:ST].\Gamma_A[\mathbf{r}_t:|\tau_x|]\end{aligned}$$

(a) Heap fragments.

To show:

$$\begin{aligned}|\Psi| \vdash F; : \Psi(F; F_1; F_2) \\ F_1; \\ F_2\end{aligned}$$

By inversion:

$$\begin{aligned}\Psi; \Delta; \Gamma; \mathcal{A}, \epsilon, \sigma_3 \vdash \text{jmp } \ell_{\text{post}}[(\Pi_1 \Delta), \epsilon, \sigma_h] \ e_1 : \tau_x \rightsquigarrow I_1; F_1 \\ \Psi; \Delta; \Gamma, x_2:\text{Dyn}; \mathcal{A}, \sigma_1, \sigma_2 \vdash \text{jmp } \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_f \circ \sigma_1, \sigma_2] \ e_2 : \tau_x \rightsquigarrow I_2; F_2 \\ \Psi; \Delta; \Gamma, x:\tau_x; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F \\ \Psi; \Delta; \Gamma \vdash e_1 : \tau_x \ \mathbf{exp} \\ \Psi; \Delta; \Gamma, x_2:\text{Dyn} \vdash e_2 : \tau_x \ \mathbf{exp} \\ \Psi; \Delta; \Gamma, x:\tau_x \vdash e : \tau \ \mathbf{exp}\end{aligned}$$

By induction:

$$\begin{aligned}|\Psi| \vdash F_1 : \Psi_1 \\ |\Psi| \vdash F_2 : \Psi_2 \\ |\Psi| \vdash F : \Psi'\end{aligned}$$

So by lemma 30:

$$|\Psi| \vdash F; F_1; F_2 : \Psi', \Psi_1, \Psi_2$$

(b) The handled code block:

To show:

$$\begin{aligned} |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash & \text{push } \mathbf{r}_e & \text{ok} \\ & \text{stackcopy } (\mathbf{Exnptr}(\sigma_2)) \triangleright_{32} \sigma_f \\ & \text{sfree } 1 \\ & \text{malloc } \mathbf{r}_e[\mathbf{Exnhndler}(\sigma_h), \sigma_h] \langle \ell_{\text{handle}}[\Pi_1 \Delta, \rho_1, \rho_2], \mathbf{sp} \rangle \\ & I_1 \end{aligned}$$

By assumption:

$$\Gamma_A(\mathbf{sp}) = \sigma_f \circ \sigma_1 \circ \sigma_2$$

$$\Gamma_A(\mathbf{r}_e) = \mathbf{Exnptr}(\sigma_2)$$

The introductory sequence of instructions only modifies the stack and the exception register, so for brevity I leave the rest of the typing context implicit and only describe the types of the stack register and the exception register after the instructions.

Instruction	$\mathbf{sp}$ type	$\mathbf{r}_e$ type	Comment
push	$\sigma_3 = (\mathbf{Exnptr}(\sigma_2)) \triangleright_{32} \sigma_f \circ \sigma_1 \circ \sigma_2$	$\mathbf{Exnptr}(\sigma_2)$	By the push rule
stackcopy	$(\mathbf{Exnptr}(\sigma_2)) \triangleright_{32} \sigma_f \circ \sigma_3$	$\mathbf{Exnptr}(\sigma_2)$	By lemma 41
sfree	$\sigma_h = \sigma_f \circ \sigma_3$	$\mathbf{Exnptr}(\sigma_2)$	By the sfree rule
malloc	σ_h	$\mathbf{Exnptr}(\sigma_h)$	By the malloc rule

So it suffices to show that:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{sp}:\sigma_h\}\{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_3)\} \vdash I_1 \text{ ok}$$

But note that :

$$|\Gamma|_{\mathcal{A}}^{\sigma_h} = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{sp}:\sigma_g\} = \Gamma_A\{\mathbf{sp}:\sigma_g\}$$

So the desired result follows by induction.

(c) The handler block:

$$\text{Let } \Gamma_H = \{\mathbf{r}_1:\text{Dyn}, \mathbf{r}_2:ns_{32}, \mathbf{r}_t:ns_{32}, \mathbf{r}_e:ns_{32}, \mathbf{sp}:\sigma_h, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}\}$$

To show:

$$\begin{aligned} |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_H \vdash & \text{srmov } \mathcal{A}(x_2), \mathbf{r}_1 \text{ ok} \\ & \text{sfree}(\text{frmsz}(\mathcal{A})) \\ & \text{pop } \mathbf{r}_e \\ & \text{junk } \mathbf{r}_1 \\ & I_2 \end{aligned}$$

Instr	machine state	Comment
srmov	$\Gamma_H\{\mathcal{A}(x_2):\text{Dyn}\}$	By lemma 36
sfree	$\Gamma_H\{\mathcal{A}(x_2):\text{Dyn}\}\{\mathbf{sp}:\sigma_3\}$	By the sfree rule
pop	$\Gamma_H\{\mathcal{A}(x_2):\text{Dyn}\}\{\mathbf{sp}:\sigma_f \circ \sigma_1 \circ \sigma_2\}\{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\}$	By the pop rule
junk	$\dots \{\mathbf{r}_1:ns_{32}\}$	By lemma 34

But note that by assumption:

$$\begin{aligned} \Gamma_A & \\ &= |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\} \\ &= \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{r}_t:ns_{32}, \mathbf{r}_e:\mathbf{Exnptr}(\sigma_2), \mathbf{sp}:\sigma_f \circ \sigma_1 \circ \sigma_2, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}\} \end{aligned}$$

And by definition:

$$\begin{aligned} & \Gamma_H\{\mathbf{sp}:\sigma_f \circ \sigma_1 \circ \sigma_2\}\{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\}\{\mathbf{r}_1:ns_{32}\} \\ &= \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{r}_t:ns_{32}, \mathbf{r}_e:\mathbf{Exnptr}(\sigma_2), \mathbf{sp}:\sigma_f \circ \sigma_1 \circ \sigma_2, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}\} \\ &= \Gamma_A \end{aligned}$$

So by the good allocator assumption:

$$\begin{aligned} & \Gamma_H\{\mathcal{A}(x_2):\mathbf{Dyn}\}\{\mathbf{sp}:\sigma_f \circ \sigma_1 \circ \sigma_2\}\{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\} \\ &= \Gamma_A\{\mathcal{A}(x_2):\mathbf{Dyn}\} \\ &= |\Gamma, x_2:\mathbf{Dyn}|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}\{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\} \end{aligned}$$

Finally, by induction:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathcal{A}(x_2):\mathbf{Dyn}\} \vdash I_2 \text{ ok}$$

(d) The postlude block:

To show:

$$\begin{aligned} & |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_B[\mathbf{r}_t:|\tau|] \vdash \mathbf{sfree}(\text{frmsz}(\mathcal{A})) \quad \text{ok} \\ & \quad \text{pop } \mathbf{r}_e \\ & \quad \text{jmp } \ell_{\text{end}}[\Pi_1 \Delta, \sigma_1, \sigma_2] \end{aligned}$$

Note that $\Gamma_B(\mathbf{sp}) = \sigma_h = \sigma_f \circ \sigma_3$.

So by the **sfree** rule:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_B[\mathbf{r}_t:|\tau|] \vdash \mathbf{sfree}(\text{frmsz}(\mathcal{A})) \Rightarrow \Gamma_B\{\mathbf{r}_t:|\tau|\}\{\mathbf{sp}:\sigma_3\}$$

By the **pop** rule:

$$\begin{aligned} & |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_B\{\mathbf{r}_t:|\tau|\}\{\mathbf{sp}:\sigma_3\} \vdash \text{pop } \mathbf{r}_e \Rightarrow \\ & \Gamma_B\{\mathbf{r}_t:|\tau|\}\{\mathbf{sp}:\sigma_f \circ \sigma_1 \circ \sigma_2\}\{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\} \end{aligned}$$

Note that:

$$\begin{aligned} & \Gamma_B\{\mathbf{r}_t:|\tau|\}\{\mathbf{sp}:\sigma_f \circ \sigma_1 \circ \sigma_2\}\{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\} \\ &= |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}\{\mathbf{r}_t:|\tau|\}\{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\} \\ &= \Gamma_A\{\mathbf{r}_t:|\tau|\} \end{aligned}$$

And by assumption:

$$\ell_{\text{end}}:\forall[|\Delta|, \rho_1:ST, \rho_2:ST]. \Gamma_A[\mathbf{r}_t:|\tau_x|]$$

So by the **jmp** rule:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_t:|\tau|\} \vdash \text{jmp } \ell_{\text{end}}[\Pi_1 \Delta, \sigma_1, \sigma_2] \text{ ok}$$

(e) The continuation block:

To show:

$$\begin{aligned} & |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_t:|\tau_x|\} \vdash \mathbf{srmov } \mathcal{A}(x), \mathbf{r}_t \text{ ok} \\ & \quad \text{junk } \mathbf{r}_t \\ & \quad I \end{aligned}$$

By lemma 36:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_t:|\tau_x|\} \vdash \mathbf{srmov } \mathcal{A}(x), \mathbf{r}_t \Rightarrow \Gamma_A\{\mathbf{r}_t:|\tau_x|\}\{\mathcal{A}(x):|\tau_x|\}$$

By lemma 34:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A\{\mathbf{r}_t:|\tau_x|\}\{\mathcal{A}(x):|\tau_x|\} \vdash \text{junk } \mathbf{r}_t \Rightarrow \Gamma_A\{\mathcal{A}(x):|\tau_x|\}$$

By the good allocator assumption:

$$\Gamma_A\{\mathcal{A}(x):|\tau_x|\} = |\Gamma, x\tau_x|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2}\{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\}$$

By inversion:

$$\begin{aligned} & \Psi; \Delta; \Gamma, x:\tau_x; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e:\tau \rightsquigarrow I; F \\ & \Psi; \Delta; \Gamma, x:\tau_x \vdash e:\tau \text{ exp} \end{aligned}$$

So by induction:

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \{ \mathcal{A}(x) : |\tau_x| \} \vdash I \text{ ok}$$

17. Suppose $\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow (S:I); F$ via the spill rule.

By inversion:

$$\begin{aligned} & \Psi; \Delta; \Gamma; \mathcal{A}', \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F \\ & \text{where } \mathcal{A}' \text{ is a good allocator for } e \\ & \Psi; \Delta; \Gamma \vdash e : \tau \text{ exp} \end{aligned}$$

By induction:

$$\begin{aligned} & |\Psi| \vdash F : \Psi_F \\ & |\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_B \vdash I \text{ ok where } \Gamma_B = |\Gamma|_{\mathcal{A}'}^{\sigma_1 \circ \sigma_2} \end{aligned}$$

By inversion:

$$|\Psi|; \Delta^{\rho_1, \rho_1}; |\Gamma|_{\mathcal{A}'}^{\sigma_1 \circ \sigma_2} \vdash S \Rightarrow |\Gamma|_{\mathcal{A}'}^{\sigma_1 \circ \sigma_2}$$

So by lemma 29 (partial instruction sequence completion):

$$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash S:I \text{ ok}$$

■

8.3.6 Heap values, heaps, and programs

Heap values

$$\Psi; \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k; x_1:\tau_1, \dots, x_m:\tau_m, z_1:\phi_1, \dots, z_n:\phi_n; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{ret}} e : \tau \rightsquigarrow I; F$$

Where $\mathcal{A}$ is a good allocator for e

$$\text{and } \sigma_1 = (\text{cont}(\overline{m+2*n})(\rho_1)(\rho_2)(|\tau|)) \triangleright_{32} \sigma_{32} \circ \sigma_{64} \circ \rho_1$$

$$\text{and } \sigma_{32} = |\tau_1| \triangleright_{32} \dots \triangleright_{32} |\tau_m| \triangleright_{32} \epsilon$$

$$\text{and } \sigma_{64} = |\phi_1| \triangleright_{64} \dots \triangleright_{64} |\phi_n| \triangleright_{64} \epsilon$$

$$\text{and } \sigma_2 = \rho_2$$

$$\text{and } l = \text{frmsz}(\mathcal{A})$$

$$\text{and } hval = \text{code}[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k, \rho_1:ST, \rho_2:ST] | \text{Code}(\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n)(\tau) |$$

$$\text{salloc } l;$$

$$\text{srmov } \mathcal{A}(x_1), \text{sp}(l+0);$$

$$\vdots$$

$$\text{srmov } \mathcal{A}(x_m), \text{sp}(l+m-1);$$

$$\text{srfmov } \mathcal{A}(z_1), \text{sp}(l+m+2*0);$$

$$\vdots$$

$$\text{srfmov } \mathcal{A}(z_n), \text{sp}(l+m+2*(n-1));$$

$$I$$

$$\begin{aligned} & \Psi \vdash_h \text{code}_\tau[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k](x_1:\tau_1, \dots, x_m:\tau_m)(z_1:\phi_1, \dots, z_n:\phi_n).e : \\ & \quad \forall[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k] \text{Code}(\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n)(\tau) \rightsquigarrow hval; F \end{aligned}$$

Since the only heap values in the **LIL** are code blocks, the heap value translation has only one inference rule, for translating code functions. Code functions are translated by translating their bodies with a good allocator, and with appropriate stack segment parameters. Recall that the stack

segment parameters describe the stack below the current frame (which is managed by the allocator). Consequently, the stack segments in the code function translation reflect the calling convention on entry to the function: the top stack segment consists of an appropriate return continuation and the 32 and 64 bit arguments pushed onto the first stack segment type variable (ρ_1). The second stack segment, describing the section below the current handler, is simply the second stack type variable (ρ_2). Notice that the parameter of occurrence on the expression translation derivation indicates that the expression occurs in a return context and hence should return its value by popping the frame and calling the return continuation from the stack.

The result of translating the body is an instruction sequence I and a heap fragment F containing additional code blocks to be allocated at the top level. In order to complete the translation, it is necessary to wrap the instruction sequence I with additional code to allocate the frame and to initialize the locations assigned to the argument variables from the locations assigned to the in-arguments by the calling convention. In principle the allocator may sometimes render this initialization code unnecessary by using variables directly from their in-argument slots, but for simplicity I do not attempt to perform this optimization here. The wrapped code sequence is then abstracted with respect to the free type variables (including the stack segment parameters) to turn it into a closed heap value.

Theorem 23 (Soundness of the heap value translation)

If $\Psi \vdash hval : \tau$ **hval**
and $\Psi \vdash_h hval : \tau \rightsquigarrow hval : F$
then
 $|\Psi| \vdash F : \Psi_F$
and $|\Psi| \vdash hval : |\tau|$ **hval**

Proof: By construction.

By inversion:

$\mathcal{A}$ is a good allocator for e

$\Psi; \alpha_1 : \kappa_1, \dots, \alpha_k : \kappa_k; x_1 : \tau_1, \dots, x_m : \tau_m, z_1 : \phi_1, \dots, z_n : \phi_n \vdash e : \tau$ **exp**

$\Psi; \alpha_1 : \kappa_1, \dots, \alpha_k : \kappa_k; x_1 : \tau_1, \dots, x_m : \tau_m, z_1 : \phi_1, \dots, z_n : \phi_n; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{ret}} e : \tau \rightsquigarrow I; F$

And by theorem 15 and construction:

$\Delta^{\rho_1, \rho_2} \vdash \sigma_1 : ST$

$\Delta^{\rho_1, \rho_2} \vdash \sigma_2 : ST$

So by theorem 22:

$|\Psi| \vdash F : \Psi_F$

$|\Psi|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_A \vdash I$ **ok**

where $\Gamma_A = |x_1 : \tau_1, \dots, x_m : \tau_m, z_1 : \phi_1, \dots, z_n : \phi_n|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e : \mathbf{Exnptr}(\sigma_2)\}$

It remains to be shown that:

$|\Psi| \vdash hval : |\forall[\alpha_1 : \kappa_1, \dots, \alpha_k : \kappa_k] \mathbf{Code}(\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n)(\tau)|$ **hval**

By lemma 29, it suffices to show that:

$$\begin{aligned}
|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_0 \vdash & \text{salloc } l; & \Rightarrow \Gamma_A \\
& \text{srmov } \mathcal{A}(x_1), \text{sp}(l+0); \\
& \vdots \\
& \text{srmov } \mathcal{A}(x_m), \text{sp}(l+m-1); \\
& \text{srfmov } \mathcal{A}(z_1), \text{sp}(l+m+2*0); \\
& \vdots \\
& \text{srfmov } \mathcal{A}(z_n), \text{sp}(l+m+2*(n-1));
\end{aligned}$$

where $\Gamma_0 = \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:\mathbf{Exnptr}(\sigma_2), \mathbf{r}_t:ns_{32}, \mathbf{sp}:\sigma_1\}$

and $\Gamma_A = |x_1:\tau_1, \dots, x_m:\tau_m, z_1:\phi_1, \dots, z_n:\phi_n|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\}$

By the salloc rule, it suffices to show that :

$$\begin{aligned}
|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_1 \vdash & \text{srmov } \mathcal{A}(x_1), \text{sp}(l+0); & \Rightarrow \Gamma_A \\
& \vdots \\
& \text{srmov } \mathcal{A}(x_m), \text{sp}(l+m-1); \\
& \text{srfmov } \mathcal{A}(z_1), \text{sp}(l+m+2*0); \\
& \vdots \\
& \text{srfmov } \mathcal{A}(z_n), \text{sp}(l+m+2*(n-1));
\end{aligned}$$

where $\Gamma_1 = \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:\mathbf{Exnptr}(\sigma_2), \mathbf{r}_t:ns_{32}, \mathbf{sp}:ns_{32}^l \triangleright_{32} \sigma_1\}$

Which follows by induction on $n+m$

- Suppose $n+m=0$

Then by assumption:

$$\Gamma_A = |\bullet|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\}$$

And by the good allocator assumption:

$$\begin{aligned}
& |\bullet|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\} \\
& = \{\mathbf{r}_1:ns_{32}, \mathbf{r}_2:ns_{32}, \mathbf{f}_1:ns_{64}, \mathbf{f}_2:ns_{64}, \mathbf{r}_e:\mathbf{Exnptr}(\sigma_2), \mathbf{r}_t:ns_{32}, \mathbf{sp}:ns_{32}^l \triangleright_{32} \sigma_1\}
\end{aligned}$$

- Suppose $m > 0$. Then

By induction:

$$\begin{aligned}
|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_1 \vdash & \text{srmov } \mathcal{A}(x_1), \text{sp}(l+0); & \Rightarrow \Gamma'_A \\
& \vdots \\
& \text{srmov } \mathcal{A}(x_m), \text{sp}(l+m-1); \\
& \text{srfmov } \mathcal{A}(z_1), \text{sp}(l+m+2*0); \\
& \vdots \\
& \text{srfmov } \mathcal{A}(z_{n-1}), \text{sp}(l+m+2*(n-2));
\end{aligned}$$

where $\Gamma'_A = |x_1:\tau_1, \dots, x_m:\tau_m, z_1:\phi_1, \dots, z_{n-1}:\phi_{n-1}|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e:\mathbf{Exnptr}(\sigma_2)\}$

So by the definition of partial instruction sequences, it suffices to show that:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma'_A \vdash \text{srfmov } \mathcal{A}(z_n), \text{sp}(l+m+2*(n-1)) \Rightarrow \Gamma_A$$

But note that by lemma 37:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma'_A \vdash \text{srfmov } \mathcal{A}(z_n), \text{sp}(l+m+2*(n-1)) \Rightarrow \Gamma'_A \{\mathcal{A}(z_n):|\phi_n|\}$$

And by the good allocator assumption:

$$\begin{aligned}\Gamma_A &= |x_1:\tau_1, \dots, x_m:\tau_m, z_1:\phi_1, \dots, z_{n-1}:\phi_{n-1}, z_n:\phi_n|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e: \mathbf{Exnptr}(\sigma_2)\} \\ &= |x_1:\tau_1, \dots, x_m:\tau_m, z_1:\phi_1, \dots, z_{n-1}:\phi_{n-1}|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e: \mathbf{Exnptr}(\sigma_2)\} \{\mathcal{A}(z_n):|\phi_n|\} \\ &= \Gamma'_A \{\mathcal{A}(z_n):|\phi_n|\}\end{aligned}$$

- Suppose $m = 0$ and $n > 0$. Then

By induction:

$$\begin{aligned}|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma_1 \vdash \mathbf{rmov} \ \mathcal{A}(x_1), \mathbf{sp}(l+0); & \Rightarrow \Gamma'_A \\ \vdots & \\ \mathbf{rmov} \ \mathcal{A}(x_{m-1}), \mathbf{sp}(l+m-2); & \end{aligned}$$

$$\text{where } \Gamma'_A = |x_1:\tau_1, \dots, x_{m-1}:\tau_{m-1}, x_m:\tau_m|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e: \mathbf{Exnptr}(\sigma_2)\}$$

So by the definition of partial instruction sequences, it suffices to show that:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma'_A \vdash \mathbf{rmov} \ \mathcal{A}(x_m), \mathbf{sp}(l+m) \Rightarrow \Gamma_A$$

But note that by lemma 36:

$$|\Psi|; \Delta^{\rho_1, \rho_2}; \Gamma'_A \vdash \mathbf{rmov} \ \mathcal{A}(x_m), \mathbf{sp}(l+m) \Rightarrow \Gamma'_A \{\mathcal{A}(x_m):|\tau_m|\}$$

And by the good allocator assumption:

$$\begin{aligned}\Gamma_A &= |x_1:\tau_1, \dots, x_{m-1}:\tau_{m-1}, x_m:\tau_m|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e: \mathbf{Exnptr}(\sigma_2)\} \\ &= |x_1:\tau_1, \dots, x_{m-1}:\tau_{m-1}|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{\mathbf{r}_e: \mathbf{Exnptr}(\sigma_2)\} \{\mathcal{A}(x_m):|\tau_m|\} \\ &= \Gamma'_A \{\mathcal{A}(x_m):|\tau_m|\}\end{aligned}$$

■

Heaps

$$\frac{}{\Psi \vdash \epsilon \rightsquigarrow \epsilon} \quad \frac{\Psi[\ell:\tau] \vdash_h \text{hval} : \tau \rightsquigarrow \text{hval}'; F \quad \Psi[\ell:\tau] \vdash d \rightsquigarrow F'}{\Psi[\ell:\tau] \vdash d, \ell:\tau \mapsto \text{hval} \rightsquigarrow \ell:|\tau|. \text{hval}', F; F'}$$

The translation of heaps simply translates the individual heap bindings to produce a heap value and a new heap fragment. The heap value can then be bound to the assigned label, and incorporated with the new heap fragment into the rest of the re-written heap.

The heap soundness theorem is stated in two parts. Since the translation returns a heap fragment for each heap value, I first show that the result of the translation is well-formed as a heap fragment. The well-formedness of the translation of a closed well-formed heap then follows almost immediately.

Theorem 24 (Soundness of the heap translation)

If $\Psi \vdash d \text{ ok}$
and $\Psi \vdash d \rightsquigarrow F$
then
 $|\Psi| \vdash F \text{ ok} \Psi_F$

Proof: (By induction on heaps)

- Suppose $d = \epsilon$.

By definition:

$$\Psi \vdash \epsilon \rightsquigarrow \epsilon$$

And by the heap fragment typing rule:

$$|\Psi| \vdash \epsilon : \epsilon$$

- Suppose $d = d', \ell : \tau \mapsto hval$

By inversion:

$$\begin{aligned} \Psi[\ell : \tau] &\vdash d' \text{ ok} \\ \Psi[\ell : \tau] &\vdash hval : \tau \text{ hval} \\ \Psi[\ell : \tau] &\vdash d \rightsquigarrow F' \\ \Psi[\ell : \tau] &\vdash_h hval : \tau \rightsquigarrow hval'; F \end{aligned}$$

By induction:

$$|\Psi[\ell : \tau]| \vdash F' \text{ ok} \Psi'_F$$

By theorem 23:

$$\begin{aligned} |\Psi[\ell : \tau]| &\vdash hval' : |\tau| \\ |\Psi[\ell : \tau]| &\vdash F \text{ ok} \Psi_F \end{aligned}$$

By lemma 30:

$$|\Psi[\ell : \tau]| \vdash F; F' \text{ ok} \Psi_F, \Psi'_F$$

By the heap fragment formation rule:

$$|\Psi[\ell : \tau]| \vdash \ell : |\tau|.hval', F; F' \text{ ok} \Psi_F, \Psi'_F$$

■

Programs

$$\Psi(d) \vdash d \rightsquigarrow H' \quad \Psi(d); \bullet; \bullet; \mathcal{A}, \epsilon, \epsilon \vdash_{\text{halt}} e : \tau \rightsquigarrow I'; F$$

Where $\mathcal{A}$ is a good allocator for e
and $H = F; H'$,

$$\begin{aligned} &\ell_{\text{ihandle}} : \text{Exnhandler}(\epsilon) \\ &\text{mov } \mathbf{r}_t, \mathbf{r}_1; \\ &\text{halt}_{\text{Dyn}} \end{aligned}$$

$$\begin{aligned} \text{and } R &= \{\mathbf{r}_1 \mapsto ns_{32}, \mathbf{r}_2 \mapsto ns, \mathbf{r}_e \mapsto ns_{32}, \mathbf{r}_t \mapsto ns, \mathbf{f}_1 \mapsto ns_{64}, \mathbf{f}_2 \mapsto ns_{64}, \mathbf{f}_t \mapsto ns_{64}, \mathbf{sp} \mapsto \epsilon\} \\ \text{and } I &= \text{malloc } \mathbf{r}_e [\text{Exnhandler}(\epsilon), \epsilon] \langle \ell_{\text{ihandle}}, \mathbf{sp} \rangle; I' \end{aligned}$$

$$\bullet \vdash \text{letrec } d \text{ in } e : \tau \rightsquigarrow H, R, I$$

The translation of a **LIL** program is obtained by translating heap to obtain a **TILTAL** heap (H'), and translating the body of the program to obtain an instruction sequence (I') and some additional heap blocks to be incorporated into the heap (F). Note that the body of the program is translated with a `halt` occurrence parameter, indicating that the instruction sequence should be terminated with a `halt` instruction. In order to satisfy the calling conventions, it is necessary to add a final exception handler block to the heap (ℓ_{ihandle}) and allocate an appropriate exception frame containing it in $\mathbf{r}_e$. The initial register file contains nonsense in all of the general purpose registers, and an empty stack in the stack register.

Theorem 25 (Soundness of the program translation)

If $\bullet \vdash \text{letrec } d \text{ in } e : \tau$
 and $\bullet \vdash \text{letrec } d \text{ in } e : \tau \rightsquigarrow (H, R, I)$
 then $\bullet \vdash (H, R, I) \text{ ok}$

Proof:

By inversion:

$$\begin{aligned} \Psi(d) &\vdash d \text{ ok} \\ \Psi(d) &\vdash d \rightsquigarrow H' \end{aligned}$$

So by theorem 24:

$$|\Psi(d)| \vdash H' \text{ ok} \Psi'_H$$

By inversion:

$$\begin{aligned} \Psi(d); \bullet; \bullet; \mathcal{A}, \epsilon, \epsilon &\vdash_{\text{halt}} e : \tau \rightsquigarrow I'; F \\ \Psi(d); \bullet; \bullet &\vdash e : \tau \text{ exp} \end{aligned}$$

And by the empty stack rule:

$$\Delta^{\rho_1, \rho_2} \vdash \epsilon : ST$$

So by theorem 22:

$$\begin{aligned} |\Psi(d)| &\vdash F : \Psi_F \\ |\Psi(d)|, \Psi_F; \Delta^{\rho_1, \rho_2}; \Gamma_0 \{ \mathbf{r}_e : \mathbf{Exnptr}(\epsilon) \} &\vdash I' \text{ ok} \\ \text{where } \Gamma_0 &= | \bullet |_{\mathcal{A}}^{\epsilon} \end{aligned}$$

By construction:

$$\begin{aligned} \ell_{\text{ihandle}} : \mathbf{Exnhandler}(\epsilon) &\vdash \ell_{\text{ihandle}} : \mathbf{Exnhandler}(\epsilon) \text{ ok} \\ &\quad \text{mov } \mathbf{r}_t, \mathbf{r}_1; \\ &\quad \text{halt}_{\text{Dyn}} \end{aligned}$$

And hence:

$$\begin{aligned} \ell_{\text{ihandle}} : \mathbf{Exnhandler}(\epsilon), \Psi'_H, \Psi_F &\vdash F; H', \text{ ok} \\ &\quad \ell_{\text{ihandle}} : \mathbf{Exnhandler}(\epsilon) \\ &\quad \text{mov } \mathbf{r}_t, \mathbf{r}_1; \\ &\quad \text{halt}_{\text{Dyn}} \end{aligned}$$

By the good allocator assumption:

$$\begin{aligned} | \bullet |_{\mathcal{A}}^{\epsilon} \{ \mathbf{r}_e : \mathbf{Exnptr}(\epsilon) \} \\ = \{ \mathbf{r}_1 : ns_{32}, \mathbf{r}_2 : ns, \mathbf{r}_e : \mathbf{Exnptr}(\epsilon), \mathbf{r}_t : ns, \mathbf{f}_1 : ns_{64}, \mathbf{f}_2 : ns_{64}, \mathbf{f}_t : ns_{64}, \mathbf{sp} : \epsilon \} \end{aligned}$$

By the malloc rule:

$$\ell_{\text{ihandle}} : \mathbf{Exnhandler}; \Delta; \Gamma_0 \vdash \text{malloc } \mathbf{r}_e [\mathbf{Exnhandler}(\epsilon), \epsilon] \langle \ell_{\text{ihandle}}, \mathbf{sp} \rangle \Rightarrow \Gamma_0 \{ \mathbf{r}_e : \mathbf{Exnptr}(\epsilon) \}$$

So by the instruction rule:

$$|\Psi(d)|, \Psi(F), \ell_{\text{ihandle}} : \mathbf{Exnhandler}; \Delta; \Gamma_0 \vdash I \text{ ok}$$

So by the program rule :

$$\vdash (H, R, I) \text{ ok}$$

■

8.4 The complete translation rules

Small values

$$\boxed{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau \rightsquigarrow sv'}$$

$$\frac{}{\Psi; \Delta; \Gamma[x:\tau]; \mathcal{A}[x \rightarrow \mathbf{r}], \sigma_1, \sigma_2 \vdash_{32} x : \tau \rightsquigarrow \mathbf{r}}$$

$$\frac{}{\Psi; \Delta; \Gamma[x:\tau]; \mathcal{A}[x \rightarrow \mathbf{sp}(i)], \sigma_2, \sigma_2 \vdash_{32} x : \tau \rightsquigarrow \mathbf{sp}(i)}$$

$$\frac{}{\Psi[\ell:\tau]; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} \ell : \tau \rightsquigarrow \ell}$$

$$\frac{}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} i : \mathbf{Int} \rightsquigarrow i}$$

$$\frac{\begin{array}{c} \Delta \vdash c \equiv \bigvee[\dots, c_i, \dots] : \mathbf{T}_{32} \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : c_i \rightsquigarrow sv' \end{array}}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} \mathbf{inj_union}_c sv : c \rightsquigarrow \mathbf{inj_union}_{|c|} sv'}$$

$$\frac{\begin{array}{c} \Delta \vdash \tau \equiv \mathbf{Rec}[\kappa](c)(c_p) : \mathbf{T}_{32} \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : c(\mathbf{Rec}[\kappa]c)c_p \rightsquigarrow sv' \end{array}}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} \mathbf{roll}_\tau sv : \tau \rightsquigarrow \mathbf{roll}_{|\tau|} sv'}$$

$$\frac{\begin{array}{c} \Delta \vdash \tau \equiv \mathbf{Rec}[\kappa](c)(c_p) : \mathbf{T}_{32} \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau \rightsquigarrow sv' \end{array}}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} \mathbf{unroll}_\tau sv : c(\mathbf{Rec}[\kappa]c)c_p \rightsquigarrow \mathbf{unroll}_{|\tau|} sv'}$$

$$\frac{\begin{array}{c} \Delta \vdash \tau \equiv \exists[\kappa](c') : \mathbf{T}_{32} \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : c'c \rightsquigarrow sv' \end{array}}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} \mathbf{pack} sv \text{ as } \tau \text{ hiding } c : \tau \rightsquigarrow (\mathbf{pack}[|\tau|][c])sv'}$$

$$\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \forall[\kappa](c') \rightsquigarrow sv'}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv[c] : c'c \rightsquigarrow sv'[|c|]}$$

Float values

$$\boxed{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{64} fv : \phi \rightsquigarrow fv'}$$

$$\frac{}{\Psi; \Delta; \Gamma[x_f:\phi]; \mathcal{A}[x_f \rightarrow \mathbf{f}], \sigma_1, \sigma_2 \vdash_{64} x_f : \phi \rightsquigarrow \mathbf{f}}$$

$$\frac{}{\Psi; \Delta; \Gamma[x_f:\phi]; \mathcal{A}[x_f \rightarrow \mathbf{sp}(i)], \sigma_2, \sigma_2 \vdash_{32} x_f : \phi \rightsquigarrow \mathbf{sp}(i)}$$

$$\frac{}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_2, \sigma_2 \vdash_{32} \mathbf{r} : \mathbf{Float} \rightsquigarrow \mathbf{r}}$$

64 bit Instructions

$$\boxed{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \text{dest}_{32} \leftarrow fopr : \phi \rightsquigarrow S}$$

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{f}_t \leftarrow fopr : \phi \rightsquigarrow S$$

$$\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{sp}(i) \leftarrow fopr : \phi \rightsquigarrow S; \quad \mathbf{fswrite} \ \mathbf{sp}(i), \mathbf{f}_t; \quad \mathbf{fjunk} \ \mathbf{f}_t}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{f} \leftarrow fv : \phi \rightsquigarrow fv'}$$

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} fv : \phi \rightsquigarrow fv'$$

$$\frac{}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{f} \leftarrow fv : \phi \rightsquigarrow \mathbf{fmov} \ \mathbf{f}_t, fv'}$$

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \mathbf{Array}_{64}(\phi) \rightsquigarrow sv'_1 \quad \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_2 : \mathbf{Int} \rightsquigarrow sv'_2$$

$$\frac{}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{f} \leftarrow \mathbf{sub}_{\phi}(sv_1, sv_2) : \phi \rightsquigarrow \mathbf{sub}_{|\phi|} \ \mathbf{f}, sv'_1, sv'_2}$$

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \mathbf{Boxed}(\phi) \rightsquigarrow sv'$$

$$\frac{}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{f} \leftarrow \mathbf{unbox} \ sv : \phi \rightsquigarrow \mathbf{floadr} \ \mathbf{f}, sv'}$$

Instructions

$$\boxed{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \text{dest}_{32} \leftarrow opr : \tau \rightsquigarrow S}$$

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r}_t \leftarrow opr : \tau \rightsquigarrow S$$

$$\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{sp}(i) \leftarrow opr : \tau \rightsquigarrow S; \quad \mathbf{swrite} \ \mathbf{sp}(i), \mathbf{r}_t; \quad \mathbf{junk} \ \mathbf{r}_t}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau \rightsquigarrow sv'}$$

$$\frac{}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow sv : \tau \rightsquigarrow \mathbf{mov} \ \mathbf{r}, sv'}$$

$$\begin{array}{c}
\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau_1 \times \dots \tau_n \rightsquigarrow sv'}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \text{select}^i sv : \tau_i \rightsquigarrow \text{mov } \mathbf{r}, sv' \quad \text{loadr } \mathbf{r}, \mathbf{r}(i)} \\
\\
\frac{}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \text{dyntag}_c : \text{Dyntag}(c) \rightsquigarrow \text{dyntag}_{|c|} \mathbf{r}} \\
\\
\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{64} fv : \text{Float} \rightsquigarrow fv'}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \text{box } fv : \text{Boxed}(\phi) \rightsquigarrow \text{malloc}_{|\phi|} \mathbf{r}, fv'} \\
\\
\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \tau_i \rightsquigarrow sv'_i}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \langle sv_1, \dots, sv_n \rangle : \tau_1 \times \dots \times \tau_n \rightsquigarrow \text{malloc } \mathbf{r}, [| \tau_1 |, \dots, | \tau_n |] \langle sv'_1, \dots, sv'_n \rangle} \\
\\
\frac{\begin{array}{l} |\Gamma|_{\mathcal{A}}^{\epsilon}(\mathbf{sp}) = \{\mathbf{r}_1 : ns_{32}, \mathbf{r}_2 : ns_{32}, \mathbf{r}_e : ns_{32}, \mathbf{r}_t : ns_{32}, \mathbf{f}_1 : ns_{64}, \mathbf{f}_2 : ns_{64}, \mathbf{sp} : \sigma_f\} \\ \Delta^{\rho_1, \rho_2} \vdash \sigma_f \equiv \underbrace{ns_{32} \triangleright_{32} \dots \triangleright_{32} ns_{32}}_m \triangleright_{32} \underbrace{ns_{64} \triangleright_{64} \dots \triangleright_{64} ns_{64}}_k \triangleright_{64} \sigma' : ST \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \text{Code}(\tau_0, \dots, \tau_{m-1})(\phi_0, \dots, \phi_{k-1}) \rightarrow \tau \rightsquigarrow sv' \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_i : \tau_i \rightsquigarrow sv'_i \quad i \in 0 \dots m-1 \\ \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} fv_i : \phi_i \rightsquigarrow fv'_i \quad i \in 0 \dots k-1 \end{array}}{\begin{array}{l} \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \text{call } sv(sv_0, \dots, sv_{m-1})(fv_0, \dots, fv_{k-1}) : \tau \rightsquigarrow \\ \quad \text{fswrite } \mathbf{sp}(m + 2 * (k - 1)), fv'_{k-1} \\ \quad \vdots \\ \quad \text{fswrite } \mathbf{sp}(m), fv'_0 \\ \quad \text{swrite } \mathbf{sp}(m - 1), sv'_{m-1} \\ \quad \vdots \\ \quad \text{swrite } \mathbf{sp}(0), sv'_0 \\ \quad \text{junkregs} \\ \quad \text{call } sv'[\sigma' \circ \sigma_1, \sigma_2] \\ \quad \text{movj } \mathbf{r}, \mathbf{r}_t \end{array}} \\
\\
\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \text{Int} \rightsquigarrow sv'_1 \quad \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_2 : \tau \rightsquigarrow sv'_2}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \text{array}_{\tau}(sv_1, sv_2) : \text{Array}_{32}(\tau) \rightsquigarrow \text{malloc}_{|\tau|} \mathbf{r}, sv'_1, sv'_2} \\
\\
\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \text{Int} \rightsquigarrow sv' \quad \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{64} fv : \phi \rightsquigarrow fv'}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \text{farray}(sv, fv) : \text{Farray}(\phi) \rightsquigarrow \text{fmalloc}_{|\phi|} \mathbf{r}, sv', fv'}
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \text{Array}_{32}(\tau) \rightsquigarrow sv'_1 \quad \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_2 : \text{Int} \rightsquigarrow sv'_2 \\
\hline
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \text{sub}_\tau(sv_1, sv_2) : \tau \rightsquigarrow \text{sub}_{|\tau|} \mathbf{r}, sv'_1, sv'_2 \\
\\
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \text{Array}_{32}(\tau) \rightsquigarrow sv'_1 \\
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_2 : \text{Int} \rightsquigarrow sv'_2 \quad \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_3 : \tau \rightsquigarrow sv'_3 \\
\hline
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \text{upd}_\tau(sv_1, sv_2, sv_3) : \text{Unit} \rightsquigarrow \text{upd } sv'_1, sv'_2, sv'_3 \\
\text{malloc } \mathbf{r}, [\langle \rangle] \\
\\
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \text{Array}_{64}(\phi) \rightsquigarrow sv'_1 \\
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_2 : \text{Int} \rightsquigarrow sv'_2 \quad \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{64} fv : \phi \rightsquigarrow fv' \\
\hline
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathbf{r} \leftarrow \text{upd}_\phi(sv_1, sv_2, fv) : \text{Unit} \rightsquigarrow \text{upd } sv'_1, sv'_2, fv' \\
\text{malloc } \mathbf{r}, [\langle \rangle]
\end{array}$$

Expressions

$$\boxed{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; H}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma; \mathcal{A}', \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F \\
|\Psi|; \Delta^{\rho_1, \rho_1}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash S \Rightarrow |\Gamma|_{\mathcal{A}'}^{\sigma_1 \circ \sigma_2} \\
\text{where } \mathcal{A}' \text{ is a good allocator for } e \\
\hline
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow (S; I); F
\end{array}$$

$$\begin{array}{c}
\Delta^{\rho_1, \rho_2} \vdash \sigma_1 \equiv (\text{cont}(\overline{m})(\sigma'_1)(\sigma_2)(|\tau_r|)) \triangleright_{32} \sigma'_1 : ST \\
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau_r \rightsquigarrow sv' \quad \text{frmsz}(\mathcal{A}) = n \\
\hline
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{ret}} sv : \tau_r \rightsquigarrow \text{mov } \mathbf{r}_t, sv'; \\
\text{sfree } n; \\
\text{junkregs}; \\
\text{junkstack } 1 \dots m; \\
\text{ret}
\end{array}$$

$$\begin{array}{c}
|\Psi|; \Delta^{\rho_1, \rho_2}; |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \vdash sv_l : \Gamma\{\mathbf{r}_t : |\tau|\} \rightarrow 0 \\
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \tau \rightsquigarrow sv' \\
\hline
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp } sv_l} sv : \tau \rightsquigarrow \text{mov } \mathbf{r}_t, sv' \\
\text{jmp } sv_l
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma; \mathcal{A}, \epsilon, \epsilon \vdash_{32} sv : \tau_r \rightsquigarrow sv' \quad \text{frmsz}(\mathcal{A}) = n \\
\hline
\Psi; \Delta; \Gamma; \mathcal{A}, \epsilon, \epsilon \vdash_{\text{halt}} sv : \tau_r \rightsquigarrow \text{mov } \mathbf{r}_t, sv' \\
\text{sfree } n \\
\text{halt}_{\tau_r}
\end{array}$$

$$\begin{array}{c}
\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathcal{A}(x) \leftarrow opr : \tau_i \rightsquigarrow S \quad \Psi; \Delta; \Gamma, x : \tau_i; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let}_\tau x = opr \text{ in } e : \tau \rightsquigarrow (S; I); F} \\
\\
\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash \mathcal{A}(x_f) \leftarrow i : \phi_i \rightsquigarrow S \quad \Psi; \Delta; \Gamma, x_f : \phi_i; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let}_\tau x_f = fopr \text{ in } e : \tau \rightsquigarrow (S; I); F} \\
\\
\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \exists[\kappa][c] \rightsquigarrow sv' \quad \Delta, \alpha : \kappa; \Gamma, x : (c\alpha); \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F}{\Psi; \Delta; \Gamma; \mathcal{A} \vdash_C \text{let}[\alpha, x] = \text{unpack } sv \text{ in } e : \tau \rightsquigarrow \text{unpack } [\alpha, \mathbf{r_t}] sv'; \\ \mathbf{srmov } \mathcal{A}(x), \mathbf{r_t}; \\ \mathbf{junk } \mathbf{r_t}; \\ I; \\ F} \\
\\
\frac{\Delta \vdash c \equiv \langle c_1, c_2 \rangle : \kappa_1 \times \kappa_2 \quad \Psi; \Delta, ; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e[c_1, c_2/\beta, \gamma] : \tau[c_1, c_2/\beta, \gamma] \rightsquigarrow I; F}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let}\langle \beta, \gamma \rangle = c \text{ in } e : \tau \rightsquigarrow I; F} \\
\\
\frac{\Delta \vdash c \equiv \alpha : \kappa_1 \times \kappa_2 \quad \left. \begin{array}{l} \Psi; \Delta, \beta : \kappa_1, \gamma : \kappa_2; \\ \Gamma[\langle \beta, \gamma \rangle / \alpha]; \\ \mathcal{A}, \sigma_1[\langle \beta, \gamma \rangle / \alpha], \sigma_2[\langle \beta, \gamma \rangle / \alpha] \end{array} \right\} \vdash_{C[\langle \beta, \gamma \rangle / \alpha]} e[\langle \beta, \gamma \rangle / \alpha] : \tau[\langle \beta, \gamma \rangle / \alpha] \rightsquigarrow I; F}{\Psi; \Delta, \alpha : \kappa_1 \times \kappa_2; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let}\langle \beta, \gamma \rangle = c \text{ in } e : \tau \rightsquigarrow \text{refine}\langle \beta, \gamma \rangle = |c|; \\ I; \\ F} \\
\\
\frac{\Delta \vdash c \equiv \text{fold}_{\mu j. \kappa} c' : \mu j. \kappa \quad \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_1 \vdash_C e[c'/\beta] : \tau \rightsquigarrow I; F}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let fold } \beta = c \text{ in } e : \tau \rightsquigarrow I; F} \\
\\
\frac{\Delta, \alpha : \mu j. \kappa, \Delta' \vdash c \equiv \alpha : \mu j. \kappa \quad \left. \begin{array}{l} \Psi; \Delta, \beta : \kappa[\mu j. \kappa / j], \Delta'; \\ \Gamma[\text{fold}_{\mu j. \kappa} \beta / \alpha]; \\ \mathcal{A}, \\ \sigma_1[\text{fold}_{\mu j. \kappa} \beta / \alpha], \\ \sigma_1[\text{fold}_{\mu j. \kappa} \beta / \alpha] \end{array} \right\} \vdash_{C[\text{fold}_{\mu j. \kappa} \beta / \alpha]} e[\text{fold}_{\mu j. \kappa} \beta / \alpha] : \tau[\text{fold}_{\mu j. \kappa} \beta / \alpha] \rightsquigarrow I; F}{\Psi; \Delta, \alpha : \mu j. \kappa, \Delta'; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let fold } \beta = c \text{ in } e : \tau \rightsquigarrow \text{refine}[\text{fold } \beta]c; \\ I; \\ F}
\end{array}$$

$$\frac{\Delta \vdash c \equiv \text{inj}_i^{+[\kappa_1, \dots, \kappa_i, \dots, \kappa_n]} c' : +[\kappa_1, \dots, \kappa_i, \dots, \kappa_n] \quad \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e[c'/\beta] : \tau \rightsquigarrow I; F}{\Psi; \Delta; \Gamma \vdash_C \text{let}_\tau \text{inj}_i \beta = (c, sv) \text{ in } e : \tau \rightsquigarrow I; F}$$

$$\frac{\begin{array}{c} \Delta, \alpha : +[\kappa_1 \dots, \kappa_n], \Delta' \vdash c \equiv \alpha : +[\kappa_1 \dots \kappa_n] \quad \beta \notin \Delta, \Delta' \\ \Psi; \Delta, \beta : \kappa_i, \Delta'; \\ \Gamma[\text{inj}_i \beta / \alpha]; \\ \mathcal{A}, \sigma_1[\text{inj}_i \beta / \alpha], \sigma_2[\text{inj}_i \beta / \alpha] \end{array} \left. \vphantom{\begin{array}{c} \Delta, \alpha : +[\kappa_1 \dots, \kappa_n], \Delta' \vdash c \equiv \alpha : +[\kappa_1 \dots \kappa_n] \quad \beta \notin \Delta, \Delta' \\ \Psi; \Delta, \beta : \kappa_i, \Delta'; \\ \Gamma[\text{inj}_i \beta / \alpha]; \\ \mathcal{A}, \sigma_1[\text{inj}_i \beta / \alpha], \sigma_2[\text{inj}_i \beta / \alpha] \end{array}} \right\} \vdash_{C[\text{inj}_i \beta / \alpha]} e[\text{inj}_i \beta / \alpha] : \tau[\text{inj}_i \beta / \alpha] \rightsquigarrow I; F \\ \Psi; \Delta, \beta : \kappa_j, \Delta'; \Gamma[\text{inj}_j \beta / \alpha]; \mathcal{A}, \sigma_1[\text{inj}_j \beta / \alpha], \sigma_2[\text{inj}_j \beta / \alpha] \vdash_{32} sv[\text{inj}_j \beta / \alpha] : \text{Void} \rightsquigarrow sv' \\ j \in 1 \dots i-1, i+1 \dots n}{\Psi; \Delta, \alpha : +[\kappa_1, \dots, \kappa_n], \Delta' \vdash_C \text{let}_\tau \text{inj}_i \beta = (c, sv) \text{ in } e : \tau \rightsquigarrow \text{refine}[\text{inj}_i \beta][c], sv'; \\ I; \\ F}$$

$$\frac{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \text{Dyn} \rightsquigarrow sv'}{\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let } x = \text{raise}_\tau sv \text{ in } e : \tau \rightsquigarrow \begin{array}{l} \text{mov } \mathbf{r}_1, sv \\ \text{loadr } \mathbf{r}_t, \mathbf{r}_e(1) \\ \text{mov } \mathbf{sp}, \mathbf{r}_t \\ \text{loadr } \mathbf{r}_t, \mathbf{r}_e(0) \\ \text{jmp } \mathbf{r}_t \end{array}}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \text{Dyn} \rightsquigarrow sv' \quad \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv_1 : \text{Dyntag}(\tau_1) \rightsquigarrow sv' \\
\Psi; \Delta; \Gamma, x_1 : \times [\text{Dyntag}(\tau_1), \tau_1]; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp } \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_1, \sigma_2]} e_1 : \tau_x \rightsquigarrow I_1; F_1 \\
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp } \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_1, \sigma_2]} e_2 : \tau_x \rightsquigarrow I_2; F_2 \\
\Psi; \Delta; \Gamma, x : \tau_x; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F \\
|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} = \Gamma_A
\end{array}

\begin{array}{c}
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let } x = \text{dyncase}(sv)(sv_1 \Rightarrow x_1.e_1, - \Rightarrow e_2) \text{ in } e : \tau \rightsquigarrow \\
\text{mov } \mathbf{r_t}, sv' \\
\text{brdyn } \mathbf{r_t}, sv'_1, \ell_1[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\
\text{junk } \mathbf{r_t} \\
I_2 \\
\ell_1 : \forall[|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A[\mathbf{r_t} : \times [\text{Dyntag}(|\tau_1|), |\tau_1|]] \\
\text{srmov } \mathcal{A}(x_1), \mathbf{r_t} \\
\text{junk } \mathbf{r_t} \\
I_1 \\
\ell_{\text{end}} : \forall[|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A[\mathbf{r_t} : |\tau_x|] \\
\text{srmov } \mathcal{A}(x), \mathbf{r_t} \\
\text{junk } \mathbf{r_t} \\
I; \\
F; \\
F_1; \\
F_2
\end{array}$$

$$\begin{aligned}
& \Psi; \Delta; \Gamma; \mathcal{A}, \epsilon, \sigma_3 \vdash_{\text{jmp}} \ell_{\text{post}}[(\Pi_1 \Delta), \epsilon, \sigma_h] e_1 : \tau_x \rightsquigarrow I_1; F_1 \\
& \Psi; \Delta; \Gamma, x_2 : \text{Dyn}; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp}} \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_f \circ \sigma_1, \sigma_2] e_2 : \tau_x \rightsquigarrow I_2; F_2 \\
& \Psi; \Delta; \Gamma, x : \tau_x; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau \rightsquigarrow I; F \\
\Gamma_A = \{ & \mathbf{r}_1 : ns_{32}, \mathbf{r}_2 : ns_{32}, \mathbf{r}_t : ns_{32}, \mathbf{r}_e : \mathbf{Exnptr}(\sigma_2), \mathbf{sp} : \sigma_f \circ \sigma_1 \circ \sigma_2, \mathbf{f}_1 : ns_{64}, \mathbf{f}_2 : ns_{64} \} \\
& \text{where } \Gamma_A = |\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} \{ \mathbf{r}_e : \mathbf{Exnptr}(\sigma_2) \\
& \text{and } \Gamma_B = |\Gamma|_{\mathcal{A}}^{\sigma_3} \{ \mathbf{r}_e : \mathbf{Exnptr}(\sigma_3) \\
& \text{and } \sigma_f = |\Gamma|_{\mathcal{A}}^{\epsilon}(\mathbf{sp}) \\
& \text{and } \sigma_3 = \mathbf{Exnptr}(\sigma_2) \triangleright_{32} \sigma_f \circ \sigma_1 \circ \sigma_2 \\
& \text{and } \sigma_h = \sigma_f \circ \sigma_3
\end{aligned}$$

$$\begin{aligned}
& \Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let } x = \text{handle}_{\tau_x}(e_1, x_2.e_2) \text{ in } e : \tau \rightsquigarrow \\
& \quad \text{push } \mathbf{r}_e \\
& \quad \text{stackcopy } (\mathbf{Exnptr}(\sigma_2)) \triangleright_{32} \sigma_f \\
& \quad \text{sfree } 1 \\
& \quad \text{malloc } \mathbf{r}_e[\mathbf{Exnhndler}(\sigma_h), \sigma_h] \langle \ell_{\text{handle}}[\Pi_1 \Delta, \rho_1, \rho_2], \mathbf{sp} \rangle \\
& \quad I_1 \\
& \quad \ell_{\text{handle}} : \forall [\Delta, \rho_1, \rho_2]. \{ \mathbf{r}_1 : \text{Dyn}, \mathbf{r}_2 : ns_{32}, \mathbf{r}_t : ns_{32}, \mathbf{r}_e : ns_{32}, \mathbf{sp} : \sigma_h, \mathbf{f}_1 : ns_{64}, \mathbf{f}_2 : ns_{64} \} \\
& \quad \text{srmov } \mathcal{A}(x_2), \mathbf{r}_1 \\
& \quad \text{sfree}(\text{frmsz}(\mathcal{A})) \\
& \quad \text{pop } \mathbf{r}_e \\
& \quad \text{junk } \mathbf{r}_1 \\
& \quad I_2 \\
& \quad \ell_{\text{post}} : \forall [|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_B[\mathbf{r}_t : |\tau_x|] \\
& \quad \text{sfree}(\text{frmsz}(\mathcal{A})) \\
& \quad \text{pop } \mathbf{r}_e \\
& \quad \text{jmp } \ell_{\text{end}}[\Pi_1 \Delta, \sigma_1, \sigma_2] \\
& \quad \ell_{\text{end}} : \forall [|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A[\mathbf{r}_t : |\tau_x|] \\
& \quad \text{srmov } \mathcal{A}(x), \mathbf{r}_t \\
& \quad \text{junk } \mathbf{r}_t \\
& \quad I; \\
& \quad F; \\
& \quad F_1; \\
& \quad F_2
\end{aligned}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{32} sv : \bigvee [\tau_1, \dots, \tau_n] \rightsquigarrow sv' \\
\Delta \vdash \tau_i \equiv \text{Tag}(i) : T_{32} \quad i \in 1 \dots (j-1) \quad \Delta \vdash \tau_i \equiv \times[\text{Tag}(i), \tau'_i] : T_{32} \quad i \in j \dots n \\
\Psi; \Delta; \Gamma, x_i : \tau_i; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{jmp } \ell_{\text{end}}[(\Pi_1 \Delta), \sigma_1, \sigma_2]} e_i : \tau \rightsquigarrow I_i; F_i \quad i \in 1 \dots n \\
\Psi; \Delta; \Gamma, x : \tau; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C e : \tau_e \rightsquigarrow I; F \\
|\Gamma|_{\mathcal{A}}^{\sigma_1 \circ \sigma_2} = \Gamma_A
\end{array}$$

$$\Psi; \Delta; \Gamma; \mathcal{A}, \sigma_1, \sigma_2 \vdash_C \text{let } x = \text{case}_{\tau}(sv)(x_0.e_0, \dots, x_n.e_n) \text{ in } e : \tau_e \rightsquigarrow$$

$$\begin{array}{l}
\text{mov } \mathbf{r}_t, sv' \\
\text{brtag}_0 \mathbf{r}_t, \ell_0[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\
\vdots \\
\text{brtag}_{k-1} \mathbf{r}_t, \ell_{k-1}[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\
\text{brtgd}_k \mathbf{r}_t, \ell_k[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\
\vdots \\
\text{brtag}_{n-1} \mathbf{r}_t, \ell_{n-1}[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\
\text{mov } \mathbf{r}_t, \text{forgetunion } \mathbf{r}_t \\
\text{jmp } \ell_n[(\Pi_1 \Delta), \sigma_1, \sigma_2] \\
\ell_0 : \forall[|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A[\mathbf{r}_t : |\tau_0|] \\
\text{srmov } \mathcal{A}(x_0), \mathbf{r}_t \\
\text{junk } \mathbf{r}_t \\
I_0 \\
\ell_1 : \forall[|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A[\mathbf{r}_t : |\tau_1|] \\
\vdots \\
\ell_n : \forall[|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A[\mathbf{r}_t : |\tau_n|] \\
\text{srmov } \mathcal{A}(x_n), \mathbf{r}_t \\
\text{junk } \mathbf{r}_t \\
I_n \\
\ell_{\text{end}} : \forall[|\Delta|, \rho_1 : ST, \rho_2 : ST]. \Gamma_A[\mathbf{r}_t : |\tau|] \\
\text{srmov } \mathcal{A}(x), \mathbf{r}_t \\
\text{junk } \mathbf{r}_t \\
I; \\
F; \\
F_1; \\
\vdots \\
F_n
\end{array}$$

Heap values

$$\boxed{\Psi \vdash_h hval : \tau \rightsquigarrow hval; H}$$

$$\begin{array}{l}
\Psi; \alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k; x_1:\tau_1, \dots, x_m:\tau_m, z_1:\phi_1, \dots, z_n:\phi_n; \mathcal{A}, \sigma_1, \sigma_2 \vdash_{\text{ret}} e : \tau \rightsquigarrow I; F \\
\text{Where } \mathcal{A} \text{ is a good allocator for e} \\
\text{and } \sigma_1 = (\text{cont}(\overline{m + 2 * n})(\rho_1)(\rho_2)(|\tau|)) \triangleright_{32} \sigma_{32} \circ \sigma_{64} \circ \rho_1 \\
\text{and } \sigma_{32} = |\tau_1| \triangleright_{32} \dots \triangleright_{32} |\tau_m| \triangleright_{32} \epsilon \\
\text{and } \sigma_{64} = |\phi_1| \triangleright_{64} \dots \triangleright_{64} |\phi_n| \triangleright_{64} \epsilon \\
\text{and } \sigma_2 = \rho_2 \\
\text{and } l = \text{frmsz}(\mathcal{A}) \\
\text{and } hval = \text{code}[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k, \rho_1:ST, \rho_2:ST][\text{Code}(\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n)(\tau)] \\
\quad \text{salloc } l; \\
\quad \text{srmov } \mathcal{A}(x_1), \text{sp}(l + 0); \\
\quad \vdots \\
\quad \text{srmov } \mathcal{A}(x_m), \text{sp}(l + m - 1); \\
\quad \text{srfmov } \mathcal{A}(z_1), \text{sp}(l + m + 2 * 0); \\
\quad \vdots \\
\quad \text{srfmov } \mathcal{A}(z_n), \text{sp}(l + m + 2 * (n - 1)); \\
I
\end{array}$$

$$\begin{array}{l}
\Psi \vdash_h \text{code}_\tau[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k](x_1:\tau_1, \dots, x_m:\tau_m)(z_1:\phi_1, \dots, z_n:\phi_n).e : \\
\quad \forall[\alpha_1:\kappa_1, \dots, \alpha_k:\kappa_k] \text{Code}(\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n)(\tau) \rightsquigarrow hval; F
\end{array}$$

Heaps

$$\boxed{\Psi \vdash d \rightsquigarrow H}$$

$$\begin{array}{c}
\overline{\Psi \vdash \epsilon \rightsquigarrow \epsilon} \\
\\
\Psi[\ell:\tau] \vdash_h hval : \tau \rightsquigarrow hval'; F \quad \Psi[\ell:\tau] \vdash d \rightsquigarrow F' \\
\hline
\Psi[\ell:\tau] \vdash d, \ell:\tau \mapsto hval \rightsquigarrow \ell:|\tau|.hval', F; F'
\end{array}$$

Programs

$$\boxed{\vdash p : \tau \rightsquigarrow P}$$

$$\begin{array}{l}
\Psi(d) \vdash d \rightsquigarrow H' \quad \Psi(d); \bullet; \bullet; \mathcal{A}, \epsilon, \epsilon \vdash_{\text{halt}} e : \tau \rightsquigarrow I'; F \\
\text{Where } \mathcal{A} \text{ is a good allocator for e} \\
\text{and } H = F; H', \\
\quad \ell_{\text{ihandle}}: \text{Exnhandler}(\epsilon) \\
\quad \text{mov } \mathbf{r}_t, \mathbf{r}_1; \\
\quad \text{halt}_{\text{dyn}} \\
\text{and } R = \{\mathbf{r}_1 \mapsto ns_{32}, \mathbf{r}_2 \mapsto ns, \mathbf{r}_e \mapsto ns_{32}, \mathbf{r}_t \mapsto ns, \mathbf{f}_1 \mapsto ns_{64}, \mathbf{f}_2 \mapsto ns_{64}, \mathbf{f}_t \mapsto ns_{64}, \mathbf{sp} \mapsto \epsilon\} \\
\text{and } I = \text{malloc } \mathbf{r}_e[\text{Exnhandler}(\epsilon), \epsilon][\ell_{\text{ihandle}}, \mathbf{sp}]; I' \\
\hline
\bullet \vdash \text{letrec } d \text{ in } e : \tau \rightsquigarrow H, R, I
\end{array}$$

Chapter 9

Implementation

In the previous chapters I developed a series of translations mapping programs in an idealized version of the original **TILT** internal language (the **MIL**) down to an idealized typed assembly language (**TILTAL**) and proved the soundness of these translations. In this chapter I describe a new backend that I have implemented in the **TILT** compiler implementing a concrete version of these translations.

The overall structure of the new backend can be seen in figure 9.1. The implementations of the translations generally follow quite closely on the formal description, and the **LIL** language as implemented is almost exactly the same as the formal version. The implementation differs most significantly from the formal presentation in that its final target is the **TALx86** language [MCG⁺99, GM00]. Unlike the RISC-like **TILTAL** language presented in chapter 7, **TALx86** is specialized to the x86 architecture. Moreover, **TALx86** uses numerous additional technologies (such as alias types [WM01]) that I do not attempt to account for in my formal presentation. The **TALx86** language was chosen as the target because of its the numerous tools already existing for it (such as an assembler, typechecker, and linker) which I was able to use with minimal modification.

In this chapter, I describe the individual passes added to the compiler as part of my implementation work and discuss briefly the important differences between the implementation and the theoretical presentation. In order to demonstrate the practicality of my approach, I present some measurements to quantify the size of the typed binaries and their runtime performance as compared to other compilers.

9.1 Singleton Elimination

A key issue in typed compilation is controlling the size of the intermediate forms of programs. Type annotations on internal representation terms quickly come to dominate the overall representation size to the point that representing and traversing the type information becomes by far the dominant performance issue. However, there is in general a great deal of redundancy in this type information which can be exploited to eliminate or reduce this problem.

A number of mechanisms have been suggested [Sha97, PCHS00, GM00] to attempt to deal with this problem. In the **TILT** compiler, the **MIL** as implemented controls type sizes by using *singleton kinds* to provide a form of type definition [PCHS00, SH99] intrinsic to the language. Types are bound to variables using a derived **let** form so that the type can be referred to subsequently via its name. Compiler passes such as common sub-expression elimination find redundant uses of types




Figure 9.1: Structure of the certifying **TILT** backend

and reduce them to a single binding.

The type theory of singleton kinds has been greatly simplified by recent work. However, rather than attempt to reconcile this type theory with the **LX** type theory used for type-analysis, I chose for the purposes of the formal translation to use a singleton-free form of the **MIL**, and consequently to rely in the implementation of the new backend on other technologies to control the size of type information. Since the **TILT** compiler produces internal representations using singletons, the first additional pass to be added to the compiler was a singleton elimination pass.

Crary [Cra00] showed that the eliminability of singleton kinds follows from the completeness of the equivalence algorithm given by Stone and Harper [SH99] and gave an algorithm for computing the singleton free version of a term. The implementation of singleton elimination in the compiler follows Crary’s algorithm almost exactly with one minor changes in order to avoid problems with type sizes, as explained below.

9.1.1 Preserving sharing of types

It would in theory have been possible to have implemented singleton elimination concurrently with the translation from the **MIL** to the **LIL** language. Since the **LIL** language has its own mechanisms for preserving sharing (discussed below in section 9.6), the result would have been a translation which neither traversed types repeatedly, nor generated internal forms with excess redundant type information. However, doing these two translations simultaneous seemed likely to be difficult and error-prone. Moreover, the singleton elimination phase introduces a large amount of additional constructor data that benefits from being exposed to the **MIL** optimization passes. Eliminating singletons concurrently with the translation to **LIL** would have prevented this.

For these reasons, singleton elimination was implemented in **TILT** in two phases. The early phase eliminates all singletons except uses which occur via the derived **let** form (an immediate application of a type lambda to a type in the underlying calculus). Since the derived elimination rule for the **let** construct is simple substitution, it is trivial to postpone the elimination of the derived **let** type construct until the translation to **LIL**, providing the **MIL** optimizer with an opportunity to improve the types emitted by the elimination phase¹. This two phase approach allows singleton elimination to be done early in the compilation process while still preserving a compact representation using the **let** definitional mechanism until translation to **LIL**.

9.1.2 Optimizations for singleton elimination

Eta reduction

The singleton elimination pass introduces many new types that turn out to be eta-expansions of variables. It is straightforward to instrument the singleton elimination pass to catch many of these eta-expansions in order to perform the eta-reduction in place. However, even with this improvement on the original algorithm, singleton elimination still produces a fair number of eta-redices. Consequently, performing eta-reduction after singleton elimination is an important optimization to avoid redundant representation of types, as well as to reduce the size of the intermediate representation.

¹An alternative way of viewing this implementation choice is as mapping into a calculus which can be given two different (but presumably equivalent) theories: one which provides a **let** form derived via singleton kinds, and one which provides a primitive **let** form with a substitution semantics. The equivalence of the two theories would have to be shown, however.




Figure 9.2: Benchmark timings with singleton elimination (normalized to timings without singleton elimination).

Beta reduction

In addition to eta-redices, singleton elimination produces numerous beta-redices in the forms of projections from known records or applications of known functions. Eliminating the former is always a desired optimization for code in named-form since it can only reduce code size and eliminates an address calculation and a memory fetch. Eliminating the latter is a speculative optimization since the function may be called more than once. Consequently, even if it were practical to do these optimizations concurrently with singleton elimination, it is almost certainly preferable to leave this up to the general optimization and inlining code which must deal with these issues in any case.

Common sub-expression elimination

Finally, singleton elimination may produce several copies of equivalent types. It is therefore valuable to do common sub-expression elimination (CSE) on the resulting code. It is particularly important after singleton elimination that CSE unify alpha-equivalent type functions, since it is not uncommon for many of the generated types to be simple alpha-variants.

9.1.3 Runtime behavior

Since types must be represented as data in the **TILT** compiler, it is important to consider the effect of singleton elimination on the runtime behavior of programs. While this information is not easily available for the certifying backend since it relies intrinsically on singleton elimination, code can be generated for the untyped Sparc backend using either normal or singleton-free internal representations. Once singletons are eliminated from the internal representations of the program,

the remaining **MIL** optimization passes can be run as usual, since they introduce no uses of singleton kinds except in the form of the derived `let`. For the most part these passes behave identically when targeting both the untyped and the typed backends, up until the point where the translation to the next intermediate language takes place.

Figure 9.2 shows that across a wide range of benchmarks, the singleton-free code generally performs similarly to the non-singleton eliminated code. There was almost no difference in object code size between the two generated executables.

I speculate that the occasional improvement in the runtime behavior can be attributed to the singleton-elimination phase acting as a rudimentary form of cross-module inlining, by extracting all possible definitional information about abstract types from their kinds, and exposing that information to the optimizer. In principle, an optimization pass could be implemented to perform a similar function without actually doing singleton elimination.

9.2 MIL to LIL

The implemented translation of **MIL** code into **LIL** code follows the formal translation from chapter 5 almost exactly in structure. The only theoretically significant change is that in addition to the `vararg` and array optimizations presented in chapter 5, the implemented version performs a related representation optimization for sums. This optimization adds no significant complication to the translation, merely requiring that the encodings of types distinguish between those types which are inhabited by heap pointers and those which are inhabited by other values. This distinction is used to take advantage of the fact that pointers are never small integers. Sums with exactly one value carrying arm can be represented specially when the carried value is a pointer: there is no need to box and tag the pointer since it can always be distinguished from small integer tags.

The **MIL** intermediate representation which is passed to the translator is singleton-free, except with respect to `let`-bound definitions as discussed above. Some care is required in the implementation to avoid expanding out type definitions indiscriminately, and also to avoid unnecessary work. For example, consider a **MIL** term of the form `let a = c in e`. The bound variable `a` may potentially occur many times in the body of `e`. Moreover, it is not syntactically apparent whether `a` is used as data (and hence requires static encoding and dynamic encoding) or as a classifier (and hence requires static encoding and interpretation), or as both, or neither. The definition of singleton elimination tells us that the correct translation of this term will be equivalent to the translation of `e[c/a]` (since this is the intended meaning of the residual `let` binding). However, simply expanding out the definition via substitution and translating the resulting terms is unacceptably inefficient, both in terms of the size of the resulting intermediate representation and in terms of the cost of repeated translation of the type. The earlier compiler passes will have ensured that any such definition remaining is used at least once: often it will be the case that the binding is used many times (since similar bindings are unified via common-sub-expression elimination).

The implementation avoids the cost of re-traversing these bindings by maintaining a mapping from **MIL** type variables to memoized thunks that when forced, compute the static or dynamic encodings of variables as needed. In this way, no type binding is ever translated more than once as a type and once as a constructor. Moreover, the memoization ensures that all of the occurrences of the variable will be replaced by the same physically shared translation, preventing an explosion in the representation size.

9.3 Closure Conversion

Functions in the **LIL** intermediate forms produced by the translation from **MIL** code are still lexically nested and may refer to variables bound in enclosing scopes. In order to provide an efficient compiled implementation, a commonly used strategy is to replace uses of functions (evaluated via substitution) with uses of closed code and environments. Environments preserve the values of free variables, and are passed as additional arguments to code at runtime. Functions in the high level language become pairs of code pointers and environments in the underlying machine.

9.3.1 Closure conversion strategies

Numerous strategies have been proposed for handling closure conversion in a typed setting [CWM98, MWCG97, MMH96, MH98]. I briefly summarize the two main axes of variation from an implementation standpoint, and refer the reader to Morrisett and Harper [MH98] for a detailed analysis of the different approaches and their typing properties.

The first source of variation has to do with whether recursion is implemented via a primitive notion of recursive code, or via the introduction of recursive closures. In the recursive code approach, code functions are permitted to recursively refer to themselves and other code functions in the same nest within their own scope. In the recursive closure approach, all code is non-recursive. Closures however, may refer back to themselves.

Many recursive closure implementations are based upon the idea of backpatching, in which the environment record is initialized to contain a pointer to the closure into which it will eventually be written. While simple and expedient to implement, this approach requires either a fairly powerful type system to type the resulting code, or else the use of ad-hoc “dummy closures” to stand in for the recursive reference in the initially allocated environment. In the interest of avoiding adding extra complexity to the problem, I chose to avoid this approach.

Another approach to implementing recursion is to parameterize code functions over their entire closure, instead of just the environment portion. Since the closure contains the code pointer to which it is being passed, this approach is sometimes referred to as the self-application approach.

Self-application is a fairly elegant approach to solving the problem of closure converting recursive code. However, it would require a fair bit of re-tooling in the existing TILT compiler in order to implement. In particular, the **MIL** type system does not support recursive types at higher kinds, which would seem to be necessary to efficiently support the self application semantics.

For the sake of simplicity in the implementation of closure conversion then, I chose to implement the relatively straightforward recursive code approach to closure conversion. Under this approach, mutually recursive functions nests become mutually recursive code nests, parameterized over a common environment. Functions become existentially packed tuples containing the code pointer and the environment of the function. At call sites, the code pointer and environment are projected out of the tuple, and the former is applied to the latter, along with the original arguments of the function.

9.3.2 Recursive code closure conversion

There are two well-known disadvantages to the recursive code approach to closure conversion.

The first is that since each function must be able to create the environment of each of its callees, the environment of each function in the nest must contain the transitive closure of all of

the free variables of all of the functions in the nest (assuming that nests have been reduced to strongly connected components). In practice, this means that all of the functions in the nest share one environment, which is consequently larger than the individual environments might be. This is somewhat offset by the fact that only one environment per nest need be allocated.

The second, and potentially more serious drawback of the recursive code approach is that an escaping occurrence of a function within its own body requires the closure to be reconstructed on each invocation. For functions implementing loops, this can obviously be problematic. It is important to note however that it is *not* necessary to re-allocate the environment on each invocation: only the two element closure. Moreover, this particular idiom is relatively rare: most recursive functions do not escape within their body. Nonetheless, this remains a serious disadvantage of this approach.

The advantage of this approach is that it simple to implement, and moreover is relatively simple to implement efficiently. Most applications of local functions (including all recursive calls) can be implemented as calls directly to code, without reference to closures. Consequently, for most non-escaping functions closures need never be allocated.

More specifically, as implemented in the **LIL** backend, every recursive call (including all calls to other functions in the same nest) is implemented as a direct code call. In addition, every call to a function defined in the same lexical scope as the call is implemented as a direct code call. It is straightforward to extend this to functions defined in any enclosing scope: however this may increase closure sizes and hence is not done in the **LIL** backend.

In addition to calling code functions directly, some additional optimizations are performed on the calling sequences of non-escaping functions (that is, functions which are only called directly as code). For such functions, the environment record is never actually constructed. Instead, the components of the environment are passed as additional arguments to the function (just as with the varargs optimization for ordinary records). Floating point (64 bit) elements of the environment are further optimized by passing them unboxed on the stack. For all functions, environments with only a single element are “unwrapped”: that is, the underlying singleton element is passed unboxed, instead of boxed in an environment. Finally, all closed closures are statically allocated as data (a closure may be closed because its environment is empty, or because all of its constituents are statically allocated data).

One of the advantages of using a type erasable language is that the treatment of types in closure conversion is much simpler than in the type passing case. Since types are used at compile time only, type environments can be passed immediately via application, instead of being packed into an existential. Implementing closure conversion for types was straightforward.

9.4 General Optimization

The process of encoding types as term data and making typecase explicit exposes new structure to the compiler which is available for optimization. Moreover, it greatly complicates other phases of the backend to be overly concerned with generating no dead, redundant, or otherwise sub-optimal code. Consequently, it seemed most expedient to port the **MIL** one pass optimizer to the **LIL** so that basic optimizations could be performed on **LIL** programs. The **LIL** optimizer performs CSE, dead code elimination, eta reductions for records and functions, projection from known records, reduction of known switches, folding of boxes and unboxes, constant folding, and reduction of coercions. Additionally, it makes some effort to recognize boolean terms which are used only as

arguments to switches, and hence which can be compiled to produce jumps to branch targets directly from comparisons, instead of evaluating the term to produce a boolean value, and then re-branching on this result.

The general optimizer is most useful to help clean up after the closure converter, since closure conversion involves a fairly substantial rewriting of programs. Many box/unbox reductions become available after closure conversion, since floating point terms are currently boxed in environments. Numerous record projections and record formations can be eliminated as well in cases where only some or none of the environment is used in a given context. Multiple calls to the same function in the same scope can be optimized so that only one projection of the code pointer and the environment need be done.

All these optimizations (and more) are performed on the closure converted code. In general however, a great deal more could be done to optimize after closure conversion. In particular, I believe that the output of the closure converter could benefit greatly from some form of partial-redundancy elimination which could move code evaluated on only some paths down into the arms of switches. The closure converter introduces numerous projections from a function environment in the function header, some of which are only used conditionally in the function. It may also be forced to re-allocate closures for recursive escaping occurrences which are again only used conditionally in the function. Such optimization is beyond the scope of this thesis.

9.5 Lil To Tal

The implemented translation from **LIL** to **TALx86** is conceptually quite similar to the formal description given in chapter 8, despite the fact that the **TALx86** language is substantially different from the idealized **TILTAL** language. While details of the implementation of particular constructs are specialized to the **TALx86** type system, the overall code generation approach is very similar.

Code generation and register allocation are done simultaneously in (notionally) a single pass. In fact, for simplicity the code generator is run twice on each function: once using a virtual frame pointer and once using stack pointer based frames. The output of the first run is ignored: it is done only to determine the frame size so that offsets from the stack pointer can be computed directly. This was done for the sake of simplicity in the coding only: it would be straightforward to implement a second pass to eliminate uses of the frame pointer.

9.5.1 Code generation

As with the formal description, code generation is done by translating terms with a pre-assigned destination in a similar fashion to that described by Dybvig et al. [DHB90]. This is particularly convenient given the register allocation strategy chosen, in which let-bound variables generally have already been assigned a location when their bindings are translated. Consequently, operations which are bound to variables can often be translated in such a way as to compute their values directly into the appropriate location.

As in the formal development from chapter 8, terms are also translated with a context of occurrence which indicates the use to which the result of a computation is to be put. For example, terms which are translated in a “return” context are expected to de-allocate the frame and return directly, instead of returning a value. Consequently, general tail-call elimination falls out very naturally: function bodies that end in function calls (even when guarded by a conditional or a

switch) simply de-allocate the frame and jump to the new callee. (Some additional adjustment of the stack is necessary when the callee has a different number of arguments than the caller).

Additional contexts of occurrence are used to generate efficient code for conditional branches, and for terms evaluated only for side-effects. A term whose value is used only to decide a branch is evaluated in a context of occurrence which provides appropriate continuations which can be called directly. A term whose value is unused is evaluated in a context of occurrence which indicates that no return value need be provided.

9.5.2 Register allocation

The assignment of variables and temporaries to registers and stack slots is performed simultaneously with code generation. Code generation is performed in a bottom up fashion, so that the live range of variables is apparent from the binding structure of a program: a variable becomes live when it is first encountered in the bottom-up pass, and becomes dead when its binding site is reached. Registers and stack slots are assigned to variables greedily, as needed by the generated code. When a variable use is encountered by the code generator, the register allocator is queried to provide the machine location in which the variable resides. The code generator dictates whether or not this location is permitted to be a stack slot, or must be a register.

Register assignment and spilling

The register allocator maintains state describing the assignment of variables to registers and stack slots and vice versa. This state encodes the assumptions which the previously emitted code makes about the state of the machine upon entry. When the register allocator is queried about a variable to which it has already assigned a machine location it returns the previously assigned location *if it is appropriate*. However, if the variable was previously assigned to a stack slot and the code generator requires that the use under consideration be assigned a register, the register allocator chooses a free register and changes the state to map the variable to this new location. Since previously emitted code assumes a different location for the variable, the register allocator is responsible for emitting code to initialize the previously assigned stack slot from the newly assigned register.

If a variable is required to be in a register and all registers are already assigned to variables, a variable must be chosen to spill to the stack. Since previously emitted code assumes that the variable was in a register, the allocator must emit code to load the register from the newly assigned stack slot. As a simple spill heuristic, the variable whose next use is farthest into the future (that is, whose next use is farthest forward in the instruction stream) is spilled. This is intended to attempt to provide more intervening instructions to hide the latency of the memory access, as well as to attempt to keep locally active variables in registers. No effort was put into attempting to validate this choice of heuristic as extensive tuning of the code generator is beyond the scope of this thesis. However, from informal inspection of the generated code, the allocator seems to usually make reasonable choices, though there are also cases where it could be improved.

If a location is requested for a previously unseen variable, it is assigned a register if possible, and otherwise a stack slot if permitted by the code generator. If neither of these cases apply, a variable must be spilled as described above to free up a register for the new variable.

At a variable binding site, the register allocator provides the code generator with the location into which the variable initialization code should write its value. Variables which have not been assigned a location when their binding site is encountered are unused, and the result of the

initialization code may be safely discarded.

Unlike many register allocation schemes in which variables are assigned a single location with spill code emitted to load them into temporary registers when necessary, this register allocation approach allows variable/location assignments to be changed dynamically. Variables may be moved between stack slots and registers multiple times over the course of a function. In this sense this approach is similar to the “second chance binpacking” allocation strategy described by Traub, et al. [THS98].

Machine state types

An important side benefit of the register allocation data structures is that they provide, at any given point during code generation, a snapshot of the machine state. The register allocator maintains a mapping from variables to registers and frame locations. Consequently, it is straightforward in this framework to define the translation of a **LIL** context Γ as described in chapter 8. The operation $|\Gamma|_A^\sigma$ mapping a stack tail σ and a **LIL** context Γ to a machine type Γ_A is trivial to implement given the data structures already used by the register allocator. The intention is that the register allocator as implemented should constitute a “good allocator” as described in chapter 8.

9.5.3 Additional TILT constructs

There are many more language constructs in the actual **LIL** language as implemented than in the formal **LIL** description. For the most part however, these are simply additional primitive operations on basic types, and are not especially interesting. However, two such constructs warrant additional comment.

Coercions

Vanderwaart et al. [VDP⁺03] describe a coercion calculus used in the **TILT** compiler to faithfully implement Standard ML datatypes in an efficient manner. Applications of coercions into and out of datatypes in this system incur zero runtime overhead. This system is quite simple, requiring only one new type constructor and three new term level constructs. However, it was not supported by the existing **TALx86** type system, and I preferred to modify this system as little as possible.

A simple technique for eliminating these coercions is to replace them with function calls. This approach has the benefit of placing no additional burden on the code generator and requiring no changes to the **TALx86** type system, but suffers from a substantial performance penalty [VDP⁺03]. As a compromise, I chose to implement coercions as functions with a highly specialized calling convention. In particular, they are implemented as functions whose argument and result are both in the same register, and for which all other registers are callee-save. Since the actual bodies of the functions do no runtime work, this means that the runtime cost of applying a coercion is reduced to two instructions: a call and an immediate return. There may be some additional cost since variables may need to be spilled or re-shuffled in order to place the argument value in the appropriate register before calling the coercion. However, in general the overhead is quite low.

External functions

In general, most programs need at some point to call external functions. At the very least, any interesting program must at some point make a system call to do input and output. These external

functions are probably not in general provably safe, and are in any case not available in a verifiable form. Consequently, these external operations are generally considered part of the trusted computing base: that is, the safety of the checked code is certified relative to the correctness and safety of the external code. However, it is important that there be some mechanism for controlling which external code the certified program has access to, both so that the certification system cannot be trivially subverted, and so that the ability of mobile code to access system resources can be controlled by the host.

One approach to this is to explicitly incorporate the external functionality needed by programs into the runtime layer upon which the program runs. Essentially, the source and target language are extended with a statically known set of additional primitives. This has the benefit of making it easy to control the set of resources available to programs, and of routing access to these resources through a central authority (leaving open the possibility of doing additional sandbox style dynamic checks). A disadvantage of this approach is that it is generally difficult to extend the set of external functions available, since the runtime itself must be modified.

The approach taken in this thesis is to push the issue of access to external resources out to the linking phases. The **TILT** compiler extends the Standard ML language with a primitive function type classifying functions obeying the C calling convention, and a primitive application form for applying such functions. Any compilation unit may declare a set of external functions against which it is compiled. The generated **TALx86** modules list these functions as typed imports. The code in the compiled module is certified to be safe under the assumption that the imported labels are safe at their declared types.

The question of how these import assumptions are discharged then becomes a policy decision on the part of the host at link time. The **TALx86** linker provides three different modes by which import assumptions can be discharged.

The first is that a typed binary may be provided to satisfy the import assumptions. In this case, these assumptions will be explicitly checked by the linker, and so long as the **TALx86** type system is correct, the linked result is guaranteed to be safe.

The second mode is that an untyped binary can be provided, along with an export interface providing types for the exported labels. In this case, the **TALx86** linker can check that the import assumptions of the client module match the exported interface of the untyped binary, but cannot check that the untyped binary actually provides the given interface. Consequently, the safety of the linked program is contingent on the untyped binary actually matching its exported interface.

The third linking option is to simply link an untyped binary against the typed client code, with no interface checking whatsoever. In this case the safety of the linked program is contingent on the untyped binary matching the declared import assumption of the typed binary.

The difference between the second and the third case is important, but subtle. The key element is that in the third case, there are no restrictions on what functions the typed binary can call, and what types it assigns to them. Consequently, it is trivial to subvert the type system simply by providing a typed binary which imports functions at incorrect types. In the second case, this is not possible, since the host provides an explicit interface against which the typed client is to be linked. So long as the host provides the correct types for the exported functions (and so long as those functions are implemented correctly and safely), the client is unable to subvert the type system since it is constrained to respect the exported interface.

In the **TILT** typed backend, a combination of the first two approaches is used to implement the runtime layer upon which programs run. Some parts of the runtime functionality are implemented in

TALx86 and hence can be checked directly. Other parts are implemented in C, with an explicitly provided interface. The runtime layer as currently implemented is very minimal and does not provide any support for system calls or any useful library functions. A version of the Standard ML Basis Library was modified slightly to provide this sort of operating system functionality. Currently, all of this code is linked using the third approach. This is satisfactory for using the compiler as a general x86 backend, but for use as part of a certifying compilation infrastructure it would be necessary to define a more limited interface to host resources, and to define explicit export interfaces against which client code could be checked.

9.6 Engineering

While extensive engineering of the **TILT** certifying backend is beyond the scope of this thesis, a certain amount of engineering work was required simply to demonstrate the feasibility of certified compilation in this setting. This section describes some of the techniques used.

9.6.1 Type representation

Unlike the **MIL**, the **LIL** lacks an internal definition mechanism for types. In the absence of any sharing mechanism whatsoever, types are observed to grow infeasibly large. The solution to this problem used in the **LIL** is to enable sharing via *hash-consing*.

Hash-consing is a popular technique for controlling representation size which works by looking up each newly allocated type into a table containing all previously allocated types. Whenever a previously allocated version is found in the table, the two copies may be represented as pointers to the same data-structure. Hashing is used to make this reasonably efficient: each node in the parse tree of a type contains the hash value of the type, which can then be used to produce hash values for parent nodes. The table of allocated nodes can then be implemented as a hash table with reasonably fast lookup times.

Hash-consing ensures that all syntactically equal types can be shared. However, types which are alpha-variants of each other will not be shared in this scheme (except accidentally) since the hash codes of different variables will usually be different. This problem can be completely avoided by switching to a representation in which variables are represented via deBruijn indices, wherein a variable use is implemented as a count of the number of binding sites between the variable use and the variable binding site. When variables are implemented in this fashion, all types which are alpha-equal will also be syntactically equal, and hence can be shared via hash-consing.

However, deBruijn indices are notoriously difficult to work with, and give no guarantee of improved sharing. Moreover, all of the existing **TILT** infrastructure was built around using variables. Consequently, I chose to use a standard representation of variables. In practice, this seems to work well so long as care is taken to avoid introducing un-necessary alpha-variants.

In addition to sharing syntactically equal types in memory, it turns out to be very important to share kinds as well. In almost all programs, the number of syntactically distinct kind nodes is smaller than one-hundred. In the absence of hash-consing for kinds however, the number of actual kind-nodes in memory is commonly three or four orders of magnitude larger than this. In general throughout this section, all of the techniques discussed in the context of sharing types should be considered to apply to kind information as well.

9.6.2 Traversing programs

Sharing types and kinds in memory is crucial in order to reduce internal representations of programs to a manageable size: that is, to reduce the space usage of the compiler. However, it is almost equally important to ensure that in addition to avoiding redundant representation of types, the compiler also avoid redundant computation on types. Hash-consing permits type parse-trees to be represented more compactly as directed acyclic graphs (DAGs). Compiler passes that rewrite types represented in this form must also be careful to traverse these programs as DAGs.

Traversing programs in DAG form is implemented by providing an abstract lookup table in which hash-consed type nodes can be inserted. Node lookup tables are used in compiler passes to keep sets of nodes already traversed: or more generally, to keep mappings from previously traversed nodes to the traversal result (or other useful information). Before attempting to rewrite a node, the compiler pass checks the table for previously computed result. If no such result is found, the result is computed and entered into the table for future use. This allows passes to avoid re-traversing types unnecessarily without impacting the structure of algorithms in any significant way.

Compiler passes that produce different results based on contextual information are slightly more complicated to implement in this way, since it is no longer necessarily the case that syntactically equal types rewrite to the same results. This situation only really arose in the **LIL** typechecker which is discussed extensively below.

9.6.3 Fast substitution

A very common operation in the compiler is the substitution of types for free type variables in other types (as well as kinds for kind variables). In practice, substitution often proves to be an efficiency bottleneck in the **TILT** compiler. In order to make substitution faster, the **LIL** backend maintains sets containing free type and kind variables at each type and kind node. This allows substitutions for variables to only traverse those nodes which actually contain free occurrences of the substituted variables.

Another important benefit of maintaining sets of free variables is that it allows the substitution code to avoid introducing unnecessary alpha-variants of types. Capture avoiding substitution must in general alpha-vary when crossing a type variable binding, since the bound variable may occur free in the type being substituted. This alpha-variation may cause an entirely new copy of the type to be created, even if the substitution only affects small sub-trees of the type. This can be avoided in the case where the bound variable does not occur free in the type being substituted. Maintaining sets of free variables makes this check efficient.

9.6.4 Fast alpha-equivalence

The hash-consed implementation allows for a fast syntactic equality check, since any two syntactically equal types will be represented by the same type node and hence can be checked for equality in constant time. In cases where the equality check fails, a normal equivalence check must be performed, since alpha-variants of the same type will not in general be represented by the same node.

9.6.5 Weak Head-Normalization

Weak head-normalization is a common operation on types, since it is the first step in checking type equality. In addition to representing types as a directed acyclic graph, it is important to normalize them as graphs as well to avoid repeated reductions of the same terms. This is implemented in the **LIL** back end by associating with every hash-consed type node an optional weak head-normal form. The reduction algorithm then checks for a previously computed normal form before reducing the term and avoids re-computing it if present. If the normal form is not present it must be computed: but the new normal form can be stored with the node for use by subsequent reductions of the same shared node. This technique (among others) has been previously described and empirically characterized by Shao et al [SLM98].

9.6.6 Engineering the LIL type checker

A valuable benefit of using a typed internal language in a compiler is that it allows the output of each compiler phase to be checked to ensure that nothing untoward has happened. Many, if not most compiler bugs will cause a compiler to generate type incorrect code at some point. An internal type checker therefore provides the valuable service of catching and pinpointing compiler bugs before runtime, when they are likely to be much harder to track down.

In order to reap these benefits however, the compiler's internal type checker must be efficient enough to quickly check anything the compiler produces. However, the simple technique described in the previous section for traversing types as directed acyclic graphs cannot be applied naively to a type checker. The reason is that type checking is a context sensitive operation: syntactically equal types may be well-formed in one typing context and not in another (for example because of the presence of free variables referring to different binding sites). In the **TILT** compiler, maintaining the simplicity of the typechecker implementation was an important goal since the correctness of the typechecker is both important and hard to test (at least in terms of soundness). However, the straightforward approach of completely ignoring sharing of types and kinds proved completely impractical. Typechecking small programs easily exhausted the resources of a Pentium 4 machine with one gigabyte of RAM. As a first step, it was necessary to find a way to avoid redundant type checking.

Exploiting physical sharing in the type checker

While the nature of typechecking is context sensitive, it is by far the most common case that checking syntactically equal terms will produce the same result (that is, that the terms will be well-typed at the same kind). Moreover, it is most frequently the case that two syntactically equal terms are judged to be well-typed for the same reasons: that is that the contexts in which they are checked are (essentially) equal. A first cut at improving the performance of the typechecker then is to keep a table mapping types to contexts in which they have already been judged equivalent (as well as the kind classifying them). Before checking a type, the compiler first checks a lookup table to see if the type under consideration has been checked before in an equivalent context (and with an equivalent kind).

Strict pointwise equality on typing contexts could in principle be somewhat expensive to compute, and moreover is vastly more restrictive than is necessary. It is frequently the case that two syntactically equal terms will be checked in two contexts which differ, but only in irrelevant ways (for example, one context may be a well-formed extension of the other). An improvement on the

simple technique described in the previous paragraph is to observe that so long as the two context are both well-formed and are pointwise equal on the free variables of the type under consideration, it is unnecessary to re-check the term in the new context.

These techniques suffice to permit traversal of types as DAGs in the **LIL** typechecker with a fairly small overhead.

Context substitutions

The technique described in the previous section avoids most redundant checking of types, but is not sufficient in and of itself to make checking large programs practical in **TILT**. The **LIL** presents an unusual engineering challenge because of the nature of its type analysis mechanism. Several of the **LIL** type checking rules induce substitutions into typing contexts as part of the type refinement mechanism. For example, in one of the rules for type checking type pair refinements, an explicit type pair is substituted for a variable in the typing context.

$$\frac{\Psi; \Delta, \beta:\kappa_1, \gamma:\kappa_2, \Delta'; \Gamma[\langle\beta, \gamma\rangle/\alpha] \vdash e[\langle\beta, \gamma\rangle/\alpha] : \tau[\langle\beta, \gamma\rangle/\alpha] \text{ exp} \quad \Delta, \alpha:\kappa_1 \times \kappa_2, \Delta' \vdash c \equiv \alpha : \kappa_1 \times \kappa_2}{\Psi; \Delta, \alpha:\kappa_1 \times \kappa_2, \Delta'; \Gamma \vdash \text{let } \langle\beta, \gamma\rangle = c \text{ in } e : \tau \text{ exp}}$$

This context substitution operation ($\Gamma[\langle\beta, \gamma\rangle/\alpha]$) turns out to be quite expensive, since implemented naively it involves duplicating the entire typing context after every type refinement. The fast substitution cut off afforded by the free variable sets on type nodes means that types in the context are not traversed unnecessarily, but in the absence of a very specialized data structure, the container implementing the context must be copied in entirety. Since type refinement operations are quite frequent in many programs, this rapidly becomes unmanageable.

There are numerous approaches that one could take to making this more efficient. The approach taken in the **LIL** backend is based primarily on two observations. Firstly, many elements of a typing context will not be affected by the context substitutions at all. Of those that are, many of them will not actually be referenced in any given term: it is often the case that the body of a type refinement construct will have at most two or three free variables. Any effort spent traversing unreferenced context elements is wasted.

Secondly, it is common to encounter several type refinement operations consecutively. Implemented naively, this means that multiple passes over the typing context may be made with no intervening variable lookups.

The **LIL** implementation of context substitution therefore attempts to do two things. In order to take advantage of the second property, substitutions are aggregated and carried out lazily. Substituting into a context simply records the substitution without carrying it out. Multiple substitutions can consequently be aggregated together. In order to take advantage of the first property, these aggregated substitutions are then only carried out on the result of a variable lookup. Full context substitutions are never performed.

There are a number of subtle points involved in this optimization. It is clearly not valid to simply maintain a substitution mapping variables to types and apply it to every type that is looked up in the context, since different substitutions will have been in place at different binding sites. In the next section, I give an informal presentation of an extension to the static semantics of the **LIL** that captures this optimization, and then briefly discuss the differences with the actual implementation.

Fast context substitution

As a first cut, I extend the notion of a typing context Γ with substitution nodes. For brevity, I elide 64 bit context entries - their addition is trivial.

$$\Gamma ::= \bullet \mid \Gamma, x:\tau \mid \alpha \rightarrow c$$

In order to ensure that substitutions are carried out correctly, I replace the implicit use of context re-ordering in the **LIL** static semantics with explicit re-ordering rules that carry out substitution nodes as needed.

$$\frac{\Psi; \Delta; \Gamma, y:\tau_y, x:\tau_x, \Gamma' \vdash sv : \tau}{\Psi; \Delta; \Gamma, x:\tau_x, y:\tau_y, \Gamma' \vdash sv : \tau} \quad \frac{\Psi; \Delta; \Gamma, \alpha \rightarrow c, x:\tau_x[c/\alpha], \Gamma' \vdash sv : \tau}{\Psi; \Delta; \Gamma, x:\tau_x, \alpha \rightarrow c, \Gamma' \vdash sv : \tau}$$

The variable lookup rule is then restricted in the usual way, requiring that variables be re-ordered to the end of the context for lookup.

$$\frac{\Delta \vdash \Gamma, x:\tau \text{ ok} \quad \vdash \Psi \text{ ok}}{\Psi; \Delta; \Gamma, x:\tau \vdash x : \tau}$$

Context substitutions can then be implemented in terms of these substitution nodes.

$$\frac{\begin{array}{c} \Psi; \Delta, \beta:\kappa_1, \gamma:\kappa_2, \Delta'; \Gamma, \alpha \rightarrow \langle \beta, \gamma \rangle \vdash e[\langle \beta, \gamma \rangle / \alpha] : \tau[\langle \beta, \gamma \rangle / \alpha] \text{ exp} \\ \Delta, \alpha:\kappa_1 \times \kappa_2, \Delta' \vdash c \equiv \alpha : \kappa_1 \times \kappa_2 \end{array}}{\Psi; \Delta, \alpha:\kappa_1 \times \kappa_2, \Delta'; \Gamma \vdash \text{let} \langle \beta, \gamma \rangle = c \text{ in } e : \tau \text{ exp}}$$

This version of the variable pair refinement rule simply adds a substitution node $(\Gamma, \alpha \rightarrow \langle \beta, \gamma \rangle)$ instead of explicitly substituting for α as in the original version $(\Gamma[\langle \beta, \gamma \rangle / \alpha])$.

This extension to typing contexts captures the notion that variable substitution need only be performed on referenced variables, and can be deferred until the point of reference. This easily extends to capture the idea of aggregating substitutions by replacing the single substitutions in the context with simultaneous substitutions for multiple variables. A simultaneous substitution Θ maps variables to types.

$$\Theta ::= \bullet \mid \alpha \rightarrow c, \Theta$$

The operation of a substitution Θ on a type c is defined by a function $c[\Theta]$ mapping types to types.

$$\begin{array}{ll} c[\bullet] & \stackrel{\text{def}}{=} c \\ \alpha[\alpha \rightarrow c, \Theta] & \stackrel{\text{def}}{=} c \\ \alpha'[\alpha \rightarrow c, \Theta] & \stackrel{\text{def}}{=} \alpha'[\Theta] \quad (\alpha \neq \alpha') \\ \dots & \stackrel{\text{def}}{=} \dots \end{array}$$

The remaining cases proceed compositionally over the structure of types exactly as with a normal substitution (taking care to avoid capture when crossing binding sites). Composition of simultaneous substitutions $\Theta \circ \Theta'$ is defined explicitly as an operation on substitutions obeying the equation $c[\Theta \circ \Theta'] = (c[\Theta])[\Theta']$.

$$\begin{array}{ll} \bullet \circ \Theta & \stackrel{\text{def}}{=} \Theta \\ (\alpha \rightarrow c, \Theta) \circ \Theta' & \stackrel{\text{def}}{=} \alpha \rightarrow (c[\Theta']), (\Theta \circ \Theta') \end{array}$$

Aggregation of substitutions in typing contexts is implemented by replacing single substitution nodes with simultaneous substitution nodes.

$$\Gamma ::= \bullet \mid \Gamma, x:\tau \mid \Gamma, \Theta$$

The re-ordering rule for substitution nodes is modified appropriately.

$$\frac{\Psi; \Delta; \Gamma, \Theta, x:(\tau_x[\Theta]), \Gamma' \vdash sv : \tau}{\Psi; \Delta; \Gamma, x:\tau_x, \Theta, \Gamma' \vdash sv : \tau}$$

A new structural rule is introduced permitting adjacent substitutions to be aggregated using substitution composition.

$$\frac{\Psi; \Delta; \Gamma, \Theta_1 \circ \Theta_2, \Gamma' \vdash sv : \tau}{\Psi; \Delta; \Gamma, \Theta_1, \Theta_2, \Gamma' \vdash sv : \tau}$$

Fast context substitution in practice

The actual implementation used in the **LIL** backend differs from this presentation slightly. While the formal system implements typing contexts as lists, the implementation uses a splay tree implementation for efficient variable lookup. In order to avoid the need to design and implement a custom balanced tree implementation supporting explicit substitution nodes, I chose instead to maintain the substitution nodes in an auxiliary data-structure, allowing the core splay tree data structure to remain unchanged.

Notionally, the implemented version can be derived from the system described above in the following manner. Number the substitution nodes in a typing context, starting from the “leftmost” substitution node. With every variable/type pair in the context, associate the number of the first substitution to its right (that is, the first substitution that applies to it). Finally, remove the substitution nodes into a separate data structure associating each substitution with its index.

Upon looking up a variable, apply the composition of all substitutions with indices greater than or equal to that associated with the variable (that is, all substitutions that were to the right of the variable in the original context). This corresponds to the sequence of substitution that would have been performed by the sequence of structural re-ordering steps required to move the variable to the end of the context. As an optimization, this composition of substitutions may be computed eagerly.

Upon inserting a variable into the context, associate with it the first index larger than the largest currently used index.

Upon substituting into the context, associate the new substitution with the first index larger than the largest currently used index. The composition of this substitution with previous substitutions may be eagerly computed if it is more efficient to do so. If no variables have been inserted into the context since the last context substitution was performed, it is also possible to simply replace the substitution associated with the largest index with the composition of the old substitution and the new substitution, avoiding the addition of a new substitution node. Intuitively, this corresponds to eagerly applying the context composition structural rule.

This algorithm provides a straightforward implementation of explicit substitution nodes in a typing context without requiring any modification to the underlying data structure. There are additional improvements possible to the algorithm. For example, in the static semantics as presented above, it is possible to (non-deterministically) choose to apply the context re-ordering rules for a variable at the first point in a derivation past which more than one immediate sub-derivation contains a lookup of that variable. In this way, the application of any given substitution to the type associated with a variable can be guaranteed to be performed at most once. The algorithm as implemented does not implement this: the substitution is re-applied to the original type upon every variable look up. In practice this does not seem to be a problem. However, if after further engineering this became a bottleneck, it would be straightforward to memoize the application of substitutions to ensure that no unnecessary traversals of types is done.

9.7 Compilation units

The definition of Standard ML [MTHM97] does not define a notion of separate compilation. In order to remedy this, the **TILT** compiler defines a notion of interface that generalizes Standard ML signatures, and uses these interfaces to mediate between separately compiled source units. Any unit of source code may be compiled in complete isolation from any other unit upon which it relies, so long as suitable interfaces are provided.

The certifying backend described in this dissertation implements the full **TILT** separate compilation system. This is done by viewing each compilation unit as a functor which maps its imported units to its exports. In this way, the linking process is modeled as function application, avoiding the need for a complicated type theory to track the initialization of globally visible locations.

More concretely, a **LIL** compilation unit may be thought of as a pair consisting of a type component and a term component. The type component is a type function mapping imported types to exported types. Similarly, the term component is a function mapping imported terms to exported terms. In actuality, the type portion is implemented not as a function, but instead by using the **TALx86** linker directly: each unit lists imported types with their kinds as explicit imports instead of parameterizing the components over them.

A **LIL** interface classifying such a compilation unit is a translucent sum: the first component of which is a kind classifying the type element of the compilation unit, and the second component of which is a type classifying the term element of the compilation unit. The pair is dependent in the usual manner: the term classifier portion may refer to the label of the type classifier portion.

The **TALx86** linker as defined and implemented by Glew et al. [GM99] does not in fact support translucent sums, and so it was necessary to extend the **TALx86** implementation with an alternative (and simpler) notion of typed linking which implements the standard translucent sum matching rules. This was the only significant change needed in the **TALx86** infrastructure.

9.7.1 Compiler generated files

The **LIL** backend emits four files for every compilation unit. The first is a typed assembly file named **asm.tal** containing the decorated assembly code for the unit: this corresponds to a **LIL** compilation unit as described above. The interface of a compilation unit (whether explicit or derived) is compiled to a term export file named **asm.e.tali** classifying the exported term portion of the unit, and a type export file named **tali** classifying the exported type portion. Finally, an additional file is emitted containing the signatures of any external C functions used by the unit.

The **TALx86** assembler type checks and assembles the typed assembly file (**asm.tal**) to produce two object files: a standard object file (**obj.o**) and a type object file (**obj.to**) containing the type annotations necessary for typechecking the object file. In the process, it verifies that the provided code matches the interfaces specified in the **asm.e.tali** and **tali** files.

After compiling each individual unit in a program, **TILT** must also produce an additional unit to serve as the “link” unit for the entire program: that is, a unit which applies each compilation unit function to its imports to produce its exported result. This exported result is then passed on to subsequent units which import it. This “link” unit implements the ML level linking. The **TALx86** linker is used to typecheck and construct the final executable program, linking together the “link” unit with each of the compilation units providing the individual unit initialization functions.

9.8 Measurements

The goal of this dissertation is to demonstrate the feasibility of certifying compilation in a type analysis framework. The bulk of the dissertation is concerned with simply developing and describing the framework necessary for this process. This provides the first argument for feasibility: that it is possible at all. This section is intended to provide some evidence that not only is it possible, but in fact that it is practical. In the following sections, I present some empirical measurements of the performance of certifying **TILT**. In particular, I present measurements quantifying the size of the generated type annotations on the emitted code, and measurements of the run time of various benchmark programs with comparisons to other compilers.

It is important to re-iterate here that engineering the compiler to improve its behavior along either of these axes is beyond the scope of this thesis. In both cases, the engineering goal was simply to develop a working system, and to measure the result. No significant effort was spent on improving the system based on these measurements, and inspection of the emitted code suggests that substantial improvements along both of these axes could be implemented without running into any fundamental limitations of the framework.

9.8.1 Benchmarks

The benchmark programs measured in the following sections were selected to be representative of a number of different sorts of programs, ranging in size from 24 to more than 2000 lines of code. Each benchmark was compiled separately to an object file, and then subsequently linked into a testing harness from which it was run. The benchmarks were also linked against the full Standard ML Basis Library which provides many of the basic data types for Standard ML, along with access to system I/O facilities. Several of the benchmarks use additional data structure implementations from the SML/NJ library. Figure 9.3 lists the benchmark programs used in this section, along with their sizes (in lines of code). Note that the size given is for each single benchmark file only: code from other compilation units (such as libraries and the test harness) are not included in this count.

9.8.2 Type size

As discussed in chapter 1.4, one of the most commonly cited applications for certifying compilation is as a security mechanism for mobile code. Certified code that is downloaded to be run from potentially un-trusted sources (or over un-trusted communication channels) can be checked for any violations of the safety properties implied by the type system. An important property of a such

Benchmark name	Description	Lines
Taku	uncurried function calls	24
Take	curried function calls	25
Fib	fib, fact with default Int	38
Fib32	fib, fact with 32 bit ints	39
PI	approximation of pi (fp)	28
Msort	Merge sort (lists)	48
ISort	Insertion sort (lists)	50
Quick	Quicksort (version 1)	130
Quick2	Quicksort (version 2)	152
Life	Game of life (lists)	205
FFT	Fast fourier transform	271
PQueens	P queens problem (arrays)	292
Frank	Small theorem prover	473
Leroy	Knuth bendix completion (exceptions)	537
BarnesHut	N-body simulation	684
Simple	Spherical fluid dynamics	860
Tyan	Grobner basis calculation	896
Boyer	Theorem proving	959
Lexgen	Lexical-analyzer generator	1178
Pia	Perspective inversion algorithm	2074

Figure 9.3: Benchmarks




Figure 9.4: Size breakdown of individual benchmark units, in kilobytes.

a certification system is that the certificate size should not be unduly large so that the additional security provided does not come at a prohibitive cost in terms of bandwidth needed. In this section, I provide measurements quantifying the size overhead of the certificates generated by **TILT**.

In section 9.7.1 I described the files generated by the **TILT** compiler and the **TALx86** assembler. At the assembly code level, there are notionally two elements of a **TILT** compiled binary: a typed assembly file containing the actual assembly code annotated with type information (**asm.tal**), and some additional files describing the type of the exported interface provided by the binary (the **asm_e.tali**, **asm_i.tali**, and **tali** files). The typed assembly file refers to the exported interface files of any units which it uses, and any units which use it will in turn refer to its exported interface files. Interface files mediate between compilation units, while the type annotations within an assembly file allow individual units to be checked.

The typed assembly file is further split by the assembler into a standard untyped object file (**obj.o**) and a type annotation file (**obj.to**) that contains sufficient information to reconstruct the annotations on the untyped object file.

A reasonable measurement of the certificate size for a compilation unit after assembly then is the sum of the sizes of its export interface files (the **asm_e.tali**, **asm_i.tali**, and **tali** files) and the type annotation file generated by the assembler (**obj.to**), since these represent the incremental contribution of each compilation unit to the overall certificate size of the entire compilation system. Note that all of the interface information is present only to support separate compilation. Once linked together, almost all of the interface information can be discarded: in this sense this measurement is an upper bound.




Figure 9.5: Size breakdown of individual benchmark units, relative.

9.8.3 Benchmark unit sizes

In this section, I give some size measurements for the **TILT** benchmark suite. Each of the units under consideration is a single compilation unit compiled from a single source file to a single object file. These object files are linked together with libraries, a test harness and the link unit to produce the final executable. Note that only the benchmark units themselves are included in this section. Measurements of the libraries and the link unit are discussed in the next section.

Figure 9.4 gives the absolute size in kilobytes of each benchmark. The columns are sorted by increasing size of the generated (untyped) object file. The segments of each column indicate the contribution of each of the different files making up the compilation unit. The size overhead of the safety certificate is everything other than the object file itself (the bottom segment). Note that in all cases, the contribution of the **asm_i.tali** file (containing the types of imported C functions) is too small to be visible.

Figure 9.5 provides a different view of this same data: in this figure each segment of each column indicates a percentage of the total size of the unit contributed by a particular source. For very small programs, the amount of type information dominates the object code size. For larger programs, the percentage of the total space usage devoted to the certificate decreases. This reflects a certain amount of fixed overhead required for each program: basic types such as exception handlers, types for printing routines, etc. As programs get larger, the cost of this fixed overhead goes down. For the largest of the benchmarks, the certificate occupies roughly sixty percent of the total size. Summed over all of the benchmarks, the certificate information occupies roughly eighty percent of the total space, indicating roughly a factor of five space penalty for certification.

Export interfaces

The topmost visible segment in both graphs represents the contribution to the total size of the exported interface of each of the object files. Generally speaking, this component is of comparable size to the type information in the object file itself. As would be expected, for bigger programs the percentage overhead of the exported interface decreases somewhat, since the overhead is amortized over a larger amount of actual object code. In fact, this overhead is probably almost entirely eliminable, for three reasons.

First, by the nature of the compilation approach taken in the **LIL** backend, each of these export interfaces describes a single function mapping each of the object files logical imports to its logical exports. Consequently, the size of this file increases significantly with the number of imported units. In addition, this implies that almost all of the type decorations in an export interface file are already present in the export interface files of its ancestors. It is very likely that this redundant information could be eliminated entirely by exporting a single canonical abbreviation for the exported type of each unit along with its exported type component.

Second, because **TILT** implements separate compilation, the interface file and the actual implementation files may be produced and used independently. As a consequence, each contains its own copy of the type of the main exported initialization function (in fact, the entire export interface file consists solely of this). It is clear that in the common case where both files are produced as part of the same compilation, this redundancy could be factored out into a common definition site. This is almost certainly expressible in the **TALx86** linking system without modification.

Finally, note that these export interfaces exist solely to mediate between compilation units. Linked as a whole program, all of this interface disappears. Moreover, in many cases (such as in the benchmark suite), almost none of the entry points described by these interfaces are exported from the local group of compilation units. Even in the Standard ML Basis Library, there are many compilation units whose logical scope is entirely local to the library. A partial linking strategy wherein groups of compilation units are closed up into a single file presenting a single export interface would almost certainly eliminate a great deal of this overhead in most cases, even when whole program linking is not possible.

9.8.4 Libraries and linking

The measurements presented in figures 9.4 and 9.5 cover only the benchmark programs and the testing harness. These files must be linked against additional libraries before running: the Standard ML Basis Library, the SML/NJ Library, and a small command line argument parsing library. It is likely that in a certified compilation system, a subset of these libraries would be provided by the client, rather than as part of the certified binary. This is particularly true for the Standard ML Basis Library, which encapsulates the operating system functionality. Nonetheless, it is useful to measure the behavior of the compiler on these libraries as well as part of the overall system.

In addition to the libraries, there is one additional compilation unit that makes up part of the final executable program. As discussed previously, the compilation strategy employed in **TILT** maps each compilation unit to a single entry point implementing a function which takes as arguments the logical imports of the unit, and computes the logical exports as a result. The final step in compilation then is to create a unit which stitches together the whole program at runtime by applying each of these functions in turn. I refer to this unit as the link unit, since it implements the logical linking of the program. Note that this is distinct from the **TALx86** notion of linking,




Figure 9.6: Total sizes of compilation groups (kilobytes).

which is used to resolve the free references within the link unit back to the exported entry points of each compilation unit.

Figure 9.6 compares the absolute sizes in kilobytes of each of the major compilation groups of the benchmark suite: the three libraries (the Basis, the SML/NJ Library, and the command line argument library), the link unit, and the benchmarks themselves (including the testing harness). As before, each column is broken into segments showing the contribution to the total of each of the constituent files. Also as before, figure 9.7 presents the same data as a percentage of the total for each group.

The Basis Library

There are several interesting points to note here. Firstly, it is clear that the Basis library (and to a lesser extent the SML/NJ library) have more certificate overhead than the benchmark suite. (This is also true of the argument library, but this is most likely because of its very small size.) While the Basis and the benchmarks viewed as a whole have similar amounts of actual object data, the certificate size for the Basis is larger by roughly a factor of four. While I have not investigated this phenomenon in detail, I conjecture that it is likely that this is partially a result of the particular architecture of the Basis library. Within the Basis, there are numerous units which consist solely of a few (or even one) functor applications, or which contain structures which aggregate numerous other structures together as sub-structures. Such units produce almost no object code, but have quite large types.

Further evidence for this can be seen in figure 9.8, which presents the contributions of each of




Figure 9.7: Total sizes of compilation groups (relative).




Figure 9.8: Breakdown of certificate overhead across individual Basis units (kilobytes)

the constituent elements for each compilation unit in the Basis library. For clarity given the large number of files, the graph is presented as a continuous area chart. As with the previous graphs, the units are sorted by object file size. For the most part, the certificate size (the sum of the top components) increases fairly smoothly with object file size. However, there are a number of very substantial spikes at points in the graph where very small files generate certificates comparable to those of the largest files in the library. All of these spikes that I have examined in detail correspond to the sort of aggregating files discussed above. Since for such files, much of the type information will have already been written out in the export files of the logically imported compilation units, it is likely that much of this overhead could be eliminated, even in a separate compilation setting (since the type decorations needed to describe the logical imports must be present in the interface file, which is in turn needed for separate compilation). However, in the current framework these small units with large types increase the aggregate certificate overhead substantially.

The link unit

A second interesting observation about figures 9.6 and 9.7 is that the size of the link unit is entirely dominated by the certificate size. The actual object code for the link unit makes up less than one percent of the total size, and the total size itself is substantial in absolute terms (significantly larger than all of the benchmarks put together). In principle this is actually somewhat understandable. The link unit refers to every compilation unit in the entire program, including all of the libraries. It therefore must be able to refer to the type of the exported entry point (and its result) for each compilation unit. In some sense then, it is not surprising that the size of the certificate for the link unit should be comparable to the sum of the sizes of the export interfaces for all of the units in the program.

However, all of the type information needed to describe the entry points of the compilation units must be present in the export interfaces of the compilation units themselves. And since these export interfaces are needed by the link unit, there is no reason that the link unit needs to contain its own copy of these types. In fact, the typed assembly code produced by the **LIL** backend takes advantage of this property by using the abbreviation mechanism provided by the **TALx86** system. The link unit code generated by the **LIL** backend simply refers to the type of a compilation unit by using a canonical abbreviation name that is given a definition by the export interface of the unit. The assembly code produced by the **LIL** backend for the link unit is consequently quite compact. Surprisingly, the **TALx86** assembler seems unable to preserve this compact representation.

Figure 9.9 compares the aggregate sizes of the compilation groups before and after assembly. For each compilation group, the first column represents the total size of the assembly (**asm.tal**) files for the unit, and the second column represents the sum of the object and typed object files (**obj.o** and **obj.to**). Note that the files from these two columns represent the same information. The object and type object files (**obj.o** and **obj.to**) are produced by assembly the assembly file (**asm.tal**), and can be subsequently dis-assembled to re-produce the original assembly representation.

In general, the **TALx86** assembler does a very good job of managing certificate sizes. In all cases except the link unit, the assembled version is noticeably smaller than the original assembly files. This is universally true for individual files within the compilation groups, as well as in aggregate. The only exception to this is the link unit: in this case the assembled version is almost forty times the size of the original assembly file (even though notionally they represent the same information).

While I have not investigated this in detail, I believe that this is almost certainly because of a failure of the **TALx86** assembler to preserve the sharing from the original assembly file. I




Figure 9.9: Assembly file size vs assembled output size (kilobytes)

conjecture that the references to abbreviated types from imported units in the original assembly file are being expanded out in the type object file. I do not believe that there is anything inherent to the structure of the link unit that prevents this sharing from being maintained: it should be possible to engineer an assembler to preserve this information.

Scalability

The data in figure 9.5 suggest strongly that the certificate overhead scales well as the size of compilation units increase. For larger units, the overall percentage overhead is substantially smaller than the overhead for smaller units. In order to demonstrate the scalability of the system further, I took additional measurements on two very large programs.

The first of these large programs consisted of a subset of the benchmarks along with all of the library code upon which they rely, concatenated into a single file. In addition to the SML/NJ libraries, this included a large portion of the Standard Basis Library as well. Unfortunately, because of a limitation with the **TILT** foreign function interface, it was not possible to concatenate the entire Standard Basis Library into a single valid Standard ML unit. Therefore, certain of the benchmarks (such as those performing file i/o) could not be included. The resulting file consisted of 8332 lines of code.

The second of these large programs consisted of approximately half of the **TILT** compiler (by lines of code) concatenated into a single file and compiled as a whole program. Note that in this case, all library code (including the Standard Basis Library) was compiled separately. The resulting source file consisted of 32555 lines of code!

The absolute size (in kilobytes) of the generated type and object files is given in figure 9.10.




Figure 9.10: Absolute sizes of large units (kilobytes).




Figure 9.11: Relative sizes of large units.

The same data is plotted in figure 9.11 as a percentage of the total, with comparisons to the totals for the separately compiled libraries and benchmarks as shown in figure 9.7. For both of these large programs, the certificate overhead is substantially smaller than the aggregate totals for the other compilation groups, representing approximately 60% of the total size. This is comparable to the overhead for the largest of the benchmark units described in figure 9.5. These two data points combined with figure 9.5 strongly suggest that while there is a noticeable initial overhead for certification, the certificate size remains a relatively constant fraction of the overall size for a large range of program sizes. That is to say, empirically speaking the certificate size seems to grow linearly with the program size.

9.8.5 Run time

The **LIL** backend is not designed to be an optimizing backend. Code generation and register allocation are both done in a single linear pass, and no optimizations are done after code generation. Nonetheless, it is useful to measure the runtime performance of the compiled code. It is important that certification not overly restrict the compilation process in such a way as to make efficient code generation impossible. In this section, I argue that even a non-optimizing certifying backend can produce reasonably efficient code.

In order to make this argument, I present comparisons with two other compilers: the Standard ML of New Jersey compiler and the MLton compiler. These comparisons are designed to give some indication of the relative performance of the certifying **TILT** backend with respect to the state of the art in Standard ML compilers. It is important to note however that each of these compilers is

fundamentally addressing a very different compilation problem from each of the others. For this reason, this comparison is only really meaningful at the most general high level.

The MLton compiler is designed to produce very efficient code. In order to do this, it compiles only complete programs. The Standard ML of NJ compiler is designed to be used as an interactive system. As such, it supports incremental (but not separate) compilation. It performs significant optimization as well, but is limited somewhat by the needs of incremental compilation and an interactive frontend. Both of these compilers implement precise garbage collection.

The **TILT** compiler on the other hand implements true separate compilation. This greatly limits its ability to optimize code that crosses compilation units. The **TALx86** backend also uses a conservative garbage collector, with a malloc based allocator.




Figure 9.12: Normalized runtime (MLton = 1)

Figure 9.12 shows the relative performance of the compiled versions of the **TILT** benchmarks as compiled by each of these three compilers. The results are normalized to the results for the MLton compiler, which is generally the fastest. In order to get some sense of the overhead of separate compilation, some of the benchmarks were compiled by **TILT** as whole programs as well (e.g. with all of the library code and the Standard Basis Library included into one file). This data is presented in the second column labelled **TILT(Whole)**. Because of limitations in the **TILT** foreign function interface however, none of the benchmarks requiring i/o could be compiled as whole programs and so certain of the data points are missing.

Overall, the certifying backend ranges from slightly faster than MLton to almost sixteen times slower, and from approximately twice as fast as SML/NJ to seven times slower. For very small benchmarks, such as the arithmetic benchmarks and the takc/taku benchmarks, **TILT** does quite well. For small benchmarks such as these, there is no penalty for the inability to optimize across compilation units, and the relatively simple register allocation in **TILT** is sufficient. The one

exception to this is the pi benchmark, on which **TILT** does quite poorly. This is almost certainly because of the lack of floating point register allocation in **TILT**, since this benchmark is dominated by floating point operations.

The performance of the larger benchmarks is much more varied. In some cases (tyan, lexgen, pqueens) **TILT** is within a factor of two of one or both of the other compilers. In the worst case (leroy) **TILT** is almost sixteen times slower than MLton, and five times slower than SML/NJ. I conjecture that part of this is likely due to separate compilation: all of the larger benchmarks cross compilation unit boundaries a fair bit, via library calls. In addition, several of the benchmarks involve floating point computation, and several of them are fairly allocation intensive. In both of these areas, **TILT** is likely to suffer: from the lack of floating point register allocation in the first case, and from a more expensive memory allocation strategy in the second case.

The whole program version of each **TILT** compiled benchmark is noticeably faster than the separately compiled version: in one particular case (PQueens), faster by a factor of approximately six. This is despite the fact that all of the **TILT** optimizations assume a separate compilation setting even when given whole programs.

Chapter 10

Conclusions and future work

10.1 Summary

In this dissertation I have shown that certified compilation is possible for full Standard ML, even in the presence of type analysis based optimizations.

10.1.1 Theory

To provide a theoretical foundation for this, I defined a series of translations mapping a polymorphically typed lambda calculus extended with type analysis primitives to a typed assembly language. These translations serve to make type analysis explicit in the intermediate representation; to translate uses of functions to uses of closures; and to make control flow and machine state explicit in an abstract assembly language.

I proved each of these translations sound in the sense that each translation maps well-typed terms to well-typed terms. In order to avoid overly constraining the implementation, I also introduced a novel approach for dealing with register allocation, allowing the typing assumptions required for the proof of soundness to be separated from the semantic behavior of the register allocator. In this way, the translation to assembly code remains parametric over the choice of register allocation algorithm, so long as the algorithm chosen satisfies certain minimal typing restrictions.

10.1.2 Practice

In order to validate the practicality of my approach, I implemented a certifying backend for the **TILT** compiler using the formal translation as a guide.

The **TILT** middle internal language corresponds closely to the polymorphically typed lambda calculus used as a starting point for the formal compilation process described in this dissertation. For each of the compiled passes described formally in this dissertation, I implemented a corresponding compiler pass in the **TILT** compiler. In addition, I implemented a one pass optimizer to perform simple optimizations to take advantage of the additional program structure exposed by the new translations. The final target of this backend is a slightly modified version of the **TALx86** framework. The final translation to the **TALx86** language follows closely the format of the formal translation in this dissertation, despite the significant syntactic differences between the idealized typed assembly language used in the formal presentation and the **TALx86** language.

Separate compilation of full Standard ML is supported by this new backend, and a large number of programs have been successfully compiled including the entire Standard ML Basis Library and the Standard ML of New Jersey library. In addition I compiled the **TILT** benchmark suite using this typed backend, and measured the results of compilation both in terms of performance as compared to other Standard ML compilers, and in terms of certificate size overhead in the compiled binaries.

10.1.3 Conclusions

Certified compilation is possible even for languages as rich as Standard ML, even in the presence of complex type based optimizations. Proving the soundness of typed compiler translations in this setting is feasible. However, there is a trade off between the closeness of the formal model to the implementation and the complexity of the proof process. In this dissertation I attempted to keep a very close correspondence between the formal model and the implementation.

In order to make this more feasible I introduced new techniques for factoring out some inessential choices of the implementation, such as the particular choice of mappings of variables to registers and the method by which this mapping is arrived at. While this particular technique should scale to more complex optimizing code generation, it is likely that maintaining the close correspondence between the formal model and the implementation would become more difficult if more complex code generation and optimization techniques were performed.

The problem of managing certificate size is shown here to be manageable. With no major engineering or tuning of the new **TILT** back end, certificate size overhead was shown to be in general quite reasonable. While the overhead for supporting separate compilation makes up almost half of the aggregate certificate size for the measured programs, simple analysis of the structure of compilation units shows that much of this overhead could be eliminated with a more sophisticated mechanism for sharing type abbreviations across interfaces. Numerous other opportunities exist for eliminating redundant type information, both between units and within units.

In the one case where the certificate overhead was observed to be drastically larger than expected (the link unit discussed in the previous section), this was shown to be due to a failure of the **TALx86** assembler (which otherwise performed admirably) rather than a structural failure of the compiler.

10.1.4 Compiler engineering

Maintaining type information on the compiler intermediate forms imposes an additional burden on a compiler writer, just as programming in a typed language imposes an additional burden on a programmer. As is the case with type safe programming languages however, it is becoming increasingly clear that the engineering benefits of typed compilation substantially outweigh the costs. A type preserving compiler provides a form of automatic self-checking that is extremely valuable to the implementer of the compiler.

Compiler bugs are notoriously difficult to track down and fix, since they often exhibit themselves only as second-order effects. That is, the compiler itself appears to run correctly: it is only in the running of the generated code that incorrect behavior appears. To make matters worse, the problem in the generated code may very well manifest itself not at the incorrect program point, but at some arbitrary later point in the program's run (e.g. because of memory or stack corruption). Matching an incorrect behavior of a generated program to the bug in the compiler that caused it often requires an extensive process of analysis and deduction. By and large, the state of the art in

untyped compiler debugging generally relies on carefully stepping through optimized (and hence often obfuscated) code in a debugger.

The great benefit of a type-preserving compiler is that it moves the point of discovery for compiler bugs out of the generated code and into the compiler itself, for a large class of bugs. In other words, many or even most compiler bugs are caught and flagged as soon as they are produced in the compiler, rather than at some arbitrary point in a future run of the generated program. Moreover, if a self-check is run between every compiler pass (at least while in development), the point at which the error is flagged indicates not only the location of the error in the intermediate code but also the particular pass of the compiler that is the most likely culprit (i.e. the pass immediately preceding the failed self-check).

Of course, type safety does not guarantee correctness. It is still possible for the compiler to generate incorrect code that nonetheless happens to be well-typed. In practice however, this seems to be relatively rare – most errors (and most of the most pernicious errors) tend to be caught. In particular, note that the entire class of errors involving memory corruption are guaranteed to be caught by the type checker! Over the course of developing the certifying backend for **TILT**, my experience was that almost all compiler bugs were caught statically. For example, while developing the register allocator (a notorious source of difficult bugs), all of the register allocation bugs that I encountered were caught by the **TALx86** typechecker.

Most of the bugs that were not caught by the typechecker arose from the more permissive nature of the **TALx86** type system as compared to the Standard ML type system. For example, **TALx86** very reasonably defines 32 bit integer arithmetic to have silent overflow semantics: the compiler is responsible for generating explicit overflow checks and raising exceptions as appropriate. If the compiler fails to emit such a check, the code is still well-typed with respect to the **TALx86** type system: however, its behavior on overflow is incorrect with respect to the semantics of Standard ML. An interesting area for future work in the area of typed assembly languages would be to provide facilities for encoding such source language constraints into the type system (without specializing the type system itself to a particular language).

10.2 Future work

This dissertation demonstrates the feasibility of performing certified compilation for Standard ML in a type analysis framework. There remain several directions in which this work could be extended in order to make this more useful and practical.

10.2.1 Optimization

The certifying backend implemented as part of this thesis is for the most part not an optimizing one. While some simple optimizations are performed on the **LIL** intermediate code, inspection of the output of the compiler suggests numerous ways in which the intermediate code could be improved (particularly after closure conversion).

Some of these improvements are as simple as extending the **LIL** language with additional constructs. A simple example of such an extension is that of heterogeneous tuples, which would allow the closure converter to avoid boxing and un-boxing 64 bit floating point numbers. Others are more complex: for example, implementing partial redundancy elimination to re-locate closure environment operations used only conditionally.

In addition to improvements in the optimization of the intermediate code, there is considerable room for improvement of the code generator itself. The register allocation technique used is quite simplistic: a more sophisticated algorithm based on graph coloring or graph fusion would produce significantly better code. No floating point register allocation is currently done at all.

The code generator itself is quite limited. No attempt is made at scheduling instructions intelligently, and the instruction selection is fairly ad-hoc. While it does a relatively good job of taking advantage of the CISC nature of the x86 instruction set, a more uniform approach would almost certainly improve the generated code.

The implementation of exceptions in the **LIL** backend is also quite inefficient, both in setting new handlers and in executing handler bodies.

Finally, I believe there is a good deal of work left to be done in tuning and improving the use of type analysis in the **TILT** compiler. No tuning has been done to adjust the parameters to the type analysis optimizations, such as the maximum width record to flatten into registers. Additionally, no work has been done to quantify the actual benefits of type analysis as currently implemented. It would be valuable to measure the effect of these optimizations in isolation, and to use this information to look for additional opportunities to take advantage of the type analysis infrastructure already in place.

10.2.2 Garbage collection

The **TALx86** infrastructure assumes the use of a conservative garbage collector. The untyped **TILT** backend makes use of type information at runtime to do precise garbage collection. An interesting topic of future research would be to replace or extend the **TALx86** infrastructure in such a way as to support precise garbage collection using the type information already kept at runtime [VC03].

10.2.3 Infrastructure improvements

The **TALx86** certification infrastructure proved impressively flexible in serving as a target for the **TILT** compiler with a minimum of modification. However, the performance of the **TALx86** type checker could be improved upon substantially. Currently, assembling and typechecking large units takes dramatically longer than the entire process of compilation within **TILT** (including numerous internal type checks). While the type system is more detailed and low level at the **TALx86** level, many of the techniques discussed in chapter 9 used to improve the performance of the **LIL** type checker are still applicable.

In addition, as discussed in section 9.8.4 there are a few important cases where the **TALx86** assembler seems to fail to preserve the physical sharing present in the original typed assembly source file which degrade its performance on certain units immensely.

10.2.4 Reducing the trusted computing base

Another concern with the **TALx86** infrastructure is that there is no proof of type safety for the language as implemented. This is of significant concern in an actual certifying compilation system, since in the absence of such a proof, even the correctness of the **TALx86** type checker does not necessarily imply the safety of the programs it certifies. And of course, there is no guarantee that the implemented type checker faithfully and soundly implements the static semantics of the type system for the language.

Attempts have recently been made to address both of these problems by providing certified code infrastructures that both implement a language with a formal proof of type safety, and which try to reduce the complexity of the checking problem as much as possible so as to minimize the amount of trusted code in the certifier [Cra03, App01]. A valuable area of future research would be to re-target the certifying **TILT** backend to such a system.

10.3 Conclusions

This dissertation has clearly shown that certifying compilation is feasible for a rich language like Standard ML, even in the presence of type analysis and other other advanced optimizations. The problem of controlling the intermediate size of programs and of controlling the certificate overhead on the generated code is almost certainly tractable. Even in the system implemented in this thesis, which made no effort to optimize for these properties beyond the minimum necessary to demonstrate feasibility, the results were easily within reach of being sufficient for a practical and usable system.

The problem of connecting formal models of compiler translations to their actual implementations remains a difficult one. As the formal models scale up to more closely model the actual languages and transformations implemented in the compiler, the syntactic overhead and complexity of proofs about the model increases significantly to the point that confidence in the correctness of the proofs must inevitably fall. New techniques for dealing with this complexity are needed, and while this dissertation attempts to address this to a certain extent, in general it remains an open problem. One promising approach to dealing with this lies in the use of logical frameworks to mechanically check proofs, or even to assist in generating them.

Certifying compilation is the natural extension of typed compilation. In this role, a certifying compiler greatly increases the ability of the compiler writer to write a correct compiler by extending the benefits of type checking to a lower level. As with all disciplines, the discipline of having always to consider the type correctness of compiler transformations is sometimes burdensome. I believe that it is also a valuable one. It should always be clear to the compiler writer why a particular transformation is safe. The type checker in a type preserving compiler enforces this discipline. It is important that the flexibility of type systems for compiler internal languages continue to be improved upon, so that the cases where the type checker must reject safe code due to its own limitations can be made increasingly rare.

Most importantly, certifying compilation provides an important tool for coming to grips with the increasing problem of providing security in a wide-open, networked universe. Delivering security along with downloaded code is an essential part of providing value to the end user. This thesis demonstrates that automatically generating such security in the form of certified code is well within our grasp.

Appendix A

MIL static semantics

Well-formed Kind

$\boxed{\vdash \kappa \text{ ok}}$

$$\frac{}{\vdash T_{32} \text{ ok}} \quad \frac{\vdash \kappa_1 \text{ ok} \quad \vdash \kappa_2 \text{ ok}}{\vdash \kappa_1 \rightarrow \kappa_2 \text{ ok}} \quad \frac{\vdash \kappa_1 \text{ ok} \quad \vdash \kappa_2 \text{ ok}}{\vdash \kappa_1 \times \kappa_2 \text{ ok}}$$

Well-formed type context

$\boxed{\vdash \Delta \text{ ok}}$

$$\frac{}{\vdash \bullet \text{ ok}} \quad \frac{\vdash \Delta \text{ ok} \quad \Delta \vdash \kappa \text{ ok}}{\vdash \Delta, \alpha : \kappa \text{ ok}} \alpha \notin \Delta$$

Well-formed Constructor

$\boxed{\Delta \vdash c : \kappa}$

$$\frac{\vdash \Delta, \alpha : \kappa, \Delta' \text{ ok}}{\Delta, \alpha : \kappa, \Delta' \vdash \alpha : \kappa} \quad \frac{\vdash \Delta \text{ ok}}{\Delta \vdash \text{Int} : T_{32}} \quad \frac{\vdash \Delta \text{ ok}}{\Delta \vdash \text{Boxf} : T_{32}} \quad \frac{\vdash \Delta \text{ ok}}{\Delta \vdash \text{Farray} : T_{32}}$$

$$\frac{\vdash \Delta \text{ ok}}{\Delta \vdash \text{Exn} : T_{32}} \quad \frac{\Delta \vdash c : T_{32}}{\Delta \vdash \text{Array}_c : T_{32}} \quad \frac{\Delta \vdash c : T_{32}}{\Delta \vdash \text{Dyntag}_c : T_{32}}$$

$$\frac{\Delta \vdash c_1 : T_{32} \quad \Delta \vdash c_2 : T_{32}}{\Delta \vdash \text{Vararg}_{c_1 \rightarrow c_2} : T_{32}} \quad \frac{\Delta, \alpha : T_{32}, \beta : T_{32} \vdash c_1 : T_{32} \quad \Delta, \alpha : T_{32}, \beta : T_{32} \vdash c_2 : T_{32}}{\Delta \vdash \mu(\alpha, \beta).(c_1, c_2) : T_{32} \times T_{32}} \alpha, \beta \notin \Delta$$

$$\frac{\Delta \vdash c_i : T_{32} \quad i \in 0 \dots j \quad \Delta \vdash c : T_{32}}{\Delta \vdash (c_0 \dots c_j) \rightarrow c : T_{32}} \quad \frac{\Delta \vdash c_i : T_{32} \quad i \in 0 \dots j}{\Delta \vdash c_1 \times \dots \times c_j : T_{32}}$$

$$\begin{array}{c}
\frac{\Delta \vdash c_i : T_{32} \quad i \in 0 \dots j}{\Delta \vdash \text{Sum}_n(c_0 \dots c_j) : T_{32}} \quad \frac{\Delta \vdash c_i : T_{32} \quad i \in 0 \dots j}{\Delta \vdash c : T_{32} \quad m \leq n + j} \\
\\
\frac{\vdash \kappa_1 \text{ ok} \quad \Delta, \alpha : \kappa_1 \vdash c : \kappa_2}{\Delta \vdash \lambda(\alpha :: \kappa).c : \kappa_1 \rightarrow \kappa_2} \alpha \notin \Delta \quad \frac{\Delta \vdash c_1 : \kappa_1 \rightarrow \kappa_2 \quad \Delta \vdash c_2 : \kappa_1}{\Delta \vdash c_1 c_2 : \kappa_2} \\
\\
\frac{\Delta \vdash c_1 : \kappa_1 \quad \Delta \vdash c_2 : \kappa_2}{\Delta \vdash \langle c_1, c_2 \rangle : \kappa_1 \times \kappa_2} \quad \frac{\Delta \vdash c : \kappa_1 \times \kappa_2}{\Delta \vdash \pi_i c : \kappa_i} (i \in \{1, 2\})
\end{array}$$

Base constructors

□

The base constructors consist of **Int**, **Boxf**, arrows, records, projections from mu types, **Vararg**, **Sum**, **Array**, **Farray**, **Exn**, and **Dyntag**. The non-record base constructors consist of any of the above except record types.

Constructor Equivalence

$$\boxed{\Delta \vdash c \equiv c' : \kappa}$$

$$\begin{array}{c}
\frac{\Delta \vdash c : \kappa}{\Delta \vdash c \equiv c : \kappa} \quad \frac{\Delta \vdash c' \equiv c : \kappa}{\Delta \vdash c \equiv c' : \kappa} \quad \frac{\Delta \vdash c_1 \equiv c_2 : \kappa \quad \Delta \vdash c_2 \equiv c_3 : \kappa}{\Delta \vdash c_1 \equiv c_3 : \kappa} \\
\\
\frac{\Delta \vdash c : \kappa_1 \rightarrow \kappa_2 \quad \alpha \notin \text{FV}(c)}{\Delta \vdash \lambda(\alpha : \kappa_1).c \alpha \equiv c : \kappa_1 \rightarrow \kappa_2} \quad \frac{\Delta \vdash c : \kappa_1 \times \kappa_2}{\Delta \vdash \langle \pi_1 c, \pi_2 c \rangle \equiv c : \kappa_1 \times \kappa_2} \\
\\
\frac{\Delta \vdash c_1 : \kappa_1 \quad \Delta \vdash c_2 : \kappa_2}{\Delta \vdash \pi_1 \langle c_1, c_2 \rangle \equiv c_1 : \kappa_1} \quad \frac{\Delta \vdash c_1 : \kappa_1 \quad \Delta \vdash c_2 : \kappa_2}{\Delta \vdash \pi_2 \langle c_1, c_2 \rangle \equiv c_2 : \kappa_2} \\
\\
\frac{\vdash \kappa_1 \text{ ok} \quad \Delta, \alpha : \kappa_1 \vdash c_1 : \kappa_2 \quad \Delta \vdash c_2 : \kappa_1}{\Delta \vdash (\lambda(\alpha : \kappa_1).c_1)c_2 \equiv c_1[c_2/\alpha] : \kappa_2} \\
\\
\frac{\vdash \kappa_1 \text{ ok} \quad \Delta, \alpha : \kappa_1 \vdash c_1 \equiv c_2 : \kappa_2}{\Delta \vdash \lambda(\alpha : \kappa_1).c_1 \equiv \lambda(\alpha : \kappa_1).c_2 : \kappa_1 \rightarrow \kappa_2} \quad \frac{\Delta \vdash c_1 \equiv c'_1 : \kappa_1 \rightarrow \kappa_2 \quad \Delta \vdash c_2 \equiv c'_2 : \kappa_1}{\Delta \vdash c_1 c_2 \equiv c'_1 c'_2 : \kappa_2}
\end{array}$$

$$\Delta \vdash c_1 \text{ ok} \quad \Delta \vdash c_2 \text{ ok} \quad \Delta \vdash c_1 \equiv \times[c'_1, \dots, c'_n] : T_{32} \quad n \leq \text{flattenlimit}$$

$$\Delta \vdash \text{Vararg}_{g_{c_1 \rightarrow c_2}} \equiv (c'_1, \dots, c'_n) \rightarrow c_2 : T_{32}$$

$$\Delta \vdash c_1 \text{ ok} \quad \Delta \vdash c_2 \text{ ok} \quad \Delta \vdash c_1 \equiv \times[c'_1, \dots, c'_n] : T_{32} \quad n > \text{flattenlimit}$$

$$\Delta \vdash \text{Vararg}_{g_{c_1 \rightarrow c_2}} \equiv c_1 \rightarrow c_2 : T_{32}$$

$$\Delta \vdash c_1 \text{ ok} \quad \Delta \vdash c_2 \text{ ok} \quad \Delta \vdash c_1 \equiv c'_1 : T_{32}$$

where c'_1 is a non-record base constructor

$$\Delta \vdash \text{Vararg}_{g_{c_1 \rightarrow c_2}} \equiv c_1 \rightarrow c_2 : T_{32}$$

$$\Delta \vdash c \equiv c' : \kappa_1 \times \kappa_2 \qquad \Delta \vdash c \equiv c' : \kappa_1 \times \kappa_2$$

$$\Delta \vdash \pi_1 c \equiv \pi_1 c' : \kappa_1 \qquad \Delta \vdash \pi_2 c \equiv \pi_2 c' : \kappa_2$$

Type equivalence

$$\boxed{\Delta \vdash \tau \equiv \tau' : T_{32}}$$

$$\Delta \vdash c \equiv c' : T_{32}$$

$$\Delta \vdash \tau_i \equiv \tau'_i : T_{32} \quad i \in 1 \dots n$$

$$\Delta \vdash T(c) \equiv T(c') : T_{32}$$

$$\Delta \vdash \tau_1 \times \dots \times \tau_n \equiv \tau'_1 \times \dots \times \tau'_n : T_{32}$$

$$\Delta \vdash \tau_i \equiv \tau'_i : T_{32} \quad i \in 1 \dots n$$

$$\Delta \vdash \tau \equiv \tau' : T_{32}$$

$$\Delta \vdash \forall[\alpha :: \kappa_1, \dots, \alpha :: \kappa_n](\tau_1, \dots, \tau_m)(k) \rightarrow \tau \equiv \forall[\alpha :: \kappa_1, \dots, \alpha :: \kappa_n](\tau'_1, \dots, \tau'_m)(k) \rightarrow \tau' : T_{32}$$

Well-formed float value

$$\boxed{\Delta; \Gamma \vdash fv : \text{Float}}$$

$$\vdash \Delta \text{ ok} \quad \Delta \vdash \Gamma \text{ ok}$$

$$\frac{\vdash \Delta \text{ ok} \quad \Delta \vdash \Gamma \text{ ok}}{\Delta; \Gamma \vdash r : \text{Float}} r$$

$$\vdash \Delta \text{ ok} \quad \Delta \vdash \Gamma[x_f] \text{ ok}$$

$$\frac{\vdash \Delta \text{ ok} \quad \Delta \vdash \Gamma[x_f] \text{ ok}}{\Delta; \Gamma[x_f] \vdash x_f : \text{Float}} fvar$$

Well-formed small value

$\Delta; \Gamma \vdash sv : \tau$

$$\frac{\Delta \vdash \Gamma[x:\tau] \text{ ok}}{\Delta; \Gamma[x:\tau] \vdash x : \tau} \text{ var}$$

$$\frac{\Delta \vdash \Gamma \text{ ok}}{\Delta; \Gamma \vdash i : \text{Int}} \text{ int}$$

$$\frac{\Delta; \Gamma \vdash sv : c_i[\pi_1 \mu(\alpha, \beta)(c_1, c_2)/\alpha, \pi_2 \mu(\alpha, \beta)(c_1, c_2)/\beta]}{\Delta; \Gamma \vdash \text{roll}_{\pi_i \mu(\alpha, \beta)(c_1, c_2)} sv : \pi_i \mu(\alpha, \beta)(c_1, c_2)} \text{ roll}$$

$$\frac{\Delta \vdash c \equiv \pi_i \mu(\alpha, \beta)(c_1, c_2) : T_{32} \quad \Delta; \Gamma \vdash sv : \pi_i \mu(\alpha, \beta)(c_1, c_2)}{\Delta; \Gamma \vdash \text{unroll}_c sv : c_i[\pi_1 \mu(\alpha, \beta)(c_1, c_2)/\alpha, \pi_2 \mu(\alpha, \beta)(c_1, c_2)/\beta]} \text{ unroll}$$

$$\frac{\Delta \vdash \text{Sum}_i^j(\vec{c}) : T_{32}}{\Delta; \Gamma \vdash \text{inj_tag}_{\text{Sum}_i(\vec{c})}^j : \text{Sum}_i(\vec{c})} \text{ inj_tag}$$

Well formed 32 bit instructions

$\Delta; \Gamma \vdash opr : \tau$

$$\frac{\Delta; \Gamma \vdash sv : \tau}{\Delta; \Gamma \vdash sv : \tau} \text{ sv}$$

$$\frac{\Delta; \Gamma \vdash fv : \text{Float}}{\Delta; \Gamma \vdash \text{boxf } fv : \text{Boxf}} \text{ boxf}$$

$$\frac{\Delta; \Gamma \vdash sv_i : \tau_i}{\Delta; \Gamma \vdash \langle sv_1, \dots, sv_n \rangle : \tau_1 \times \dots \times \tau_n} \text{ tuple}$$

$$\frac{\Delta; \Gamma \vdash sv : c_1 \rightarrow c_2}{\Delta; \Gamma \vdash \text{vararg}_{c_1 \rightarrow c_2} sv : \text{Vararg}_{c_1 \rightarrow c_2}} \text{ vararg}$$

$$\begin{array}{c}
\frac{\Delta; \Gamma \vdash sv : \mathbf{Vararg}_{c_1 \rightarrow c_2}}{\Delta; \Gamma \vdash \mathbf{onearg}_{c_1 \rightarrow c_2} sv : c_1 \rightarrow c_2} \text{onearg} \\
\\
\frac{\begin{array}{c} \Delta; \Gamma \vdash sv : \forall[\alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n](\tau_1, \dots, \tau_m)(k) \rightarrow \tau \\ \Delta; \Gamma \vdash sv_i : \tau_i[c_1/\alpha_1, \dots, c_n/\alpha_n] \\ \Delta; \Gamma \vdash fv_i : \mathbf{Float} \end{array}}{\Delta; \Gamma \vdash sv[c_1, \dots, c_n](sv_1, \dots, sv_m)(fv_1, \dots, fv_k) : \tau[c_1/\alpha_1, \dots, c_n/\alpha_n]} \text{app} \\
\\
\frac{\begin{array}{c} \Delta \vdash \mathbf{Sum}_i^j(\vec{c}) : \mathbf{T}_{32} \\ \Delta; \Gamma \vdash sv : \mathbf{Sum}_i^j(\vec{c}) \end{array}}{\Delta; \Gamma \vdash \mathbf{proj}_{\mathbf{Sum}_i(\vec{c})}^j sv : c_j} \text{proj} \\
\\
\frac{\Delta \vdash \mathbf{Sum}_i^j(\vec{c}) : \mathbf{T}_{32} \quad \Delta; \Gamma \vdash sv : c_j}{\Delta; \Gamma \vdash \mathbf{inj}_{\mathbf{Sum}_i(\vec{c})}^j sv : \mathbf{Sum}_i(\vec{c})} \text{inj} \\
\\
\frac{\Delta; \Gamma \vdash sv : \tau_1 \times \dots \times \tau_n}{\Delta; \Gamma \vdash \mathbf{select}^i sv : \tau_i} \text{select} \\
\\
\frac{\begin{array}{c} \Delta \vdash \tau : \mathbf{T}_{32} \\ \Delta; \Gamma \vdash sv : \mathbf{Sum}_i(\vec{c}) \quad \Delta; \Gamma[x_j : \mathbf{Sum}_i^j(\vec{c})] \vdash e_j : \tau \end{array}}{\Delta; \Gamma \vdash \mathbf{case}_\tau(sv)(x_1.e_1, \dots, x_n.e_n) : \tau} \text{case} \\
\\
\frac{\begin{array}{c} \Delta; \Gamma \vdash sv : \mathbf{Exn} \Delta \vdash \tau : \mathbf{T}_{32} \\ \Delta; \Gamma \vdash sv_1 : \mathbf{Dyntag}_{c_1} \quad \Delta; \Gamma[x_1 : c_1] \vdash e_1 : \tau \quad \Delta; \Gamma \vdash e_2 : \tau \end{array}}{\Delta; \Gamma \vdash \mathbf{exncase}_\tau(sv)(sv_1 \Rightarrow x_1.e_1, - \Rightarrow e_2) : \tau} \text{exncase} \\
\\
\frac{\Delta; \Gamma \vdash e_1 : \tau \quad \Delta; \Gamma[x : \mathbf{Exn}] \vdash e_2 : \tau}{\Delta; \Gamma \vdash \mathbf{handle}_\tau(e_1, x.e_2) : \tau} \text{handle} \\
\\
\frac{\Delta; \Gamma \vdash sv_1 : \mathbf{Dyntag}_c \quad \Delta; \Gamma \vdash sv_2 : c}{\Delta; \Gamma \vdash \mathbf{inj_dyn}_c(sv_1, sv_2) : \mathbf{Exn}} \text{inj_exn}
\end{array}$$

$$\frac{\Delta \vdash c : T_{32} \quad \Delta; \Gamma \vdash sv_1 : \text{Array}_c \quad \Delta; \Gamma \vdash sv_2 : \text{Int}}{\Delta; \Gamma \vdash \text{sub}_c(sv_1, sv_2) : c} \text{sub}$$

$$\frac{\Delta; \Gamma \vdash sv_1 : \text{Int} \quad \Delta; \Gamma \vdash sv_2 : c}{\Delta; \Gamma \vdash \text{array}_c(sv_1, sv_2) : \text{Array}_c} \text{array}$$

$$\frac{\Delta; \Gamma \vdash sv : \text{Int} \quad \Delta \vdash fv : \text{Float}}{\Delta; \Gamma \vdash \text{farray}(sv, fv) : \text{Farray}} \text{farray}$$

$$\frac{\Delta \vdash \tau : T_{32} \quad \Delta; \Gamma \vdash sv : \text{Exn}}{\Delta; \Gamma \vdash \text{raise}_\tau sv : \tau} \text{raise}$$

$$\frac{\Delta \vdash c : T_{32} \quad \Delta \vdash \Gamma \text{ ok}}{\Delta; \Gamma \vdash \text{mkexntag}_c : \text{Dyntag}_c} \text{mkexntag}$$

Well formed float instructions

$\Delta; \Gamma \vdash \text{opr} : \text{Float}$

$$\frac{\Delta; \Gamma \vdash sv : \text{Boxf}}{\Delta; \Gamma \vdash \text{unboxf } sv : \text{Float}} \text{unbox}$$

$$\frac{\Delta; \Gamma \vdash sv_1 : \text{Farray} \quad \Delta; \Gamma \vdash sv_2 : \text{Int}}{\Delta; \Gamma \vdash \text{fsub}(sv_1, sv_2) : \text{Float}} \text{fsub}$$

$$\frac{\Delta; \Gamma \vdash fv : \text{Float}}{\Delta; \Gamma \vdash fv : \text{Float}} \text{fv}$$

Well-formed Expression

$$\boxed{\Delta; \Gamma \vdash e : \tau}$$

$$\frac{\Delta; \Gamma \vdash sv : \tau}{\Delta; \Gamma \vdash sv : \tau} \text{sv}$$

$$\frac{\Delta; \Gamma \vdash i : \tau' \quad \Delta; \Gamma[x : \tau'] \vdash e : \tau}{\Delta; \Gamma \vdash \text{let}_\tau x = i \text{ in } e : \tau} i$$

$$\frac{\Delta; \Gamma \vdash i : \text{Float} \quad \Delta; \Gamma[x_f] \vdash e : \tau}{\Delta; \Gamma \vdash \text{let}_\tau x_f = i \text{ in } e : \tau} i64$$

$$\frac{\begin{array}{c} \Delta; \Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r, \alpha::\kappa, x::\tau, \vec{x}_f] \vdash e_f : \tau_r \\ \Delta; \Gamma[f : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau_r] \vdash e : \tau \end{array}}{\Delta; \Gamma \vdash \text{let}_\tau \text{rec}_{\tau_r} f[\alpha::\kappa](x::\tau)(\vec{x}_f).e_f \text{ in } e : \forall[\alpha::\kappa](\vec{\tau})(|x_f|) \rightarrow \tau} \text{rec}$$

Well-formed Context

$$\boxed{\Delta \vdash \Gamma \text{ ok}}$$

$$\frac{\vdash \Delta \text{ ok}}{\Delta \vdash \bullet \text{ ok}} \quad \frac{\Delta \vdash \Gamma \text{ ok} \quad \Delta \vdash \tau : T_{32}}{\Delta \vdash \Gamma, x : \tau \text{ ok}} x \notin \Gamma \quad \frac{\Delta \vdash \Gamma \text{ ok}}{\Delta \vdash \Gamma, x_{64} : \phi \text{ ok}} x_{64} \notin \Gamma$$

Appendix B

LIL static semantics

Definitions

□

$$\kappa \text{ list} \stackrel{\text{def}}{=} \mu j.1 + \kappa \times j$$

$$\text{nat} \stackrel{\text{def}}{=} \mu j.1 + j$$

$$\begin{aligned} \bar{0} &\stackrel{\text{def}}{=} \text{fold}_{\text{nat}} \text{inj}_1^{1+\text{nat}} * \\ \overline{n+1} &\stackrel{\text{def}}{=} \text{fold}_{\text{nat}} \text{inj}_2^{1+\text{nat}}(\bar{n}) \end{aligned}$$

Well-formed Kind

$\Delta \vdash \kappa \text{ ok}$

$$\begin{array}{c} \frac{\vdash \Delta \text{ ok}}{\Delta \vdash \text{T}_{32} \text{ ok}} \quad \frac{\vdash \Delta \text{ ok}}{\Delta \vdash \text{T}_{64} \text{ ok}} \quad \frac{\vdash \Delta \text{ ok}}{\Delta \vdash 1 \text{ ok}} \quad \frac{\vdash \Delta, j, \Delta' \text{ ok}}{\Delta, j, \Delta' \vdash j \text{ ok}} \\[10pt] \frac{\Delta \vdash \kappa_1 \text{ ok} \quad \Delta \vdash \kappa_2 \text{ ok}}{\Delta \vdash \kappa_1 \rightarrow \kappa_2 \text{ ok}} \quad \frac{\Delta \vdash \kappa_1 \text{ ok} \quad \Delta \vdash \kappa_2 \text{ ok}}{\Delta \vdash \kappa_1 \times \kappa_2 \text{ ok}} \quad \frac{\Delta \vdash \kappa_1 \text{ ok} \quad \Delta \vdash \kappa_2 \text{ ok}}{\Delta \vdash \kappa_1 + \kappa_2 \text{ ok}} \\[10pt] \frac{\Delta, j \vdash \kappa \text{ ok}}{\Delta \vdash \mu j. \kappa \text{ ok}} (j \notin \Delta, j \text{ only positive in } \kappa) \quad \frac{\Delta, j \vdash \kappa \text{ ok}}{\Delta \vdash \forall j. \kappa \text{ ok}} (j \notin \Delta) \end{array}$$

Well-formed type and kind ontext

$\vdash \Delta \text{ ok}$

$$\frac{}{\vdash \bullet \text{ ok}} \quad \frac{\vdash \Delta \text{ ok}}{\vdash \Delta, j \text{ ok}} j \notin \Delta \quad \frac{\vdash \Delta \text{ ok} \quad \Delta \vdash \kappa \text{ ok}}{\vdash \Delta, \alpha : \kappa \text{ ok}} \alpha \notin \Delta$$

Constants and their kinds		
$\text{Float} : \mathsf{T}_{64}$	$\text{Int} : \mathsf{T}_{32}$	$\text{Void} : \mathsf{T}_{32}$
$\text{Array}_{32} : \mathsf{T}_{32} \rightarrow \mathsf{T}_{32}$	$\text{Array}_{64} : \mathsf{T}_{64} \rightarrow \mathsf{T}_{32}$	$\text{Boxed} : \mathsf{T}_{64} \rightarrow \mathsf{T}_{32}$
$\text{Tag} : \mathsf{nat} \rightarrow \mathsf{T}_{32}$	$\text{Dyntag} : \mathsf{T}_{32} \rightarrow \mathsf{T}_{32}$	$\text{Dyn} : \mathsf{T}_{32}$
$\times : \mathsf{T}_{32} \text{list} \rightarrow \mathsf{T}_{32}$	$\rightarrow : \mathsf{T}_{32} \text{list} \rightarrow \mathsf{T}_{64} \text{list} \rightarrow \mathsf{T}_{32} \rightarrow \mathsf{T}_{32}$	
$\vee : \mathsf{T}_{32} \text{list} \rightarrow \mathsf{T}_{32}$	$\text{Code} : \mathsf{T}_{32} \text{list} \rightarrow \mathsf{T}_{64} \text{list} \rightarrow \mathsf{T}_{32} \rightarrow \mathsf{T}_{32}$	
$\forall : \forall j. (j \rightarrow \mathsf{T}_{32}) \rightarrow \mathsf{T}_{32}$	$\exists : \forall j. (j \rightarrow \mathsf{T}_{32}) \rightarrow \mathsf{T}_{32}$	
$\text{Rec} : \forall j. ((j \rightarrow \mathsf{T}_{32}) \rightarrow (j \rightarrow \mathsf{T}_{32})) \rightarrow j \rightarrow \mathsf{T}_{32}$		

$$\frac{\vdash \Delta \text{ ok}}{\Delta \vdash * : 1} \quad \frac{\vdash \Delta, \alpha : \kappa, \Delta' \text{ ok}}{\Delta, \alpha : \kappa, \Delta' \vdash \alpha : \kappa}$$

$$\frac{\Delta \vdash \kappa_1 \text{ ok} \quad \Delta, \alpha : \kappa_1 \vdash c : \kappa_2}{\Delta \vdash \lambda(\alpha :: \kappa).c : \kappa_1 \rightarrow \kappa_2} \alpha \notin \Delta \quad \frac{\Delta \vdash c_1 : \kappa_1 \rightarrow \kappa_2 \quad \Delta \vdash c_2 : \kappa_1}{\Delta \vdash c_1 c_2 : \kappa_2}$$

$$\frac{\Delta \vdash c_1 : \kappa_1 \quad \Delta \vdash c_2 : \kappa_2}{\langle c_1, c_2 \rangle \vdash \kappa_1 \times \kappa_2 :} \quad \frac{\Delta \vdash c : \kappa_1 \times \kappa_2}{\Delta \vdash \pi_i c : \kappa_i} (i \in \{1, 2\})$$

$$\frac{\Delta \vdash c : \kappa_i \quad (i \in \{1 \dots n\})}{\Delta \vdash \text{inj}_i^{+[\kappa_1, \dots, \kappa_n]} c : +[\kappa_1, \dots, \kappa_n]} \quad \frac{\Delta \vdash c : +[\kappa_1, \dots, \kappa_n] \quad \Delta, \alpha_i : \kappa_i \vdash c_i : \kappa \quad i \in 1 \dots n}{\Delta \vdash \text{case } c[(\alpha_1.c_1, \dots, \alpha_n.c_n)] : \kappa} \alpha_i \notin \Delta$$

$$\frac{\Delta \vdash \mu j. \kappa \text{ ok} \quad \Delta \vdash c : \kappa[\mu j. \kappa / j]}{\Delta \vdash \text{fold}_{\mu j. \kappa} c : \mu j. \kappa}$$

$$\frac{\Delta \vdash \mu j. \kappa \text{ ok} \quad \Delta \vdash \kappa' \text{ ok} \quad j, \alpha, \rho, \notin \Delta \quad \Delta, j, \alpha : \kappa, \rho : (j \rightarrow \kappa'), \vdash c : \kappa'}{\Delta \vdash \text{pr}(j, \alpha : \kappa, \rho : (j \rightarrow \kappa'), \text{in } c) : \mu j. \kappa \rightarrow \kappa'}$$

$$\frac{\Delta \vdash c : \forall j. \kappa' \quad \Delta \vdash \kappa \text{ ok}}{\Delta \vdash c[\kappa] : \kappa'[\kappa / j]} \quad \frac{\Delta, j \vdash c : \kappa}{\Delta \vdash \Lambda j. c : \forall j. \kappa} j \notin \Delta$$

Constructor Equivalence

$$\Delta \vdash c \equiv c' : \kappa$$

$$\frac{\Delta \vdash c : \kappa}{\Delta \vdash c \equiv c : \kappa} \quad \frac{\Delta \vdash c' \equiv c : \kappa}{\Delta \vdash c \equiv c' : \kappa} \quad \frac{\Delta \vdash c_1 \equiv c_2 : \kappa \quad \Delta \vdash c_2 \equiv c_3 : \kappa}{\Delta \vdash c_1 \equiv c_3 : \kappa}$$

$$\frac{\Delta \vdash c : \kappa_1 \rightarrow \kappa_2 \quad \alpha \notin \text{FV}(c)}{\Delta \vdash \lambda(\alpha : \kappa_1).c \alpha \equiv c : \kappa_1 \rightarrow \kappa_2} \quad \frac{\Delta \vdash c : \kappa_1 \times \kappa_2}{\Delta \vdash \langle \pi_1 c, \pi_2 c \rangle \equiv c : \kappa_1 \times \kappa_2}$$

$$\Delta \vdash c : +[\kappa_1, \dots, \kappa_n]$$

$$\Delta \vdash \text{case}(c, [\alpha_1. \text{inj}_1^{+[\kappa_1, \dots, \kappa_n]} \alpha_1, \dots, \alpha_n. \text{inj}_n^{+[\kappa_1, \dots, \kappa_n]} \alpha_n]) \equiv c : +[\kappa_1, \dots, \kappa_n]$$

$$\frac{\Delta \vdash c_1 : \kappa_1 \quad \Delta \vdash c_2 : \kappa_2}{\Delta \vdash \pi_1 \langle c_1, c_2 \rangle \equiv c_1 : \kappa_1} \quad \frac{\Delta \vdash c_1 : \kappa_1 \quad \Delta \vdash c_2 : \kappa_2}{\Delta \vdash \pi_2 \langle c_1, c_2 \rangle \equiv c_2 : \kappa_2}$$

$$\frac{\Delta \vdash \kappa_1 \text{ ok} \quad \Delta, \alpha : \kappa_1 \vdash c_1 : \kappa_2 \quad \Delta \vdash c_2 : \kappa_1}{\Delta \vdash (\lambda(\alpha : \kappa_1).c_1).c_2 \equiv c_1[c_2/\alpha] : \kappa_2} \quad \frac{\Delta \vdash \kappa \text{ ok} \quad \Delta, j \vdash c : \kappa'}{\Delta \vdash (\Lambda j.c)[\kappa] \equiv c[\kappa/j] : \kappa'[\kappa/j]}$$

$$\frac{\Delta \vdash c : \kappa_i \quad \Delta \vdash \kappa_j \text{ ok} \quad j \in 1 \dots n \quad \Delta, \alpha_j : \kappa_j \vdash c_j : \kappa \quad j \in 1 \dots n}{\Delta \vdash \text{case}(\text{inj}_i^{+[\kappa_1, \dots, \kappa_n]} c, [\dots, \alpha_i.c_i, \dots]) \equiv c_i[c/\alpha_i] : \kappa}$$

$$\frac{\Delta \vdash c' : \kappa[\mu j. \kappa/j] \quad \Delta \vdash \kappa' \text{ ok} \quad (j, \alpha, \rho, \notin \Delta) \quad \Delta, j, \alpha : \kappa, \rho : (j \rightarrow \kappa'), \vdash c : \kappa' \quad \Delta \vdash \mu j. \kappa \text{ ok}}{\Delta \vdash \text{pr}(j, \alpha : \kappa, \rho : (j \rightarrow \kappa'), \text{in } c) \text{ fold}_{\mu j. \kappa} c' \equiv c[\mu j. \kappa, c', \text{pr}(j, \alpha : \kappa, \rho : (j \rightarrow \kappa') \text{ in } c), /j, \alpha, \rho] : \mu j. \kappa \rightarrow \kappa'}$$

$$\frac{\Delta \vdash \kappa_1 \text{ ok} \quad \Delta, \alpha : \kappa_1 \vdash c_1 \equiv c_2 : \kappa_2}{\Delta \vdash \lambda(\alpha : \kappa_1).c_1 \equiv \lambda(\alpha : \kappa_1).c_2 : \kappa_1 \rightarrow \kappa_2} \quad \frac{\Delta \vdash c_1 \equiv c'_1 : \kappa_1 \rightarrow \kappa_2 \quad \Delta \vdash c_2 \equiv c'_2 : \kappa_1}{\Delta \vdash c_1 c_2 \equiv c'_1 c'_2 : \kappa_2}$$

$$\frac{\Delta \vdash c \equiv c' : \kappa_1 \times \kappa_2}{\Delta \vdash \pi_1 c \equiv \pi_1 c' : \kappa_1} \quad \frac{\Delta \vdash c \equiv c' : \kappa_1 \times \kappa_2}{\Delta \vdash \pi_2 c \equiv \pi_2 c' : \kappa_2}$$

$$\begin{array}{c}
\Delta \vdash c \equiv c' : \kappa_i \quad \Delta \vdash \kappa_j \text{ ok} \quad j \in 1 \dots n \\
\hline
\Delta \vdash \text{inj}_i^{+[\kappa_1, \dots, \kappa_n]} c \equiv \text{inj}_i^{+[\kappa_1, \dots, \kappa_n]} c' : +[\kappa_1, \dots, \kappa_n] \\
\\
\Delta \vdash c \equiv c' : +[\kappa_1, \dots, \kappa_n] \\
\Delta, \alpha_i : \kappa_i \vdash c_i \equiv c'_i : \kappa \quad i \in 1 \dots n \\
\hline
\Delta \vdash \text{case}(c, [\alpha_1.c_1, \dots, \alpha_n.c_n]) \equiv \text{case}(c', [\alpha_1.c'_1, \dots, \alpha_n.c'_n]) : \kappa \\
\\
\Delta \vdash c \equiv c' : \kappa[\mu j.\kappa/j] \\
\hline
\Delta \vdash \text{fold}_{\mu j.\kappa} c \equiv \text{fold}_{\mu j.\kappa} c' : \mu j.\kappa \\
\\
j, \alpha, \rho, \notin \Delta \\
\Delta \vdash \mu j.\kappa \text{ ok} \quad \Delta \vdash \kappa' \text{ ok} \quad \Delta, j, \alpha : \kappa, \rho : (j \rightarrow \kappa'), \vdash c \equiv c' : \kappa' \\
\hline
\Delta \vdash \text{pr}(j, \alpha : \kappa, \rho : j \rightarrow \kappa', \text{in } c) \\
\equiv \text{pr}(j, \alpha : \kappa, \rho : j \rightarrow \kappa', \text{in } c') : \mu j.\kappa \rightarrow \kappa' \\
\\
\Delta, j \vdash c \equiv c' : \kappa' \quad j \notin \Delta \quad \Delta \vdash c \equiv c' : \forall j.\kappa' \quad \Delta \vdash \kappa \text{ ok} \\
\hline
\Delta \vdash \Lambda j.c \equiv \Lambda j.c' : \forall j.\kappa' \quad \Delta \vdash c[\kappa] \equiv c'[\kappa] : \kappa'[\kappa/j]
\end{array}$$

Well-formed term context

$\Delta \vdash \Gamma \text{ ok}$

$$\frac{\vdash \Delta \text{ ok}}{\Delta \vdash \bullet \text{ ok}} \quad \frac{\Delta \vdash \Gamma \text{ ok} \quad \Delta \vdash \tau : T_{32}}{\Delta \vdash \Gamma, x : \tau \text{ ok}} \quad x \notin \Gamma \quad \frac{\Delta \vdash \Gamma \text{ ok} \quad \Delta \vdash \phi : T_{64}}{\Delta \vdash \Gamma, x_{64} : \phi \text{ ok}} \quad x_{64} \notin \Gamma$$

Well-formed heap context

$\vdash \Psi \text{ ok}$

$$\frac{}{\vdash \bullet \text{ ok}} \quad \frac{\vdash \Psi \text{ ok} \quad \bullet \vdash \tau : T_{32}}{\vdash \Psi, \ell : \tau \text{ ok}} \quad \ell \notin \Gamma$$

Well-formed 64 bit value

$\Delta; \Gamma \vdash fv : \phi$

$$\frac{\Delta \vdash \Gamma \text{ ok} \quad \vdash \Psi \text{ ok}}{\Psi; \Delta; \Gamma \vdash \mathbf{r} : \text{Float}} \quad \frac{\Delta \vdash \Gamma, x_{64} : \phi, \Gamma' \text{ ok} \quad \vdash \Psi \text{ ok}}{\Psi; \Delta; \Gamma, x_{64} : \phi, \Gamma' \vdash x_{64} : \phi}$$

Well-formed 32 bit Value

$$\boxed{\Psi; \Delta; \Gamma \vdash sv : \tau}$$

$$\begin{array}{c}
\frac{\Delta \vdash \Gamma, x:\tau, \Gamma' \text{ ok} \quad \vdash \Psi \text{ ok}}{\Psi; \Delta; \Gamma, x:\tau, \Gamma' \vdash x : \tau} \quad \frac{\Delta \vdash \Gamma \text{ ok} \quad \vdash \Psi \text{ ok}}{\Psi; \Delta; \Gamma \vdash i : \text{Int}} \quad \frac{\Delta \vdash \Gamma \text{ ok} \quad \vdash \Psi_1, \ell:\tau, \Psi_2 \text{ ok}}{\Psi_1, \ell:\tau, \Psi_2; \Delta; \Gamma \vdash \ell : \tau} \\
\\
\frac{\Delta \vdash \tau \equiv \text{Rec}[\kappa](c)(c_p) : T_{32} \quad \Psi; \Delta; \Gamma \vdash sv : c(\text{Rec}[\kappa]c)c_p}{\Psi; \Delta; \Gamma \vdash \text{roll}_\tau sv : \tau} \quad \frac{\Delta \vdash \tau \equiv \text{Rec}[\kappa](c)(c_p) : T_{32} \quad \Psi; \Delta; \Gamma \vdash sv : \tau}{\Psi; \Delta; \Gamma \vdash \text{unroll}_\tau sv : c(\text{Rec}[\kappa]c)c_p} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv_1 : \text{DynTag } \tau \quad \Psi; \Delta; \Gamma \vdash sv_2 : \tau}{\Psi; \Delta; \Gamma \vdash \text{inj_dyn}_\tau(sv_1, sv_2) : \text{Dyn}} \quad \frac{\Delta \vdash c \equiv \bigvee[\dots, c_i, \dots] : T_{32} \quad \Psi; \Delta; \Gamma \vdash sv : c_i}{\Psi; \Delta; \Gamma \vdash \text{inj_union}_c sv : c} \\
\\
\frac{\Delta \vdash \tau \equiv \exists[\kappa](c') : T_{32} \quad \Delta \vdash c : \kappa \quad \Psi; \Delta; \Gamma \vdash sv : c'c]}{\Psi; \Delta; \Gamma \vdash \text{pack } sv \text{ as } \tau \text{ hiding } c : \tau} \quad \frac{\Psi; \Delta; \Gamma \vdash sv : \forall[\kappa](c') \quad \Delta \vdash c : \kappa}{\Psi; \Delta; \Gamma \vdash sv[c] : c'c} \\
\\
\frac{\Delta \vdash \Gamma \text{ ok} \quad \vdash \Psi \text{ ok}}{\Psi; \Delta; \Gamma \vdash \text{tag}_i : \text{Tag}(\bar{i})}
\end{array}$$

Well formed 64 bit operations

$$\boxed{\Psi; \Delta; \Gamma \vdash fopr : \phi \text{ opr}_{64}}$$

$$\begin{array}{c}
\frac{\Psi; \Delta; \Gamma \vdash fv : \phi}{\Psi; \Delta; \Gamma \vdash fv : \phi \text{ opr}_{64}} \quad \frac{\Psi; \Delta; \Gamma \vdash sv : \text{Boxed } \phi}{\Psi; \Delta; \Gamma \vdash \text{unbox } sv : \phi \text{ opr}_{64}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv_1 : \text{Array}_{64} \phi \quad \Psi; \Delta; \Gamma \vdash sv_2 : \text{Int}}{\Psi; \Delta; \Gamma \vdash \text{sub}_\phi(sv_1, sv_2) : \phi \text{ opr}_{64}}
\end{array}$$

Well-formed 32 bit operation

$$\boxed{\Psi; \Delta; \Gamma \vdash opr : \tau \text{ opr}_{32}}$$

$$\begin{array}{c}
\frac{\Psi; \Delta; \Gamma \vdash sv : \tau}{\Psi; \Delta; \Gamma \vdash sv : \tau \text{ opr}_{32}} \quad \frac{\Psi; \Delta; \Gamma \vdash sv : \times (\tau_0 :: \dots :: \tau_i :: c')}{\Psi; \Delta; \Gamma \vdash \text{select}^i sv : \tau_i \text{ opr}_{32}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash fv : \phi}{\Psi; \Delta; \Gamma \vdash \text{box } fv : \text{Boxed } \phi} \quad \frac{\Psi; \Delta; \Gamma \vdash sv_i : \tau_i \quad i \in 0, \dots, n}{\Psi; \Delta; \Gamma \vdash \langle sv_0, \dots, sv_n \rangle : \times [\tau_0, \dots, \tau_n]}
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash sv : \bigvee[\tau_0, \dots, \tau_k] \quad \Psi; \Delta; \Gamma, x_i : \tau_i \vdash e_i : \tau \textbf{exp} \\
\Delta \vdash \tau_i \equiv \textbf{Tag}(i) : T_{32} \quad i \in 0 \dots (j-1) \\
\Delta \vdash \tau_i \equiv \times[\textbf{Tag}(i), \tau'_i] : T_{32} \quad i \in j \dots k \\
\hline
\Psi; \Delta; \Gamma \vdash \textbf{case}(sv)(x.e_0, \dots, x.e_k) : \tau \textbf{opr}_{32}
\end{array}$$

$$\begin{array}{c}
\Delta \vdash \tau : T_{32} \\
\hline
\Psi; \Delta; \Gamma \vdash \textbf{dyntag}_{\tau} : \textbf{Dyntag} \tau \textbf{opr}_{32}
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash sv : \textbf{Dyn} \quad \Psi; \Delta; \Gamma \vdash sv_1 : \textbf{Dyntag} \tau_1 \\
\Psi; \Delta; \Gamma, x_1 : \tau_1 \vdash e_1 : \tau \textbf{exp} \quad \Psi; \Delta; \Gamma \vdash e : \tau \textbf{exp} \\
\hline
\Psi; \Delta; \Gamma \vdash \textbf{dyncase}(sv)(sv_1 \Rightarrow x_1.e_1, - \Rightarrow e) : \tau \textbf{opr}_{32}
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash sv : \textbf{Dyn} \quad \Delta \vdash \tau : T_{32} \quad \Psi; \Delta; \Gamma \vdash e_1 : \tau \textbf{exp} \\
\hline
\Psi; \Delta; \Gamma \vdash \textbf{raise}_{\tau} sv : \tau \textbf{opr}_{32}
\end{array}
\quad
\begin{array}{c}
\Psi; \Delta; \Gamma, x : \textbf{Dyn} \vdash e_2 : \tau \textbf{exp} \\
\hline
\Psi; \Delta; \Gamma \vdash \textbf{handle}_{\tau}(e_1, x.e_2) : \tau \textbf{opr}_{32}
\end{array}
\quad x \notin \Gamma$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash sv : \rightarrow([\tau_0, \dots, \tau_n])([\phi_0, \dots, \phi_k])(\tau) \\
\Psi; \Delta; \Gamma \vdash sv_i : \tau_i \quad \Psi; \Delta; \Gamma \vdash fv_i : \phi_i \\
\hline
\Psi; \Delta; \Gamma \vdash sv(sv_0, \dots, sv_n)(fv_0, \dots, fv_k) : \tau \textbf{opr}_{32}
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash sv : \textbf{Code}[\tau_0, \dots, \tau_n][\phi_0, \dots, \phi_k](\tau) \\
\Psi; \Delta; \Gamma \vdash sv_i : \tau_i \quad \Psi; \Delta; \Gamma \vdash fv_i : \phi_i \\
\hline
\Psi; \Delta; \Gamma \vdash \textbf{call} sv(sv_0, \dots, sv_n)(fv_0, \dots, fv_k) : \tau \textbf{opr}_{32}
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash sv_1 : \textbf{Int} \quad \Psi; \Delta; \Gamma \vdash sv_2 : \tau \\
\hline
\Psi; \Delta; \Gamma \vdash \textbf{array}_{\tau}(sv_1, sv_2) : \textbf{Array}_{32}(\tau) \textbf{opr}_{32}
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash sv_1 : \textbf{Array}_{32}(\tau) \quad \Psi; \Delta; \Gamma \vdash sv_2 : \textbf{Int} \\
\hline
\Psi; \Delta; \Gamma \vdash \textbf{sub}_{\tau}(sv_1, sv_2) : \tau \textbf{opr}_{32}
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash sv_1 : \textbf{Array}_{32}(\tau) \quad \Psi; \Delta; \Gamma \vdash sv_2 : \textbf{Int} \quad \Psi; \Delta; \Gamma \vdash sv_3 : \tau \\
\hline
\Psi; \Delta; \Gamma \vdash \textbf{upd}_{\tau}(sv_1, sv_2, sv_3) : \textbf{Unit} \textbf{opr}_{32}
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash sv : \text{Int} \quad \Psi; \Delta; \Gamma \vdash fv : \phi \\
\hline
\Psi; \Delta; \Gamma \vdash \text{array}_\phi(sv, fv) : \text{Array}_{64}\phi \text{ opr}_{32} \\
\\
\Psi; \Delta; \Gamma \vdash sv_1 : \text{Array}_{64}\phi \quad \Psi; \Delta; \Gamma \vdash sv_2 : \text{Int} \quad \Psi; \Delta; \Gamma \vdash fv : \phi \\
\hline
\Psi; \Delta; \Gamma \vdash \text{upd}_\phi(sv_1, sv_2, fv) : \text{Unit opr}_{64}
\end{array}$$

Well-formed Expression

$$\boxed{\Psi; \Delta; \Gamma \vdash e : \tau \text{ exp}}$$

$$\begin{array}{c}
\Delta \vdash \kappa_i \text{ ok} \quad \Delta, \vec{\alpha} : \vec{\kappa} \vdash \tau_i : \text{T}_{32} \quad \Delta, \vec{\alpha} : \vec{\kappa} \vdash \phi_i : \text{T}_{64} \\
\Psi; \Delta, \vec{\alpha} : \vec{\kappa}; \Gamma, f : \forall[\vec{\alpha} : \vec{\kappa}](\vec{\tau})(\vec{\phi}) \rightarrow \tau, \vec{x} : \vec{\tau}, x_{64} : \vec{\phi} \vdash e : \tau \text{ exp} \\
\Psi; \Delta; \Gamma, f : \forall[\vec{\alpha} : \vec{\kappa}](\vec{\tau})(\vec{\phi}) \rightarrow \tau \vdash e' : \tau' \text{ exp} \\
\hline
\Psi; \Delta; \Gamma \vdash \text{let rec}_\tau f[\vec{\alpha} : \vec{\kappa}](\vec{x} : \vec{\tau})(x_{64} : \vec{\phi}).e \text{ in } e' : \tau' \text{ exp} \quad \vec{\alpha}, \vec{x}, x_{64}, f \notin \Delta, \Gamma
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash sv : \tau \\
\hline
\Psi; \Delta; \Gamma \vdash sv : \tau \text{ exp}
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash \text{opr} : \tau \text{ opr}_{32} \quad \Psi; \Delta; \Gamma \vdash \text{opr} : \phi \text{ opr}_{32} \\
\Psi; \Delta; \Gamma, x : \tau \vdash e : \tau' \text{ exp} \quad \Psi; \Delta; \Gamma, x_{64} : \phi \vdash e : \tau' \text{ exp} \\
\hline
\Psi; \Delta; \Gamma \vdash \text{let } x = \text{opr in } e : \tau' \text{ exp} \quad \Psi; \Delta; \Gamma \vdash \text{let } x_{64} = \text{fopr in } e : \tau' \text{ exp}
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash sv : \exists[\kappa](c) \\
\Psi; \Delta, \alpha : \kappa; \Gamma, x : (c\alpha) \vdash e : \tau' \text{ exp} \quad \alpha \notin \text{fv}(\tau') \\
\hline
\Psi; \Delta; \Gamma \vdash \text{let}[\alpha, x] = \text{unpack } sv \text{ in } e : \tau' \text{ exp} \quad \alpha, x \notin \Delta, \Gamma
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta, \beta : \kappa_1, \gamma : \kappa_2, \Delta'; \Gamma[\langle \beta, \gamma \rangle / \alpha] \vdash e[\langle \beta, \gamma \rangle / \alpha] : \tau[\langle \beta, \gamma \rangle / \alpha] \text{ exp} \\
\Delta, \alpha : \kappa_1 \times \kappa_2, \Delta' \vdash c \equiv \alpha : \kappa_1 \times \kappa_2 \\
\hline
\Psi; \Delta, \alpha : \kappa_1 \times \kappa_2, \Delta'; \Gamma \vdash \text{let} \langle \beta, \gamma \rangle = c \text{ in } e : \tau \text{ exp} \quad \beta, \gamma \notin \Delta
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta; \Gamma \vdash e[c_1, c_2 / \beta, \gamma] : \tau \text{ exp} \\
\Delta \vdash c \equiv \langle c_1, c_2 \rangle : \kappa_1 \times \kappa_2 \\
\hline
\Psi; \Delta; \Gamma \vdash \text{let} \langle \beta, \gamma \rangle = c \text{ in } e : \tau \text{ exp} \quad \beta \notin \Delta
\end{array}$$

$$\begin{array}{c}
\Psi; \Delta, \beta : \kappa[\mu j. \kappa / j], \Delta'; \Gamma[\text{fold}_{\mu j. \kappa} \beta / \alpha] \vdash e[\text{fold}_{\mu j. \kappa} \beta / \alpha] : \tau[\text{fold}_{\mu j. \kappa} \beta / \alpha] \text{ exp} \\
\Delta, \alpha : \mu j. \kappa, \Delta' \vdash c \equiv \alpha : \mu j. \kappa \\
\hline
\Psi; \Delta, \alpha : \mu j. \kappa, \Delta'; \Gamma \vdash \text{let}(\text{fold } \beta) = c \text{ in } e : \tau \text{ exp}
\end{array}$$

$$\frac{\Psi; \Delta; \Gamma \vdash e[c'/\beta] : \tau \text{ exp} \quad \Delta \vdash c \equiv \text{fold}_{\mu j. \kappa} c' : \mu j. \kappa}{\Psi; \Delta; \Gamma \vdash \text{let}(\text{fold } \beta) = c \text{ in } e : \tau \text{ exp}}$$

$$\frac{\begin{array}{l} \Delta, \alpha : \kappa_1 + \kappa_2, \Delta' \vdash c \equiv \alpha : \kappa_1 + \kappa_2 \quad \alpha_1, \alpha_2 \notin \Delta, \Delta' \\ \Psi; \Delta, \beta : \kappa_i, \Delta'; \Gamma[\text{inj}_i \beta / \alpha] \vdash e[\text{inj}_i \beta / \alpha] : \tau[\text{inj}_i \beta / \alpha] \text{ exp} \\ \Psi; \Delta, \beta : \kappa_j, \Delta'; \Gamma[\text{inj}_j \beta / \alpha] \vdash sv[\text{inj}_j \beta / \alpha] : \text{Void} \quad j \in 1 \dots i-1, i+1 \dots n \end{array}}{\Psi; \Delta, \alpha : +[\kappa_1, \dots, \kappa_n], \Delta' \vdash \text{let}_\tau \text{inj}_i \beta = (c, sv) \text{ in } e : \tau \text{ exp}}$$

$$\frac{\Delta \vdash c \equiv \text{inj}_i^{+[\kappa_1, \dots, \kappa_i, \dots, \kappa_n]} c' : +[\kappa_1, \dots, \kappa_i, \dots, \kappa_n] \quad \Psi; \Delta; \Gamma \vdash e[c'/\beta] : \tau \text{ exp}}{\Psi; \Delta; \Gamma \vdash \text{let}_\tau \text{inj}_i \beta = (c, sv) \text{ in } e : \tau \text{ exp}}$$

Well-formed *hval*

$$\boxed{\Psi \vdash hval : \tau \text{ hval}}$$

$$\frac{\begin{array}{l} \bullet \vdash \kappa_i \text{ ok} \quad (i \in 1 \dots k) \\ \alpha_1 : \kappa_1, \dots, \alpha_k : \kappa_k \vdash \tau_i : \text{T}_{32} \quad (i \in 1 \dots m) \\ \alpha_1 : \kappa_1, \dots, \alpha_k : \kappa_k \vdash \phi_i : \text{T}_{64} \quad (i \in 1 \dots n) \\ \alpha_1 : \kappa_1, \dots, \alpha_k : \kappa_k \vdash \tau : \text{T}_{32} \end{array} \quad \Psi; \alpha_1 : \kappa_1, \dots, \alpha_k : \kappa_k; x_1 : \tau_1, \dots, x_m : \tau_m, z_1 : \phi_1, \dots, z_n : \phi_n \vdash e : \tau \text{ exp}}{\Psi \vdash \text{code}_\tau[\alpha_1 : \kappa_1, \dots, \alpha_k : \kappa_k](x_1 : \tau_1, \dots, x_m : \tau_m)(z_1 : \phi_1, \dots, z_n : \phi_n).e : \forall[\alpha_1 : \kappa_1, \dots, \alpha_k : \kappa_k] \text{Code}(\tau_1, \dots, \tau_m)(\phi_1, \dots, \phi_n)(\tau) \text{ hval}}$$

Well-formed heap

$$\boxed{\Psi \vdash d \text{ ok}}$$

$$\vdash \Psi \text{ ok}$$

$$\Psi \vdash \epsilon \text{ ok}$$

$$\frac{\begin{array}{l} \Psi[\ell : \tau] \vdash hval : \tau \text{ hval} \\ \Psi[\ell : \tau] \vdash d \text{ ok} \end{array}}{\Psi[\ell : \tau] \vdash d, \ell : \tau \mapsto hval \text{ ok}}$$

Well-formed program

$$\boxed{\vdash p : \tau}$$

$$\begin{array}{l} \Psi(\epsilon) \stackrel{\text{def}}{=} \bullet \\ \Psi(d, \ell : \tau \mapsto hval) \stackrel{\text{def}}{=} \Psi(d), \ell : \tau \end{array}$$

$$\frac{\Psi, \Psi(d) \vdash d \text{ ok} \quad \Psi, \Psi(d); \bullet; \bullet \vdash e : \tau \text{ exp}}{\Psi \vdash \text{letrec } d \text{ in } e : \tau}$$

Appendix C

TILTAL static and dynamic semantics

C.1 TILTAL Static semantics

Notation

□

$$\begin{array}{ll}
 (\tau \triangleright_{32} \sigma)[0]_{32} \stackrel{\text{def}}{=} \tau & (\phi \triangleright_{64} \sigma)[0]_{64} \stackrel{\text{def}}{=} \phi \\
 (\tau \triangleright_{32} \sigma)[n+1]_{32} \stackrel{\text{def}}{=} \sigma[n]_{32} & (\tau \triangleright_{32} \sigma)[n+1]_{64} \stackrel{\text{def}}{=} \sigma[n]_{64} \\
 (\phi \triangleright_{64} \sigma)[n+2]_{32} \stackrel{\text{def}}{=} \sigma[n]_{32} & (\phi \triangleright_{64} \sigma)[n+2]_{64} \stackrel{\text{def}}{=} \sigma[n]_{64}
 \end{array}$$

$$\begin{array}{ll}
 (\tau \triangleright_{32} \sigma)[0]_{32} \leftarrow \tau' & \stackrel{\text{def}}{=} \tau' \triangleright_{32} \sigma \\
 (\phi \triangleright_{64} \sigma)[0]_{32} \leftarrow \tau' & \stackrel{\text{def}}{=} \tau' \triangleright_{32} ns_{32} \triangleright_{32} \sigma \\
 (\phi \triangleright_{64} \sigma)[1]_{32} \leftarrow \tau' & \stackrel{\text{def}}{=} ns_{32} \triangleright_{32} \tau' \triangleright_{32} \sigma \\
 (\tau \triangleright_{32} \sigma)[n+1]_{32} \leftarrow \tau' & \stackrel{\text{def}}{=} (\sigma)[n]_{32} \leftarrow \tau' \\
 (\phi \triangleright_{64} \sigma)[n+2]_{32} \leftarrow \tau' & \stackrel{\text{def}}{=} (\sigma)[n]_{32} \leftarrow \tau'
 \end{array}$$

$$\begin{array}{ll}
 (\tau_1 \triangleright_{32} \tau_2 \triangleright_{32} \sigma)[0]_{64} \leftarrow \phi' & \stackrel{\text{def}}{=} \phi' \triangleright_{64} \sigma \\
 (\tau_1 \triangleright_{32} \phi \triangleright_{64} \sigma)[0]_{64} \leftarrow \phi' & \stackrel{\text{def}}{=} \phi' \triangleright_{64} ns_{32} \triangleright_{32} \sigma \\
 (\phi \triangleright_{64} \sigma)[0]_{64} \leftarrow \phi' & \stackrel{\text{def}}{=} \phi' \triangleright_{64} \sigma \\
 (\phi \triangleright_{64} \sigma)[1]_{64} \leftarrow \phi' & \stackrel{\text{def}}{=} ns_{32} \triangleright_{32} (ns_{32} \triangleright_{32} \sigma)[0]_{64} \leftarrow \phi' \\
 (\tau \triangleright_{32} \sigma)[n+1]_{64} \leftarrow \phi' & \stackrel{\text{def}}{=} (\sigma)[n]_{64} \leftarrow \phi' \\
 (\phi \triangleright_{64} \sigma)[n+2]_{64} \leftarrow \phi' & \stackrel{\text{def}}{=} (\sigma)[n]_{64} \leftarrow \phi'
 \end{array}$$

$$\begin{array}{ll}
 |\epsilon| & \stackrel{\text{def}}{=} 0 \\
 |\tau \triangleright_{32} \sigma| & \stackrel{\text{def}}{=} 1 + |\sigma| \\
 |\phi \triangleright_{64} \sigma| & \stackrel{\text{def}}{=} 2 + |\sigma|
 \end{array}$$

New well-formed kind rules

$\Delta \vdash \kappa \text{ ok}$

$$\frac{}{\Delta \vdash ST \text{ ok}}$$

Well-formed Constructor

$\Delta \vdash c : \kappa$

Constants and their kinds					
$\epsilon : ST$	$\triangleright_{32} : T_{32} \rightarrow ST \rightarrow ST$	$\triangleright_{64} : T_{64} \rightarrow ST \rightarrow ST$			
	$\circ : ST \rightarrow ST \rightarrow ST$	$\text{sptr} : ST \rightarrow T_{32}$			
$ns_{64} : T_{64}$	$ns_{32} : T_{32}$	$\bigvee : T_{32} \text{list} \rightarrow T_{32}$			

$$\Delta \vdash \Gamma \text{ ok}$$

$$\frac{}{\Delta \vdash \Gamma \rightarrow 0 : T_{32}}$$

Well-formed coercion

$\Delta \vdash q : \tau \Rightarrow \tau'$

$$\Delta \vdash \tau \equiv \text{rec}[\kappa](c)(c_p) : T_{32}$$

$$\frac{}{\Delta \vdash \text{roll}_\tau : (c(\text{Rec}[\kappa]c)c_p) \Rightarrow \tau}$$

$$\Delta \vdash \tau \equiv \text{rec}[\kappa](c)(c_p) : T_{32}$$

$$\frac{}{\Delta \vdash \text{unroll}_\tau w : (c(\text{Rec}[\kappa]c)c_p) \Rightarrow \tau}$$

$$\Delta \vdash c \equiv \bigvee(\tau_0 :: \dots :: \tau_i :: \dots :: \text{nil}) : T_{32} \text{list}$$

$$\frac{}{\Delta \vdash \text{inj_union}_{(i,c)} : \tau_i \Rightarrow c}$$

$$\Delta \vdash \tau \equiv \exists[\kappa](c') : T_{32} \quad \Delta \vdash c : \kappa$$

$$\frac{}{\Delta \vdash \text{pack}[\tau]c : (c' c) \Rightarrow \tau}$$

$$\Delta \vdash \tau : T_{32}$$

$$\frac{}{\Delta \vdash \text{inj_dyn}_\tau : (\times [\text{Dyntag}(\tau), \tau]) \Rightarrow \text{Dyn}}$$

Well-formed stack

$\Psi; \Delta \vdash s : \sigma$

$\vdash \Delta \text{ ok}$	$\Psi; \Delta \vdash w : \tau$	$\Psi; \Delta \vdash l : \phi$
	$\Psi; \Delta \vdash s : \sigma$	$\Psi; \Delta \vdash s : \sigma$
$\Psi; \Delta \vdash \epsilon : \epsilon$	$\Psi; \Delta \vdash w \triangleright_{32} s : \tau \triangleright_{32} \sigma$	$\Psi; \Delta \vdash l \triangleright_{32} s : \phi \triangleright_{64} \sigma$

Well-formed 64-bit value

$\Psi; \Delta \vdash l : \phi$

$$\frac{\vdash \Psi \text{ ok} \quad \vdash \Delta \text{ ok}}{\Psi; \Delta \vdash ns_{64} : ns_{64}}$$

$$\frac{\vdash \Psi \text{ ok} \quad \vdash \Delta \text{ ok}}{\Psi; \Delta \vdash \mathbf{r} : \text{Float}} \quad \frac{\vdash \Psi, \ell : \phi, \Psi' \text{ ok} \quad \vdash \Delta \text{ ok}}{\Psi, \ell : \phi, \Psi'; \Delta \vdash \ell : \phi}$$

Well-formed 32-bit value

$\Psi; \Delta \vdash w : \tau$

$$\frac{\vdash \Psi, \ell : \tau, \Psi' \text{ ok} \quad \vdash \Delta \text{ ok}}{\Psi, \ell : \tau \Psi'; \Delta \vdash \ell : \tau} \quad \frac{\vdash \Psi \text{ ok} \quad \vdash \Delta \text{ ok}}{\Psi; \Delta \vdash i : \text{Int}} \quad \frac{\vdash \Psi \text{ ok} \quad \vdash \Delta \text{ ok}}{\Psi; \Delta \vdash ns_{32} : ns_{32}}$$

$$\frac{\vdash \Psi \text{ ok} \quad \vdash \Delta \text{ ok}}{\Psi; \Delta \vdash \text{tag}_i : \text{Tag}(\bar{i})} \quad \frac{\Psi; \Delta \vdash w : \forall[\kappa](c') \quad \Delta \vdash c : \kappa}{\Psi; \Delta \vdash w[c] : c'c}$$

$$\frac{\Psi; \Delta \vdash \sigma \equiv \sigma' : ST \quad (|\sigma'| = i)}{\Psi; \Delta \vdash \text{sptr}(i) : \text{sptr}(\sigma)} \quad \frac{\Delta \vdash q : \tau_1 \Rightarrow \tau_2 \quad \Psi; \Delta \vdash w : \tau_1}{\Psi; \Delta \vdash q w : \tau_2}$$

Well-formed 64-bit operand

$\Psi; \Delta; \Gamma \vdash fv : \phi$

$$\frac{\Gamma(f) = \phi}{\Psi; \Delta; \Gamma \vdash f : \phi} \quad \frac{\Psi; \Delta \vdash l : \phi}{\Psi; \Delta; \Gamma \vdash l : \phi} \quad \frac{\Gamma(\mathbf{sp}) = \sigma \quad \sigma[i]_{64} = \phi}{\Psi; \Delta; \Gamma \vdash \mathbf{sp}(i) : \phi}$$

Well-formed 32-bit operand

$\Psi; \Delta; \Gamma \vdash sv : \tau$

$$\frac{\Gamma(r) = \tau}{\Psi; \Delta; \Gamma \vdash r : \tau} \quad \frac{\Psi; \Delta \vdash w : \tau}{\Psi; \Delta; \Gamma \vdash w : \tau} \quad \frac{\Gamma(\mathbf{sp}) = \sigma \quad \sigma[i]_{32} = \tau}{\Psi; \Delta; \Gamma \vdash \mathbf{sp}(i) : \tau}$$

$$\frac{\Psi; \Delta; \Gamma \vdash sv : \forall[\kappa](c') \quad \Delta \vdash c : \kappa}{\Psi; \Delta; \Gamma \vdash sv[c] : c'c} \quad \frac{\Delta \vdash q : \tau_1 \Rightarrow \tau_2 \quad \Psi; \Delta; \Gamma \vdash sv : \tau_1}{\Psi; \Delta; \Gamma \vdash q sv : \tau_2}$$

$$\begin{array}{c}
\frac{\Psi; \Delta; \Gamma \vdash sv : \tau}{\Psi; \Delta; \Gamma \vdash \text{mov } \mathbf{r}, sv \Rightarrow \Gamma\{\mathbf{r}:\tau\}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash \mathbf{r}_s : \times [\tau_0, \dots, \tau_i, \dots, \tau_n]}{\Psi; \Delta; \Gamma \vdash \text{loadr } \mathbf{r}_d, \mathbf{r}_s(i) \Rightarrow \Gamma\{\mathbf{r}_d:\tau_i\}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv : \tau_i \quad \Psi; \Delta; \Gamma \vdash \mathbf{r}_s : \times [\tau_0, \dots, \tau_i, \dots, \tau_n]}{\Psi; \Delta; \Gamma \vdash \text{store } \mathbf{r}_s(i), sv \Rightarrow \Gamma} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv_i : \tau_i}{\Psi; \Delta; \Gamma \vdash \text{mallocr}[\tau_1, \dots, \tau_n] \langle sv_1, \dots, sv_n \rangle \Rightarrow \Gamma\{\mathbf{r} : \times [\tau_1, \dots, \tau_n]\}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv_1 : \text{Int} \quad \Psi; \Delta; \Gamma \vdash sv_2 : \tau}{\Psi; \Delta; \Gamma \vdash \text{malloc}_\tau \mathbf{r}(sv_1, sv_2) \Rightarrow \Gamma\{\mathbf{r}:\text{Array}_{32}(\tau)\}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv : \text{Int} \quad \Psi; \Delta; \Gamma \vdash fv : \phi}{\Psi; \Delta; \Gamma \vdash \text{malloc}_\phi \mathbf{r}(sv, fv) \Rightarrow \Gamma\{\mathbf{r}:\text{Array}_{64}(\phi)\}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash fv : \phi}{\Psi; \Delta; \Gamma \vdash \text{malloc}_\phi \mathbf{r}, fv \Rightarrow \Gamma\{\mathbf{r}:\text{Boxed}(\phi)\}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv : \Gamma\{\mathbf{sp}:(\Gamma_{\text{ret}} \rightarrow 0) \triangleright_{32} \sigma\} \rightarrow 0 \quad (\Gamma(\mathbf{sp}) = \sigma)}{\Psi; \Delta; \Gamma \vdash \text{call } sv \Rightarrow \Gamma_{\text{ret}}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv : \Gamma\{\mathbf{r}:\text{Tag}(\bar{i})\} \rightarrow 0 \quad \Delta \vdash \Gamma(\mathbf{r}) \equiv \bigvee (\tau_0 :: \dots :: \tau_{k-1} :: \text{Tag}(\bar{i}) :: c) : \text{T}_{32}}{\Psi; \Delta; \Gamma \vdash \text{brtag}_i \mathbf{r}, sv \Rightarrow \Gamma\{\mathbf{r}:\bigvee (\tau_0 :: \dots :: \tau_{k-1} :: c)\}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv : \Gamma\{\mathbf{r}:\times [\text{Tag}(\bar{i}), \tau]\} \rightarrow 0 \quad \Delta \vdash \Gamma(\mathbf{r}) \equiv \bigvee (\tau_0 :: \dots :: \tau_{k-1} :: \tau_k :: \dots \tau_n :: (\times [\text{Tag}(\bar{i}), \tau]) :: c) : \text{T}_{32}}{\Psi; \Delta; \Gamma \vdash \text{brtgd}_i \mathbf{r}, sv \Rightarrow \Gamma\{\mathbf{r}:\bigvee (\tau_0 :: \dots :: \tau_n :: c)\}}
\end{array}$$

$$\frac{\Gamma(wreg) = \text{Dyn} \quad \Psi; \Delta; \Gamma \vdash sv_1 : \text{Dyn}(\tau) \quad \Psi; \Delta; \Gamma \vdash sv_2 : \Gamma\{\mathbf{r} : \times [\text{Dyn}(\tau), \tau]\} \rightarrow 0}{\Psi; \Delta; \Gamma \vdash \text{brdyn } \mathbf{r}, sv_1, sv_2 \Rightarrow \Gamma}$$

$$\frac{\vdash \Psi \text{ ok} \quad \Delta \vdash \tau : T_{32}}{\Psi; \Delta; \Gamma \vdash \text{dyntag}_\tau \mathbf{r} \Rightarrow \Gamma\{\mathbf{r} : \text{Dyn}(\tau)\}}$$

$$\frac{\Psi; \Delta; \Gamma \vdash sv : \tau \quad \Gamma(\mathbf{sp}) = \sigma \quad \Delta \vdash \sigma \equiv \sigma' : ST \quad (\sigma')[i]_{32} \leftarrow \tau = \sigma''}{\Psi; \Delta; \Gamma \vdash \text{swrite } \mathbf{sp}(i), sv \Rightarrow \Gamma\{\mathbf{sp} : \sigma''\}}$$

$$\frac{\vdash \Psi \text{ ok} \quad \Delta \vdash \Gamma \text{ ok}}{\Psi; \Delta; \Gamma \vdash \text{salloc } n \Rightarrow \Gamma\{\mathbf{sp} : \underbrace{ns_{32} \triangleright_{32} \cdots \triangleright_{32} ns_{32}}_n \triangleright_{32} \Gamma(\mathbf{sp})\}}$$

$$\frac{\Gamma(\mathbf{sp}) = \sigma \quad \Delta \vdash \sigma \equiv \sigma_1 \circ \sigma_2 : ST \quad |\sigma_1| = n}{\Psi; \Delta; \Gamma \vdash \text{sfree } n \Rightarrow \Gamma\{\mathbf{sp} : \sigma_2\}}$$

$$\frac{\Gamma(\mathbf{sp}) = \sigma \quad \Delta \vdash \sigma \equiv \sigma_1 \circ (\phi \triangleright_{64} \sigma_2) : ST \quad |\sigma_1| = (n-1)}{\Psi; \Delta; \Gamma \vdash \text{sfree } n \Rightarrow \Gamma\{\mathbf{sp} : ns_{32} \triangleright_{32} \sigma_2\}}$$

$$\frac{\Gamma(\mathbf{sp}) = \sigma}{\Psi; \Delta; \Gamma \vdash \text{mov } \mathbf{r}, \mathbf{sp} \Rightarrow \Gamma\{\mathbf{r} : \text{sptr}(\sigma)\}}$$

$$\frac{\Psi; \Delta; \Gamma \vdash sv : \text{sptr}(\sigma_1) \quad \Delta \vdash \Gamma(\mathbf{sp}) \equiv \sigma_2 \circ \sigma_2 : ST}{\Psi; \Delta; \Gamma \vdash \text{mov } \mathbf{sp}, sv \Rightarrow \Gamma\{\mathbf{sp} : \sigma_2\}}$$

$$\frac{\Psi; \Delta; \Gamma \vdash sv_1 : \text{Array}_{32}(\tau) \quad \Psi; \Delta; \Gamma \vdash sv_2 : \text{Int}}{\Psi; \Delta; \Gamma \vdash \text{sub}_\tau \mathbf{r}, sv_1, sv_2 \Rightarrow \Gamma\{\mathbf{r} : \tau\}}$$

$$\frac{\Psi; \Delta; \Gamma \vdash sv_1 : \text{Array}_{32}(\tau) \quad \Psi; \Delta; \Gamma \vdash sv_2 : \text{Int} \quad \Psi; \Delta; \Gamma \vdash sv_3 : \tau}{\Psi; \Delta; \Gamma \vdash \text{upd}_\tau sv_1, sv_2, sv_3 \Rightarrow \Gamma}$$

$$\begin{array}{c}
\frac{\Psi; \Delta; \Gamma \vdash fv : \phi}{\Psi; \Delta; \Gamma \vdash \text{fmov } \mathbf{f}, fv \Rightarrow \Gamma\{\mathbf{f}:\phi\}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash \mathbf{r}_s : \text{Boxed}(\phi)}{\Psi; \Delta; \Gamma \vdash \text{floadr } \mathbf{f}_d, \mathbf{r}_s \Rightarrow \Gamma\{\mathbf{f}_d:\phi\}} \quad \frac{\Psi; \Delta; \Gamma \vdash fv : \phi \quad \Psi; \Delta; \Gamma \vdash \mathbf{r} : \text{Boxed}(\phi)}{\Psi; \Delta; \Gamma \vdash \text{store } \mathbf{r}, fv \Rightarrow \Gamma} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash fv : \phi \quad \Gamma(\mathbf{sp}) = \sigma \quad \Delta \vdash \sigma \equiv \sigma' : ST \quad (\sigma')[i]_{64} \leftarrow \phi = \sigma''}{\Psi; \Delta; \Gamma \vdash \text{fswrite } \mathbf{sp}(i), fv \Rightarrow \Gamma\{\mathbf{sp}:\sigma''\}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv_1 : \text{Array}_{64}(\tau) \quad \Psi; \Delta; \Gamma \vdash sv_2 : \text{Int}}{\Psi; \Delta; \Gamma \vdash \text{sub}_\phi \mathbf{f}, sv_1, sv_2 \Rightarrow \Gamma\{\mathbf{f}:\phi\}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv_1 : \text{Array}_{64}(\tau) \quad \Psi; \Delta; \Gamma \vdash sv_2 : \text{Int} \quad \Psi; \Delta; \Gamma \vdash fv : \phi}{\Psi; \Delta; \Gamma \vdash \text{upd}_\phi sv_1, sv_2, fv \Rightarrow \Gamma}
\end{array}$$

Well-formed Instruction Sequence

$\Psi; \Delta; \Gamma \vdash I : \tau$

$$\begin{array}{c}
\frac{\Gamma(\mathbf{sp}) = (\Gamma_{\text{ret}} \rightarrow 0) \triangleright_{32} \sigma \quad \Gamma_{\text{ret}} = \Gamma\{\mathbf{sp}:\sigma\}}{\Psi; \Delta; \Gamma \vdash \text{ret } \mathbf{ok}} \quad \frac{\Psi; \Delta; \Gamma \vdash sv : \Gamma \rightarrow 0}{\Psi; \Delta; \Gamma \vdash \text{jmp } sv \mathbf{ok}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash \mathbf{r}_t : \tau}{\Psi; \Delta; \Gamma \vdash \text{halt}_\tau \mathbf{ok}} \quad \frac{\Psi; \Delta; \Gamma \vdash i \Rightarrow \Gamma' \quad \Psi; \Delta; \Gamma' \vdash I \mathbf{ok}}{\Psi; \Delta; \Gamma \vdash i; I \mathbf{ok}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash sv : \exists[\kappa](c) \quad \Psi; \Delta, \alpha:\kappa; \Gamma\{\mathbf{r}:(c\alpha)\} \vdash I \mathbf{ok}}{\Psi; \Delta; \Gamma \vdash \text{unpack}[\alpha, \mathbf{r}], sv; I \mathbf{ok}} \quad \alpha \notin \Delta \\
\\
\frac{\Psi[\text{inj}_1^{\kappa_1+\kappa_2} \alpha_1/\alpha]; \Delta, \alpha_1:\kappa_1, \Delta'; \Gamma[\text{inj}_1^{\kappa_1+\kappa_2} \alpha_1/\alpha] \vdash sv[\text{inj}_1^{\kappa_1+\kappa_2} \alpha_1/\alpha] : \text{Void} \quad \Psi[\text{inj}_2^{\kappa_1+\kappa_2} \alpha_2/\alpha]; \Delta, \alpha_2:\kappa_2, \Delta'; \Gamma[\text{inj}_2^{\kappa_1+\kappa_2} \alpha_2/\alpha] \vdash I[\text{inj}_2^{\kappa_1+\kappa_2} \alpha_2/\alpha] \mathbf{ok} \quad \Delta, \alpha:\kappa_1 + \kappa_2, \Delta'; \Gamma \vdash c \equiv \alpha : \kappa_1 + \kappa_2 \quad \alpha_1, \alpha_2 \notin \Delta, \Delta'}{\Psi; \Delta, \alpha:\kappa_1 + \kappa_2, \Delta'; \Gamma \vdash \text{vcase}[\alpha_1. \text{dead } sv, \alpha_2]c; I \mathbf{ok}}
\end{array}$$

$$\frac{\Psi[\text{inj}_1^{\kappa_1+\kappa_2} \alpha_1/\alpha]; \Delta, \alpha_1:\kappa_1, \Delta'; \Gamma[\text{inj}_1^{\kappa_1+\kappa_2} \alpha_1/\alpha] \vdash I[\text{inj}_1^{\kappa_1+\kappa_2} \alpha_1/\alpha] \text{ ok} \quad \Psi[\text{inj}_1^{\kappa_1+\kappa_2} \alpha_1/\alpha]; \Delta, \alpha_2:\kappa_2, \Delta'; \Gamma[\text{inj}_2^{\kappa_1+\kappa_2} \alpha_2/\alpha] \vdash sv[\text{inj}_2^{\kappa_1+\kappa_2} \alpha_2/\alpha] : \text{Void} \quad \Delta, \alpha:\kappa_1 + \kappa_2, \Delta' \vdash c \equiv \alpha : \kappa_1 + \kappa_2 \quad \alpha_1, \alpha_2 \notin \Delta, \Delta'}{\Psi; \Delta, \alpha:\kappa_1 + \kappa_2, \Delta' \vdash \text{vcase}[\alpha_1, \alpha_2. \text{dead } sv]c; I \text{ ok}}$$

$$\Delta \vdash c \equiv \text{inj}_2^{\kappa_1+\kappa_2} c' : \kappa_1 + \kappa_2 \quad \Psi; \Delta; \Gamma \vdash I[c'/\alpha_2] \text{ ok}$$

$$\Psi; \Delta; \Gamma \vdash \text{vcase}[\alpha_1. \text{dead } sv, \alpha_2]c; I) \text{ ok}$$

$$\Delta \vdash c \equiv \text{inj}_1^{\kappa_1+\kappa_2} c' : \kappa_1 + \kappa_2 \quad \Psi; \Delta; \Gamma \vdash I[c'/\alpha_1] \text{ ok}$$

$$\Psi; \Delta; \Gamma \vdash \text{vcase}[\alpha_1, \alpha_2. sv]c; I \text{ ok}$$

$$\frac{\Psi[\langle \beta, \gamma \rangle / \alpha]; \Delta, \beta:\kappa_1, \gamma:\kappa_2, \Delta'; \Gamma[\langle \beta, \gamma \rangle / \alpha] \vdash I[\langle \beta, \gamma \rangle / \alpha] \text{ ok} \quad \Delta, \alpha:\kappa_1 \times \kappa_2, \Delta' \vdash c \equiv \alpha : \kappa_1 \times \kappa_2}{\Psi; \Delta, \alpha:\kappa_1 \times \kappa_2, \Delta'; \Gamma \vdash \text{refine}[\langle \beta, \gamma \rangle] c; I \text{ ok}}$$

$$\frac{\Psi; \Delta; \Gamma \vdash I[c_1, c_2/\beta, \gamma] \text{ ok} \quad \Delta \vdash c \equiv \langle c_1, c_2 \rangle : \kappa_1 \times \kappa_2}{\Psi; \Delta; \Gamma \vdash \text{refine}[\langle \beta, \gamma \rangle] c; I \text{ ok}}$$

$$\frac{\Psi[\text{fold}_{\mu j. \kappa} \beta/\alpha]; \Delta, \beta:\kappa[\mu j. \kappa/j], \Delta'; \Gamma[\text{fold}_{\mu j. \kappa} \beta/\alpha] \vdash I[\text{fold}_{\mu j. \kappa} \beta/\alpha] \text{ ok} \quad \Delta, \alpha:\mu j. \kappa, \Delta' \vdash c \equiv \alpha : \mu j. \kappa}{\Psi; \Delta, \alpha:\mu j. \kappa, \Delta'; \Gamma \vdash \text{refine}[\text{fold } \beta] c; I \text{ ok}}$$

$$\frac{\Psi; \Delta; \Gamma \vdash I[c'/\beta] \text{ ok} \quad \Delta \vdash c \equiv \text{fold}_{\mu j. \kappa} c' : \mu j. \kappa}{\Psi; \Delta; \Gamma \vdash \text{refine}[\text{fold } \beta] c; I \text{ ok}}$$

Well-formed *hval*

$\Psi \vdash hval : \tau \text{ hval}$

$$\frac{\Psi; \bullet \vdash w_i : \tau_i \quad i = 1 \dots n}{\Psi \vdash \langle w_1, \dots, w_n \rangle : \times [\tau_1, \dots, \tau_n] \text{ hval}}$$

$$\frac{\Psi; \bullet \vdash w_i : \tau \quad i = 1 \dots n}{\Psi \vdash [w_1, \dots, w_n] : \text{Array}_{32} \tau \text{ hval}}$$

$$\frac{\Psi; \bullet \vdash l : \phi}{\Psi \vdash l : \text{Boxed } \phi \text{ hval}} \quad \frac{\Psi; \bullet \vdash l_i : \phi \quad i = 1 \dots n}{\Psi \vdash [l_1, \dots, l_n] : \text{Array}_{64} \phi \text{ hval}}$$

$$\frac{\bullet \vdash \tau : \text{T}_{32}}{\Psi; \bullet \vdash \text{dtag} : \text{Dyntag}_\tau \text{ hval}}$$

$$\frac{\alpha_0 : \kappa_0, \dots, \alpha_{i-1} : \kappa_{i-1} \vdash \kappa_i \text{ ok} \quad \vec{\alpha} : \vec{\kappa} \vdash \Gamma \text{ ok} \quad \Psi; \vec{\alpha} : \vec{\kappa}; \Gamma \vdash I \text{ ok}}{\Psi \vdash \text{code}_\tau[\vec{\alpha} : \vec{\kappa}] \Gamma. I : \forall[\vec{\alpha} : \vec{\kappa}] \Gamma \rightarrow 0 \text{ hval}}$$

Well-formed Heap

$$\boxed{\vdash H : \Psi}$$

$$\frac{\vdash \Psi \text{ ok} \quad \Psi \vdash \text{hval}_i : \Psi(\ell_i) \text{ hval}}{\Psi \vdash \{\ell_1 : \tau_1 \mapsto \text{hval}_1, \dots, \ell_n : \tau_n \mapsto \text{hval}_n\} \text{ ok}} \quad \Psi = \{\ell_1 : \tau_1, \dots, \ell_n : \tau_n\}$$

Well-formed Register File

$$\boxed{\Psi \vdash R : \Gamma}$$

$$\frac{\begin{array}{c} \Psi; \bullet \vdash w_1 : \tau_1 \quad \Psi; \bullet \vdash w_2 : \tau_2 \quad \Psi; \bullet \vdash w_e : \tau_e \quad \Psi; \bullet \vdash w_t : \tau_t \\ \Psi; \bullet \vdash l_1 : \phi_1 \quad \Psi; \bullet \vdash l_2 : \phi_2 \quad \Psi; \bullet \vdash s : \sigma \end{array}}{\Psi \vdash \{\mathbf{r}_1 \mapsto w_1, \mathbf{r}_2 \mapsto w_2, \mathbf{r}_e \mapsto w_e, \mathbf{r}_t \mapsto w_t, \mathbf{f}_1 \mapsto l_1, \mathbf{f}_2 \mapsto l_2 : \mathbf{sp} \mapsto s\} : \{\mathbf{r}_1 : \tau_1, \mathbf{r}_2 : \tau_2, \mathbf{r}_e : \tau_e, \mathbf{r}_t : \tau_t, \mathbf{f}_1 : \phi_1, \mathbf{f}_2 : \phi_2, \mathbf{sp} : \sigma\}}$$

Well-formed Program

$$\boxed{\vdash (H, R, I) \text{ ok}}$$

$$\frac{\vdash H : \Psi \quad \Psi \vdash R : \Gamma \quad \Psi(H; \bullet; \Gamma \vdash I \text{ ok})}{\vdash (H, R, I) \text{ ok}}$$

C.2 TILTAL dynamic semantics

C.2.1 Definitions

$$\begin{array}{lll} (w \triangleright_{32} s)[0]_{32} & \stackrel{\text{def}}{=} w & (l \triangleright_{64} s)[0]_{64} \stackrel{\text{def}}{=} l \\ (w \triangleright_{32} s)[n+1]_{32} & \stackrel{\text{def}}{=} s[n]_{32} & (w \triangleright_{32} s)[n+1]_{64} \stackrel{\text{def}}{=} s[n]_{64} \\ (l \triangleright_{64} s)[n+2]_{32} & \stackrel{\text{def}}{=} s[n]_{32} & (l \triangleright_{64} s)[n+2]_{64} \stackrel{\text{def}}{=} s[n]_{64} \end{array}$$

$$\begin{aligned}
(w \triangleright_{32} s)[0]_{32} \leftarrow w' & \stackrel{\text{def}}{=} w' \triangleright_{32} s \\
(l \triangleright_{64} s)[0]_{32} \leftarrow w' & \stackrel{\text{def}}{=} w' \triangleright_{32} ns_{32} \triangleright_{32} s \\
(l \triangleright_{64} s)[1]_{32} \leftarrow w' & \stackrel{\text{def}}{=} ns_{32} \triangleright_{32} w' \triangleright_{32} s \\
(w \triangleright_{32} s)[n+1]_{32} \leftarrow w' & \stackrel{\text{def}}{=} (s)[n]_{32} \leftarrow w' \\
(l \triangleright_{64} s)[n+2]_{32} \leftarrow w' & \stackrel{\text{def}}{=} (s)[n]_{32} \leftarrow w'
\end{aligned}$$

$$\begin{aligned}
(w_1 \triangleright_{32} w_2 \triangleright_{32} s)[0]_{64} \leftarrow l' & \stackrel{\text{def}}{=} l' \triangleright_{64} s \\
(w_1 \triangleright_{32} l \triangleright_{64} s)[0]_{64} \leftarrow l' & \stackrel{\text{def}}{=} l' \triangleright_{64} ns_{32} \triangleright_{32} s \\
(l \triangleright_{64} s)[0]_{64} \leftarrow l' & \stackrel{\text{def}}{=} l' \triangleright_{64} s \\
(l \triangleright_{64} s)[1]_{64} \leftarrow l' & \stackrel{\text{def}}{=} ns_{32} \triangleright_{32} (ns_{32} \triangleright_{32} s)[0]_{64} \leftarrow l' \\
(w \triangleright_{32} s)[n+1]_{64} \leftarrow l' & \stackrel{\text{def}}{=} (s)[n]_{64} \leftarrow l' \\
(l \triangleright_{64} s)[n+2]_{64} \leftarrow l' & \stackrel{\text{def}}{=} (s)[n]_{64} \leftarrow l'
\end{aligned}$$

$$\hat{R}(sv) \stackrel{\text{def}}{=} \begin{cases} R(r) & \text{when } sv = r \\ R(\mathbf{sp})[i]_{32} & \text{when } sv = \mathbf{sp}(i) \\ w & \text{when } sv = w \\ \mathbf{tag}_i & \text{when } sv = \mathbf{tag}_i \\ \hat{R}(sv')[\vec{c}] & \text{when } sv = sv'[\vec{c}] \\ q \hat{R}(sv') & \text{when } sv = q \ sv' \\ \mathbf{pack} \hat{R}(sv') \text{ as } \tau \text{ hiding } c & \text{when } sv = \mathbf{pack} \ sv' \text{ as } \tau \text{ hiding } c \end{cases}$$

$$\hat{R}(fv) \stackrel{\text{def}}{=} \begin{cases} R(f) & \text{when } fv = f \\ R(\mathbf{sp})[i]_{64} & \text{when } fv = \mathbf{sp}(i) \end{cases}$$

C.2.2 Transitions

$(H, R, I) \mapsto P$ where	
if $I =$	then $P =$
mov $\mathbf{r}, sv; I'$	$(H, R\{\mathbf{r} \mapsto \hat{R}(sv)\}, I')$
loadr $\mathbf{r}_d, \mathbf{r}_s(i); I'$	$(H, R\{\mathbf{r} \mapsto w_i\}, I')$ where $R(\mathbf{r}_d) = \ell$ and $H(\ell) = \langle w_0, \dots, w_{n-1} \rangle$
store $\mathbf{r}_d(i), sv; I'$	$(H\{\ell \mapsto \langle w_0, \dots, \hat{R}(sv), \dots, w_{n-1} \rangle\}, R, I')$ where $R(\mathbf{r}_d) = \ell$ and $H(\ell) = \langle w_0, \dots, w_i, \dots, w_{n-1} \rangle$
mallocr $[\tau_1, \dots, \tau_n] \langle sv_0, \dots, sv_{n-1} \rangle; I'$	$(H\{\ell \mapsto \langle \hat{R}(sv_0), \dots, \hat{R}(sv_{n-1}) \rangle\}, R\{\mathbf{r}_d \mapsto \ell\}, I')$ where $\ell \notin H$
malloc_{τ} $\mathbf{r}_d(sv_1, sv_2); I'$	$(H\{\ell \mapsto \underbrace{[\hat{R}(sv_2), \dots, \hat{R}(sv_2)]}_{\hat{R}(sv_1)}\}, R\{\mathbf{r}_d \mapsto \ell\}, I')$ where $\ell \notin H$
malloc_{ϕ} $\mathbf{r}, fv; I'$	$(H\{\ell \mapsto \hat{R}(fv)\}, R\{\mathbf{r}_d \mapsto \ell\}, I')$ where $\ell \notin H$
malloc_{ϕ} $\mathbf{r}(sv, fv); I'$	$(H\{\ell \mapsto \underbrace{[\hat{R}(fv), \dots, \hat{R}(fv)]}_{\hat{R}(sv)}\}, R\{\mathbf{r}_d \mapsto \ell\}, I')$ where $\ell \notin H$
dyntag_c $\mathbf{r}; I'$	$(H\{\ell \mapsto \mathbf{exn_tag}\}, R\{\mathbf{r} \mapsto \ell\}, I')$ where $\ell \notin H$
swrite $\mathbf{sp}(i), sv; I'$	$(H, R\{\mathbf{sp} \mapsto s'\}, I')$ where $s' = (R(\mathbf{sp}))[i]_{32} \leftarrow \hat{R}(sv)$
salloc $n; I'$	$(H, R\{\mathbf{sp} \mapsto \underbrace{ns_{32} \triangleright_{32} \dots \triangleright_{32} ns_{32}}_n \triangleright_{32} R(\mathbf{sp})\}, I')$
sfree $n; I'$	$(H, R\{\mathbf{sp} \mapsto s'\}, I')$ where $\mathbf{pop}(R(\mathbf{sp}), n) = s'$
mov $\mathbf{r}, \mathbf{sp}; I'$	$(H, R\{\mathbf{r} \mapsto \mathbf{sptr}(R(\mathbf{sp}))\}, I')$
mov $\mathbf{sp}, sv; I'$	$(H, R\{\mathbf{sp} \mapsto s'\}, I')$ where $R(\mathbf{sp}) = s$ and $\hat{R}(sv) = \mathbf{sptr}(j) = s' $ and s' is a suffix of s

Figure C.1: TILTAL transitions (part I)

$(H, R, I) \mapsto P$ where	
call_Γ $sv; I'$	$(H\{\ell \mapsto \text{code}[\Gamma.I']\}, R\{\mathbf{sp} \mapsto \ell \triangleright_{32} \hat{R}(\mathbf{sp})\}, I_d[\vec{c}/\Delta])$ where $\hat{R}(sv) = \ell_d[\vec{c}]$ and $H(\ell_d) = \text{code}[\Delta]\Gamma.I_d$ and $\ell \notin H$
brtag_i $\mathbf{r}, sv; I'$ when $R(r) = \text{inj_union}_{(c_1, c_2)}^i \text{tag}_i$	$(H, R\{\mathbf{r} \mapsto \text{inj_union}_{([\text{Tag}(\bar{i})], \text{nil})}^i \text{tag}_i\}, I_d)$ where $\hat{R}(sv) = \ell$ and $H(\ell) = \text{code}[\Gamma].I_d$
brtag_i $\mathbf{r}, sv; I'$ when $R(r) \neq \text{inj_union}_{(c_1, c_2)}^i \text{tag}_i$	$(H, R\{\mathbf{r} \mapsto \text{inj_union}_{(c'_1, c_2)}^i \text{tag}_i\}, I')$ where c_1 normalizes to $\tau_0 :: \dots :: \tau_k :: \dots \tau_n$ and $\tau_k = \text{Tag}(\bar{i})$ and $c'_1 = \tau_0 :: \dots :: \tau_{k-1} :: \tau_{k+1} :: \dots \tau_n$
brtgd_i $\mathbf{r}, sv; I'$ when $R(r) = \text{inj_union}_{(c_1, c_2)}^i \ell_1$ and $H(\ell_1) = \langle \text{tag}_i, w \rangle$	$(H, R\{\mathbf{r} \mapsto \text{inj_union}_{(\text{nil}, [\times[\text{Tag}(\bar{i}), \tau_w]])}^i \ell_1\}, I_d)$ where $\hat{R}(sv) = \ell$ and $H(\ell) = \text{code}[\Gamma].I_d$
brtgd_i $\mathbf{r}, sv; I'$ when $R(r) = \text{inj_union}_{(c_1, c_2)}^i \text{tag}_k$	$(H, R\{\mathbf{r} \mapsto \text{inj_union}_{(c_1, c'_2)}^i \text{tag}_k\}, I_d)$ where c_2 normalizes to $\tau_0 :: \dots :: \tau_k :: \dots \tau_n$ and $\tau_k = \times[\text{Tag}(\bar{i}), \tau_w]$ and $c'_2 = \tau_0 :: \dots :: \tau_{k-1} :: \tau_{k+1} :: \dots \tau_n$
brtgd_i $\mathbf{r}, sv; I'$ when $R(r) = \text{inj_union}_{(c_1, c_2)}^i \ell$ and $H(\ell_1) = \langle \text{tag}_k, w \rangle$ ($i \neq k$)	$(H, R\{\mathbf{r} \mapsto \text{inj_union}_{(c_1, c'_2)}^i \ell\}, I_d)$ where c_2 normalizes to $\tau_0 :: \dots :: \tau_k :: \dots \tau_n$ and $\tau_k = \times[\text{Tag}(\bar{i}), \tau_w]$ and $c'_2 = \tau_0 :: \dots :: \tau_{k-1} :: \tau_{k+1} :: \dots \tau_n$
ret	$(H, R\{\mathbf{sp} \mapsto s'\}, I[\vec{c}/\Delta])$ where $R(\mathbf{sp}) = \ell[\vec{c}] \triangleright_{32} s'$ and $H(\ell) = \text{code}[\Delta]\Gamma.I$
jmp sv	$(H, R, I[\vec{c}/\Delta])$ where $\hat{R}(sv) = \ell[\vec{c}]$ and $H(\ell) = \text{code}[\Delta]\Gamma.I$

Figure C.2: TILTAL transitions (part II)

$(H, R, I) \mapsto P$ where	
$\text{unpack}[\alpha, \mathbf{r}], sv; I'$	$(H, R\{\mathbf{r} \mapsto w\}, I'[c/\alpha])$ where $\hat{R}(sv) = \text{pack } w \text{ as } \tau \text{ hiding } c$
$\text{vcase}(c, \alpha_1. \text{dead } sv, \alpha_2); I'$	$(H, R, I'[c'/\alpha_2])$ where c normalizes to $\text{inj}_2 c'$
$\text{vcase}(c, \alpha_1, \alpha_2. \text{dead } sv); I'$	$(H, R, I'[c'/\alpha_1])$ where c normalizes to $\text{inj}_1 c'$
$\langle \beta, \gamma \rangle = c; I'$	$(H, R, I'[c_1, c_2/\beta, \gamma])$ where c normalizes to $\langle c_1, c_2 \rangle$
$(\text{fold } \beta) = c; I'$	$(H, R, I'[c'/\beta])$ where c normalizes to $\text{fold}_{\mu_j.\kappa} c'$
$\text{sub}_\tau \mathbf{r}, sv_1, sv_2; I'$	$(H, R\{\mathbf{r} \mapsto w_i\}, I')$ where $\hat{R}(sv_1) = \ell$ and $\hat{R}(sv_2) = i$ and $H(\ell) = [w_0, \dots, w_{n-1}]$ and $0 \leq i < n$
$\text{upd}_\tau sv_1, sv_2, sv_3; I'$	$(H\{\ell \mapsto [w_0, \dots, w_{i-1}, \hat{R}(sv_3), w_{i+1}, \dots, w_{n-1}]\}, R, I')$ where $\hat{R}(sv_1) = \ell$ and $\hat{R}(sv_2) = i$ and $H(\ell) = [w_0, \dots, w_{n-1}]$ and $0 \leq i < n$
$\text{fmov } \mathbf{f}, fv; I'$	$(H, R\{\mathbf{f} \mapsto \hat{R}(fv)\})$
$\text{floadr } \mathbf{f}, \mathbf{r}; I'$	$(H, R\{\mathbf{f} \mapsto l\}, I')$ where $R(\mathbf{r}) = \ell$ and $H(\ell) = l$
$\text{fstorer } \mathbf{r}, fv; I'$	$(H\{\ell \mapsto l\}, R\{\mathbf{r} \mapsto \ell\}, I')$ where $\hat{R}(fv) = l$ and $\ell \notin H$
$\text{fswritesp}(i), fv; I'$	$(H, R\{\mathbf{sp} \mapsto s'\}, I')$ where $s' = (R(\mathbf{sp}))[i]_{64} \leftarrow \hat{R}(fv)$
$\text{sub}_\phi \mathbf{f}, sv_1, sv_2; I'$	$(H, R\{\mathbf{f} \mapsto l_i\}, I')$ where $\hat{R}(sv_1) = \ell$ and $\hat{R}(sv_2) = i$ and $H(\ell) = [l_0, \dots, l_{n-1}]$ and $0 \leq i < n$
$\text{upd}_\phi sv_1, sv_2, fv; I'$	$(H\{\ell \mapsto [l_0, \dots, l_{i-1}, \hat{R}(fv), l_{i+1}, \dots, l_{n-1}]\}, R, I')$ where $\hat{R}(sv_1) = \ell$ and $\hat{R}(sv_2) = i$ and $H(\ell) = [l_0, \dots, l_{n-1}]$ and $0 \leq i < n$

Figure C.3: TILTAL transitions (part III)

Bibliography

- [App01] Andrew W. Appel. Foundational proof-carrying code. In *Proceedings of the 16th Annual Symposium on Logic in Computer Science (LICS'01)*. IEEE Computer Society Press, June 2001.
- [CLN⁺00] Christopher Colby, Peter Lee, George C. Necula, Fred Blau, Mark Plesko, and Kenneth Cline. A certifying compiler for Java. In *Proceedings of the Conference on Programming Language Design and Implementation (PLDI'00)*, pages 95–107, Vancouver, Canada, June 2000. ACM Press.
- [Cra00] Karl Crary. Sound and complete elimination of singleton kinds. Technical Report CMU-CS-00-104, School of Computer Science, Carnegie Mellon University, 2000.
- [Cra03] Karl Crary. Toward a foundational typed assembly language. In *POPL '03: Proceedings of the 30th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 198–212, New York, NY, USA, 2003. ACM Press.
- [CW99] Karl Crary and Stephanie Weirich. Flexible type analysis. In *1999 ACM International Conference on Functional Programming*, Paris, France, September 1999. ACM Press.
- [CWM98] Karl Crary, Stephanie Weirich, and Greg Morrisett. Intensional polymorphism in type-erasure semantics. In *1998 ACM International Conference on Functional Programming*, pages 301–312, Baltimore, September 1998. Extended version published as Cornell University technical report TR98-1721.
- [DHB90] R. Kent Dybvig, Robert Hieb, and Tom Butler. Destination-driven code generation. Technical Report 302, Indiana University Computer Science Department, Bloomington, IN, February 1990.
- [FSDF93] Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. The essence of compiling with continuations. In *Proceedings of the ACM SIGPLAN 1993 conference on Programming language design and implementation*, pages 237–247. ACM Press, 1993.
- [GM99] Neal Glew and Greg Morrisett. Type-safe linking and modular assembly language. In *Conference Record of the 26th Symposium on Principles of Programming Languages (POPL'99)*, pages 262–275. ACM Press, January 1999.
- [GM00] Dan Grossman and Greg Morrisett. Scalable certification for typed assembly language. In *2000 ACM SIGPLAN Workshop on Types in Compilation '00*, Montreal, Canada, September 2000.

- [HM95] Robert Harper and Greg Morrisett. Compiling polymorphism using intensional type analysis. In *Twenty-Second ACM Symposium on Principles of Programming Languages*, pages 130–141, San Francisco, CA, January 1995.
- [HMM90] Robert Harper, John C. Mitchell, and Eugenio Moggi. Higher-order modules and the phase distinction. In *Seventeenth ACM Symposium on Principles of Programming Languages*, San Francisco, CA, January 1990.
- [MCG⁺99] Greg Morrisett, Karl Crary, Neal Glew, Dan Grossman, Richard Samuels, Frederick Smith, David Walker, Stephanie Weirich, and Steve Zdancewic. TALx86: A realistic typed assembly language. In *Second Workshop on Compiler Support for System Software*, pages 25–35, Atlanta, Georgia, May 1999.
- [MCGW02] Greg Morrisett, Karl Crary, Neal Glew, and David Walker. Stack-based typed assembly language. *Journal of Functional Programming*, 12(1):43–88, January 2002.
- [MH98] Greg Morrisett and Robert Harper. Typed closure conversion for recursively-defined functions (extended abstract). In Andrew Gordon, Andrew Pitts, and Carolyn Talcott, editors, *Second Workshop on Higher-Order Techniques in Operational Semantics (HOOTS-II)*, volume 10 of *Electronic Notes in Theoretical Computer Science*, Stanford University, Stanford, California, December 1998. Elsevier. URL: <http://www.elsevier.nl/locate/entcs/volume10.html>.
- [MMH96] Yasuhiko Minamide, Greg Morrisett, and Robert Harper. Typed closure conversion. In *Twenty-Third ACM Symposium on Principles of Programming Languages*, pages 271–283, St. Petersburg, FL, January 1996.
- [Mog89] Eugenio Moggi. Computational lambda-calculus and monads. In *Proceedings 4th Annual IEEE Symp. on Logic in Computer Science, LICS'89, Pacific Grove, CA, USA, 5–8 June 1989*, pages 14–23. IEEE Computer Society Press, Washington, DC, 1989.
- [MTHM97] Robin Milner, Mads Tofte, Robert Harper, and Dave MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997.
- [MWCG97] Greg Morrisett, David Walker, Karl Crary, and Neal Glew. From System F to Typed Assembly Language. Technical Report TR97-1651, Department of Computer Science, Cornell University, 1997.
- [MWCG98] Greg Morrisett, David Walker, Karl Crary, and Neal Glew. From System F to typed assembly language. In *Twenty-Fifth ACM Symposium on Principles of Programming Languages*, pages 85–97, San Diego, January 1998. Extended version published as Cornell University technical report TR97-1651.
- [Nec98] George C. Necula. *Compiling with Proofs*. PhD thesis, Carnegie Mellon University, October 1998. Available as Technical Report CMU-CS-98-154.
- [NL98] George C. Necula and Peter Lee. The design and implementation of a certifying compiler. In Keith D. Cooper, editor, *Proceedings of the Conference on Programming Language Design and Implementation (PLDI'98)*, pages 333–344, Montreal, Canada, June 1998. ACM Press.

- [PCHS00] Leaf Petersen, Perry Cheng, Robert Harper, and Chris Stone. Implementing the TILT internal language. Technical Report CMU-CS-00-180, School of Computer Science, Carnegie Mellon University, December 2000.
- [SH99] Christopher A. Stone and Robert Harper. Deciding Type Equivalence in a Language with Singleton Kinds. Technical Report CMU-CS-99-155, Department of Computer Science, Carnegie Mellon University, 1999.
- [Sha97] Zhong Shao. An overview of the FLINT/ML compiler. In *1997 Workshop on Types in Compilation*, Amsterdam, June 1997. ACM SIGPLAN. Published as Boston College Computer Science Department Technical Report BCCS-97-03.
- [SLM98] Zhong Shao, Christopher League, and Stefan Monnier. Implementing typed intermediate languages. In *ICFP '98: Proceedings of the third ACM SIGPLAN international conference on Functional programming*, pages 313–323, New York, NY, USA, 1998. ACM Press.
- [THS98] Omri Traub, Glenn H. Holloway, and Michael D. Smith. Quality and speed in linear-scan register allocation. In *SIGPLAN Conference on Programming Language Design and Implementation*, pages 142–151, 1998.
- [TMC⁺96] David Tarditi, Greg Morrisett, Perry Cheng, Chris Stone, Robert Harper, and Peter Lee. TIL: A type-directed optimizing compiler for ML. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 181–192, Philadelphia, PA, May 1996.
- [VC03] Joseph C. Vanderwaart and Karl Crary. A typed interface for garbage collection. In *TLDI '03: Proceedings of the 2003 ACM SIGPLAN international workshop on Types in languages design and implementation*, pages 109–122, New York, NY, USA, 2003. ACM Press.
- [VDP⁺03] Joseph C. Vanderwaart, Derek R. Dreyer, Leaf Petersen, Karl Crary, and Robert Harper. Typed compilation of recursive datatypes. In *Proceedings of the TLDI 2003: ACM SIGPLAN International Workshop on Types in Language Design and Implementation*, pages 98–108, New Orleans, LA, January 2003.
- [WM01] David Walker and Greg Morrisett. Alias types for recursive data structures. *Lecture Notes in Computer Science*, 2071, 2001.

