

REASONING ABOUT SOFTWARE IN THE  
PRESENCE OF TRANSIENT FAULTS

FRANCES JANE PERRY

A DISSERTATION  
PRESENTED TO THE FACULTY  
OF PRINCETON UNIVERSITY  
IN CANDIDACY FOR THE DEGREE  
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE  
BY THE DEPARTMENT OF  
COMPUTER SCIENCE

ADVISOR: DAVID P. WALKER

SEPTEMBER 2008

© Copyright by Frances Jane Perry, 2008. All rights reserved.

## Abstract

A transient fault occurs when an energetic particle strikes a chip and causes a change in state in the processor. Although there is no permanent damage, the current computation may become corrupt. Transient faults have been shown to be the cause of crashes at major companies, and current technology trends will make future processors more susceptible to them. Researchers have already developed a number of solutions using combinations of software and hardware where the general approach is to duplicate computations and check for consistency between the copies. Unfortunately, generating correct fault-tolerant code is difficult, and there has been little research on proving the correctness of these techniques. Reasoning formally about fault tolerance is challenging, because invariants that hold at compile time, such as standard type safety, may not actually hold at runtime. Previous work on formalizing fault tolerance has been at the high level, including proofs about fault-tolerant algorithms and the definition of a fault-tolerant lambda calculus.

This dissertation presents the first set of techniques for statically proving fault-tolerance properties of actual executable code. To address this challenge, we develop modifications to the general methodology for typed assembly languages. An assembly-level type system incorporates invariants about fault tolerance that are strong enough to prove that all well-typed programs have the desired behaviors. All that is required to guarantee that a specific piece of code is fault-tolerant is to type-check the code at the conclusion of compilation.

More specifically, we introduce a family of three type systems.  $TAL_{FT}$  is a core language with simple instructions that guarantees that well-typed programs implementing a hybrid fault-tolerance scheme will always detect a single fault before the fault can affect the observable behavior of the program.  $ETAL_{FT}$  extends  $TAL_{FT}$  with features necessary to support realistic compilation, including stack activation frames and dynamic memory

allocation. The third language,  $TAL_{CF}$ , precisely captures the behavior of software solutions for control flow faults and can provably detect any fault that causes incorrect control transfers between basic blocks before control exits that first incorrect block. Although each typed assembly language is designed for a specific hardware context, the type systems and proof methods use similar designs, allowing us to demonstrate general approaches needed for reasoning in the presence of transient faults. For example, to statically reason about values that may be corrupted at runtime, we generalize the “color systems” of previous work into a framework for classifying values with related reliability properties.

As well as being the first to prove that executable code is fault-tolerant, this dissertation makes three additional contributions. By including a language of static expressions within the type system, we can verify low-level fault-tolerant code independently of the compilation process. We show to apply fault-tolerant typed assembly languages to two different fault models: a hybrid system to detect data faults and a software-only system to detect control-flow faults. Finally, we investigate how to generate fault-tolerant typed assembly language for a realistic compiler.

## Acknowledgments

This dissertation owes the most to my advisor David Walker. He is an extremely knowledgeable and patient teacher who taught me to be clear and rigorous in my work and confident in myself. I clearly remember the day I first openly disagreed with him, and he responded, “That’s it! Now you are becoming a researcher.”

The members of my thesis committee were extremely helpful in improving the quality of this dissertation, as well as my job talk. I am very grateful for their time and feedback. Andrew Appel helped me develop clear explanations and pull out the high-level ideas. David Tarditi inspired me to work that bit harder to become a better writer. Margaret Martonosi and David August’s different backgrounds helped me target my work to the broader community.

Additionally, this dissertation benefited from the support of many others. George Reis and Neil Vachharajani were a great help with the architecture side of things. I am grateful for the wonderful administrative help and friendship of the staff at Princeton, especially Melissa Lawson and Donna O’Leary. This work was funded by the National Science Foundation award CNS-0627650 and a Microsoft fellowship. Any opinions, findings, and conclusions or recommendations expressed in this material are mine and do not necessarily reflect the views of the NSF or Microsoft.

During internships at Microsoft, I was lucky to work with a number of amazing people. Manuvir Das, Jason Yang, and the PAG group taught me that sometimes it’s okay to sacrifice soundness, as long as you can find real bugs in real code. Juan Chen and Chris Hawblitzel showed me the joy of straddling the boundary between theory and practice to design elegant, yet practical, systems.

My fellow students were always there to advise me, commiserate with me, and support me. (And, of course, procrastinate with me...) Limin Jia was both a fantastic

colleague and a close friend, and our afternoon tea times were always a bright spot in my day. Ananya Misra's sarcastic emails and dry wit kept me laughing. Siggi Cherem and Sharon Betz kept me sane through the home stretch. In addition, my survival of grad school is due in part to Melissa Carroll, Dan Dantas, Haakon Larson, Guilherme Ottoni, Shirley Gaw, Henning Rohde, Zach Anderson, Akash Lal, Dan Wang, Jay Liggatti, the pizza crowd, my officemates, and many others.

My family's continued love and support means everything to me. My father Jeff, though he never did teach me how to throw a ball, encouraged me through my disgruntled phases and kept me focused on the final goal; my mother Liz reminded me to take care of myself and was always willing to proofread my papers for punctuation, even if the words themselves were gobbledygook. As we grew up, my sister Julia and I realized we loved each other after all, and our differences turned out to be quite complementary – I fix her computer and she picks out my interview suits. Erik, Isabelle, Becky, Ben, and Hunter have become my second family. They were always there, kept Dan and me well-fed, and rushed to our aid during the occasional emergency. I will greatly miss having them so close by. My "Aunt" Marilyn is always proud of me and reminds me to be proud of myself.

And finally, there is Dan. Although his continual jokes about "Frances the 12th year grad student" did get a bit old, he is, and always will be, my best friend. I am thrilled to have him by my side as we enter this next phase in our lives – moving cross country, becoming homeowners, making new friends, and who knows what else!

# Contents

Abstract . . . . .	iii
<b>1 Introduction and Background</b>	<b>1</b>
1.1 Transient Faults . . . . .	1
1.1.1 Issues Caused by Transient Faults . . . . .	1
1.1.2 Transient Fault Trends . . . . .	3
1.2 Existing Transient Fault Solutions . . . . .	4
1.2.1 An Example Solution: SWIFT . . . . .	5
1.2.2 Do Existing Solutions Work? . . . . .	6
1.3 Proof-Carrying Code . . . . .	7
1.3.1 Typed Assembly Languages . . . . .	7
1.3.2 Using Typed Assembly Languages . . . . .	8
1.4 Thesis Scope . . . . .	10
1.4.1 Modeling Transient Faults . . . . .	10
1.4.2 Defining Fault Tolerance . . . . .	12
1.4.3 Invariants of Fault Tolerance . . . . .	13
1.4.4 Proving Fault Tolerance . . . . .	13
1.4.5 Compiling for Fault Tolerance . . . . .	14
1.5 Thesis Organization . . . . .	14

<b>2</b>	<b>TAL<sub>FT</sub> : Fault-tolerant Typed Assembly Language</b>	<b>15</b>
2.1	The Fault-Tolerant Hardware . . . . .	16
2.1.1	The Fault Model . . . . .	18
2.1.2	Instruction Semantics . . . . .	20
2.2	Typing . . . . .	28
2.2.1	Static Expressions . . . . .	30
2.2.2	Value Typing . . . . .	33
2.2.3	Instruction Typing . . . . .	35
2.2.4	Machine State Typing . . . . .	44
2.3	Formal Results . . . . .	47
2.3.1	Type Safety . . . . .	48
2.3.2	Fault Tolerance . . . . .	49
2.4	Performance . . . . .	51
2.5	Summary . . . . .	53
<b>3</b>	<b>ETAL<sub>FT</sub> : Generating fault-tolerant assembly code</b>	<b>54</b>
3.1	MiniC . . . . .	55
3.2	Extending TAL <sub>FT</sub> to ETAL <sub>FT</sub> . . . . .	57
3.2.1	Memory Layout . . . . .	57
3.2.2	Stacks . . . . .	59
3.2.3	Memory Allocation and Initialization . . . . .	69
3.2.4	Removal of Color Tags . . . . .	74
3.2.5	Other Changes . . . . .	76
3.3	ETAL <sub>FT</sub> Formal Results . . . . .	77
3.3.1	Type Safety . . . . .	77



3.3.2	Fault Tolerance . . . . .	78
3.4	Translation from MiniC to ETAL <sub>FT</sub> . . . . .	82
3.4.1	Translation Introduction . . . . .	82
3.4.2	Translation Details . . . . .	83
3.4.3	Translation Formal Results . . . . .	85
3.5	Translation Example . . . . .	86
3.6	Type-preserving Optimizations . . . . .	92
3.6.1	General Considerations . . . . .	93
3.6.2	Removal of Redundant Moves . . . . .	94
3.6.3	Register Allocation . . . . .	94
3.6.4	Common Subexpression Elimination . . . . .	94
3.6.5	Dead Code Elimination . . . . .	95
3.6.6	Constant Folding and Propagation . . . . .	95
3.6.7	Stack Packing . . . . .	95
3.6.8	Instruction Scheduling . . . . .	96
3.7	Summary . . . . .	97
<b>4</b>	<b>TAL<sub>CF</sub> : Reasoning about Control Flow</b>	<b>98</b>
4.1	Informal Overview . . . . .	100
4.2	The Control-Flow Machine . . . . .	106
4.2.1	Dynamic Semantics . . . . .	107
4.3	Typing . . . . .	111
4.3.1	Value Typing . . . . .	117
4.3.2	Instruction Typing . . . . .	121
4.3.3	Block Typing . . . . .	123

4.3.4	Machine State Typing . . . . .	126
4.4	Formal Results . . . . .	128
4.5	Translation . . . . .	130
4.6	Summary and Future Work . . . . .	131
<b>5</b>	<b>Related Work and Conclusion</b>	<b>133</b>
5.1	PCC and TAL . . . . .	133
5.2	Research in Transient Fault Tolerance . . . . .	135
5.2.1	Hardware-Based Solutions . . . . .	135
5.2.2	Software-Based Solutions . . . . .	136
5.2.3	Hybrid Solutions . . . . .	140
5.3	Formal Reasoning about Faults . . . . .	140
5.4	Concluding Remarks . . . . .	142
<b>A</b>	<b>TAL<sub>FT</sub> Proof Details</b>	<b>145</b>
A.1	Lemmas . . . . .	145
A.1.1	Properties of Static Expressions . . . . .	146
A.1.2	Properties of Well-Typed Values . . . . .	147
A.1.3	Properties of Well-typed Memories . . . . .	149
A.1.4	Properties of Well-typed Queues . . . . .	150
A.2	Type Safety . . . . .	151
A.3	Multistep Transitions . . . . .	154
A.3.1	No False Positives . . . . .	154
A.3.2	Multistep Split and Combine . . . . .	158
A.4	Fault Tolerance . . . . .	158
A.4.1	Simulation Relation . . . . .	159

A.4.2	Singlestep Fault Detection . . . . .	160
A.4.3	Multistep Fault Detection . . . . .	161
A.4.4	Fault Similarity . . . . .	163
A.4.5	Fault Tolerance Theorem . . . . .	164
<b>B</b>	<b>MiniC Typing Rules</b>	<b>166</b>
<b>C</b>	<b>Complete Rules for <math>\text{ETAL}_{\text{FT}}</math></b>	<b>169</b>
C.1	Syntax of Machine States . . . . .	169
C.2	Dynamic Semantics . . . . .	170
C.2.1	Fault Rules . . . . .	170
C.2.2	Normal Execution Rules . . . . .	170
C.3	Static Semantics . . . . .	174
C.3.1	Syntax . . . . .	174
C.3.2	Properties of Static Expressions . . . . .	174
C.3.3	Value Typing . . . . .	176
C.3.4	Subtyping . . . . .	177
C.3.5	Stack Typing Judgments . . . . .	178
C.3.6	Instruction Typing Rules . . . . .	179
C.3.7	Machine State Typing . . . . .	183
C.4	Simulation of Machine States . . . . .	185
C.4.1	Color Extraction . . . . .	185
C.4.2	Simulation . . . . .	186
<b>D</b>	<b><math>\text{ETAL}_{\text{FT}}</math> Formal Results</b>	<b>188</b>
D.1	Modified Lemmas . . . . .	188

D.2	New Lemmas For Stacks . . . . .	191
D.3	New Lemmas For Dynamic Memory Allocation . . . . .	192
D.4	Type Safety . . . . .	194
D.5	Fault Tolerance . . . . .	195
<b>E</b>	<b>MiniC to ETAL<sub>FT</sub> Translation</b>	<b>198</b>
E.1	Overview . . . . .	198
E.2	Translating Typing Information . . . . .	200
E.2.1	Type Translation . . . . .	200
E.2.2	Context Generation . . . . .	200
E.3	Extending Code Memory . . . . .	203
E.3.1	Invariants Between Functions . . . . .	204
E.3.2	Partial Code Memories . . . . .	205
E.3.3	Invariants within a Function . . . . .	205
E.3.4	Adding Instructions to Code Memory . . . . .	207
E.4	Value Translation . . . . .	208
E.5	Value List Translation . . . . .	210
E.6	Local Declaration Translation . . . . .	211
E.7	Statement Translation . . . . .	215
E.8	Function Translation . . . . .	216
E.8.1	Prologue and Epilogue . . . . .	216
E.8.2	Function Translation Judgment . . . . .	217
E.9	Program Translation . . . . .	219
<b>F</b>	<b>TAL<sub>CF</sub> Formal Results</b>	<b>221</b>
F.1	Type Safety . . . . .	221

F.1.1	Typing Lemmas . . . . .	221
F.1.2	Type Safety . . . . .	224
F.2	Fault Tolerance Results . . . . .	226
F.2.1	Machine State Simulation . . . . .	226
F.2.2	Block Execution . . . . .	226
F.2.3	Fault Recovery . . . . .	228
F.2.4	Block Transitions . . . . .	229
F.2.5	Program Execution . . . . .	230
F.2.6	The Fault Tolerance Theorem . . . . .	231
<b>G</b>	<b>TAL<sub>CF</sub> Translation</b> . . . . .	<b>234</b>
G.1	A Simple While Loop Language . . . . .	234
G.2	Checking Code and Exit Code Macros . . . . .	235
G.3	Translating Variable Context $V$ . . . . .	235
G.4	Partial Translations . . . . .	237
G.5	Translating Statements . . . . .	239
G.6	The Translation Theorem . . . . .	241
	<b>Bibliography</b> . . . . .	<b>243</b>

# List of Figures

1.1	Effects of Altitude [76] and Feature Size [11] on Transient Faults . . . . .	3
1.2	Proof-Carrying Code [42] and Typed Assembly Language [39]. . . . .	8
2.1	Syntax of instructions and machine states. . . . .	18
2.2	Operational rules for basic instructions. . . . .	21
2.3	Selected operational rules for memory instructions. . . . .	23
2.4	Operational rules for control flow instructions. . . . .	27
2.5	TAL <sub>FT</sub> type syntax. . . . .	29
2.6	Semantics of Static Expressions, Part 1. . . . .	31
2.7	Semantics of Static Expressions, Part 2. . . . .	32
2.8	Value Typing. . . . .	34
2.9	Subtyping. . . . .	36
2.10	Instruction Typing Rules for Basic Instructions and Memory Instructions. . . . .	39
2.11	Instruction Typing Rules for Jump Instructions. . . . .	41
2.12	Instruction Typing Rules for Branch Instructions. . . . .	42
2.13	Machine State Element Typing. . . . .	45
2.14	Machine State Typing. . . . .	46
2.15	Similarity of Machine States. . . . .	50

2.16	Performance Normalized to Unprotected Version. . . . .	51
3.1	MiniC Example Program. . . . .	55
3.2	MiniC Syntax. . . . .	56
3.3	New Instruction Syntax and Operation Semantics to Support Stacks. . . . .	60
3.4	Modifications to Value Typing, Subtyping, and Well-Formed Judgments. . . . .	64
3.5	Stack Lookup and Update. . . . .	65
3.6	Typing Rules for the New Stack Instructions. . . . .	67
3.7	Stack Typing. . . . .	68
3.8	New Operation Semantics to Support Dynamic Memory Allocation. . . . .	70
3.9	Additions to Subtyping to Support Reference Initialization. . . . .	71
3.10	Instruction Typing Rules to Support Dynamic Memory Allocation. . . . .	73
3.11	Changes to Memory Typing Judgments. . . . .	75
3.12	Changes to the Machine State Typing Judgment. . . . .	76
3.13	Similarity of Machine States. . . . .	80
3.14	Summary of Translation Judgments. . . . .	84
3.15	Summary of Type Translation Judgments. . . . .	85
3.16	Example Translation from MiniC to ETAL <sub>FT</sub> . . . . .	87
4.1	Machine State Syntax. . . . .	106
4.2	Operational Semantics. . . . .	110
4.3	Properties of colored values and zap tags. . . . .	114
4.4	Example: Protocol Stages. . . . .	115
4.5	Typing Syntax. . . . .	116
4.6	Static Expression Judgments, Part 1. . . . .	118
4.7	Static Expression Judgments, Part 2. . . . .	119

4.8	Value Typing Judgment and Subtyping Judgment. . . . .	120
4.9	Instruction Typing Judgment. . . . .	122
4.10	Block Typing Judgment. . . . .	124
4.11	Machine State Typing. . . . .	127
A.1	Example Cases of Theorem 25 (Progress). . . . .	153
A.2	Example Case from Theorem 26 (Preservation) Part 1. . . . .	155
A.3	Example Case from Theorem 26 (Preservation) Part 1. . . . .	156
A.4	Example Case of Theorem 26 (Preservation) Part 2. . . . .	157
A.5	Similarity of Machine States. . . . .	159
A.6	Example Case of Lemma 30. . . . .	162
E.1	Type Translation. . . . .	201
E.2	Statement Translation - Basic Statements. . . . .	212
E.3	While Statement Translation. . . . .	213
E.4	If Statement Translation. . . . .	213
E.5	Function Call Translation. . . . .	214
G.1	Translation of While Programs. . . . .	240



# Chapter 1

## Introduction and Background

### 1.1 Transient Faults

A *transient fault* occurs when an energetic particle strikes a transistor or wire in a processor and causes a change in state. These particles do not permanently damage the hardware, but they may corrupt the computation that is currently executing by depositing a charge that alters stored values and signals.

#### 1.1.1 Issues Caused by Transient Faults

As one might expect, transient faults can cause a host of issues. In particular, as these anecdotes show, systems can crash due to transient faults.

- In 2000, Sun Microsystems acknowledged that transient faults interfered with cache memories and caused crashes in server systems at major customer sites, including AOL, eBay, and dozens of others [6].

- Cypress Semiconductor has stated “The wake-up call came in the end of 2001 with a major customer reporting havoc at a large telephone company. Technically it was found that a single soft fail... was causing an interleaved system farm to crash” [76].
- At Los Alamos in 2003, the ASC Q supercomputer crashed regularly due to soft errors [37].

Of course, and perhaps even worse, there is always the possibility that the result of a program is incorrect. Such an outcome is referred to as *silent data corruption*. Although these issues with crashing and correctness are bad enough, transient faults can also cause some more surprising problems.

Researchers [24] have shown how an attacker can exploit transient faults to take over a Java Virtual Machine. In a Java Virtual Machine, the virtual machine runs in the same address space as the untrusted program, relying on type safety to keep the untrusted program from touching its address space. Essentially, the attacker can craft a program that waits until a fault results in a pointer with a runtime type that does not match its static type and then uses this mismatch to execute arbitrary code.

Cryptographic protocols can be broken using transient faults [7, 8, 19, 56, 4]. For example, certain implementations of RSA based on the Chinese Remainder Theorem are vulnerable to a single fault [9]. RSA relies on the inability to factor a large number  $N$  into two prime numbers  $p$  and  $q$ . A signature computed from a message  $m$  and a private key  $d$  has the form  $a * s_p + b * s_q$  where  $a$  and  $b$  are precomputed values and  $s_n$  is a function of  $m$ ,  $d$  and  $n$ . If an attacker obtains two signatures (one correct, one faulty) of the same message, he can take the difference of the two signatures. One of the two terms cancels

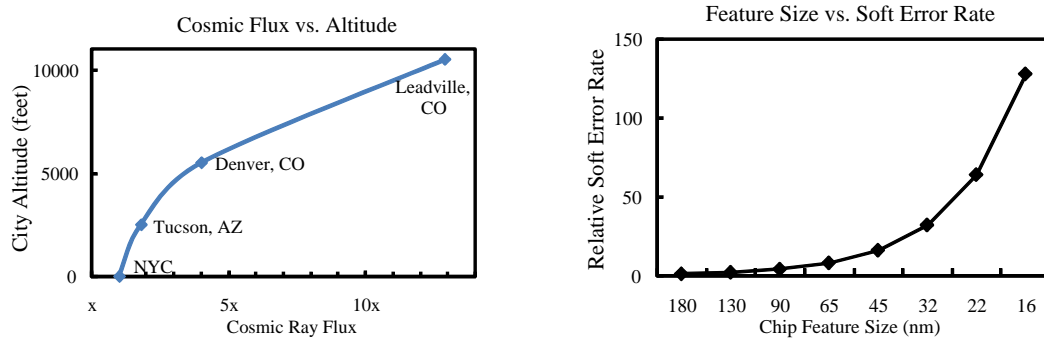


Figure 1.1: Effects of Altitude [76] and Feature Size [11] on Transient Faults

out, leaving the difference of the other. Then by factoring out the precomputed value and calculating a greatest common divisor, the attacker can determine  $p$  and  $q$ .

## 1.1.2 Transient Fault Trends

It is difficult to get precise numbers on transient fault rates, but as a benchmark, in 2004, a typical laptop with 1GB DRAM had approximately 1 soft fail per year [76].

The particles that cause transient faults are mostly high-energy neutrons in cosmic-ray radiation [76], and so fault rates increase with altitude. The graph in Figure 1.1 shows that the amount of cosmic ray flux in Denver, Colorado is about four times that of New York City. The data was gathered over three years beginning in 1986, with New York City averaging one failure every 45 days, and Leadville averaging one failure every two days.

These results are for transient faults that affect memory. Transient faults can also affect the latches in the processor itself, and these faults are much harder to detect. Unfortunately, trends in processor manufacturing are causing these fault rates to increase over time. Faster clock rates, increasing transistor density, decreasing voltages and smaller feature sizes are resulting in processors that are more susceptible to transient

faults [5, 46, 63]. For example, decreasing voltages reduce the critical voltage of each transistor, making it more likely that adding the extra charge from a fault will change a value. On the other hand, decreasing feature size makes it less likely that a fault will affect a given transistor. As a result, the fault rate per bit is expected to remain roughly constant for the next few generations [26, 31]. However, the number of transistors per chip is increasing, and so the overall result is an increase in the fault rate per processor. Figure 1.1 shows how the transient fault rate increases as the feature size decreases in a processor. Because of the combination of these factors, fault rates are increasing approximately 8% per processor generation [28].

So although the average computer user may not be aware of transient faults today, transient faults are likely to become a significant problem in the near future.

## 1.2 Existing Transient Fault Solutions

The computer architecture and compiler communities are well aware of the issues surrounding transient faults, and dealing with transient faults has been an active area of research for many years. At the high level, all the solutions work by adding redundancy. Redundancy can be added in time (by computing a result and then computing it again), in space (by storing values in two locations), or in information (using techniques such as checksums or error-correcting codes). The many solutions can be divided into two main categories: hardware-based solutions that duplicate hardware structures in the processor or exploit multiple cores, and software-based solutions that insert additional instructions to verify that values are not corrupted.

Chapter 5 goes into more detail on the existing solutions, but as one would expect, there are trade-offs to each type. It is standard to evaluate a system in terms of how it

performs along three axes: cost, performance, and power. Transient faults add a fourth axis to be considered – reliability. Hardware solutions are generally more efficient in terms of performance, but more expensive in terms of power and cost. In addition, once a hardware solution is deployed, it cannot be changed. Software solutions, on the other hand, are more flexible. Because they are controlled by a compiler, they can be deployed when, where, and to the degree necessary. Software solutions have no additional hardware cost, but they often have a noticeable performance overhead. The reliability added by each solution varies, but it is safe to say that although hardware solutions can achieve very high rates of fault detection and recovery, software-solutions are fundamentally limited without hardware support. More recently, researchers have begun proposing hybrid solutions that use small amounts of additional hardware that is controlled by the software with the hopes that such solutions can combine the best of both worlds.

### **1.2.1 An Example Solution: SWIFT**

As an example, let us look more closely at one specific solution. SWIFT (SoftWare Implemented Fault Tolerance) [60] is a software-based solution to transient faults. During compilation, the SWIFT compiler duplicates the original computation, interleaves the two copies in some manner, and inserts comparisons before storing values to memory to ensure that the two versions of the computation agree. Because the two computations are completely independent, a single transient fault will likely result in a difference between the two computations. If this corruption spreads to a value that will be written to memory, the inserted comparison will fail, and control is transferred to a designated location containing error handling or recovery code. If a transient fault corrupts a value

but does not result in a noticeable difference in the values written to memory, then it does not need to be detected.

The researchers who developed SWIFT evaluated it by randomly injecting faults and looking at the resulting performance and detection rate. At first, the detection rates were not nearly as good as they were expecting. After more investigation, they discovered that the compiler, having added the redundant computation, had been doing what compilers do best – optimizing the code to remove redundancy. Optimizations such as common subexpression elimination were resulting in a single computation, with additional code before store instructions that duplicated the value to be stored and then compared the value and its duplicate. Clearly, faults affecting the computation itself were not being detected. The solution was to turn off or reorder the optimizations that the compiler performed until the results were more in line with what was expected.<sup>1</sup>

This leaves an interesting question. How do we know that the code that SWIFT generates now is as reliable as intended? Any realistic optimizing compiler is extremely complicated, and it is very difficult to understand how all the phases may interact. The high-level intuition used in SWIFT is simple to understand, but how do we ensure that the compiled code correctly implements this intuition?

### **1.2.2 Do Existing Solutions Work?**

Most other transient fault solutions are evaluated in similar ways: a promising idea is presented, a system is implemented, and experimental results show an increase in fault detection. None of these solutions provide rigorous proofs of their correctness. In fact, many do not even precisely define which faults they can handle.

---

<sup>1</sup>Personal communication with G.A. Reis, October 2006.

The goal of this thesis is provide an approach to formally reason about the behavior of transient fault solutions, focusing on those approaches that make use of software. The initial task is to determine the correct level of abstraction and best proof methodology for reasoning about these situations. As the example with SWIFT showed, it is not enough to simply reason about the *algorithm*; we need to reason about the *implementation* as well. Transient faults affect the hardware, so the reasoning technique needs to model primitive instructions, memory, registers, and any other relevant hardware structures. The addition of software redundancy occurs as a compiler phase, just like optimizations and register allocation. Because compiler phases may have unexpected interactions, we want to reason about the final assembly code generated by the compiler.

## 1.3 Proof-Carrying Code

Proof-Carrying Code [43, 42] is a technique for verifying properties of untrusted, low-level code. In a proof-carrying code system, the compiler is responsible for generating two things: the low-level code and a safety proof that the low-level code obeys a predefined set of properties. Then anyone who wishes to run the code can first verify the proof to ensure that the code will not behave unexpectedly. Traditionally, these safety proofs have included guarantees of memory safety and type safety.

### 1.3.1 Typed Assembly Languages

One standard way to represent these safety proofs is using typed assembly languages [39]. To generate a typed assembly language, the compiler begins with a type-safe source program. Instead of type checking this program and then discarding the typing information, the compiler preserves types through every level of intermediate representation. Each

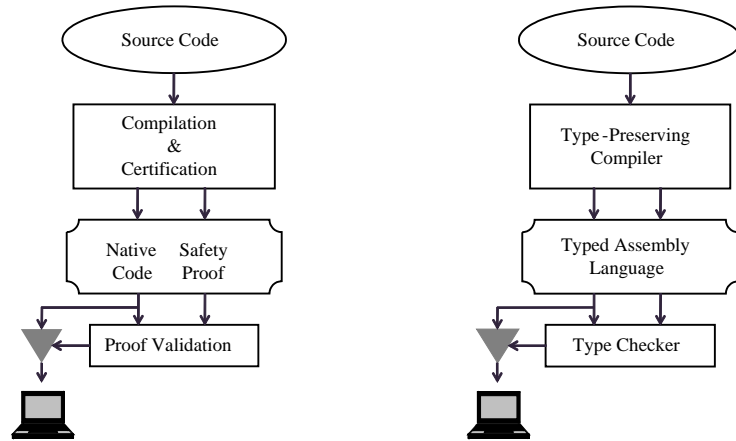


Figure 1.2: Proof-Carrying Code [42] and Typed Assembly Language [39].

intermediate representation has its own associated type system, and as the representations become closer to the machine level, the type systems become more complicated in order to maintain information about the program.

Once the compiler generates the final assembly code, type checking the code guarantees that the code obeys the properties encapsulated by the type system. Nothing is required to be known about the correctness of the compiler. If the generated code obeys the safety proof, then that proves that it behaves according to the properties captured by the type system.

### 1.3.2 Using Typed Assembly Languages

Typed assembly languages are an active area of research, and Chapter 5 gives some more background in this area. For now, we will explain the general methodology for using a typed assembly language.

1. **Model machine execution.** The goal of typed assembly languages is to reason about the execution behavior of a program, so the first step is to develop a model



of how programs execute. One common way to do this is by defining a small-step operational semantics. Imagine that we could pause execution and look at the current state of an abstract processor. We could inspect the values in each register and in memory, see which instruction will be executed next, and so on. The small-step operational semantics takes a snapshot of the machine like this and shows how the snapshot is modified by executing one single instruction. By sequencing many of these small steps together, we can reason about the execution of an entire program.

2. **Specify the desired execution behavior.** Given the model of execution, we can define which executions have the behavior that we desire.
3. **Design the type system.** The type system encapsulates the desired invariants that the program should maintain. Whereas high-level type systems track the types of program variables, assembly-level type systems track the types of the values in the machine state, including the register file and memory.
4. **Prove Soundness.** Once we have these three definitions, the next step is to prove that the type system is sound with respect to the machine model. In other words, any well-typed program is guaranteed to execute as desired on the machine model.
5. **Show that the type system is expressive.** Finally, the last step is to show that the typed assembly language is expressive. We want to rule out the possibility that we have designed an overly restrictive type system. If there are no interesting programs that type-check, then it means nothing to prove properties about the behavior of all well-typed programs. We can show that this is not the case by taking a high-level language and showing how any well-typed source program can be compiled into a well-typed assembly program.

## 1.4 Thesis Scope

This thesis investigates the use of typed assembly languages as a method for verifying software-based transient fault solutions.

Typed assembly languages are especially well-suited for this because of the sheer complexity of reasoning about transient faults. Just as the many possible interleavings of threads make it difficult to reason about concurrent programs, the many possible transient faults make it difficult to reason about fault tolerance. For this reason, testing is not sufficient. To truly test that the code a compiler generates can detect all possible faults, it would be necessary to test all combinations of features in the compiler in conjunction with all possible faults. The drastic explosion in the number of test cases makes this infeasible. By using a typed assembly language, we can statically guarantee that a program will achieve perfect fault coverage relative to the fault model.

In particular, this thesis defines two families of typed assembly languages, each designed for a different type of solution.  $TAL_{FT}$  and its close cousin  $ETAL_{FT}$  are designed for a hybrid solution that uses a mix of hardware and software to detect faults.  $TAL_{CF}$  is a separate language that takes the first steps towards reasoning about software-only solutions to control-flow faults. In order to use these typed assembly languages for verifying the behavior of programs in the presence of transient faults, we make some slight modifications to the usual methodology for typed assembly languages.

### 1.4.1 Modeling Transient Faults

The machine model needs to represent hardware that may be affected by transient faults.

We will assume that memory is protected by error-correcting codes (ECC) [25]. Unlike simple parity which only detects single-bit errors, ECC is capable of both correcting

single-bit errors and detecting multi-bit errors. Each piece of data has a corresponding ECC code. The size of the code depends on the size of the data to be protected. Protecting 32 bits of data requires a 7 bit code, and protecting 64 bits of data requires an 8 bit code. When data is written to memory, the memory controller encodes the data using a device-specific algorithm and updates the corresponding ECC code with the result. When data is read from memory (or during periodic consistency checks), the read value is encoded and the result is compared to the stored code. If the two codes do not match, the memory controller corrects a single-bit error or reports a multi-bit error to the operating system. ECC has been shown to be extremely effective for protecting memory [12] and has been used since the 1950s [29, 64].

However, applying ECC to the register file is too costly in terms of both performance [68] and power [55] because of the frequency of accesses. We assume that the register file contents are vulnerable to transient faults and specifically model transient faults to registers in between the execution of two instructions. As the model machine executes its small-step semantics, it may nondeterministically insert an additional step rule that corrupts a random register value.

Transient faults may actually affect any part of the processor datapath or control. Many “inter-instruction” faults can be modeled as a correct instruction execution followed by a fault to the destination register. For example, a transient fault that strikes the ALU during the execution of an add instruction can be modeled as a correct add and then a register fault to the destination register. A fault that causes a multiply to be performed instead of add can be handled in the same way. However, this is not the case for all faults. The model does not account for faults that affect instructions with side effects beyond the register file, instruction decoding, the virtual memory page table, memory buses, and so on. (Section 2.1.2 gives a concrete example of a fault is not modeled.) This is not to

say that solutions such as SWIFT would fail to detect all these unmodeled faults, in fact many such faults would likely result in a difference between the two computations and be detected, but only that all such faults are not guaranteed to be handled in such a system.

By specifying the machine with precise operational rules, we clearly identify which transient faults are under consideration and what hardware behavior we rely on. If a fault can modify the assumed hardware behavior and cannot be modeled as a register fault, then this identifies a vulnerability that needs to be addressed with additional hardware or software techniques.

As is standard in the literature [57, 61], we will work under the assumption that only one fault occurs during the execution of a program. This does not mean that multiple faults would not be caught, only that with multiple faults there is a minuscule chance that they may occur in a way that tricks the fault detection mechanisms. For example, the first fault may corrupt a value, while the second affects the checking code that would have otherwise detected the error. Multiple faults might also cause both computations to calculate the same incorrect result.

### 1.4.2 Defining Fault Tolerance

The goal is to prove that programs are fault-tolerant, but first we need to formally define what this means. Abstractly, a program is fault-tolerant if no fault can change the observable behavior of a program. The definition of “observable” may change, but one example would be to assume the system operates in the presence of a memory-mapped IO device. In this case, the program is fault-tolerant if a fault does not cause a change the sequence of values written to memory up until the point where a fault is detected.

We will focus only on formalizing *fault detection*, and not specify *fault recovery*. There are a number of known recovery techniques applicable to this problem domain.

The simplest way to recover from a transient error is to restart the computation. For example, Google’s MapReduce implementation distributes work over thousands of machines. When a machine fails, its task is reassigned to another machine [18]. By running three copies of a computation, it is possible to use majority voting to recover from a single error [59, 73], though this requires consistently paying a high performance overhead to recover from a rare event. Checkpointing and rollback recovery are commonly used to recover from faults in databases and parallel systems [34, 67]. In our current setting, we already assume that memory is protected, so this can easily be used as stable storage for the checkpointing. Because transient faults are rare, it is not unreasonable to pay a high performance cost for rolling back to the most recent checkpoint.

### 1.4.3 Invariants of Fault Tolerance

The type systems are designed specifically to capture invariants about the fault-tolerance solution. For example, when using two computations to detect data faults, there are three main invariants that must be maintained. The copies of the computation need to be independent, so that a fault to one computation cannot affect the other computation. The computations need to be redundant, so that in the absence of faults they calculate equal values. In addition, any action that may affect the observable behavior of the program must be guarded by a comparison of the two computations.

### 1.4.4 Proving Fault Tolerance

We formalize the behavior of well-typed programs as a mathematical theorem that relates faulty and non-faulty executions of a program. First, we define a simulation relationship between a program snapshot and the equivalent faulty snapshot that essentially means that

the uncorrupted computations are identical, although the corrupted computations may differ. Then we show that continuing to execute the two snapshots will always maintain the simulation relationship up until the point where the fault is detected. This is used to guarantee that a well-typed program behaves equivalently whether or not it is affected by a single fault. In addition, we show that well-typed code only claims to have detected a fault when a fault has actually occurred.

### 1.4.5 Compiling for Fault Tolerance

Finally, we show how to take high-level (unduplicated) programs and translate them into well-typed fault-tolerant assembly language programs. We include realistic low-level features such as memory management and discuss the effect optimizations may have on our typed assembly language.

## 1.5 Thesis Organization

The remainder of this dissertation is organized as follows. Chapters 2 and 3 reason about a hybrid transient fault solution modeled after the CRAFT system [61]. Chapter 2 begins the process with a core assembly language called  $TAL_{FT}$  and walks through the first four stages of using typed assembly languages. In order to show expressiveness, Chapter 3 defines  $ETAL_{FT}$ , an extension to  $TAL_{FT}$ , and discusses how to use  $ETAL_{FT}$  as the target language of an optimizing compiler. Chapter 4 takes the first steps toward reasoning about software-only solutions to control-flow faults using a language called  $TAL_{CF}$ . Chapter 5 goes into more detail on related work and ends with some final remarks.

## Chapter 2

# TAL<sub>FT</sub> : Fault-tolerant Typed Assembly Language

This chapter presents TAL<sub>FT</sub>, a typed assembly language designed for reasoning formally about a hybrid transient fault solution. Our system is abstractly modeled after CRAFT [61], a solution similar to SWIFT. Being a software-only solution, SWIFT has some fundamental limitations. For example, before storing a value to memory, SWIFT performs a comparison between the two computations. However, a fault that occurs after this comparison, but before the store instruction, will not be caught. CRAFT introduces a little bit of extra hardware to help close this window of vulnerability.

The rest of this chapter presents the details of the abstracted hybrid fault-tolerance technique. Section 2.1 presents the syntax and operational semantics of the machine model. It is a RISC-based architecture with special instructions to facilitate reliable communication with memory and to detect control-flow faults. Section 2.2 presents the

---

*Material in this chapter is joint work with Lester Mackey, George A. Reis, Jay Ligatti, David I. August, and David Walker and has been previously published as [51] and [52].*

key principles and formal definitions for the type system. Although the typing rules are specific to the particular setting, the underlying principles are more general; we believe many of these principles will apply to reasoning about related fault-tolerant systems. The combination of a TAL-like type-theory with concepts from classical Hoare Logics is a particularly general and important technical contribution. Section 2.3 describes the key theorems we have proven including Progress, Preservation, “No False Positives,” and Fault Tolerance. Section 2.4 provides empirical evidence that our new hybrid solution to fault tolerance is feasible for many applications by measuring performance results on simulated hardware. Finally, Section 2.5 summarizes the contributions of  $TAL_{FT}$ .

## 2.1 The Fault-Tolerant Hardware

The first task is to develop the machine model so that we can reason about the execution of a program. The model hardware is based on a simple RISC architecture, extended with features to support detection of control-flow faults and safe interaction with memory-mapped output devices. Correct use of these features makes it possible to detect all faults that might change a program’s observable behavior. Most practical systems also need a fault recovery mechanism of some kind. However, because recovery is largely a secondary issue to detection, we omit recovery for now.

The general strategy is to maintain two redundant and independent threads of computation, which we name the *green* ( $G$ ) computation and the *blue* ( $B$ ) computation. The green computation generally leads slightly, and the blue computation generally trails, though there is a fair amount of flexibility in how the instructions in each computation may be interleaved. Prior to writing data out to a memory-mapped output device, the results of the two computations are checked for equivalence. If the results are not



equivalent, the machine will signal that a fault has been detected. The arguments to any control-flow transfer must also be checked for faults.

The execution of assembly programs is specified using a small-step operational semantics that maps *machine states* ( $\Sigma$ ) to other machine states. These machine states are made up of a number of components. The first component is the machine's *register bank*  $R$ , which is a total function that maps register names to the values contained therein. The metavariable  $a$  ranges over all registers, and metavariable  $r$  ranges only over general-purpose registers ( $r_1, r_2, \dots$ ). In addition to general-purpose registers, there are two program counter registers ( $pc_G$  and  $pc_B$ ), which contain the same value unless there has been the fault. There is one additional special register, the *green destination register*,  $gd$ . Its role in control-flow checking will be explained later.

To facilitate proofs of certain theorems, the value in each register is tagged with the color (either green or blue) of the computation to which it belongs. However, these tags have no effect on the run-time behavior of programs. Section 3.2.4 discusses how color tags can be removed by doing more work in the proofs. In contrast, the tags on instruction opcodes, to be introduced momentarily, *do* have an effect on evaluation.

In addition to a register bank, the machine state includes a *code memory*  $C$ , which we model as a function mapping integer addresses  $n$  to instructions. Address 0 is not considered a valid code address. The machine also has a *value memory*  $M$ , which maps addresses to integer values. In between the value memory and the processor is a special *store queue*,  $Q$ , which is used to detect faults before data is written to a memory-mapped output device. The store queue is a queue of address-value pairs. We will discuss the role of the queue in greater detail later.

Overall, an abstract machine state ( $\Sigma$ ) may have the form *fault*, indicating the hardware has detected a transient fault, or the ordinary state  $(R, C, M, Q, ir)$ , where the first

<i>colors</i>	$c$	$::= G \mid B$
<i>integers</i>	$n$	$::= \dots \mid -1 \mid 0 \mid 1 \mid \dots$
<i>colored values</i>	$v$	$::= c \ n$
<i>general registers</i>	$r$	$::= r_n$
<i>all regs</i>	$a$	$::= r \mid gd \mid pc_c$
<i>register file</i>	$R$	$::= \cdot \mid R, a \rightarrow v$
<i>code memory</i>	$C$	$::= \cdot \mid C, n \rightarrow i$
<i>value memory</i>	$M$	$::= \cdot \mid M, n \rightarrow n$
<i>store queue</i>	$Q$	$::= \overline{(n, n)}$
<i>ALU ops</i>	$op$	$::= add \mid sub \mid mul$
<i>instructions</i>	$i$	$::= op \ r_d, r_s, r_t \mid op \ r_d, r_s, v$ $\quad \mid ld_c \ r_d, r_s \mid st_c \ r_d, r_s \mid mov \ r_d, v$ $\quad \mid bz_c \ r_z, r_d \mid jmp_c \ r_d$
<i>inst register</i>	$ir$	$::= i \mid \cdot$
<i>state</i>	$\Sigma$	$::= (R, C, M, Q, ir) \mid fault$

Figure 2.1: Syntax of instructions and machine states.

four components are as discussed above, and  $ir$  is either an instruction  $i$  to be executed, or “.” indicating the next instruction should be fetched from code memory. Figure 2.1 summarizes the syntax of machine states. We use overbar notation to indicate a sequence of objects.

### 2.1.1 The Fault Model

The operational semantics is designed both to model proper execution of machine instructions and to make explicit, precise, and transparent all of the assumptions about when and where faults may occur. The central operational judgment has the form  $\Sigma_1 \xrightarrow{k}^s \Sigma_2$ , which expresses a single step transition from state  $\Sigma_1$  to state  $\Sigma_2$  while incurring  $k$  faults and writing data  $s$  to a memory-mapped output device. We will work under the standard assumption of a single upset event [57, 61] and hence  $k$  will always be either 0 or 1. The data  $s$  is a (possibly empty) sequence of address-value pairs. While the operational semantics models the internal workings of the machine, the only externally observable

behavior of the machine is the sequence of writes  $s$  to the output device or the signaling of a hardware-detected fault. If faults cause the processor to have drastically different internal behavior, but the externally observable sequence  $s$  is unchanged, we consider the program to have executed successfully.

Different fault-tolerance techniques protect different components of machines. In the literature, the protected areas are usually inside the *Sphere of Replication* (SoR) [57]. In our case, we target faults that may occur in data manipulated within the processor. We assume that both code memory  $C$  and value memory  $M$  are fully protected. This is often the case since error-correcting codes can very efficiently protect memory. To make these assumptions explicit, the following three operational rules specify exactly how faults may occur within our system.

$$\frac{R(a) = c n}{(R, C, M, Q, ir) \longrightarrow_1 (R[a \mapsto c n'], C, M, Q, ir)} \text{ (reg-zap)}$$

$$\frac{Q_1 = \overline{(n_1, n'_1)}, (m_1, m'), \overline{(n_2, n'_2)}}{Q_2 = \overline{(n_1, n'_1)}, (m_2, m'), \overline{(n_2, n'_2)}} \text{ (Q-zap1)}$$

$$(R, C, M, Q_1, ir) \longrightarrow_1 (R, C, M, Q_2, ir)$$

$$\frac{Q_1 = \overline{(n_1, n'_1)}, (m, m'_1), \overline{(n_2, n'_2)}}{Q_2 = \overline{(n_1, n'_1)}, (m, m'_2), \overline{(n_2, n'_2)}} \text{ (Q-zap2)}$$

$$(R, C, M, Q_1, ir) \longrightarrow_1 (R, C, M, Q_2, ir)$$

Rule *reg-zap* nondeterministically introduces a fault into any register by replacing the value in that register with some other arbitrary value. There are no restrictions on how the underlying value might be changed. For instance, code pointers can be changed to arbitrary integer values; references may no longer be in bounds. However, the color tag is preserved to facilitate fault-tolerance proofs. Because the color tag is fictional (has no

effect on run-time behavior), this poses no limitation on the fault model. Rules  $Q_1\text{-zap}$  and  $Q_2\text{-zap}$  alter the contents of the store queue in similar ways.

Formally, these are the only faults that can occur. However, notice that since the program counters and targets of indirect jumps are susceptible to the  $reg\text{-zap}$  rule, we effectively capture many forms of “control-flow faults” studied previously. Notice also that we do not explicitly consider faults that occur *during* execution of an instruction. However, many such faults may easily be shown to be equivalent to correct execution of an instruction composed with a fault either immediately before or afterward. For example, consider a simple register-to-register add instruction. Any fault within the adder hardware during execution of the add is equivalent to a correct add followed by a fault in the destination register. On the other hand, as we will see later, a fault during the  $st_B$  instruction cannot be modeled in our system.

An important benefit of our formal model is that there is actually a precise, concrete specification to analyze. Moreover, if a researcher wants to reason about the consequences of some fault that lives outside the formal model, this may be done by adding a new operational rule to the system and studying its semantic effect.

### 2.1.2 Instruction Semantics

The syntax of machine instructions was presented along with the rest of the components of our abstract machine in Figure 2.1. The semantics is described formally by the inference rules in Figures 2.2, 2.3, and 2.4, and explained informally below. The formal rules use several notational conventions. For instance, if  $R$  is a register file then  $R(a)$  is the contents of register  $a$  and  $R[a \mapsto v]$  is the updated register file with register  $a$  mapped to  $v$ .  $R++$  is the register file that results from incrementing both  $pc_G$  and  $pc_B$  by 1. If  $R(a)$  is the colored value  $c\ n$ , we write  $R_{val}(a)$  to denote  $n$  and  $R_{col}(a)$  to denote  $c$ . The function

*Instruction Fetch:*

$$\frac{R_{val}(pc_G) = R_{val}(pc_B) \quad R_{val}(pc_G) \in Dom(C)}{(R, C, M, Q, \cdot) \longrightarrow_0 (R, C, M, Q, C(R_{val}(pc_G)))} \text{ (fetch)}$$

$$\frac{R_{val}(pc_G) \neq R_{val}(pc_B)}{(R, C, M, Q, \cdot) \longrightarrow_0 \text{ fault}} \text{ (fetch-fail)}$$

*Basic Instructions:*

$$\frac{R' = R_{++}[r_d \mapsto R_{col}(r_t) (R_{val}(r_s) \text{ op } R_{val}(r_t))]}{(R, C, M, Q, \text{op } r_d, r_s, r_t) \longrightarrow_0 (R', C, M, Q, \cdot)} \text{ (op2r)}$$

$$\frac{R' = R_{++}[r_d \mapsto c (R_{val}(r_s) \text{ op } n)]}{(R, C, M, Q, \text{op } r_d, r_s, c n) \longrightarrow_0 (R', C, M, Q, \cdot)} \text{ (op1r)}$$

$$\frac{}{(R, C, M, Q, \text{mov } r_d, v) \longrightarrow_0 (R_{++}[r_d \mapsto v], C, M, Q, \cdot)} \text{ (mov)}$$

Figure 2.2: Operational rules for basic instructions.

$find(Q, n)$  produces the first pair  $(n, n')$  that appears in  $Q$ , or  $()$  if no pair  $(n, n')$  appears in  $Q$ .

**Instruction Fetch.** The machine operates by fetching an instruction from code memory and executing that instruction. The corresponding operational semantics rules are shown in Figure 2.2. When there is no current instruction to execute (i.e.  $ir = \cdot$ ), the *fetch* rule should fire. This rule tests for equality of the two program counters to check for faults and loads the appropriate instruction from code memory. If  $pc_G$  and  $pc_B$  are the same but  $R_{val}(pc_G)$  is not a valid address in code memory, execution “gets stuck” (no rule fires). Well-typed programs never get stuck, even when a single fault occurs. On the other hand, a fault can render the two program counters inequivalent. In this case, rule *fetch-fail* fires and causes a transition to the fault state. Abstractly, this transition represents hardware

detection of a transient fault. Controlled program termination or perhaps recovery may follow. The fault model does not allow for the instruction itself to be corrupted.

**Basic Instructions.** The arithmetic and move instructions (rules *op2r*, *op1r*, and *mov*) shown in Figure 2.2 are completely standard. The first arithmetic operation *op r<sub>d</sub>, r<sub>s</sub>, r<sub>t</sub>* performs *op* on the values in *r<sub>s</sub>* and *r<sub>t</sub>*, storing the result in *r<sub>d</sub>*. The second arithmetic operation uses a constant operand *v* in addition to *r<sub>s</sub>* and *r<sub>d</sub>*. All constants are annotated with the color of the computation they belong to. Likewise, the *mov* instruction loads an annotated constant into a register.

**Memory Instructions.** Transient faults are problematic only when they change the results of computations and those results are *observed* by an external user. In this model, the only way a result can be observed is for a program to write it to memory, where a memory-mapped output device may read and process it.

Without special hardware it appears *impossible* to guarantee that storage operations guard access to memory properly. No matter what sophisticated software checking is performed just before a conventional store instruction, it will be undone if a fault strikes between the check and execution of the store instruction.

To address this vulnerability, the machine possesses a modified store buffer (the queue *Q*), which is similar to the store buffer used in previous hardware [57] and hybrid [61] fault tolerant systems. In addition, there are two special storage instructions, each tagged with a color. The green store instruction *st<sub>G</sub> r<sub>d</sub>, r<sub>s</sub>* places the address-value pair  $(R_{val}(r_d), R_{val}(r_s))$  on the front of the queue (rule *st<sub>G</sub>-queue* in Figure 2.3). The blue store instruction *st<sub>B</sub> r<sub>d</sub>, r<sub>s</sub>* retrieves the pair  $(n_l, n'_l)$  on the back of the queue, checks that it equals  $(R_{val}(r_d), R_{val}(r_s))$ , and then stores it in memory (rule *st<sub>B</sub>-mem*). If the pairs are different, the hardware signals a fault (rule *st<sub>B</sub>-mem-fail*). Because green stores must

$$\boxed{\Sigma \xrightarrow[k]{s} \Sigma'}$$

$$\overline{(R, C, M, Q, st_G \ r_d, r_s) \longrightarrow_0 (R_{++}, C, M, ((R_{val}(r_d), R_{val}(r_s)), Q), \cdot)} \quad (st_G\text{-queue})$$

$$\frac{R_{val}(r_d) = n_l \quad R_{val}(r_s) = n'_l}{(R, C, M, ((\overline{(n, n')}, (n_l, n'_l)), st_B \ r_d, r_s) \longrightarrow_0^{(n_l, n'_l)} (R_{++}, C, M[n_l \mapsto n'_l], \overline{(n, n')}, \cdot)} \quad (st_B\text{-mem})$$

$$\frac{R_{val}(r_d) \neq n_l \text{ or } R_{val}(r_s) \neq n'_l}{(R, C, M, ((\overline{(n, n')}, (n_l, n'_l)), st_B \ r_d, r_s) \longrightarrow_0 \text{ fault}} \quad (st_B\text{-mem-fail})$$

$$\frac{\text{find}(Q, R_{val}(r_s)) = (R_{val}(r_s), n)}{(R, C, M, Q, ld_G \ r_d, r_s) \longrightarrow_0 (R_{++}[r_d \mapsto G \ n], C, M, Q, \cdot)} \quad (ld_G\text{-queue})$$

$$\frac{\text{find}(Q, R_{val}(r_s)) = () \quad R_{val}(r_s) \in \text{Dom}(M)}{(R, C, M, Q, ld_G \ r_d, r_s) \longrightarrow_0 (R_{++}[r_d \mapsto G \ M(R_{val}(r_s))], C, M, Q, \cdot)} \quad (ld_G\text{-mem})$$

$$\frac{\text{find}(Q, R_{val}(r_s)) = () \quad R_{val}(r_s) \notin \text{Dom}(M)}{(R, C, M, Q, ld_G \ r_d, r_s) \longrightarrow_0 \text{ fault}} \quad (ld_G\text{-fail})$$

$$\frac{R_{val}(r_s) \in \text{Dom}(M)}{(R, C, M, Q, ld_B \ r_d, r_s) \longrightarrow_0 (R_{++}[r_d \mapsto B \ M(R_{val}(r_s))], C, M, Q, \cdot)} \quad (ld_B\text{-mem})$$

$$\frac{R_{val}(r_s) \notin \text{Dom}(M)}{(R, C, M, Q, ld_B \ r_d, r_s) \longrightarrow_0 \text{ fault}} \quad (ld_B\text{-fail})$$

$$\frac{\text{find}(Q, R_{val}(r_s)) = () \quad R_{val}(r_s) \notin \text{Dom}(M)}{(R, C, M, Q, ld_G \ r_d, r_s) \longrightarrow_0 (R_{++}[r_d \mapsto G \ n], C, M, Q, \cdot)} \quad (ld_G\text{-rand})$$

$$\frac{R_{val}(r_s) \notin \text{Dom}(M)}{(R, C, M, Q, ld_B \ r_d, r_s) \longrightarrow_0 (R_{++}[r_d \mapsto B \ n], C, M, Q, \cdot)} \quad (ld_B\text{-rand})$$

Figure 2.3: Selected operational rules for memory instructions.

always come before blue stores, instruction scheduling is somewhat constrained. As we will show later in Section 2.4, we have evaluated the performance with and without this scheduling constraint and show that its performance impact is negligible.

As an example, consider the following straight-line sequence:

```
1 mov r1, G 5
2 mov r2, G 256
3 stG r2, r1
4 mov r3, B 5
5 mov r4, B 256
6 stB r4, r3
```

These six instructions have the effect of storing 5 into memory address 256. Moreover, a fault at any point in execution, to either blue or green values or addresses, will be caught by the hardware when the blue store (instruction 6) compares its operands to those in the queue. In addition, our instruction set gives a compiler the freedom to allocate registers however it chooses (*e.g.*, reusing registers 1 and 2 in instructions 4-6 instead of registers 3 and 4) and to change the instruction schedule in various ways (*e.g.*, moving instruction 3 to a position between instructions 5 and 6).

Interestingly, however, not all conventional optimizations are sound. This is why type checking generated code can be so helpful in detecting compiler errors. For example,



common subexpression elimination might result in the following code:

```

1 mov r1, G 5
2 mov r2, G 256
3 stG r2, r1
4 stB r2, r1

```

In this case, a fault in  $r_1$  after instruction 1, or a fault in  $r_2$  after instruction 2 will cause both instructions 3 and 4 to manipulate the same, but incorrect, address-value pair. The result would be to store an incorrect value at the correct location or a correct value at an incorrect location. Fortunately, the  $TAL_{FT}$  type system catches reliability errors like this one.

As mentioned in Section 2.1.1, many “intra-instruction” faults can be modeled by modifying the register file before or after the instruction. However, this is not the case for a fault that occurs during the execution of the  $st_B$ -mem rule in between the comparisons and the store. The hardware designer must implement structures that detect or mask any faults that occur here. If the hardware designer cannot meet the specification given by the operational semantics, he acknowledges there may be a vulnerability.

The load instructions also come in pairs:  $ld_B$  and  $ld_G$ . As Figure 2.3 shows, the only difference in their semantics is that  $ld_G$  checks for a pending store in the queue before loading its value from memory, whereas  $ld_B$  goes directly to memory, ignoring the queue. This wrinkle increases the freedom in instruction scheduling by allowing the green computation to load a value it may have recently stored before the blue computation has necessarily committed the store. Rules  $ld_G$ -queue,  $ld_G$ -mem, and  $ld_B$ -mem specify these behaviors.

Notice that there is no mechanism for verifying the address used in loads. Hence, a fault can result in an invalid address. In practice such a load might induce a hardware exception such as a segmentation fault (rules *ld<sub>G</sub>-fail* and *ld<sub>B</sub>-fail*) or might result in loading some arbitrary value (rules *ld<sub>G</sub>-rand* and *ld<sub>B</sub>-rand*).

**Control-Flow Instructions.** Any change in the control-flow of a program may cause a different sequence of values to be stored and observed by an external user. Consequently, the hardware contains mechanisms to detect faults in addresses that serve as jump targets. Intuitively, these mechanisms mirror the solution to faults in stored data in that execution of a control-flow transfer is accomplished through two instructions. Our solution uses a combination of software and hardware control-flow protection that is similar to watchdog processors [35], but that makes both versions of the control flow explicit as in software-only control flow protection [47, 60]. The operational rules that implement this solution are shown in Figure 2.4.

To achieve an unconditional jump, one executes a *jmp<sub>G</sub>* instruction first and a related *jmp<sub>B</sub>* instruction at some point in the future. A *jmp<sub>G</sub> r<sub>1</sub>* moves the destination address from *r<sub>1</sub>* into the special destination register *gd* (rule *jmp<sub>G</sub>*). Like the store queue, the destination register stores a programmer intention, initiated by the green computation. Later, when the blue computation attempts to commit the jump by executing a *jmp<sub>B</sub> r<sub>2</sub>* instruction, the contents of *r<sub>2</sub>* are compared to the contents of the destination register and if they are equal, control jumps to that location (rule *jmp<sub>B</sub>*). If the addresses are different, the hardware detects a fault (see rule *jmp<sub>B</sub>-fail*). Similar to the constraint for the store queue, forcing green control flow instructions to be executed before the corresponding blue version constrains the instruction schedule. Section 2.4 will show that this scheduling constraint has only a minimal performance impact.

$$\boxed{\Sigma \xrightarrow[k]{s} \Sigma'}$$

$$\frac{R_{val}(gd) = 0}{(R, C, M, Q, jmp_G r_d) \longrightarrow_0 (R_{++}[gd \mapsto R(r_d)], C, M, Q, \cdot)} \quad (jmp_G)$$

$$\frac{R_{val}(gd) \neq 0}{(R, C, M, Q, jmp_G r_d) \longrightarrow_0 \text{fault}} \quad (jmp_G\text{-fail})$$

$$\frac{R_{val}(gd) \neq 0 \quad R_{val}(r_d) = R_{val}(gd) \quad R' = R[pc_G \mapsto R(gd)][pc_B \mapsto R(r_d)][d \mapsto G 0]}{(R, C, M, Q, jmp_B r_d) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (jmp_B)$$

$$\frac{R_{val}(r_d) \neq R_{val}(gd) \text{ or } R_{val}(gd) = 0}{(R, C, M, Q, jmp_B r_d) \longrightarrow_0 \text{fault}} \quad (jmp_B\text{-fail})$$

$$\frac{R_{val}(gd) = 0 \quad R_{val}(r_z) \neq 0}{(R, C, M, Q, bz_c r_z, r_d) \longrightarrow_0 (R_{++}, C, M, Q, \cdot)} \quad (bz\text{-untaken})$$

$$\frac{R_{val}(gd) \neq 0 \quad R_{val}(r_z) \neq 0}{(R, C, M, Q, bz_c r_z, r_d) \longrightarrow_0 \text{fault}} \quad (bz\text{-untaken-fail})$$

$$\frac{R_{val}(gd) = 0 \quad R_{val}(r_z) = 0}{(R, C, M, Q, bz_G r_z, r_d) \longrightarrow_0 (R_{++}[gd \mapsto R(r_d)], C, M, Q, \cdot)} \quad (bz_G\text{-taken})$$

$$\frac{R_{val}(gd) \neq 0 \quad R_{val}(r_z) = 0}{(R, C, M, Q, bz_G r_z, r_d) \longrightarrow_0 \text{fault}} \quad (bz_G\text{-taken-fail})$$

$$\frac{R_{val}(gd) \neq 0 \quad R_{val}(r_z) = 0 \quad R_{val}(r_d) = R_{val}(gd) \quad R' = R[pc_G \mapsto R(gd)][pc_B \mapsto R(r_d)][gd \mapsto G 0]}{(R, C, M, Q, bz_B r_z, r_d) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (bz_B\text{-taken})$$

$$\frac{R_{val}(r_z) = 0 \quad (R_{val}(r_d) \neq R_{val}(gd) \text{ or } R_{val}(gd) = 0)}{(R, C, M, Q, bz_B r_z, r_d) \longrightarrow_0 \text{fault}} \quad (bz_B\text{-taken-fail})$$

Figure 2.4: Operational rules for control flow instructions.

The following code illustrates a typical control-flow transfer.

```

1 ldG r1, r2
2 jmpG r1
3 ldB r3, r4
4 jmpB r3

```

Initially, registers  $r_2$  and  $r_4$  should point to the same memory location, which contains a code pointer to jump to. The example illustrates some of the flexibility in scheduling jump instructions.

Conditional jumps are more complex, but follow the same principles. The green conditional  $bz_G r_z, r_d$  tests  $r_z$  and if it is 0, moves the contents of  $r_d$  into destination register  $gd$  (rules  $bz$ -*untaken* and  $bz_G$ -*taken*). No control-flow transfer occurs until a blue conditional  $bz_B r'_z, r'_d$  tests the contents of its  $r'_z$  register. If  $r'_z$  is 0 then  $r'_d$  must equal the contents of  $gd$ , and if so, the control flow transfer occurs (rule  $bz_B$ -*taken*). If  $r'_z$  is not 0, it is not good enough merely to fall through — the contents of  $r'_z$  might be faulty. To avoid this possibility, the instruction examines the destination register. If it is 0 (and hence a prior  $bz_G$  instruction did not store an address), the fall-through occurs (rule  $bz$ -*untaken*). Our metatheory will show that this mechanism suffices to detect faults either in the green computation (registers  $r_z$  and  $r_d$ ) or the blue computation (registers  $r'_z$  and  $r'_d$ ).

## 2.2 Typing

The primary goal of the  $TAL_{FT}$  type system is to ensure that well-typed programs exhibit fail-safe behavior in the presence of transient faults. In other words, well-typed programs must guarantee that even when a single fault occurs, a memory-mapped output device

## Static Expressions

<i>exp kinds</i>	$\kappa ::= \kappa_{int} \mid \kappa_{mem}$
<i>exp contexts</i>	$\Delta ::= \cdot \mid \Delta, x : \kappa$
<i>exps</i>	$E ::= x \mid n \mid E \text{ op } E \mid \text{sel } E_m E_n \mid \text{emp} \mid \text{upd } E_m E_{n_1} E_{n_2}$
<i>substitutions</i>	$S ::= \cdot \mid S, E/x$

## Types

<i>zap tags</i>	$Z ::= \cdot \mid c$
<i>basic types</i>	$b ::= int \mid \Theta \rightarrow void \mid b \text{ ref}$
<i>reg types</i>	$t ::= \langle c, b, E \rangle \mid E' = 0 \Rightarrow \langle c, b, E \rangle$
<i>reg file types</i>	$\Gamma ::= \cdot \mid \Gamma, a \rightarrow t$
<i>result types</i>	$RT ::= \Theta \mid void$

## Contexts

<i>heap typing</i>	$\Psi ::= \cdot \mid \Psi, n : b$
<i>static context</i>	$\Theta ::= \Delta; \Gamma; \overline{(E_d, E_s)}; E_m$

Figure 2.5: TAL<sub>FT</sub> type syntax.

will never read a corrupt value and make it visible to a user. We call this property “fault tolerance.”

The syntax of the type system is presented in Figure 2.5. In the following sections, we explain the intuitions and principles behind the various elements of the type system. Throughout the discussion, the reader will notice that our typing rules are not syntax-directed. In other words, there may be situations when more than one rule can be applied at a given point. Of course, as with other sorts of typed assembly language or proof-carrying code, this presents no particular difficulty in practice — during the translation, the compiler knows the correct typing derivation. The compiler can either produce this full derivation or, for efficiency’s sake, can produce sufficient “typing hints” to make type reconstruction trivial.

### 2.2.1 Static Expressions

Our “type system” is actually a combination of two theories, one being a relatively simple type theory for assembly language, inspired by previous work on TAL [40], and the second being a Hoare Logic, designed to enforce the more precise invariants required for fault tolerance. The latter component requires we define a language of *static expressions* for reasoning about values and storage.

The static expressions are drawn from the standard theory of arithmetic and arrays used in many classical Hoare Logics (*c.f.*, Necula’s thesis [45]). These static expressions are classified as either integers (kind  $\kappa_{int}$ ) or memories (kind  $\kappa_{mem}$ ). The integer expressions include variables, constants, simple arithmetic operations, and values from a memory (*sel*  $E_m E_n$  is the integer located at address  $E_n$  in  $E_m$ ). The memory expressions include variables, the empty memory (*emp*), and memory updates (*upd*  $E_m E_{n_1} E_{n_2}$  is a memory  $E_m$  updated so that address  $E_{n_1}$  stores value  $E_{n_2}$ ).

The context  $\Delta$  is a mapping from variables to kinds and contains the free variables used in the static expressions at compile time. The judgment  $\Delta \vdash E : \kappa$  classifies expression  $E$  as having kind  $\kappa$ . At runtime, these variables will have known values, so in order to reason about program execution we use a substitution  $S$  which maps expression variables to expressions. The judgment  $\Delta \vdash S : \Delta'$  holds when the substitution  $S$  maps variables in  $Dom(\Delta')$  to expressions well-formed in  $\Delta$  with types in  $Rng(\Delta')$ . The judgment  $\Delta \vdash E_1 = E_2$  is valid when  $E_1$  and  $E_2$  are equal objects in the standard model. The function  $\llbracket E \rrbracket$  supplies the denotation of the closed static expression  $E$  as either an integer or a memory, depending on its kind. The judgments are defined in Figures 2.6 and 2.7, and the denotation function is defined in Figure 2.7.

$$\boxed{\Delta \vdash E : \kappa}$$

$$\frac{x \in \text{Dom}(\Delta)}{\Delta \vdash x : \Delta(x)} \text{ (E-var-t)}$$

$$\frac{}{\Delta \vdash n : \kappa_{int}} \text{ (E-int-t)} \quad \frac{\Delta \vdash E_1 : \kappa_{int} \quad \Delta \vdash E_2 : \kappa_{int}}{\Delta \vdash E_1 \text{ op } E_2 : \kappa_{int}} \text{ (E-op-t)}$$

$$\frac{}{\Delta \vdash \text{emp} : \kappa_{mem}} \text{ (E-emp-t)} \quad \frac{\Delta \vdash E_m : \kappa_{mem} \quad \Delta \vdash E_n : \kappa_{int}}{\Delta \vdash \text{sel } E_m E_n : \kappa_{int}} \text{ (E-sel-t)}$$

$$\frac{\Delta \vdash E_m : \kappa_{mem} \quad \Delta \vdash E_{n_1} : \kappa_{int} \quad \Delta \vdash E_{n_2} : \kappa_{int}}{\Delta \vdash \text{upd } E_m E_{n_1} E_{n_2} : \kappa_{mem}} \text{ (E-upd-t)}$$

$$\boxed{\Delta \vdash S : \Delta'}$$

$$\frac{}{\Delta \vdash \cdot : \cdot} \text{ (sub-emp-t)}$$

$$\frac{\Delta \vdash S' : \Delta'' \quad \Delta \vdash E : \kappa \quad x \notin \text{Dom}(\Delta) \cup \text{Dom}(\Delta'')}{\Delta \vdash S', E/x : \Delta'', x : \kappa} \text{ (sub-t)}$$

Figure 2.6: Semantics of Static Expressions, Part 1.

$\llbracket E \rrbracket$ 

$$\begin{aligned}
 \llbracket n \rrbracket &= n \\
 \llbracket emp \rrbracket &= \cdot \\
 \llbracket E_1 \text{ op } E_2 \rrbracket &= \llbracket E_1 \rrbracket \text{ op } \llbracket E_2 \rrbracket \\
 \llbracket sel E_m E_n \rrbracket &= \llbracket E_m \rrbracket (\llbracket E_n \rrbracket) \\
 \llbracket upd E_m E_1 E_2 \rrbracket &= \llbracket E_m \rrbracket [ \llbracket E_1 \rrbracket \mapsto \llbracket E_2 \rrbracket ]
 \end{aligned}$$

 $\Delta \vdash E_1 = E_2$ 

$$\frac{\Delta \vdash E_1 : \kappa_{int} \quad \Delta \vdash E_2 : \kappa_{int} \quad \forall S. \cdot \vdash S : \Delta \implies \llbracket S(E_1) \rrbracket = \llbracket S(E_2) \rrbracket}{\Delta \vdash E_1 = E_2} \text{ (E-eq)}$$

$$\frac{\Delta \vdash E_1 : \kappa_{int} \quad \Delta \vdash E_2 : \kappa_{int} \quad \forall S. \cdot \vdash S : \Delta \implies \llbracket S(E_1) \rrbracket \neq \llbracket S(E_2) \rrbracket}{\Delta \vdash E_1 \neq E_2} \text{ (E-neq)}$$

$$\frac{\Delta \vdash E_1 : \kappa_{mem} \quad \Delta \vdash E_2 : \kappa_{mem} \quad \forall \ell \in \text{Dom}(\llbracket S(E_1) \rrbracket) \cup \text{Dom}(\llbracket S(E_2) \rrbracket). \llbracket S(E_1) \rrbracket(\ell) = \llbracket S(E_2) \rrbracket(\ell)}{\Delta \vdash E_1 = E_2} \text{ (E-mem-eq)}$$

Figure 2.7: Semantics of Static Expressions, Part 2.



### 2.2.2 Value Typing

Because faults strike values, corrupting their bit patterns in arbitrary ways, the subtleties of value typing are a key concern. Informally, the type system maintains three key pieces of information about every value:

1. *A color (green or blue).* The type system is organized to ensure that when a value is known to be green, its contents can only depend on the contents of other green values not blue ones, and likewise, blue can only depend upon blue. Hence, while a fault in a green value can eventually corrupt arbitrarily many other green values, it cannot corrupt any blue values, and vice versa.
2. *A “basic type”.* When no fault has occurred in the value’s color, the value’s basic type describes its shape. Values with type *int* may have any bit pattern. Values with type  $\Theta \rightarrow \text{void}$  are pointers to code (continuations). One must satisfy the precondition  $\Theta$  before jumping to them. Values with type *b ref* are pointers to values with type *b*.
3. *A static expression.* When there has been no fault in a value’s color, the value exactly equals the static expression. Static expressions are used to guarantee that in the absence of faults, the green and blue computations produce equal values, and hence, dynamic fault detection checks always succeed.

To summarize, every value is typed using a triple  $\langle c, b, E \rangle$ , where *c* is a color, *b* is a basic type, and *E* is a static expression. The presence of the static expression makes this type a sort of singleton type.

**Value Typing Judgment.** The value typing judgment in Figure 2.8 has the form  $\Psi; \Delta \vdash^Z v : t$ , where  $\Psi$  maps heap addresses to basic types, and  $\Delta$  contains the free expression

$$\boxed{\Psi \vdash n : b}$$

$$\overline{\Psi \vdash n : int} \quad (int-t) \qquad \overline{\Psi \vdash n : \Psi(n)} \quad (base-t)$$

$$\boxed{\Psi; \Delta \vdash^Z v : t}$$

$$\frac{\Psi \vdash n : b \quad \Delta \vdash E = n}{\Psi; \Delta \vdash^Z c n : \langle c, b, E \rangle} \quad (val-t)$$

$$\frac{n \neq 0 \quad \Psi; \Delta \vdash^Z c n : \langle c, b, E \rangle \quad \Delta \vdash E' = 0}{\Psi; \Delta \vdash^Z c n : E' = 0 \Rightarrow \langle c, b, E \rangle} \quad (cond-t)$$

$$\frac{\Delta \vdash E' \neq 0}{\Psi; \Delta \vdash^Z c 0 : E' = 0 \Rightarrow \langle c, b, E \rangle} \quad (cond-t-n0)$$

$$\frac{\Delta \vdash E : \kappa_{int}}{\Psi; \Delta \vdash^c c n : \langle c, b, E \rangle} \quad (val-zap-t)$$

$$\frac{\Delta \vdash E : \kappa_{int}}{\Psi; \Delta \vdash^c c n : E' = 0 \Rightarrow \langle c, b, E \rangle} \quad (val-zap-cond)$$

Figure 2.8: Value Typing.

variables. In the rule *val-t*, a colored value  $c\ n$  is given the type  $\langle c, b, E \rangle$  when the static expression  $E$  is equal to  $n$ , and  $\Psi \vdash n : b$ . The judgment  $\Psi \vdash n : b$  allows  $n$  to be given either the basic type *int* or the type of the address  $n$  in memory.

The two rules *cond-t* and *cond-t-n0* are used to type the conditional type  $(E' = 0 \Rightarrow \langle G, \Theta \rightarrow void, E'_r \rangle)$ . When the static expression  $E'$  is equal to zero, values of this type also have type  $\langle G, \Theta \rightarrow void, E'_r \rangle$ . When  $E'$  is not equal to zero, values with this type must be 0.

The final two rules for  $\Psi; \Delta \vdash^Z v : t$  make use of the *zap tag*  $Z$ , which is either empty or a color  $c$ . If the zap tag is a color  $c$ , then there may have been a fault affecting data of that color. Data colored the same as the zap tag can be given any type, as it may have been arbitrarily corrupted. The static expression used in this type may not contain any free expression variables.

**Value Subtyping.** There is also a subtyping relation  $\Delta \vdash t \leq t'$  that allows all types  $\langle c, b, E_1 \rangle$  to be subtypes of  $\langle c, int, E_2 \rangle$  when  $\Delta \vdash E_1 = E_2$ . This relation is extended to register file subtyping  $\Delta \vdash \Gamma_1 \leq \Gamma_2$ , by requiring that the type of each general-purpose register in  $\Gamma_2$  be a supertype of the corresponding register in  $\Gamma_1$ . Note that here is no required relationship between the special registers *gd*, *pc<sub>G</sub>*, and *pc<sub>B</sub>*. The rules for these judgments appear in Figure 2.9.

### 2.2.3 Instruction Typing

While many of the instruction typing rules are quite complex, the essential principles guiding their construction may be summarized as follows.

1. *In the absence of faults, standard type theoretic principles should be valid.* In order to guarantee basic safety properties, the type system checks standard properties in

$$\boxed{\Delta \vdash t \leq t'}$$

$$\frac{\Delta \vdash E_1 = E_2}{\Delta \vdash \langle c, b, E_1 \rangle \leq \langle c, b, E_2 \rangle} \text{ (subtp-triple)}$$

$$\frac{\Delta \vdash E_1 = E_2}{\Delta \vdash \langle c, b, E_1 \rangle \leq \langle c, \text{int}, E_2 \rangle} \text{ (subtp-int)}$$

$$\frac{\Delta \vdash t \leq t' \quad \Delta \vdash E = E'}{\Delta \vdash (E = 0 \Rightarrow t) \leq (E' = 0 \Rightarrow t')} \text{ (subtp-cond)}$$

$$\boxed{\Delta \vdash \Gamma_1 \leq \Gamma_2}$$

$$\frac{\forall r \in \text{Dom}(\Gamma_2). \Gamma_1(r) \leq \Gamma_2(r)}{\Delta \vdash \Gamma_1 \leq \Gamma_2} \text{ (reg-file-comp)}$$

Figure 2.9: Subtyping.

much the same manner as previous typed assembly languages [40]. For example, jump targets must have code types, while loads and stores must operate over values with reference types.

2. *Green values only depend on other green values, and blue values only depend on blue values.* When this invariant is maintained, a fault in a blue value can never corrupt a green value and vice versa.
3. *Both green and blue computations have equal say in any dangerous actions.* Dangerous actions include storing values to memory-mapped output devices and executing control-flow operations. When both blue and green computations are involved, a fault in just one color is insufficient to deceive the hardware fault detection mechanisms.

4. *In the absence of faults, green and blue computations must compute identical values.* To be more precise, green and blue computations must store identical values to identical storage locations and must issue orders to transfer control to identical addresses. If not, the hardware will claim to detect faults when there have been none, or alternatively, might exhibit incorrect behaviors when there is a fault.

The first three principles are relatively straightforward to enforce. The fourth principle leads to the most technical challenges as it requires we check equality constraints between values. Moreover, since construction of these values depends on storage, the type system must maintain a relatively accurate static representation of storage. We accomplish this latter challenge using techniques drawn from Hoare Logics. The former challenge (testing values for equality) is achieved through the use of the singleton types described earlier.

**The Instruction Typing Judgment.** The judgment for typing instructions has the form  $\Psi; \Theta \vdash ir \Rightarrow RT$ . Unlike the context  $\Psi$ , which only contains invariant heap typing assumptions,  $\Theta$  contains fine-grained context-sensitive information about the current state of memory and the register file. More specifically,  $\Theta$  consists of the following sub-contexts: (1)  $\Delta$ , which describes the free expression variables appearing in the other context-sensitive objects, (2)  $\Gamma$ , which describes the mapping of register names to types for register values, (3)  $\overline{(E_d, E_s)}$ , which describes the values in the queue, and (4)  $E_m$ , which describes memory, as one does in Hoare Logic.

The “result” of checking an instruction is a result type  $RT$ . A result type may either be *void*, indicating control does not proceed past the instruction (it is a jump), or a postcondition  $\Theta'$ , which describes the state of memory and the register file after execution of the instruction.

The typing rules are defined using several notational abbreviations. The notation  $\Gamma_{++}$  adds one to the static expression associated with each program counter register in  $\Gamma$ . The expression  $\overline{upd} E_m \overline{(E_d, E_s)}$  is  $(upd \dots (upd E_m E_{d_k} E_{s_k}) \dots) E_{d_1} E_{s_1}$  when  $\overline{(E_d, E_s)} = ((E_{d_1}, E_{s_1}), \dots, (E_{d_k}, E_{s_k}))$ . Figures 2.10, 2.11, and 2.12 presents the typing rules for instructions, and the following paragraphs explain the main points of interest.

**Typing Basic Instructions.** Basic arithmetic operations are not “dangerous” to execute, so the definitions of their typing rules are driven by principles 1 and 2, mentioned earlier. Consider, for example, rule  $op2r-t$  for an arithmetic operation  $op$ . This rule requires that the operand registers contain integers with the same color  $c$  in accordance with principal 2 (green depends on green, blue depends on blue). The result register  $r_d$  has a type colored  $c$  as well. In accordance with principle 1, the result has integer type. The rule also states that the static expression describing the result register is  $E'_s op E'_t$  and that the state of the queue and memory are unchanged by evaluation of the instruction.

**Typing Memory Instructions.** Store operations are “dangerous” — they make computed values observable by the outside world — so we must be particularly careful in the formulation of their typing rules. In accordance with principle 1, both green and blue store instructions (rules  $st_G-t$  and  $st_B-t$ ) require that the address register has the basic type  $b ref$  and the value register has the corresponding basic type  $b$ . Intuitively, the store queue is a green object, and in accordance with principle 2, the green store instruction may push an address-value pair onto the front the queue as long as the address and the value are green. In accordance with principle 4, the rule for the blue store checks that the address-value pair to be stored is exactly equal to the address-value pair at the end of the queue. Since the arguments to the blue store have a blue type and the queue always

$$\boxed{\Psi; \Theta \vdash ir \Rightarrow RT}$$

$$\frac{}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash \cdot \Rightarrow (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m)} (\cdot\text{-}t)$$

$$\frac{\Gamma(r_s) = \langle c, \text{int}, E'_s \rangle \quad \Gamma(r_t) = \langle c, \text{int}, E'_t \rangle}{\Gamma' = \Gamma_{++}[r_d \mapsto \langle c, \text{int}, E'_s \text{ op } E'_t \rangle]} (op2r\text{-}t)$$

$$\frac{}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash \text{op } r_d, r_s, r_t \Rightarrow (\Delta; \Gamma'; \overline{(E_d, E_s)}; E_m)}$$

$$\frac{\Gamma(r_s) = \langle c, \text{int}, E'_s \rangle}{\Gamma' = \Gamma_{++}[r_d \mapsto \langle c, \text{int}, E'_s \text{ op } n \rangle]} (op1r\text{-}t)$$

$$\frac{}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash \text{op } r_d, r_s, c \ n \Rightarrow (\Delta; \Gamma'; \overline{(E_d, E_s)}; E_m)}$$

$$\frac{\Psi; \Delta \vdash v : t}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash \text{mov } r_d, v \Rightarrow (\Delta; \Gamma_{++}[r_d \mapsto t]; \overline{(E_d, E_s)}; E_m)} (mov\text{-}t)$$

$$\frac{\Gamma(r_s) = \langle G, b \text{ ref}, E'_s \rangle \quad E = \text{sel } (\overline{\text{upd}} E_m \overline{(E_d, E_s)}) E'_s}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash \text{ld}_G r_d r_s \Rightarrow (\Delta; \Gamma_{++}[r_d \mapsto \langle G, b, E \rangle]; \overline{(E_d, E_s)}; E_m)} (ld_G\text{-}t)$$

$$\frac{\Gamma(r_s) = \langle B, b \text{ ref}, E'_s \rangle \quad E = \text{sel } E_m E'_s}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash \text{ld}_B r_d r_s \Rightarrow (\Delta; \Gamma_{++}[r_d \mapsto \langle B, b, E \rangle]; \overline{(E_d, E_s)}; E_m)} (ld_B\text{-}t)$$

$$\frac{\Gamma(r_d) = \langle G, b \text{ ref}, E'_d \rangle \quad \Gamma(r_s) = \langle G, b, E'_s \rangle}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash \text{st}_G r_d r_s \Rightarrow (\Delta; \Gamma_{++}; (E'_d, E'_s), \overline{(E_d, E_s)}; E_m)} (st_G\text{-}t)$$

$$\frac{\Gamma(r_d) = \langle B, b \text{ ref}, E''_d \rangle \quad \Gamma(r_s) = \langle B, b, E''_s \rangle}{\Delta \vdash E'_d = E''_d \quad \Delta \vdash E'_s = E''_s} (st_B\text{-}t)$$

$$\frac{}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}, (E'_d, E'_s); E_m) \vdash \text{st}_B r_d r_s \Rightarrow (\Delta; \Gamma_{++}; \overline{(E_d, E_s)}; \text{upd } E_m E'_d E'_s)}$$

Figure 2.10: Instruction Typing Rules for Basic Instructions and Memory Instructions.

contains green objects, both blue and green computations contribute to the actual storage operation (in accordance with principle 3).

The load operations are somewhat simpler than the store instructions since they are not “dangerous” in our model. However, like the store instructions, the operands of blue loads must be blue and the operands of green loads must be green. Once again, in accordance with principle 2, the result of a blue load is value with a blue type and likewise for a green load. Other than the colors, the main difference between two load typing rules is the way they obtain the expression that described the loaded value. The blue load (rule  $ld_B-t$ ) simply uses  $sel\ E_m\ E'_s$ , which is the expression that describes the contents of the address described by  $E'_s$  in the memory description  $E_m$ . However, the operational rule for  $ld_G$  first checks the queue for any pending stores before loading the value from memory, and so the typing rule  $ld_G-t$  uses the expression  $sel\ (\overline{upd}\ E_m\ \overline{(E_d, E_s)})\ E'_s$ . In other words, before taking the corresponding expression from  $E_m$ , all the pending changes in the queue (described by  $\overline{(E_d, E_s)}$ ) are applied. This way, if the loaded value also has a pending update, the selected expression will describe the updated value.

**Typing Control-Flow Instructions.** The rules for the green jump and branch instructions are relatively simple, as they do not actually involve a control flow transfer.

The simplest situation involves the green jump (rule  $jmp_G-t$ ). This instruction is just a move from register  $r_d$  to the special destination register  $gd$ . The type of register  $gd$  is updated to the type of  $r_d$  (obeying both principles 1 and 2). The rule contains constraints that  $gd$  must be equal to 0 in both  $\Gamma$  and  $\Gamma'$  since the hardware resets the destination register to 0 after a jump.

The typing of the green conditional branch (rule  $bz_G-t$ ) is quite similar to that of the green jump. One difference is that the  $bz_G$  instruction is now a conditional move as



$$\boxed{\Psi; \Theta \vdash ir \Rightarrow RT}$$

$$\begin{array}{c}
 \Gamma(gd) = \langle G, int, 0 \rangle \\
 \Gamma(r_d) = \langle G, \Theta \rightarrow void, E_{rd'} \rangle \\
 \Theta = (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m) \\
 \Gamma'(gd) = \langle G, int, 0 \rangle \\
 \hline
 \Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash jmp_G r_d \\
 \Rightarrow (\Delta; \Gamma_{++}[gd \mapsto \langle G, \Theta \rightarrow void, E_{rd'} \rangle]; \overline{(E_d, E_s)}; E_m)
 \end{array} \quad (jmp_G-t)$$

$$\begin{array}{c}
 \Gamma(gd) = \langle G, (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m) \rightarrow void, E'_r \rangle \\
 \Gamma(r_d) = \langle B, (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m) \rightarrow void, E_r \rangle \\
 \Delta \vdash E_r = E'_r \\
 \exists S. \Delta \vdash S : \Delta' \\
 S(\Gamma')(gd) = \langle G, int, 0 \rangle \\
 S(\Gamma')(pc_G) = \langle G, int, E'_r \rangle \\
 S(\Gamma')(pc_B) = \langle B, int, E_r \rangle \\
 \Delta \vdash \Gamma \leq S(\Gamma') \\
 \Delta \vdash \overline{(E_d, E_s)} = S(\overline{(E'_d, E'_s)}) \\
 \Delta \vdash E_m = S(E'_m) \\
 \hline
 \Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash jmp_B r_d \Rightarrow void
 \end{array} \quad (jmp_B-t)$$

Figure 2.11: Instruction Typing Rules for Jump Instructions.

$$\boxed{\Psi; \Theta \vdash ir \Rightarrow RT}$$

$$\begin{array}{c}
 \Gamma(gd) = \langle G, int, 0 \rangle \\
 \Gamma(r_z) = \langle G, int, E_z \rangle \\
 \Gamma(r_d) = \langle G, \Theta \rightarrow void, E'_d \rangle \\
 \Theta = (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m) \\
 \Gamma'(gd) = \langle G, int, 0 \rangle \\
 \hline
 \Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash bz_G r_z r_d \quad (bz_G-t) \\
 \Rightarrow (\Delta; \Gamma_{++}[gd \mapsto E_z = 0 \Rightarrow \langle G, \Theta \rightarrow void, E'_d \rangle]; \overline{(E_d, E_s)}; E_m)
 \end{array}$$

$$\begin{array}{c}
 \Gamma(r_z) = \langle B, int, E_z \rangle \\
 \Gamma(r_d) = \langle B, (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m) \rightarrow void, E_r \rangle \\
 \Gamma(gd) = E'_z = 0 \Rightarrow \langle G, (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m) \rightarrow void, E'_r \rangle \\
 \Delta \vdash E_z = E'_z \\
 \Delta \vdash E_r = E'_r \\
 \exists S. \Delta \vdash S : \Delta' \\
 S(\Gamma')(gd) = \langle G, int, 0 \rangle \\
 S(\Gamma')(pc_G) = \langle G, int, E'_r \rangle \\
 S(\Gamma')(pc_B) = \langle B, int, E_r \rangle \\
 \Delta \vdash \Gamma \leq S(\Gamma') \\
 \Delta \vdash \overline{(E_d, E_s)} = S(\overline{(E'_d, E'_s)}) \\
 \Delta \vdash E_m = S(E'_m) \\
 \hline
 \Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash bz_B r_z r_d \Rightarrow (\Delta; \Gamma_{++}; \overline{(E_d, E_s)}; E_m) \quad (bz_B-t)
 \end{array}$$

Figure 2.12: Instruction Typing Rules for Branch Instructions.

opposed to an unconditional move. Hence, to represent the result of the move (unknown at compile time) the conditional type  $(E'_z = 0 \Rightarrow \langle G, \Theta \rightarrow void, E'_r \rangle)$  is used. In addition, since the conditional branch may fall through, the result of typing the  $bz_G$  instruction is a proper postcondition as opposed to  $void$ , like  $jmp_G$ .

The typing rules for the blue control-flow instructions continue to follow the same four principles as the other instructions. Much of the complexity is inherently due to principle 1, which mandates checking all the usual constraints associated with jumps in any typed assembly language. For comparison, a jump rule in a non-fault-tolerant typed assembly language requires that the target register contains a code pointer, and that the current register file is a subtype of that expected by the target code. (To be consistent with TAL<sub>FT</sub>, the subtyping relationship below does not include the program counter register, and the requirement for the program counter is written separately.)

$$\frac{\begin{array}{l} \Gamma(r_d) = \Gamma' \rightarrow void \\ \Gamma'(pc) = int \\ \Delta \vdash \Gamma \leq \Gamma' \end{array}}{\Psi; (\Delta; \Gamma) \vdash jmp\ r_d \Rightarrow void} \text{ (conventional-jmp-t)}$$

The blue jump is a true jump. According to principle 1, the typing rule  $jmp_{B-t}$  checks the standard typing invariants needed to ensure safety in any typed assembly language, including (1) that both copies of the jump target have a code type (see the first two premises), and (2) that the current state, including register file, memory, and queue, matches the expected state at the jump target, modulo some substitution  $S$  of static expressions for universally quantified variables  $\Delta$  from the code type (see the final seven premises).

The typing of the blue conditional branch (rule  $bz_{B-t}$ ) is similar to the jump. The conditional type of the destination register  $gd$  must depend on an expression that is equivalent to the expression used in the blue branch. And again, the instruction may fall through, so the rule has a proper postcondition.

## 2.2.4 Machine State Typing

In order to prove various properties of the type system, we need to specify the invariants of machine states that are preserved during execution. The judgments for typing a machine state  $\Sigma$  are shown in Figures 2.13 and 2.14 and explained below.

**Register File Typing.** The judgment  $\Psi \vdash^Z R : \Gamma$  states that the register file  $R$  has the register file type  $\Gamma$  under heap typing  $\Psi$  and a zap tag  $Z$ . The contents of each register must have the type given to that register by  $\Gamma$ . Each program counter must have the appropriate color, and the program counters must contain equal values. (In the case where one program counter is corrupted, the zap tag  $Z$  in the first premise allows its actual value to differ from the expected computed value.)

**Code Typing.** The judgment  $\Psi \vdash C$  states that code memory  $C$  is well-formed with respect to heap typing  $\Psi$ . The address 0 is not a valid code address. Each address must have a code type, and the code type must contain the precondition for the instruction at that address. If the instruction typing results in a postcondition  $\Theta'$  (meaning that control may fall through to the next instruction) then the subsequent instruction must be well typed using  $\Theta'$  as its precondition.

**Memory Typing.** The judgment  $\Psi \vdash M : E_m$  states that given heap typing  $\Psi$  the value memory  $M$  is well-formed and can be described by the static expression  $E_m$ . The static

$$\boxed{\Psi \vdash^Z R : \Gamma}$$

$$\frac{\forall a. \Psi; \cdot \vdash^Z R(a) : \Gamma(a) \quad \cdot \vdash \Gamma(pc_G) \leq \langle G, int, E_G \rangle \quad \cdot \vdash \Gamma(pc_B) \leq \langle B, int, E_B \rangle \quad \cdot \vdash E_G = E_B}{\Psi \vdash^Z R : \Gamma} \text{ (R-t)}$$

$$\boxed{\Psi \vdash C}$$

$$\frac{\begin{array}{l} 0 \notin Dom(C) \\ \forall n \in Dom(C). \Psi(n) = \Theta \rightarrow void \wedge \Psi; \Theta \vdash C(n) \Rightarrow RT \\ \wedge (RT = \Theta' \text{ implies } \Psi(n+1) = \Theta' \rightarrow void) \end{array}}{\Psi \vdash C} \text{ (C-t)}$$

$$\boxed{\Psi \vdash M : E_m}$$

$$\frac{\cdot \vdash E_m : \kappa_{mem} \quad \llbracket E_m \rrbracket = M \quad \forall \ell \in Dom(M). \Psi \vdash \ell : b \text{ ref} \wedge \Psi \vdash M(\ell) : b}{\Psi \vdash M : E_m} \text{ (M-t)}$$

$$\boxed{\Psi \vdash^Z Q : \overline{(E_d, E_s)}}$$

$$\frac{\overline{\Psi \vdash^Z () : ()}}{\Psi \vdash^Z () : ()} \text{ (Q-emp-t)}$$

$$\frac{\begin{array}{l} Z \neq G \\ \Psi \vdash n_1 : b \text{ ref} \quad \Psi \vdash n_2 : b \quad \cdot \vdash E_d = n_1 \quad \cdot \vdash E_s = n_2 \\ \Psi \vdash^Z \overline{(n'_1, n'_2)} : \overline{(E'_d, E'_s)} \end{array}}{\Psi \vdash^Z (n_1, n_2), \overline{(n'_1, n'_2)} : (E_d, E_s), \overline{(E'_d, E'_s)}} \text{ (Q-t)}$$

$$\frac{\cdot \vdash E_d : \kappa_{int} \quad \cdot \vdash E_s : \kappa_{int} \quad \Psi \vdash^G \overline{(n'_1, n'_2)} : \overline{(E'_d, E'_s)}}{\Psi \vdash^G (n_1, n_2), \overline{(n'_1, n'_2)} : (E_d, E_s), \overline{(E'_d, E'_s)}} \text{ (Q-zap-t)}$$

Figure 2.13: Machine State Element Typing.

$$\boxed{\vdash^Z (R, C, M, Q, ir)}$$

$$\begin{array}{c}
 \text{Dom}(\Psi) = \text{Dom}(C) \cup \text{Dom}(M) \\
 Z \neq G \implies \text{Dom}(Q) \subseteq \text{Dom}(M) \\
 \Psi \vdash C \\
 \forall c \neq Z. ir \neq \cdot \implies C(R_{val}(pc_c)) = ir \\
 \forall c \neq Z. \Psi(R_{val}(pc_c)) = (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \rightarrow \text{void} \\
 \exists S. \cdot \vdash S : \Delta \\
 \Psi \vdash M : S(E_m) \\
 \Psi \vdash^Z Q : \overline{S(E_d, E_s)} \\
 \Psi \vdash^Z R : S(\Gamma) \\
 \hline
 \vdash^Z (R, C, M, Q, ir) \quad (\Sigma\text{-}t)
 \end{array}$$

Figure 2.14: Machine State Typing.

expression  $E_m$  must have kind  $\kappa_{mem}$ , and  $M$  must be the denotation of  $E_m$ . Each location in the domain of  $M$  must have a type  $b \text{ ref}$  and the contents of that location must have type  $b$ .

**Queue Typing.** The judgment  $\Psi \vdash^Z Q : \overline{(E_d, E_s)}$  means that queue  $Q$  can be described by the sequence of static expressions  $\overline{(E_d, E_s)}$  given heap typing  $\Psi$  and zap tag  $Z$ . When the queue is empty, it is described by the empty sequence. When the zap tag  $Z$  is not  $G$ , the first pair  $(n_1, n_2)$  must consist of an address  $n_1$  with type  $b \text{ ref}$  and a value  $n_2$  with type  $b$ . This pair is described by the static expression pair  $(E_d, E_s)$  when  $E_d$  evaluates to  $n_1$  and  $E_s$  evaluates to  $n_2$ . The remainder of the queue must be described by the remainder of the static expression sequence. All values in the queue are considered to be green, so when the zap tag is  $G$ , these values may have been arbitrarily corrupted. Accordingly in this case, the only requirements are that each static expression must have kind  $\kappa_{int}$  and the length of the queue must be the same as the length of the static expression sequence.

**Machine State Typing.** The judgment  $\vdash^Z \Sigma$  states that a machine state  $\Sigma$  is well-typed under zap tag  $Z$ . This judgment holds when  $\Sigma$  is a five-tuple  $(R, C, M, Q, ir)$ , and these elements are each well-typed and consistent with each other. Note that  $\Sigma$  is not well-typed when it is the fault state *fault*.

The domain of the heap typing  $\Psi$  must be the union of the domains of the code memory  $C$  and the value memory  $M$ . When  $Z$  is not equal to  $G$ , the queue has not been corrupted, and so each address in  $Q$  is also a valid address in  $M$ . The code memory  $C$  must be well-formed with respect to the heap typing  $\Psi$ .

The zap tag  $Z$  must be either empty or colored blue or green. At least one of the program counters, and possibly both, will not be colored by  $Z$ , and therefore must not be corrupted. If  $ir$  is an instruction, these correct program counters must point to an address that contains  $ir$ . The heap typing  $\Psi$  gives this address a code type  $(\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \rightarrow \text{void}$ .

The context  $\Delta$  contains the free variables in  $\Gamma$ ,  $\overline{(E_d, E_s)}$ , and  $E_m$ . There must be some substitution  $S$  that gives values to each of the variables in  $\Delta$ . Value memory  $M$  must be well-formed and described by the static expression  $S(E_m)$ . The queue  $Q$  must be described by  $S(\overline{(E_d, E_s)})$ . The register file  $R$  must have type  $S(\Gamma)$ .

## 2.3 Formal Results

Now that we have defined the machine model and the type system, we need to prove that the type system is sound with respect to the machine model. In other words, if a program is well-typed, we can guarantee certain properties about all possible executions of that program. This section provides a high-level overview of the formal results for  $TAL_{FT}$ . A more in-depth discussion appears in Appendix A.

In order to prove properties about program execution, we extend our single-step transition  $\Sigma_1 \xrightarrow{k}^s \Sigma_2$  from Section 2.1 to a sequence of  $n$  transitions containing exactly  $k$  faults  $\Sigma_1 \xrightarrow{n}^s \Sigma_2$ , where  $n$  is greater than or equal to zero, and  $k$  is still either 0 or 1.

### 2.3.1 Type Safety

Progress states that well-typed states can take a step. In particular, a machine state that is well-typed under the empty zap tag can take a step to another machine state. A machine state that is well-typed under a zap tag of color  $c$  can take a step, but the result of that step may either be another machine state or the *fault* state. (Recall that  $\vdash^Z \Sigma$  only holds when  $\Sigma$  is a five-tuple  $(R, C, M, Q, ir)$  and not the fault state *fault*.)

#### Theorem 1 (Progress)

1. If  $\vdash \Sigma$  then  $\Sigma \xrightarrow{0}^s \Sigma'$  and  $\Sigma' \neq \text{fault}$ .
2. If  $\vdash^c \Sigma$  then  $\Sigma \xrightarrow{0}^s \Sigma$ .

According to Preservation, if a machine state is well-typed under a zap tag  $Z$ , and it takes a non-faulty step to another machine state, then that resulting state will also be well-typed under  $Z$ . Additionally, if a state is well-typed under the empty zap tag, and it takes a faulty step, then there is some color  $c$  such that the resulting state is well-typed under  $c$ .

#### Theorem 2 (Preservation)

1. If  $\vdash^Z \Sigma$  and  $\Sigma \xrightarrow{0}^s \Sigma'$  and  $\Sigma' \neq \text{fault}$  then  $\vdash^Z \Sigma'$ .
2. If  $\vdash \Sigma$  and  $\Sigma \xrightarrow{1}^s \Sigma'$  then  $\exists c. \vdash^c \Sigma'$ .

Progress and Preservation define the usual notion of type safety. In addition, part one of Progress, together with part one of Preservation entail the following important



corollary: The hardware never claims to have detected a fault when no fault has occurred during execution of a well-typed program.

**Corollary 3 (No False Positives)**

If  $\vdash \Sigma$  then  $\forall n. \Sigma \xrightarrow[n]{s} \Sigma'$  and  $\vdash \Sigma'$ .

### 2.3.2 Fault Tolerance

A program is fault tolerant when all the faulty executions of that program *simulate* fault-free executions of the program. More precisely, the sequence of outputs from the faulty execution is required either to be identical to the fault-free execution or, in the case the hardware detects the fault, a prefix of the fault-free execution.

In order to reason about pairs of faulty and fault-free executions, we define similarity relations between values, register files, queues and machine states. Each of these relations is defined relative to the zap tag  $Z$ . Intuitively, if  $Z$  is empty, the related objects must be identical. If  $Z$  is a color  $c$ , the objects must be identical modulo values colored  $c$ . In the latter case, values colored  $c$  may be corrupted, and there is no hope they satisfy any particular relation. The formal definitions of these relations are shown in Figure 2.15.

Using the simulation relation, we can state and prove the fault tolerance theorem for well-typed programs precisely. Assume that machine state  $\Sigma$  is well-typed under the empty zap tag, and non-faulty execution of  $\Sigma$  for  $n$  steps results in a state  $\Sigma'$  and outputs a sequence of value-address pairs  $s$ . If somewhere during that execution a single fault is encountered, the faulty execution will either run for  $n + 1$  steps or terminate in the fault state during that time. If the faulty execution takes  $n + 1$  steps and reaches the non-faulty state  $\Sigma'_f$ , then  $\Sigma'$  simulates  $\Sigma'_f$  and the sequence of output pairs is identical the original execution. Alternatively, if the faulty execution reaches the fault state then the output pairs will be a prefix of the non-faulty output pairs.

$$\begin{array}{c}
 \boxed{v_1 \text{ sim}^Z v_2} \\
 \\
 \frac{}{C \ n \ \text{sim}^Z \ C \ n} \text{ (sim-val)} \quad \frac{}{C \ n \ \text{sim}^C \ C \ n'} \text{ (sim-val-zap)} \\
 \\
 \boxed{R \ \text{sim}^Z \ R'} \\
 \\
 \frac{\forall a. R(a) \ \text{sim}^Z \ R'(a)}{R \ \text{sim}^Z \ R'} \text{ (sim-R)} \\
 \\
 \boxed{Q \ \text{sim}^Z \ Q'} \\
 \\
 \frac{}{\cdot \ \text{sim}^Z \ \cdot} \text{ (sim-Q-empty)} \\
 \\
 \frac{G \ n_1 \ \text{sim}^Z \ G \ n'_1 \quad G \ n_2 \ \text{sim}^Z \ G \ n'_2 \quad Q \ \text{sim}^Z \ Q'}{((n_1, n_2), Q) \ \text{sim}^Z \ ((n'_1, n'_2), Q')} \text{ (sim-Q)} \\
 \\
 \boxed{\Sigma_1 \ \text{sim}^Z \ \Sigma_2} \\
 \\
 \frac{R \ \text{sim}^Z \ R' \quad Q \ \text{sim}^Z \ Q'}{(R, C, M, Q, ir) \ \text{sim}^Z \ (R', C, M, Q', ir)} \text{ (sim-}\Sigma\text{)}
 \end{array}$$

Figure 2.15: Similarity of Machine States.




Figure 2.16: Performance Normalized to Unprotected Version.

#### Theorem 4 (Fault Tolerance)

If  $\vdash \Sigma$  and  $\Sigma \xrightarrow{0}^s \Sigma'$  then either  $\Sigma \xrightarrow{1}^{s'} \Sigma'_f$

or  $\exists m \leq (n+1) . \Sigma \xrightarrow{1}^m \text{fault}$ , and

1. For all derivations  $\Sigma \xrightarrow{1}^{s'} \Sigma'_f$  where  $\Sigma'_f \neq \text{fault}$ .

$s' = s$  and  $\exists c . \Sigma' \text{ sim}^c \Sigma'_f$ .

2. For all derivations  $\Sigma \xrightarrow{1}^m \text{fault}$  where  $m \leq (n+1)$ .

$s'$  is a prefix of  $s$ .

## 2.4 Performance<sup>1</sup>

To better understand how TAL<sub>FT</sub> can be applied to real world situations, Reis [51] simulated the TAL<sub>FT</sub> hardware in the framework of a current computer architecture, the Intel Itanium 2 ISA. The instruction set of the Itanium 2 contains many more types of instructions than those specified in TAL<sub>FT</sub>. While not an exact representation of the performance of TAL<sub>FT</sub>, simulating the performance of TAL<sub>FT</sub> applied to this architecture gives guidance as to the feasibility of this system in a real architecture.

<sup>1</sup>This subsection is the work of G. A. Reis [51] and is included in this thesis for completeness.

To evaluate the performance impact of the techniques, a version of the VELOCITY compiler [69] was modified to add the reliability techniques of  $TAL_{FT}$  and was used to compile the SPEC CINT2000 and MediaBench benchmark suites. These executions were compared against binaries generated by the original VELOCITY compiler, which have no fault detection. The reliability transformation was compiled into the low-level code immediately before register allocation and scheduling. To simulate the new hardware structures of  $TAL_{FT}$ , extra instructions were inserted to emulate the timing and dependences of the hardware structure accesses. Only optimizations that were  $TAL_{FT}$  aware were used on the reliable code.

Performance metrics were obtained by running the resulting binaries with reference inputs on an HP workstation zx6000 with 2 900Mhz Intel Itanium 2 processors running Redhat Advanced Workstation 2.1 with 4Gb of memory. The `perfmon` utility was used to measure the CPU time.

Figure 2.16 presents the execution time of the fault-tolerant code relative to baseline binaries with no fault detection. Naïvely, one might expect the fault-tolerant code to run twice as slowly as the fault intolerant code since the number of instructions is essentially doubled. However, smart instruction scheduling and efficient allocation of resources reduces the execution time to only 34% more than the fault-intolerant baseline average. These simulations are in line with previously published software-only reliability performance experiments [60] that show the degradation due to redundant code to be less than double.

As alluded to in Section 2.1.2, Figure 2.16 compares the performance degradation with and without the scheduling constraint that green memory and control flow instructions must be executed before the corresponding blue versions. In order to perform the second set of experiments, the compiler was modified to produce code that had more

flexibility in the scheduling of the green and blue versions. A more aggressive hardware implementation that could correlate the original and redundant memory operations regardless of the executed order was then simulated. As expected, this version has better performance (in most cases) than the unconstrained code. Although the colored ordering restriction of  $TAL_{FT}$  may seem costly, removing this restriction provides only a small improvement. Comparing both to the unprotected code, the version without the ordering constraint increases execution time by 30% while the version with the ordering increases execution time by 34%.

## 2.5 Summary

This chapter introduced  $TAL_{FT}$ , a typed assembly language designed to verify a hybrid fault-tolerance technique.  $TAL_{FT}$  is the first technique for reasoning about fault-tolerance properties of executable code. In addition, we identify four general principles for verifying correctness of fault tolerant systems and capture these in the  $TAL_{FT}$  type system. The two main formal results show that a single fault affecting observable behavior in a well-typed  $TAL_{FT}$  program will always be detected, and that the system will not claim to have detected a fault when none has occurred. Despite the fact that well-typed programs essentially duplicate all computation, we provide simulation results showing a performance overhead of 1.34x. The next chapter investigates using a similar language as the target of a realistic optimizing compiler.

## Chapter 3

# ETAL<sub>FT</sub> : Generating fault-tolerant assembly code

The previous chapter introduced the TAL<sub>FT</sub> language and formally proved that all well-typed programs have certain behaviors, namely that they detect any single fault before it can affect the observable behavior of the program. TAL<sub>FT</sub> forms the core of a fault-tolerant typed assembly language, but is only a starting point for developing a realistic, usable language. This chapter shows how to extend TAL<sub>FT</sub> with additional features to serve as the target of a type-directed compiler.

In addition, these changes allow us to investigate the interactions between fault-tolerant typed assembly languages and other current research areas in typed assembly languages, namely typing the control stack and typing memory allocation and initialization.

The rest of the chapter is organized as follows. Section 3.1 defines a simple source-level imperative language called MiniC. Section 3.2 introduces the extended typed assembly language ETAL<sub>FT</sub> and summarizes its differences from TAL<sub>FT</sub>. Section 3.3

```

int double(x:int ref) {
  int v = 0;
  v = !x;
  v = v + v;
  x := v;
  return 0;
}

int ref y = ref 3;
int result = 0;

result = double(y);
return result;

```

Figure 3.1: MiniC Example Program.

updates the formal results from the previous chapter for ETAL<sub>FT</sub>. Section 3.4 gives a naive translation from MiniC into ETAL<sub>FT</sub> and sketches a proof that the result is well-typed. Section 3.5 gives an example of the translation. Section 3.6 discusses how standard compiler optimizations can be applied to well-typed ETAL<sub>FT</sub> programs. Finally, Section 3.7 summarizes the contributions of ETAL<sub>FT</sub>.

## 3.1 MiniC

MiniC is a simple imperative language with basic arithmetic operations, references, and functions. An example program is shown in Figure 3.1, and the complete abstract syntax is shown in Figure 3.2. The notation *ref* *v* allocates a new heap location and initializes it with value *v*. Heap updates and accesses are notated by *:=* and *!* respectively. We will assume the existence of a garbage collector to run in the background and handle deallocation. Though the syntax is extremely simplified in order to facilitate the proofs later in this chapter, it is still quite expressive.

<i>variables</i>	$x, f$	
<i>type</i>	$\tau$	$::= \text{int} \mid \tau \text{ ref} \mid X \rightarrow \tau$
<i>typing context</i>	$X, F, A, L$	$::= \cdot \mid x : \tau, X$
<i>value</i>	$v$	$::= n \mid x \mid \text{ref } v$
<i>value list</i>	$vs$	$::= \cdot \mid v, vs$
<i>binaryop</i>	$op$	$::= + \mid - \mid *$
<i>statement</i>	$s$	$::= x = v \mid x = v \text{ op } v$ $\mid x = !v \mid v := v$ $\mid x = f(vs)$ $\mid \text{if } v \text{ then } s \text{ else } s \mid \text{while } v \text{ do } s$ $\mid s; s$
<i>function declarations</i>	$fds$	$::= \cdot \mid \tau f(X) \{lds; s; \text{return } v\} fds$
<i>local declarations</i>	$lds$	$::= \cdot \mid \tau x = v; lds$
<i>program</i>	$p$	$::= fds; lds; s; \text{return } v;$

Figure 3.2: MiniC Syntax.

The main typing judgments for MiniC are summarized below, and the complete typing rules are provided in Appendix B. The typing rules are completely standard. Essentially they enforce that variables are defined before their use and that variables of a given type (integers, references, and functions) are used appropriately. We use the metavariable  $X$  to refer to a generic typing context and metavariables  $F$ ,  $A$ , and  $L$  when referring to typing contexts for function variables, argument variables, and local variables respectively.

$X \vdash v : \tau$       *Value  $v$  has type  $\tau$  in typing context  $X$ .*

$X \vdash vs : X'$       *Value list  $vs$  has type  $X'$  in typing context  $X$ .*

$F; A; L \vdash s \text{ wf}$       *Statement  $s$  is well-formed given function context  $F$ , argument context  $A$ , and local variable context  $L$ .*



$F;A;L \vdash lds : L'$      *Given function context  $F$ , argument context  $A$ , and current local declaration  $L$ , after adding local declarations  $lds$ , the new set of local declarations is  $L'$*

$F \vdash fds : F'$      *Given function context  $F$ , after adding the well-formed functions defined in  $fds$ , the new function context is  $F'$ .*

$\vdash p \text{ wf}$      *Program  $p$  is well-typed.*

We have defined MiniC in order to show how to compile MiniC programs into TAL<sub>FT</sub> programs. However, there are two things that make TAL<sub>FT</sub> unsuitable to be the target language for MiniC: (1) TAL<sub>FT</sub> has only a limited number of registers so it can't handle the unlimited number of arguments or local variables allowed by MiniC, and (2) TAL<sub>FT</sub> has no support for dynamic memory allocation.

## 3.2 Extending TAL<sub>FT</sub> to ETAL<sub>FT</sub>


In order to compile MiniC into fault-tolerant assembly code, we need to make a couple of extensions to TAL<sub>FT</sub> to address the issues with limited registers and lack of dynamic memory allocation. This new language is called ETAL<sub>FT</sub>, and this section focuses on only on how it differs from TAL<sub>FT</sub>. The complete specification of ETAL<sub>FT</sub> is provided as a reference in [Appendix C](#).

### 3.2.1 Memory Layout

In TAL<sub>FT</sub>, memory is divided into two pieces: a code memory  $C$  and a value memory  $M$ . Both of these are of a fixed size throughout the execution of a program. In addition, the abstract machine is able to tell the difference between these two memories, and many of the operational rules test to see if a location is only one of these two. (For example, rule *fetch-fail* applies when the address is not in  $C$ , even if it is in  $M$ .) It is reasonable to

assume that the abstracted hardware can make this distinction, as real machines are often able to do the same using an execute bit.

In ETAL<sub>FT</sub> we will conceptually further divide the value memory  $M$  into two pieces: a heap and a stack. The goal is to use the stack as scratch space and consider the heap to be the portion of memory that is "observable". Both the heap and stack may grow during execution, and the stack may shrink as well. To avoid dealing with the possibility of running out of memory, we will assume an infinite sized address space with code memory  $C$  in the middle, the stack in the lower addresses growing downward, and the heap in the higher addresses growing upwards.



We will continue to assume that the machine can differentiate between code memory and the value memory, but we do not assume it can distinguish between the locations in  $M$  that represent the heap and those that represent the stack.

The following judgment is used to break a memory into two pieces, by using a set of addresses  $\mathcal{L}$  to define the division. By using the domain of  $C$  as the set of addresses, we can divide the value memory  $M$  into the stack and the heap.

$$\boxed{M = M_1 \overset{\mathcal{L}}{\#} M_2}$$

$$\frac{\begin{array}{l} \text{Dom}(M) = \text{Dom}(M_1) \cup \text{Dom}(M_2) \\ \text{Dom}(M_1) \cap \text{Dom}(M_2) = \emptyset \\ \forall \ell_1 \in \text{Dom}(M_1). \forall \ell_{\mathcal{L}} \in \mathcal{L}. \forall \ell_2 \in \text{Dom}(M_2). \ell_1 < \ell_{\mathcal{L}} < \ell_2 \end{array}}{M = M_1 \overset{\mathcal{L}}{\#} M_2} \quad (\#-def)$$

### 3.2.2 Stacks

Real machines use the stack to pass arguments and return values, as temporary space when the number of live values exceeds the number of registers, for local variables whose address may be taken, and to allocate local storage without the overhead of the garbage collector.

For simplicity, we will only investigate the first two of these uses. Even so, we need to immediately think about how the stack will interact with our goal of fault tolerance. In TAL<sub>FT</sub>, values are only committed to memory when both the green and blue computations agree. However, the goal of the ETAL<sub>FT</sub> stack is to act as temporary spill space that is not externally visible. We do not want to require that the two computations synchronize when spilling values to the stack. For example, consider a program which needs many registers to perform some arithmetic computation and then stores the final result to memory. In this case, it may well make sense for the blue computation to spill all its values to the stack, reducing the register pressure for the green computation. Once the green computation has completed, it can spill the result and allow the blue computation to use all registers for its version of the task. Therefore the stack may contain a mix of green and blue values. It will be up to our type system to track which is which and to enforce the values of one color continue to only depend on the values of the same color.

The stack grows down towards lower addresses, and we will refer to the lowest location on the stack as the “top” of the stack and the highest location as the “base” of the stack.

#### **New Instruction Syntax and Operational Semantics**

For manipulating the stack, there are two new registers and four new instructions. Each of these new instructions has corresponding operational rules, shown in Figure 3.3.

*general regs*  $a ::= \dots \mid sp_c$   
*instructions*  $i ::= \dots \mid salloc\ n \mid sfree\ n \mid sld_c\ r_d\ n \mid sst\ n\ r_v$

$$\frac{
 \begin{array}{l}
 R' = R++[sp_G \mapsto R(sp_G) - n][sp_B \mapsto R(sp_B) - n] \\
 m = \min(Dom(M)) \\
 M' = (M, m - 1 \mapsto 0, \dots, m - n \mapsto 0)
 \end{array}
 }{
 (R, C, M, Q, salloc\ n) \longrightarrow_0 (R', C, M', Q, \cdot) \quad (salloc)
 }$$

$$\frac{
 \begin{array}{l}
 R' = R++[sp_G \mapsto R(sp_G) + n][sp_B \mapsto R(sp_B) + n] \\
 m = \min(Dom(M)) \\
 M = M', m \mapsto v_m, \dots, (m + n - 1) \mapsto v_1
 \end{array}
 }{
 (R, C, M, Q, sfree\ n) \longrightarrow_0 (R', C, M', Q, \cdot) \quad (sfree)
 }$$

$$\frac{
 \begin{array}{l}
 R(sp_c) + n \in Dom(M) \\
 R' = R++[r_d \mapsto M(R(sp_c) + n)]
 \end{array}
 }{
 (R, C, M, Q, sld_c\ r_d, n) \longrightarrow_0 (R', C, M, Q, \cdot) \quad (sld_c)
 }$$

$$\frac{
 R(sp_c) + n \notin Dom(M)
 }{
 (R, C, M, Q, sld_c\ r_d, n) \longrightarrow_0 \text{fault} \quad (sld_c\text{-fail})
 }$$

$$\frac{
 \begin{array}{l}
 R(sp_B) = R(sp_G) \quad R(sp_G) + n \in Dom(M)
 \end{array}
 }{
 (R, C, M, Q, sst\ n, r_v) \xrightarrow{0}^{(R(sp_G)+n, R(r_v))} (R++, C, M[R(sp_G) + n \mapsto R(r_v)], Q, \cdot) \quad (sst)
 }$$

$$\frac{
 R(sp_G) \neq R(sp_B) \text{ or } R(sp_B) + n \notin Dom(M)
 }{
 (R, C, M, Q, sst\ n, r_v) \longrightarrow_0 \text{fault} \quad (sst\text{-fail})
 }$$

Figure 3.3: New Instruction Syntax and Operation Semantics to Support Stacks.

Just as there is a program counter for each computation, there is a green stack pointer  $sp_G$  and a blue stack pointer  $sp_B$ . Both stack pointers should point to the location on the top of the single stack.

The instructions  $salloc\ n$  and  $sfree\ n$  are used to allocate and deallocate space on the stack. The rules for  $salloc$  and  $sfree$  do this by either subtracting or adding  $n$  to both stack pointers. Our abstract semantics for these instructions also modify the set of locations in memory. Stacks grow downwards, so the lowest address in  $M$  is currently the top of the stack. When allocating,  $n$  new locations are added below the original top of the stack. For convenience, we assume these locations contain the value 0, but do consider them initialized yet. The type system will prevent a programmer from loading from uninitialized locations. When deallocating, the bottom  $n$  addresses are removed.

The instructions  $sld_c\ r_d\ n$  and  $sst\ n\ r_v$  are used to load from and store to the stack at offset  $n$  above the stack pointers. The rules for  $sld_c\ n\ r_v$  are very similar to the rules for the basic load instructions. Instead of taking the address as an argument,  $sld_c$  takes an offset  $n$  and adds that to the stack pointer colored  $c$ . If that address is not in  $M$ , the machine detects a fault, otherwise the value is loaded. If the stack pointer colored  $c$  is corrupted, the loaded value may actually be a heap value or a value of the other color from the stack. But since the computation is already corrupted, no harm is done. As we will see, stores to the stack commit immediately, so there is no reason to check the queue for pending stores when loading from the stack.

In TAL<sub>FT</sub>, the goal was to detect a fault before allowing a store to corrupt  $M$  (through either a corrupt address or a corrupt value). In ETAL<sub>FT</sub>, that restriction is loosened slightly. If there has been a fault to the blue computation, that corruption may affect blue registers which are later spilled to the stack. We do not need to detect this unless it will affect the heap portion of memory. So when the zap tag is  $c$ , we will allow corrupt

values in the locations in  $M$  corresponding to the  $c$ -colored stack locations. However, we still need to be extremely careful about the addresses used in storing to the stack. Like the program counters, the stack pointers can be corrupted. Because there is only one stack containing both green and blue values, we need to make sure that we do not store a blue value at an offset off of a corrupt stack pointer, as that might actually store into a green stack location or into the heap. The following ill-typed sequence has exactly the behavior we desire:

$$\text{add } r_G \text{ sp}_G \ n; \text{ add } r_B \text{ sp}_B \ n; \text{ st}_G \text{ sp}_G \ r_v; \text{ st}_B \text{ sp}_B \ r_v$$

It uses an equivalent address from each computation, so the store will only proceed if both computations agree on the address. However, both store instructions store the same (potentially corrupt) value  $r_v$ . The instruction  $\text{sst } n \ r_v$  can be thought of as a macro for this sequence. Accordingly, it adds address-value pair written to  $M$  to the sequence of observable values, just as  $\text{st}_B$  does.

### Typing the Stack

Typing program stacks is an active area of research in typed assembly languages [39, 30, 50]. Since stack locations are reused during the life of a program to hold values of different types, the type system needs to carefully track these changes. If stack locations may be aliased, tracking these changes soundly becomes difficult. Luckily, we do not require the ability to alias stack locations in ETAL<sub>FT</sub>, and so we can use a simple stack typing method.

In ETAL<sub>FT</sub>, a stack is conceptually an ordered list of location labels and type triples. We extend a number of definitions from TAL<sub>FT</sub> and add two mutually recursive definitions for the stack type itself.

<i>exp kinds</i>	$\kappa ::= \dots \mid \kappa_\sigma$
<i>base types</i>	$b ::= \dots \mid sptr$
<i>reg types</i>	$t ::= \dots \mid ns$
<i>unlabeled stack</i>	$\sigma ::= sbase \mid \rho \mid t :: \zeta$
<i>labeled stack</i>	$\zeta ::= E : \sigma$
<i>static context</i>	$\Theta ::= \Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta$

The kinds for expression variables are extended to include a kind  $\kappa_\sigma$  for the unlabeled stack type. Stack pointers are given base type  $sptr$  and their static expression is used to refer to the exact location on the stack. The nonsense type  $ns$  is given to newly allocated stack locations. A labeled stack type  $\zeta$  consists of a static expression  $E$  that represents the address on the top of unlabeled stack  $\sigma$ . An unlabeled stack type  $\sigma$  is either the empty stack  $sbase$ , a stack variable  $\rho$  or a type  $t$  on top of a labeled stack  $\zeta$ . Stack variables like  $\rho$  have kind  $\kappa_\sigma$  and are used to abstract parts of the caller's stack from the callee. The stack type needs to be tracked at each program point, so it is added to context  $\Theta$  along with the existing register file type, queue description, and so on. The snippet below contains an example of a stack type:

$$\begin{aligned} \zeta &= E_t : \langle G, int, E \rangle :: E_t + 1 : \langle B, int, E \rangle :: E_t + 2 : \rho \\ \Gamma(sp_G) &= \langle G, sptr, E_t \rangle \end{aligned}$$

The stack type  $\zeta$  says that the top location on the stack is  $E_t$  and that location contains a green integer. The next location contains a blue integer. The rest of the stack starting at location  $E_t + 2$  is unknown and abstracted into variable  $\rho$ . In addition, register  $sp_G$  points to the top of the stack.

$$\boxed{\Psi \vdash n : b}$$

$$\dots \quad \overline{\Psi \vdash n : sptr} \quad (\text{addr-stack-t})$$

$$\boxed{\Psi; \Delta \vdash^Z n : t}$$

$$\dots \quad \overline{\Psi; \Delta \vdash^Z n : ns} \quad (ns-t)$$

$$\boxed{\Delta \vdash t \leq t'}$$

$$\dots \quad \overline{\Delta \vdash t \leq ns} \quad (\text{subtp-t-ns})$$

$$\boxed{\Delta \vdash \zeta \leq \zeta'}$$

$$\frac{\Delta \vdash E = E'}{\Delta \vdash E : sbase \leq E' : sbase} \quad (\text{subtp-}\zeta\text{-base})$$

$$\frac{\Delta \vdash E = E'}{\Delta \vdash E : \rho \leq E' : \rho} \quad (\text{subtp-}\zeta\text{-var})$$

$$\frac{\Delta \vdash E = E' \quad \Delta \vdash t \leq t' \quad \Delta \vdash \zeta \leq \zeta'}{\Delta \vdash E : (t :: \zeta) \leq E' : (t' :: \zeta')} \quad (\text{subtp-}\zeta\text{-cons})$$

$$\boxed{\Delta \vdash \zeta \text{ wf}}$$

$$\frac{\Delta \vdash E : \kappa_{int}}{\Delta \vdash E : sbase \text{ wf}} \quad (\zeta\text{-wf-base}) \quad \frac{\Delta(\rho) = \kappa_{\sigma} \quad \Delta \vdash E : \kappa_{int}}{\Delta \vdash E : \rho \text{ wf}} \quad (\zeta\text{-wf-var})$$

$$\frac{\Delta \vdash E + 1 = E' \quad \Delta \vdash t \text{ wf} \quad \Delta \vdash (E' : \sigma') \text{ wf}}{\Delta \vdash E : (t :: (E' : \sigma')) \text{ wf}} \quad (\zeta\text{-wf-cons})$$

Figure 3.4: Modifications to Value Typing, Subtyping, and Well-Formed Judgments.



$$\boxed{\Delta; \zeta \vdash E : t}$$

$$\frac{\Delta \vdash E_s = E}{\Delta; E_s : (t :: \zeta') \vdash E : t} (\zeta\text{-lookup-top}) \quad \frac{\Delta \vdash E_s \neq E \quad \Delta; \zeta' \vdash E : t}{\Delta; E_s : (t :: \zeta') \vdash E : t} (\zeta\text{-lookup-tail})$$

$$\boxed{\Delta \vdash \zeta[E \mapsto t] = \zeta'}$$

$$\frac{\Delta \vdash E_s = E}{\Delta \vdash (E_s : (t_s :: \zeta))[E \mapsto t] = E_s : (t :: \zeta)} (\zeta\text{-update-top})$$

$$\frac{\Delta \vdash E_s \neq E \quad \Delta \vdash \zeta[E \mapsto t] = \zeta'}{\Delta \vdash (E_s : (t_s :: \zeta))[E \mapsto t] = E_s : (t_s :: \zeta')} (\zeta\text{-update-tail})$$

Figure 3.5: Stack Lookup and Update.

Figure 3.4 shows the minor changes required to the value typing, subtyping judgments, and well-formed expression judgments. Any integer can be given the base type *sptr*. In other words, dangling pointers into the stack may exist, as long as we do not dereference them. Any integer can also given the type *ns*. Since values of type *ns* have no color, neither computation can manipulate them. Instead, all that can be done is to overwrite the value. The subtyping judgments are extended so that all types are subtypes of *ns* and one stack is a subtype of another if the locations are equivalent and corresponding locations have types which are subtypes of each other. In order for a stack type to be well formed, the unabstracted portion of the stack must describe sequential locations.

In addition, there are two new judgments for looking up the type of a stack location and updating the type of a stack location shown in Figure 3.5. Judgment  $\Delta; \zeta \vdash E : t$  states that location  $E$  in stack type  $\zeta$  has type  $t$ . Judgment  $\Delta \vdash \zeta[E \mapsto t] = \zeta'$  Says that the result of updating location  $E$  in stack type  $\zeta$  to contain type  $t$  is a new stack type  $\zeta'$ . All stack

updates are strong updates. In other words, it does not matter what type is currently in a location when we go to update it.

The instruction typing rules for the four new stack instructions are shown in Figure 3.6. When allocating or deallocating space on the stack, we require that both stack pointers and the location on the top of the stack be described by equivalent expressions. When allocating space, new locations with the nonsense type  $ns$  are added to the top of the stack. When deallocating, the stack must contain at least as many locations as are being freed. Both instructions update the types of the two stack pointers to refer to the new top of the stack.

To load from the stack at offset  $n$ , we must be able to look up the type of the location at that offset. The destination register is updated to contain that type. When storing to the stack, the two stack pointers must be equivalent, and the stack type is updated with the type of the value being stored. Notice that the type of the value being overwritten is irrelevant.

The rules for all the other instructions are modified to propagate the stack type as part of  $\Theta$ . The rules for  $jmp_B$  and  $brz_B$  are extended to require that the current stack type is a subtype of the target's stack type.

In terms of top-level judgments, the major change is a new judgment  $\Psi \vdash^Z M : \zeta$  shown in Figure 3.7. The stack typing judgment states that a memory  $M$  can be described by a stack type  $\zeta$  under a zap tag  $Z$ . Unlike the judgment about memory in TAL<sub>FT</sub>, this judgment is parameterized by a zap tag which is used to call the standard value typing judgment on each location and its corresponding type.

The machine state type judgment  $\vdash^Z (R, C, M, Q, ir)$  is updated to divide memory  $M$  into a stack portion  $M_s$  and a heap portion  $M_h$ . The stack is typed with  $\Psi \vdash^Z M_s : \zeta$ , and  $M_h$  is typed by a new heap typing judgment described next.

$$\boxed{\Psi; \Theta \vdash ir \Rightarrow RT}$$

...

$$\frac{\begin{array}{l} \Gamma(sp_G) = \langle G, sptr, E_g \rangle \quad \Gamma(sp_B) = \langle B, sptr, E_b \rangle \\ \Delta \vdash E_g = E_b \quad \Delta \vdash E_g = E_t \\ \Gamma' = \Gamma_{++}[sp_G \mapsto \langle G, sptr, (E_g - n) \rangle][sp_B \mapsto \langle B, sptr, (E_b - n) \rangle] \\ \zeta' = (E_t - n) : ns :: (E_t - (n + 1)) : ns :: \dots :: E_t : \sigma \end{array}}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; (E_t : \sigma)) \vdash salloc\ n \Rightarrow (\Delta; \Gamma'; \overline{(E_d, E_s)}; E_m; \zeta')} \text{ (salloc-t)}$$

$$\frac{\begin{array}{l} \Gamma(sp_G) = \langle G, sptr, E_g \rangle \quad \Gamma(sp_B) = \langle B, sptr, E_b \rangle \\ \Delta \vdash E_g = E_b \quad \Delta \vdash E_g = E_t \\ \zeta = E_t : t :: \dots :: E_f : \sigma \quad \Delta \vdash E_f = E_g + n \\ \Gamma' = \Gamma_{++}[sp_G \mapsto \langle G, sptr, (E_g + n) \rangle][sp_B \mapsto \langle B, sptr, (E_b + n) \rangle] \end{array}}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash sfree\ n \Rightarrow (\Delta; \Gamma'; \overline{(E_d, E_s)}; E_m; E_f : \sigma)} \text{ (sfree-t)}$$

$$\frac{\begin{array}{l} \Gamma(sp_c) = \langle c, sptr, E_c \rangle \quad \Delta \vdash E_c + n = E_n \quad \Delta; \zeta \vdash E_n : \langle c, b, E \rangle \end{array}}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash sld_c\ r_d\ n \Rightarrow (\Delta; \Gamma_{++}[r_d \mapsto \langle c, b, E \rangle]; \overline{(E_d, E_s)}; E_m; \zeta)} \text{ (sld_c-t)}$$

$$\frac{\begin{array}{l} \Gamma(sp_G) = \langle G, sptr, E_g \rangle \quad \Gamma(sp_B) = \langle B, sptr, E_b \rangle \quad \Delta \vdash E_g = E_b \\ \Delta \vdash E_g + n = E_n \quad \Gamma(r_v) = \langle c, b, E_v \rangle \quad \Delta \vdash \zeta[E_n \mapsto \langle c, b, E_v \rangle] = \zeta' \end{array}}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash sst\ n\ r_v \Rightarrow (\Delta; \Gamma_{++}; \overline{(E_d, E_s)}; E_m; \zeta')} \text{ (sst-t)}$$

Figure 3.6: Typing Rules for the New Stack Instructions.

$$\boxed{\Psi \vdash^Z M : \zeta}$$

$$\frac{\cdot \vdash E = \ell \quad \text{Dom}(M) = \{\ell\}}{\Psi \vdash^Z M : (E : sbase)} \quad (\zeta\text{-}t\text{-}base)$$

$$\frac{\begin{array}{l} \cdot \vdash (E : t :: \zeta') \text{ wf} \quad \cdot \vdash E = \ell \\ M = \{\ell \rightarrow n\} \# M' \\ \Psi; \cdot \vdash^Z n : t \quad \Psi \vdash^Z M' : \zeta' \end{array}}{\Psi \vdash^Z M : (E : t :: \zeta')} \quad (\zeta\text{-}t\text{-}cons)$$

Figure 3.7: Stack Typing.

### 3.2.3 Memory Allocation and Initialization

When designing type systems for assembly languages with dynamic allocation, we have to be able to separate allocation and initialization. One common approach is to use initialization flags to track whether a location has been initialized [39]. Newly allocated locations are flagged with a 0, which is changed to 1 the first time something is written to the location. Values can only be loaded from initialized locations.

We will use a similar idea, but must make some modifications for this to work in the context of fault tolerance. The issue occurs because the green and blue computations may have different views of the same location. The store queue allows a lot of flexibility in instruction scheduling, including situations where the green computation has initialized a newly allocated location, but the blue computation has not. The stored value is pending in the queue, and so a green load, which always checks the queue first, may load from the location, but a blue load may not. To solve this, ETAL<sub>FT</sub> has three initialization flags instead of the standard two. The added flag  $\frac{1}{2}$  means that only the green computation has done the initialization.

#### New Instruction Syntax and Operational Semantics

To allocate new locations, we add a new instruction  $malloc[b] r_g r_b$  shown in Figure 3.8. It allocates a new location in the heap portion of memory and puts the address of that location in a green register  $r_g$  and a blue register  $r_b$ . In addition, memory  $M$  is extended with the new location. The instruction is also annotated with the basic type  $b$  that the location will contain, but this information has no operational effect and is only used by the type system.

The operational rule currently allocates memory sequentially, but one can imagine modifying this rule to implement other allocation algorithms. In all likelihood,  $malloc$

$$\text{instructions } i ::= \dots \mid \text{malloc}[b] r_g r_b$$

$$\frac{\begin{array}{l} n = \max(\text{Dom}(M)) + 1 \\ R' = R++[r_g \mapsto n][r_b \mapsto n] \end{array}}{(R, C, M, Q, \text{malloc}[b] r_g r_b) \longrightarrow_0 (R', C, (M, n \mapsto 0), Q, \cdot)} \text{ (malloc)}$$

Figure 3.8: New Operation Semantics to Support Dynamic Memory Allocation.

would actually be implemented with a sequence of instructions or a function call with corresponding behavior.

As in MiniC, we will assume the existence of a garbage collector to handle heap deallocation.

### Typing Memory Allocation

In order to track the initialization state of heap locations, reference types are updated to contain an initialization flag. As mentioned earlier, 0 means completely uninitialized, and 1 means initialized, and  $\frac{1}{2}$  means that only the green computation has performed its initialization.

$$\begin{array}{l} \text{initialization flags } \varphi ::= 1 \mid \frac{1}{2} \mid 0 \\ \text{base types } b ::= \dots \mid b \text{ ref}^\varphi \end{array}$$

There is a subtyping relationship between the different initialization flags. A half initialized location containing a value of type  $b$  is a subtype of an uninitialized location with a value of type  $b$ . Similarly, a fully initialized location is a subtype of the corresponding half initialized location. Figure 3.9 shows the additional rules added to support these relationships.

$$\boxed{\Psi \vdash n : b}$$

$$\dots \frac{\Psi \vdash n : b \text{ ref}^\varphi \quad \varphi \leq \varphi'}{\Psi \vdash n : b \text{ ref}^{\varphi'}} \text{ (addr-subtp-t)}$$

$$\boxed{\Psi; \Delta \vdash^Z n : t}$$

$$\dots \frac{\Psi; \Delta \vdash^Z n : t' \quad \Delta \vdash t' \leq t}{\Psi; \Delta \vdash^Z n : t} \text{ (val-subtp-t)}$$

$$\boxed{\varphi \leq \varphi'}$$

$$1 \leq \frac{1}{2} \leq 0 \quad \varphi \leq \varphi$$

$$\boxed{b \leq b'}$$

$$\dots \frac{\varphi \leq \varphi'}{b \text{ ref}^\varphi \leq b \text{ ref}^{\varphi'}} \text{ (subtp-b-ref)}$$

Figure 3.9: Additions to Subtyping to Support Reference Initialization.

Figure 3.10 shows the modifications to the instruction typing judgment. When a new location is allocated to contain a value of type  $b$ , it is given type  $b \text{ ref}^0$ , and a fresh expression variable  $x$  is selected to represent the newly allocated address. The two destination registers are updated to contain green and blue type triples with the base type  $b$  and expression  $x$ . The expression  $E_m$  that describes memory is updated to map  $x$  to 0. In addition to the new rule for typing *malloc*, the typing rules for the existing load and store instructions are also modified. When the green computation wishes to load a value, the location must be at least half initialized. For the blue computation, the location must be fully initialized. The typing rules for the two store instructions are modified to update the initialization flags. When the green computation stores to a location, it uses the function  $\varphi \uparrow$  to update the flag. This function changes the uninitialized flag to be half initialized, but does not affect the other flags. The blue computation modifies the initialization flag to be fully initialized. In addition, when one of the store instructions changes the initialization flag on a value, it needs to update the type for that location used by the corresponding computation.

Now that we have initialization flags, we can no longer type of heap portion of memory by itself. We also need information about the queue to check for consistency in the initialization states of each location. The new and modified rules are shown in Figure 3.11.

Judgment  $\Psi; M; Q \vdash^Z \ell : b \text{ ref}^\varphi$  holds for location  $\ell$  when  $Q$  and  $M$  are consistent with the initialization flag  $\varphi$ . If the initialization flag is 1, then  $M(\ell)$  needs to have type  $b$ . If the initialization flag is  $\frac{1}{2}$ , and the zap tag is not green, there must be a pending store to  $\ell$  in the queue. If the initialization flag is 0, nothing is required of  $M$  or  $Q$ .



$\boxed{\Phi \uparrow}$ 

$$0 \uparrow = \frac{1}{2} \quad \frac{1}{2} \uparrow = \frac{1}{2} \quad 1 \uparrow = 1$$

 $\boxed{\Psi; \Theta \vdash ir \Rightarrow RT}$ 

...

$$\frac{\begin{array}{c} x \notin \Delta \\ \Gamma' = \Gamma_{++}[r_g \mapsto \langle G, b \text{ ref}^0, x \rangle][r_b \mapsto \langle B, b \text{ ref}^0, x \rangle] \\ E'_m = \text{upd } E_m \ x \ 0 \end{array}}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{malloc}[b] \ r_g \ r_b \Rightarrow (\Delta, x : \kappa_{\text{int}}; \Gamma'; \overline{(E_d, E_s)}; E'_m; \zeta)} \quad (\text{malloc-t})$$

$$\frac{\Delta \vdash \Gamma(r_s) \leq \langle G, b \text{ ref}^{\frac{1}{2}}, E'_s \rangle \quad E = \text{sel } (\overline{\text{upd } E_m \ \overline{(E_d, E_s)}}) \ E'_s}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{ld}_G \ r_d \ r_s \Rightarrow (\Delta; \Gamma_{++}[r_d \mapsto \langle G, b, E \rangle]; \overline{(E_d, E_s)}; E_m; \zeta)} \quad (\text{ld}_G\text{-t})$$

$$\frac{\Gamma(r_s) = \langle B, b \text{ ref}^1, E'_s \rangle \quad E = \text{sel } E_m \ E'_s}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{ld}_B \ r_d \ r_s \Rightarrow (\Delta; \Gamma_{++}[r_d \mapsto \langle B, b, E \rangle]; \overline{(E_d, E_s)}; E_m; \zeta)} \quad (\text{ld}_B\text{-t})$$

$$\frac{\begin{array}{c} \Gamma(r_d) = \langle G, b \text{ ref}^\Phi, E'_d \rangle \quad \Gamma(r_s) = \langle G, b, E'_s \rangle \\ \Gamma' = \Gamma_{++} \text{ except } \forall r \text{ where } \Gamma(r) = \langle c_r, b \text{ ref}^\Phi, E_r \rangle \text{ and } \Delta \vdash E_r = E'_d . \\ \Gamma'(r) = \langle c_r, b \text{ ref}^{\Phi \uparrow}, E_r \rangle \end{array}}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{st}_G \ r_d \ r_s \Rightarrow (\Delta; \Gamma'; (E'_d, E'_s), \overline{(E_d, E_s)}; E_m; \zeta)} \quad (\text{st}_G\text{-t})$$

$$\frac{\begin{array}{c} \Delta \vdash \Gamma(r_d) \leq \langle B, b \text{ ref}^{\frac{1}{2}}, E''_d \rangle \quad \Gamma(r_s) = \langle B, b, E''_s \rangle \\ \Delta \vdash E'_s = E''_s \quad \Delta \vdash E'_d = E''_d \\ \Gamma' = \Gamma_{++} \text{ except } \forall r \text{ where } \Gamma(r) = \langle c_r, b \text{ ref}^{\frac{1}{2}}, E_r \rangle \text{ and } \Delta \vdash E_r = E'_d . \\ \Gamma'(r) = \langle c_r, b \text{ ref}^1, E_r \rangle \end{array}}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}, (E'_d, E'_s); E_m; \zeta) \vdash \text{st}_B \ r_d \ r_s \Rightarrow (\Delta; \Gamma'; \overline{(E_d, E_s)}; \text{upd } E_m \ E'_d \ E'_s; \zeta)} \quad (\text{st}_B\text{-t})$$

Figure 3.10: Instruction Typing Rules to Support Dynamic Memory Allocation.

The queue typing judgment  $\Psi \vdash^Z Q : \overline{(E_d, E_s)}$  in TAL<sub>FT</sub> requires that all addresses in the queue have type *b ref* for some *b* unless the zap tag is green. In addition in ETAL<sub>FT</sub>, the base type of these locations must be at least half initialized.

The heap typing judgment  $\Psi \vdash^Z (M, Q) : (E_m, \overline{(E_d, E_s)})$  requires all locations to be consistent with *M* and *Q*. In addition,  $\overline{(E_d, E_s)}$  describes *Q*, and *E<sub>m</sub>* describes *M*.

Finally, the machine state typing judgment is modified so that it divides memory *M* into the stack portion *M<sub>s</sub>* and the heap portion *M<sub>m</sub>* and calls the new stack typing judgment on *M<sub>s</sub>* and the new heap typing judgment on *M<sub>m</sub>* and *Q*. The final judgment is shown in Figure 3.12.

### 3.2.4 Removal of Color Tags

In TAL<sub>FT</sub>, register values are tagged with the color of the computation to which they belong. The main use of this color tag is in the relation  $\Sigma \text{sim}^Z \Sigma'$ . Because each value is tagged with a color, we can compare two machine states directly without relying on any extra information. This simplifies the judgments quite a bit. However in ETAL<sub>FT</sub>, we don't have that luxury. Values on the stack are conceptually colored, but those in the heap are not. Adding color tags to memory doesn't make sense for all locations, and in any case, the colors already appear in the typing information. Instead we change the simulation relationship to depend on color information extracted from the typing information. We go ahead and make this change throughout the language, and remove the color tags from all values. This also enforces our earlier claim that the color tags in TAL<sub>FT</sub> were only for convenience.

This change requires removing the color tags from all values in the existing rules. The dynamic semantics only propagated the tags, so their removal has no effect. In the static semantics, the color tags are duplicated in the typing information, and so again their

$$\boxed{\Psi; M; Q \vdash^Z \ell : b \text{ ref}^\varphi}$$

$$\frac{\Psi(\ell) = b \text{ ref}^1 \quad \Psi \vdash M(\ell) : b}{\Psi; M; Q \vdash^Z \ell : b \text{ ref}^1} \text{ (init-t)} \quad \frac{\Psi(\ell) = b \text{ ref}^0}{\Psi; M; Q \vdash^Z \ell : b \text{ ref}^0} \text{ (uninit-t)}$$

$$\frac{\Psi(\ell) = b \text{ ref}^{\frac{1}{2}} \quad Z \neq G \implies \exists n. (\ell, n) \in Q}{\Psi; M; Q \vdash^Z \ell : b \text{ ref}^{\frac{1}{2}}} \text{ (halfinit-t)}$$

$$\boxed{\Psi \vdash^Z Q : \overline{(E_d, E_s)}}$$

...

$$\frac{\Psi \vdash^Z \overline{(n'_1, n'_2) : (E'_d, E'_s)} \quad \cdot \vdash E_d = n_1 \quad \cdot \vdash E_s = n_2 \quad \Psi \vdash n1 : b \text{ ref}^\varphi \quad \varphi \leq \frac{1}{2} \quad \Psi \vdash n2 : b}{\Psi \vdash^Z \overline{(n_1, n_2), (n'_1, n'_2) : (E_d, E_s), (E'_d, E'_s)}}} \text{ (Q-t)}$$

$$\boxed{\Psi \vdash^Z (M, Q) : (E_m, \overline{(E_d, E_s)})}$$

$$\frac{\forall \ell \in \text{Dom}(M). \exists \varphi. \Psi; M; Q \vdash^Z \ell : b \text{ ref}^\varphi \quad \llbracket E_m \rrbracket = M \quad \Psi \vdash^Z Q : \overline{(E_d, E_s)}}{\Psi \vdash^Z (M, Q) : (E_m, \overline{(E'_d, E'_s)})} \text{ (heap-t)}$$

Figure 3.11: Changes to Memory Typing Judgments.

$$\boxed{\vdash^Z (R, C, M, Q, ir)}$$

$$\begin{array}{l}
 \text{Dom}(\Psi) = \text{Dom}(C) \cup \text{Dom}(M_m) \\
 M = M_s \overset{\text{Dom}(C)}{\#} M_m \\
 \Psi \vdash C \\
 \forall c \neq Z. ir \neq \cdot \implies C(R(\underline{pc_c})) = ir \\
 \forall c \neq Z. \Psi(R(\underline{pc_c})) = (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \rightarrow \text{void} \\
 \exists S. \cdot \vdash S : \Delta \\
 \Psi \vdash^Z M_s : S(\zeta) \\
 \Psi \vdash^Z (M_m, Q) : (S(E_m), S(\overline{(E_d, E_s)})) \\
 \Psi \vdash^Z R : S(\Gamma) \\
 \hline
 \vdash^Z (R, C, M, Q, ir) \quad (\Sigma\text{-}t)
 \end{array}$$

Figure 3.12: Changes to the Machine State Typing Judgment.

removal has no effect. Section 3.3.2 will discuss how the simulation relation and formal results are modified.

### 3.2.5 Other Changes

In addition to the changes above, we added a new move instruction to move the contents of one register into another register. This instruction is not technically necessary, but is helpful to have when generating assembly code.

$$\text{instructions } i ::= \dots \mid \text{mov } r_d, r_s$$

The TAL<sub>FT</sub> rule for Register File Typing required specifying a type for every register. Again, for simplicity of the translation, we drop this requirement. A register file  $R$  has type  $\Gamma$  when all types in  $\Gamma$  hold true for  $R$ .

$$\boxed{\Psi \vdash^Z R : \Gamma}$$

$$\frac{\begin{array}{l} \forall a \in \text{Dom}(\Gamma). \Psi; \cdot \vdash^Z R(a) : \Gamma(a) \\ \cdot \vdash \Gamma(pc_G) \leq \langle G, \text{int}, E_G \rangle \\ \cdot \vdash \Gamma(pc_B) \leq \langle B, \text{int}, E_B \rangle \\ \cdot \vdash E_G = E_B \end{array}}{\Psi \vdash^Z R : \Gamma} \text{ (reg-file-t)}$$

The complete specification of ETAL<sub>FT</sub> is provided as a reference in Appendix C. Next, we discuss the formal results for these modified rules.

### 3.3 ETAL<sub>FT</sub> Formal Results

Because ETAL<sub>FT</sub> contains modifications to TAL<sub>FT</sub>, all the formal results need to be repeated. The results are summarized below, and Appendix D contains proof sketches of the main lemmas and theorems.

#### 3.3.1 Type Safety

The statements of Progress and Preservation do not change, although many of the cases must be modified. A number of the lemmas used in these proofs are modified, and there are also a number of new lemmas for reasoning about the stack and dynamic memory allocation.

##### Theorem 5 (Progress)

1. If  $\vdash \Sigma$  then  $\Sigma \longrightarrow_0^s \Sigma'$  and  $\Sigma' \neq \text{fault}$ .
2. If  $\vdash^c \Sigma$  then  $\Sigma \longrightarrow_0^s \Sigma$ .

**Theorem 6 (Preservation)**

1. If  $\vdash^Z \Sigma$  and  $\Sigma \longrightarrow_0^s \Sigma'$  and  $\Sigma' \neq \text{fault}$  then  $\vdash^Z \Sigma'$ .
2. If  $\vdash \Sigma$  and  $\Sigma \longrightarrow_1^s \Sigma'$  then  $\exists c. \vdash^c \Sigma'$ .

**3.3.2 Fault Tolerance**

Due to the removal of color tags, there are more significant changes needed in the formal Fault Tolerance results.

**Similarity of Machine States**

Now that values do not contain color tags, we need to extract the color of values from the typing information. In order to make as few changes as possible, we divide the judgments into two phases.

First we extract the color of the value in each register and memory location from the typing information and to store this information in a mapping  $K$ . The range of  $K$  is an *extended color*  $k$  which can either be a color  $c$ ,  $GB$ , or *none*. Registers that do not have a corresponding type in  $\Gamma$  are mapped to *none*. Locations in the heap are considered to be  $GB$ , as both the green and blue computations have agreed on their values.

$$\text{extended color } k ::= c \mid GB \mid \text{none}$$

$$\text{coloring } K ::= \cdot \mid a \mapsto k \mid \ell \mapsto k$$

Extracting the color of a type  $t$  is simple. Type triples and conditional types clearly specify the color, and the  $ns$  type has no color.

$$\begin{aligned} \text{extract}K_t(\langle c, b, E \rangle) &= c \\ \text{extract}K_t(E_z = 0 \Rightarrow \langle c, b, E \rangle) &= c \\ \text{extract}K_t(ns) &= \text{none} \end{aligned}$$

For the register file, those registers given a type by  $\Gamma$  extract the color from that type, and those registers with no typing information have no color.

$$\text{extract}K(R, \Gamma) = \forall a \in R. a \in \text{Dom}(\Gamma) \ ? \ a \mapsto \text{extract}K_t(\Gamma(a)) \ : \ a \mapsto \text{none}$$

Similarly for the stack, the color for each location is extracted from the type of that location.

$$\text{extract}K_s(M_s, \zeta) = \forall \ell \in \text{Dom}(M_s). \ .; \zeta \vdash \ell : t \Rightarrow \ell \mapsto \text{extract}K_t(t)$$

All locations in the heap have extended color  $GB$ .

$$\text{extract}K_h(M_h) = \forall \ell \in \text{Dom}(M_h). \ell \mapsto GB$$

The top-level typing judgment  $\vdash^Z \Sigma$  is modified slightly to  $\vdash^Z \Sigma : K$ , which generates the coloring calculated using the typing information in the premises of the judgment.

Given the coloring information, we modify the TAL<sub>FT</sub> simulation relations. Figure 3.13 gives the new judgments. The base judgment  $k \ n_1 \ \text{sim}^Z \ k \ n_2$  is very similar to the simulation of colored values in TAL<sub>FT</sub>. The main difference is the addition of a rule *sim-no-color* that allows any two values with no color to simulate each other regardless

$$\boxed{k n_1 \text{ sim}^Z k n_2}$$

$$\frac{}{k n \text{ sim}^Z k n} \text{ (sim-val)} \quad \frac{}{c n \text{ sim}^c c n'} \text{ (sim-val-zap)}$$

$$\frac{}{\text{none } n \text{ sim}^Z \text{ none } n'} \text{ (sim-val-no-color)}$$

$$\boxed{K \vdash R \text{ sim}^Z R'}$$

$$\frac{\forall a. K(a) R(a) \text{ sim}^Z K(a) R'(a)}{K \vdash R \text{ sim}^Z R'} \text{ (sim-R)}$$

$$\boxed{K \vdash M \text{ sim}^Z M'}$$

$$\frac{\text{Dom}(M) = \text{Dom}(M') \quad \forall \ell \in \text{Dom}(M). K(\ell) M(\ell) \text{ sim}^Z K(\ell) M'(\ell)}{K \vdash M \text{ sim}^Z M'} \text{ (sim-M)}$$

$$\boxed{Q \text{ sim}^Z Q'}$$

$$\frac{}{\cdot \text{ sim}^Z \cdot} \text{ (sim-Q-empty)}$$

$$\frac{G n_1 \text{ sim}^Z G n'_1 \quad G n_2 \text{ sim}^Z G n'_2 \quad Q \text{ sim}^Z Q'}{((n_1, n_2), Q) \text{ sim}^Z ((n'_1, n'_2), Q')} \text{ (sim-Q)}$$

$$\boxed{\Sigma_1 \text{ sim}^Z \Sigma_2}$$

$$\frac{\begin{array}{l} \vdash (R, C, M, Q, ir) : K \\ \vdash^Z (R', C, M', Q', ir) : K \\ K \vdash R \text{ sim}^Z R' \\ K \vdash M \text{ sim}^Z M' \\ Q \text{ sim}^Z Q' \end{array}}{(R, C, M, Q, ir) \text{ sim}^Z (R', C, M', Q', ir)} \text{ (sim-}\Sigma\text{)}$$

Figure 3.13: Similarity of Machine States.



of the zap tag. This rule is used when a fault affects a register that has no associated type in  $\Gamma$ .

The judgment for register files is now parameterized by a coloring which it uses to determine the color for each register. In addition there is a new rule for memories. Since the coloring is defined differently for stack and heap locations, two memories simulate each other when their heap locations are identical and their stack locations simulate each other under the zap tag. The simulation relationship for queues is unmodified.

Finally, two machine states simulate each other under  $Z$  if the first is well-typed under the empty zap tag and the second is well-typed under  $Z$ . Both typing judgments must generate the same coloring  $K$ , and all elements must simulate each other given  $K$ .

### **Fault Tolerance Theorem**

Using these modified similarity relations, we can state and prove the fault tolerance theorem for well-typed programs.

The main difference comes from the modified definition of “observable” behavior that requires the heap to be identical, but allows different values in the stack. We formalize this with the new relationships  $s' \stackrel{\ell}{\simeq} s$  and  $s' \stackrel{\ell}{\preceq} s$  that are used to relate two output sequences of address-value pairs. Because the faulty computation may store faulty values into the stack portion of memory, we can no longer use simple equality to compare address-value pairs. The judgment  $\stackrel{\ell}{\simeq}$  says that the addresses in the two sequences are equal, and for all addresses greater than  $\ell$ , the values are also equal. The judgment  $s' \stackrel{\ell}{\preceq} s$  is similar, but only requires that the locations in the first sequence are a subsequence of those in the second sequence. In other words, when given  $\max(\text{Dom}(C))$  as  $\ell$ , these judgments check that the stores committed to the heap are identical and the stores committed to the stack are to the same locations, though the values may differ.

**Theorem 7 (Fault Tolerance)**

If  $\vdash \Sigma$  and  $\Sigma \xrightarrow{n \rightarrow_0^s} \Sigma'$  then either  $\Sigma \xrightarrow{(n+1) \rightarrow_1^{s'}} \Sigma'_f$

or  $\exists m \leq (n+1) . \Sigma \xrightarrow{m \rightarrow_1^{s'}} \text{fault}$ , and

1. For all derivations  $\Sigma \xrightarrow{(n+1) \rightarrow_1^{s'}} \Sigma'_f$  where  $\Sigma'_f \neq \text{fault}$ .

$$s' \stackrel{\max(\text{Dom}(C))}{\simeq} s \text{ and } \exists c. \Sigma' \text{ sim}^c \Sigma'_f.$$

2. For all derivations  $\Sigma \xrightarrow{m \rightarrow_1^{s'}} \text{fault}$  where  $m \leq (n+1)$ .

$$s' \stackrel{\max(\text{Dom}(C))}{\preceq} s.$$

### 3.4 Translation from MiniC to ETAL<sub>FT</sub>

When designing a sound type system, we need to be careful that the restrictions imposed by the type system are not so stringent as to rule out all interesting programs. This section gives an algorithm for translating all well-typed MiniC programs into ETAL<sub>FT</sub>. By doing so, we show that the ETAL<sub>FT</sub> type system is expressive enough to be of interest.

Optimizing compilers generally generate code using a simple translation that may produce inefficient code, and then they apply optimizations to this code. Similarly, this section gives only a naive translation from MiniC to ETAL<sub>FT</sub>, and Section 3.6 will sketch the interactions between ETAL<sub>FT</sub> and some common optimizations. The rest of this chapter gives an overview of the translation. Appendix E provides the complete rules for the translation and a proof sketch of the Translation Theorem.

#### 3.4.1 Translation Introduction

For simplicity, the translation uses the designated registers  $pc_G$ ,  $pc_B$ ,  $sp_G$ ,  $sp_B$ , and  $gd$  and then as many fresh temporary registers  $t_1, t_2, \dots$  as needed. Many of these temporary

registers can be easily removed by coalescing move instructions. Section 3.6 discusses how to support register allocation if the number of temporary registers is greater than the number of actual registers.

The translation uses a simplistic calling convention. In order to support fault tolerance, all function arguments and return values need to be duplicated. Arguments are passed on the stack, with the last argument pushed on first. There will be two assembly-level arguments for each MiniC argument, and the green argument of each pair is always below the corresponding blue argument on the stack. Below the arguments are the blue and green copies of the return address. When a function returns, it pops the return address and all the arguments and pushes two copies of the return value.

On function entry, each function loads all the arguments into temporary registers. When making a function call, all temporary registers that correspond to local variables are spilled to the stack before the arguments and return address are pushed. After the call returns, the registers for the local variables are restored and then the return values are moved to their destinations.

### 3.4.2 Translation Details

At a high level, the translation works by passing around a code memory  $C$  and continually accumulating new instructions onto the end. In addition, many judgments track the following additional information:

$n$  the number of temporary registers required so far.

$V$  a mapping from MiniC variables  $x$  to pairs of registers  $(r_g, r_b)$ .

$B$  a mapping from function variables to addresses in code memory.

$\llbracket X \vdash v : \tau \rrbracket C n V = C' n' r r'$	<i>Given a value <math>v</math>, generate code to move the translation of <math>v</math> into registers <math>r</math> and <math>r'</math>.</i>
$\llbracket X \vdash vs : X \rrbracket C n V n_p = C' n'$	<i>Given a list of values <math>vs</math> to be passed to a function, generate code to push them onto the stack in reverse order.</i>
$\llbracket F; A; L \vdash s \text{ wf} \rrbracket_f C n V B = C' n'$	<i>Given a statement <math>s</math> in function <math>f</math>, generate code for the translation of <math>s</math>.</i>
$\llbracket F; A; L \vdash lds : L' \rrbracket C n V = C' n' V'$	<i>Given a list of local declarations <math>lds</math>, generate code to initialize them and modify the variable map <math>V</math> to contain the new mappings.</i>
$\llbracket F \vdash fds : F' \rrbracket C B = C' B' n$	<i>Given a list of function declarations <math>fds</math>, generate code for each function and modify the function mapping <math>B</math> to contain the starting address for each function. <math>n</math> is the maximum number of temporary registers used by any one function.</i>
$\llbracket \vdash p \text{ wf} \rrbracket = C n l$	<i>Given a program <math>p</math>, generate code for <math>p</math>, the maximum number of registers <math>n</math> required by <math>p</math>, and the starting address <math>l</math>.</i>

Figure 3.14: Summary of Translation Judgments.

The translation is defined over the typing rules for MiniC. Figure 3.14 summarizes the program translation judgments. The eventual goal is to show that all well-typed MiniC programs can be translated into well-typed ETAL<sub>FT</sub> programs, so the translation uses MiniC typing information to generate the corresponding ETAL<sub>FT</sub> typing information. The assorted type translation judgments are shown in Figure 3.15.

$\llbracket \tau \rrbracket = b$	<i>The translation of MiniC type <math>\tau</math> is ETAL<sub>FT</sub> basic type <math>b</math>.</i>
$\llbracket A \rrbracket = \zeta$	<i>The MiniC arguments <math>A</math> correspond to a stack type <math>\zeta</math>.</i>
$\llbracket A; \zeta \rrbracket_V = \Gamma$	<i>Given MiniC arguments <math>A</math> and their translation into stack type <math>\zeta</math>, calculate the register file type <math>\Gamma</math> that results after loading each argument into its corresponding register.</i>
$\llbracket X \rrbracket_V = \Gamma$	<i>Given a variable context <math>X</math> and variable mapping <math>V</math>, generate the corresponding register file type <math>\Gamma</math>.</i>
$\llbracket A \rightarrow \tau; L \rrbracket_V = \Theta$	<i>Given a function type <math>A \rightarrow \tau</math>, local variables <math>L</math>, and variable mapping <math>V</math>, generate the current precondition <math>\Theta</math>.</i>

Figure 3.15: Summary of Type Translation Judgments.

### 3.4.3 Translation Formal Results

The Translation Theorem states that the result of translating a well-typed program  $p$  can be used to create a well-typed ETAL<sub>FT</sub> machine state. The code memory in this machine state is just the code memory returned by the translation. Because ETAL<sub>FT</sub> has no true notion of termination, programs “complete” by jumping to a designated address  $l_{halt}$  which contains a short code snippet for an infinite loop. The heap is empty, and the stack contains two pointers to  $l_{halt}$ . The register file is built by calling the function  $buildR(n)$  to generate a blank register with  $n$  temporary registers. The two program counters are set to the start address generated by the translation, and the stack pointers are set to the top of the stack.

#### Theorem 8 (Translation)

If  $\llbracket \vdash p \text{ wf} \rrbracket = C \ n \ l_{start}$  then  $\vdash (R, C, M, (), \cdot)$

where  $l_\zeta = \min(\text{Dom}(C) - 3)$

$$R = buildR(n), \ pc_G \mapsto l_{start}, \ pc_B \mapsto l_{start}, \ sp_G \mapsto l_\zeta, \ sp_B \mapsto l_\zeta, \ gd \mapsto 0$$

$$M = l_\zeta \mapsto l_{halt}, \ l_\zeta + 1 \mapsto l_{halt}, \ l_\zeta + 1 \mapsto 0$$

Again, the complete translation and corresponding proof sketches are provided in Appendix E.

### 3.5 Translation Example

Figure 3.16 shows the result of applying the naive translation to the MiniC example program from the beginning of this chapter.

Clearly, this translation is not efficient. But again, that is not yet the point. The generated code is well-typed. Let us walk through the translation of the `double()` function to get an idea how this works.

We use the type translation  $\llbracket \tau \rrbracket$  (defined in Appendix E.2.1) to generate the type for the first instruction in the translation of a function. Function `double()` has type  $x : int\ ref \rightarrow int$ , and  $\llbracket x : int\ ref \rightarrow int \rrbracket$  gives us the type below. Though it looks complicated, it is just laying out the calling convention discussed earlier. When entering this function, the stack should have two copies of the argument and then two copies of the return address. The rest of the stack is unknown. The stack pointers should point to the top of the stack, the program counters are equal, and the special destination register  $gd$  must be zero (meaning that no control flow transfers are in progress). When the function returns, it should leave things in a state satisfying  $\Theta_r$ . In other words, it needs to remove the return address and arguments from the stack and push two copies of the return value.

CHAPTER 3. ETAL<sub>FT</sub> : GENERATING FAULT-TOLERANT ASSEMBLY CODE 87

```

// original MiniC program
int double(x:int ref) {
    int v = 0;
    v = !x;
    v = v + v;
    x := v;
    return 0;
}

int ref y = ref 3;
int result = 0;

result = double(y);
return result;

// infinite loop
50 mov t0 1
51 mov t1 1
52 jmpG t1
53 jmpB t1

// double()
54 sldG 2 t0 // prologue
55 sldB 3 t1
56 mov t2 0 // int v = 0
57 mov t3 0
58 mov t4 t2
59 mov t5 t3
60 ldG t4 t0 // v = !x
61 ldB t5 t1
62 add t4 t4 t4 // v = v + v
63 add t5 t5 t5
64 stG t0 t4 // x := v
65 stB t1 t5
66 mov t6 0 // get ret val
67 mov t7 0
68 sldG 0 t8 // epilogue
69 sldB 1 t9
70 sfree 4
71 salloc 2
72 sst 0 t6
73 sst 1 t7
74 jmpG t8
75 jmpB t9

// main()
76 mov t0 3 // int ref y = ref 3
77 mov t1 3
78 malloc[int] t2 t3
79 stG t2 t0
80 stB t3 t1
81 mov t4 t2
82 mov t5 t3
83 mov t6 0 // int result = 0
84 mov t7 0
85 mov t8 t6
86 mov t9 t7
87 salloc 8 // result = double(y)
88 sst 4 t4 - st local vars/args
89 sst 5 t5
90 sst 6 t8
91 sst 7 t9
92 sst 2 t4 - st args
93 sst 3 t5
94 mov t10 50 - st ret addr
95 sst 0 t10
96 mov t11 50
97 sst 1 t11
98 mov t12 54 - jmp to double()
99 mov t13 54
100 jmpG t12
101 jmpB t13
102 sldG 2 t4 - reload local vars
103 sldB 3 t5
104 sldG 4 t6
105 sldB 5 t7
106 sldG 0 t8 - set result
107 sldB 1 t9
108 sfree 6 - free call space
109 mov t14 0 // get ret val
110 mov t15 0
111 sldG 0 t16 // epilogue
112 sldB 1 t17
113 sfree 2
114 salloc 2
115 sst 0 t14
116 sst 0 t15
117 jmpG t16
118 jmpB t17

```

Figure 3.16: Example Translation from MiniC to ETAL<sub>FT</sub>.

$$\Psi(54) = (\Delta_d, \Gamma_d, (), \alpha_m, \zeta_d) \rightarrow \text{void}$$

where:

$$\begin{aligned} \Delta_d &= \alpha_p : \kappa_{int}, \alpha_m : \kappa_{mem}, \alpha_\ell : \kappa_{int}, \alpha_\sigma : \kappa_\sigma, \alpha_r : \kappa_{int}, \alpha_x : \kappa_{int} \\ \Gamma_d &= pc_G \mapsto \langle G, int, \alpha_p \rangle, pc_B \mapsto \langle B, int, \alpha_p \rangle, \\ &\quad sp_G \mapsto \langle G, sptr, \alpha_\ell - 4 \rangle, sp_B \mapsto \langle B, sptr, \alpha_\ell - 4 \rangle, \\ &\quad gd \mapsto \langle G, int, 0 \rangle \\ \zeta_d &= \alpha_\ell - 4 : \langle G, \Theta_r \rightarrow void, \alpha_r \rangle :: \alpha_\ell - 3 : \langle B, \Theta_r \rightarrow void, \alpha_r \rangle \\ &\quad :: \alpha_\ell - 2 : \langle G, int\ ref^1, \alpha_x \rangle :: \alpha_\ell - 1 : \langle B, int\ ref^1, \alpha_x \rangle \\ &\quad :: \alpha_\ell : \alpha_\sigma \end{aligned}$$

and:

$$\begin{aligned} \Theta_r &= (\Delta_r, \Gamma_r, (), \alpha'_m, \zeta_r) \\ \Delta_r &= \alpha'_m : \kappa_{mem}, \alpha_\tau : \kappa_{int} \\ \Gamma_r &= pc_G \mapsto \langle G, int, \alpha_r \rangle, pc_B \mapsto \langle B, int, \alpha_r \rangle, \\ &\quad sp_G \mapsto \langle G, sptr, \alpha_\ell - 2 \rangle, sp_B \mapsto \langle B, sptr, \alpha_\ell - 2 \rangle, \\ &\quad gd \mapsto \langle G, int, 0 \rangle \\ \zeta_d &= \alpha_\ell - 2 : \langle G, int, \alpha_\tau \rangle :: \alpha_\ell - 1 : \langle B, int, \alpha_\tau \rangle \\ &\quad :: \alpha_\ell : \alpha_\sigma \end{aligned}$$

54 sld<sub>G</sub> 2 t<sub>0</sub>

After executing the first stack load, the next instruction has a similar type except that the type of the destination register used in the load is updated to contain the same type as the stack slots that the value was loaded from.

$$\Psi(55) = (\Delta_d, \Gamma_d[t_0 \mapsto \langle G, int\ ref^1, \alpha_x \rangle], (), \alpha_m, \zeta_d) \rightarrow \text{void}$$

55 sld<sub>B</sub> 3 t<sub>1</sub>

Executing the second load has a similar effect.

$$\Psi(56) = (\Delta_d, \Gamma_d[t_0 \mapsto \langle G, int\ ref^1, \alpha_x \rangle][t_1 \mapsto \langle B, int\ ref^1, \alpha_x \rangle], (), \alpha_m, \zeta_d) \rightarrow \text{void}$$

56 mov t<sub>2</sub> 0  
 57 mov t<sub>3</sub> 0  
 58 mov t<sub>4</sub> t<sub>2</sub>  
 59 mov t<sub>5</sub> t<sub>3</sub>

This sequence of move instructions adds typing information about the temporary registers to the register file type.



$$\Psi(60) = (\Delta_d, \Gamma_{60}, (), \alpha_m, \zeta_d) \rightarrow void$$

where:  $\Gamma_{60} = \Gamma_d$  [  $t_0 \mapsto \langle G, int\ ref^1, \alpha_x \rangle$  ]  
 [  $t_1 \mapsto \langle B, int\ ref^1, \alpha_x \rangle$  ]  
 [  $t_2 \mapsto \langle G, int, 0 \rangle$  ]  
 [  $t_3 \mapsto \langle B, int, 0 \rangle$  ]  
 [  $t_4 \mapsto \langle G, int, 0 \rangle$  ]  
 [  $t_5 \mapsto \langle B, int, 0 \rangle$  ]

```

60  ldG t4 t0
61  ldB t5 t1
    
```

Each of these load instructions loads from some address  $\alpha_x$  in memory  $\alpha_m$ , so the type of the loaded values can be described by the static expression  $sel\ \alpha_m\ \alpha_x$ .

$$\Psi(62) = (\Delta_d, \Gamma_{62}, (), \alpha_m, \zeta_d) \rightarrow void$$

where:  $\Gamma_{62} = \Gamma_{60}$  [  $t_4 \mapsto \langle G, int, sel\ \alpha_m\ \alpha_x \rangle$  ]  
 [  $t_5 \mapsto \langle B, int, sel\ \alpha_m\ \alpha_x \rangle$  ]

```

62  add t4 t4 t4
63  add t5 t5 t5
    
```

When the value is added to itself, the static expressions are modified accordingly. At this point, we know that  $t_4$  and  $t_5$  both contain double the value in address  $\alpha_x$ , whatever that value may be.

$$\Psi(64) = (\Delta_d, \Gamma_{64}, (), \alpha_m, \zeta_d) \rightarrow void$$

where:  $\Gamma_{64} = \Gamma_{62}$  [  $t_4 \mapsto \langle G, int, (sel\ \alpha_m\ \alpha_x) + (sel\ \alpha_m\ \alpha_x) \rangle$  ]  
 [  $t_5 \mapsto \langle B, int, (sel\ \alpha_m\ \alpha_x) + (sel\ \alpha_m\ \alpha_x) \rangle$  ]

```

64  stG t0 t4
65  stB t1 t5
    
```

The green store typing rule will add the pair  $(\alpha_x, (sel\ \alpha_m\ \alpha_x) + (sel\ \alpha_m\ \alpha_x))$  to the queue. The blue store typing rule removes this pair from the queue description, enforces that it is equal to the expressions describing the operands, and updates the type of memory.

$$\Psi(66) = (\Delta_d, \Gamma_{64}, (), E'_m, \zeta_d) \rightarrow void$$

where:  $E'_m = upd\ \alpha_m\ \alpha_x\ ((sel\ \alpha_m\ \alpha_x) + (sel\ \alpha_m\ \alpha_x))$

```

66  mov t6 0
67  mov t7 0
68  sldG 0 t8
69  sldB 1 t9

```

Again, these instructions add more temporary register typing information to the register file type.

$$\Psi(70) = (\Delta_d, \Gamma_{70}, (), E'_m, \zeta_d) \rightarrow \text{void}$$

where:  $\Gamma_{70} = \Gamma_{64} [ t_6 \mapsto \langle G, \text{int}, 0 \rangle ]$   
 $[ t_7 \mapsto \langle B, \text{int}, 0 \rangle ]$   
 $[ t_8 \mapsto \langle G, \Theta_r \rightarrow \text{void}, \alpha_r \rangle ]$   
 $[ t_9 \mapsto \langle B, \Theta_r \rightarrow \text{void}, \alpha_r \rangle ]$

```

70  sfree 4
71  salloc 2

```

The *sfree* instruction modifies the stack type to be  $\alpha_\ell : \alpha_\sigma$  and updates the stack pointer types appropriately. Then the *salloc* instruction pushes two nonsense types onto the stack type and again updates the stack pointer types.

$$\Psi(72) = (\Delta_d, \Gamma_{72}, (), E'_m, \zeta_{72}) \rightarrow \text{void}$$

where:  $\zeta_{72} = \alpha_\ell - 2 : ns :: \alpha_\ell - 1 : ns :: \alpha_\ell : \alpha_\sigma$   
 $\Gamma_{72} = \Gamma_{70} [ sp_G \mapsto \langle G, \text{sptr}, \alpha_\ell - 2 \rangle ]$   
 $[ sp_B \mapsto \langle B, \text{sptr}, \alpha_\ell - 2 \rangle ]$

```

72  sst 0 t6
73  sst 1 t7

```

The two stack store instructions overwrite the nonsense types in the stack type with the types of registers  $t_6$  and  $t_7$ .

$$\Psi(74) = (\Delta_d, \Gamma_{70}, (), E'_m, \zeta_{74}) \rightarrow \text{void}$$

where:  $\zeta_{74} = \alpha_\ell - 2 : \langle G, \text{int}, 0 \rangle :: \alpha_\ell - 1 : \langle B, \text{int}, 0 \rangle :: \alpha_\ell : \alpha_\sigma$

```

74  jmpG t8

```

The green jump instruction updates the type of the destination register  $gd$ . For clarity, the type of location 75 is fully expanded. The subscript  $c$  appears on the current contexts, and the subscript  $r$  appears on those describing the type of the return address.

$$\begin{aligned}
 \Psi(75) &= (\Delta_c, \Gamma_c, (), E_{mc}, \zeta_c) \rightarrow \text{void} \\
 \text{where: } \Delta_c &= \alpha_p : \kappa_{int}, \alpha_m : \kappa_{mem}, \alpha_\ell : \kappa_{int}, \alpha_\sigma : \kappa_\sigma, \alpha_r : \kappa_{int}, \alpha_x : \kappa_{int} \\
 \Gamma_c &= pc_G \mapsto \langle G, int, \alpha_p \rangle, pc_B \mapsto \langle B, int, \alpha_p \rangle, \\
 &sp_G \mapsto \langle G, sptr, \alpha_\ell - 2 \rangle, sp_B \mapsto \langle B, sptr, \alpha_\ell - 2 \rangle, \\
 &gd \mapsto \langle G, \Theta_r \rightarrow \text{void}, \alpha_r \rangle, \\
 &t_0 \mapsto \langle G, int \text{ ref}^1, \alpha_x \rangle, t_1 \mapsto \langle B, int \text{ ref}^1, \alpha_x \rangle, \\
 &t_2 \mapsto \langle G, int, 0 \rangle, t_3 \mapsto \langle B, int, 0 \rangle, \\
 &t_4 \mapsto \langle G, int, (sel \alpha_m \alpha_x) + (sel \alpha_m \alpha_x) \rangle, \\
 &t_5 \mapsto \langle B, int, (sel \alpha_m \alpha_x) + (sel \alpha_m \alpha_x) \rangle, \\
 &t_6 \mapsto \langle G, int, 0 \rangle, t_7 \mapsto \langle B, int, 0 \rangle, \\
 &t_8 \mapsto \langle G, \Theta_r \rightarrow \text{void}, \alpha_r \rangle, t_9 \mapsto \langle B, \Theta_r \rightarrow \text{void}, \alpha_r \rangle, \\
 \zeta_c &= \alpha_\ell - 2 : \langle G, int, 0 \rangle :: \alpha_\ell - 1 : \langle B, int, 0 \rangle :: \alpha_\ell : \alpha_\sigma \\
 E_{mc} &= \text{upd } \alpha_m \alpha_x ((sel \alpha_m \alpha_x) + (sel \alpha_m \alpha_x)) \\
 \\
 \text{and: } \Theta_r &= (\Delta_r, \Gamma_r, (), \alpha'_m, \zeta_r) \\
 \Delta_r &= \alpha'_m : \kappa_{mem}, \alpha_\tau : \kappa_{int} \\
 \Gamma_r &= pc_G \mapsto \langle G, int, \alpha_r \rangle, pc_B \mapsto \langle B, int, \alpha_r \rangle, \\
 &sp_G \mapsto \langle G, sptr, \alpha_\ell - 2 \rangle, sp_B \mapsto \langle B, sptr, \alpha_\ell - 2 \rangle, \\
 &gd \mapsto \langle G, int, 0 \rangle \\
 \zeta_d &= \alpha_\ell - 2 : \langle G, int, \alpha_\tau \rangle :: \alpha_\ell - 1 : \langle B, int, \alpha_\tau \rangle \\
 &:: \alpha_\ell : \alpha_\sigma \\
 \\
 75 \quad \text{jmp}_B \quad t_9
 \end{aligned}$$

Type checking the blue jump instruction is more involved. There are 11 premises which must be satisfied. The first three require that the destination register and the jump target both have the same code type and are described by equivalent expressions.

1.  $\Gamma_c(gd) = \langle G, (\Delta_r; \Gamma_r; ()); \alpha'_m, \zeta_r \rangle \rightarrow \text{void}, \alpha_r \rangle$
2.  $\Gamma_c(rd) = \langle B, (\Delta_r; \Gamma_r; ()); \alpha'_m, \zeta_r \rangle \rightarrow \text{void}, \alpha_r \rangle$
3.  $\Delta_c \vdash \alpha_r = \alpha_r$

There exists a substitution  $S$  that gives expressions to substitute in for the expression variables in  $\Delta_r$ . These expressions may contain free variables from  $\Delta_c$ . We can construct such a substitution by inspecting the current typing information and the typing information required by the return address.

$$4. \quad \exists S. \Delta_c \vdash S : \Delta_r \quad (\text{let } S = E_{mc}/\alpha'_m, 0/\alpha_\tau)$$

The return address has appropriate requirements for the designated registers  $gd$ ,  $pc_G$  and  $pc_B$ .

$$\begin{aligned} 5. \quad S(\Gamma_r)(gd) &= \langle G, int, 0 \rangle \\ 6. \quad S(\Gamma_r)(pc_G) &= \langle G, int, \alpha_r \rangle \\ 7. \quad S(\Gamma_r)(pc_B) &= \langle B, int, \alpha_r \rangle \end{aligned}$$

And finally, the current register file type, queue description, memory description, and stack type are subtypes of the substituted equivalents for the return address.

$$\begin{aligned} 8. \quad \Delta_c \vdash \Gamma_c &\leq S(\Gamma_r) \\ 9. \quad \Delta_c \vdash () &= S(()) \\ 10. \quad \Delta_c \vdash E_{mc} &= S(\alpha'_m) \\ 11. \quad \Delta_c \vdash \zeta_c &\leq S(\zeta_r) \end{aligned}$$

At this point, we have shown that each instruction in the translation of the `double()` function results in a type that can be used to type check the next instruction. In other words, this sequence of instructions can be added to a well-typed code memory, and the resulting code memory will also be well-typed.

### 3.6 Type-preserving Optimizations

In this section, we discuss the effects of the type system on various common optimizations. These are the sorts of optimizations a compiler may apply to a low-level interme-

diate representation, which is very close to executable code. So by discussing how to support these optimizations, we are arguing that ETAL<sub>FT</sub> can be used for the final result of compilation.

Previous work [66, 62, 14] has demonstrated how to implement many common optimizations for a non-fault-tolerant typed assembly language. Below, we will sketch the effects of the additional typing infrastructure for supporting fault tolerance on some common optimizations. However, a full analysis of the impact of the type system on various optimizations is beyond the scope of this thesis.

### 3.6.1 General Considerations

Unfortunately for optimizations, ETAL<sub>FT</sub> works with actual code addresses instead of symbolic labels. Therefore inserting or removing instructions requires incrementing or decrementing all uses of subsequent code addresses. Though cumbersome, it is possible to deal with this. Alternatively, one could make minor modifications to ETAL<sub>FT</sub> in order to use a different code representation more similar to a control flow graph.

One important invariant that must be maintained is the equivalence of related static expressions. This equivalence is abstracted in the typing rules by the judgement  $\Delta \vdash E_1 = E_2$ . In the translation in Section 3.4, the two computations are exactly the same, and so syntactic equality is all that is required. Using a more powerful theorem prover to compare expressions allows flexibility in how an optimization is applied. For example, assume we have code to multiply a value by two. It is possible to optimize this code to use addition instead of costly multiplication. If this optimization is only applied to the green computation, the end result will be two values with type  $\langle G, int, E + E \rangle$  and  $\langle B, int, 2 * E \rangle$ . The type checker will allow this as long as the theorem prover used can determine that  $E + E = 2 * E$ .

### 3.6.2 Removal of Redundant Moves

If a move exists between two registers in well-typed code, then we know the registers must belong to the same computation. Therefore redundant moves will only occur within a single computation, and so removing them will continue to maintain the separation between colored computations.

### 3.6.3 Register Allocation

The most obvious issue with the ETAL<sub>FT</sub> translation is that it is constantly generating fresh temporary registers. Removing redundant moves will decrease the number of temporary registers, but for some programs the number of live temporary registers will exceed the number of actual registers.

We have already shown how to add a stack into ETAL<sub>FT</sub>. Once other optimizations have been completed, and compiler can determine a graph coloring and modify the code to spill the selected values to the stack.

In the overly simplistic translation, the stack is only accessed upon entering and exiting functions and before and after function calls. However, there is nothing in the typing rules to prevent stack operations at any point in the code, so the implementation of register allocation should be able to proceed in the normal way.

### 3.6.4 Common Subexpression Elimination

Common subexpressions elimination is one of the optimizations that caused issues in the initial implementation of the SWIFT compiler. Naively performing CSE on ETAL<sub>FT</sub> code would likely cause dependencies between the two computations, resulting in code that would fail to type check. However, it is simple to solve this by modifying CSE to be

aware of the color tags. As long as the optimization only applies within a computation of a single color, there appear to be no additional issues.

### 3.6.5 Dead Code Elimination

In general, dead code elimination is not affected by the addition of fault tolerance to the type system.

However, if dead stores are removed, it is imperative that these same modifications are made to each computation. Otherwise, the queue will be inconsistent, and so the code will fail to type check. (The same holds for the removal of silent stores, which are stores that overwrite a value with the same value.)

### 3.6.6 Constant Folding and Propagation

In constant folding and propagation, the constant has the same type as the expression it is replacing, and so these optimizations do not cause any issues in ETAL<sub>FT</sub> as long as the theorem prover can prove the equivalence. It is worth noting that a number of the instructions only take registers as operands (for example *add r<sub>1</sub> r<sub>2</sub> r<sub>3</sub>*), but adding other addressing modes for instructions would only require simple changes.

### 3.6.7 Stack Packing

In a stack packing optimization, the compiler reuses the same stack slot for two values with disjoint lifetimes. Because ETAL<sub>FT</sub> uses strong updates when storing to the stack, this can be supported. The information about where a stack pointer points is separate from the information about what each stack location contains. So even if one register is

used to perform a strong update on a stack pointer, any other registers that also point to this location will be aware of the change.

### 3.6.8 Instruction Scheduling

As with TAL<sub>FT</sub>, ETAL<sub>FT</sub> allows a lot of flexibility in the instruction scheduling. Within a computation, all the usual reorderings are allowed, given the standard constraints about instruction dependencies.

There is a great deal of flexibility between the green and blue computations as well. Essentially, computations are synchronized at the block level. In other words, the green computation may not execute past a control flow transfer until the blue computation has also executed the transfer. Within a block, the main restriction is that each green store instruction comes before the corresponding blue store instruction. Though our machine model assumes an infinite length store queue, in reality this would be implemented as a fixed-length buffer. This will require the compiler to limit the number of green store instructions that can come before the first of the corresponding blue instructions.

The new instructions in ETAL<sub>FT</sub> such as *malloc* and *sst* would likely be implemented with sequences of instructions, which may cause unanticipated interactions with scheduling. For example, if *sst* is implemented using *st<sub>G</sub>* and *st<sub>B</sub>*, then it cannot occur while there is a green store pending in the queue. Previous work on LTAL [16] develops a standard typed assembly language without macros to avoid interference with scheduling. Investigating the interaction of equivalent ideas with fault tolerance is left as future work.



## 3.7 Summary

This chapter has presented  $ETAL_{FT}$ , a version of  $TAL_{FT}$  extended with support for stacks and dynamic memory allocation. We show how to compile a simple imperative source language into fault-tolerant assembly code. This demonstrates that the restrictions imposed by the type system do not affect our ability to express interesting programs. In addition, we discuss how  $ETAL_{FT}$  can support common low-level optimizations, thereby showing that it can be used as the final language in a realistic compiler.

# Chapter 4

## TAL<sub>CF</sub> : Reasoning about Control Flow

The previous two chapters focused on a *hybrid* transient fault solution that included specialized hardware to help detect faults. Although this type of the solution is extremely promising, there are many cases in which additional hardware is not available. Fortunately, techniques based only on software can detect a large number of faults on off-the-shelf processors.

Chapter 5 covers related work on existing software-only techniques, but the general methodology is similar to what we have already seen, though the comparisons are performed by additional instructions in software. Doing so sometimes leaves a *window of vulnerability* in which faults may not be detected before affecting the observable behavior of the program.

In particular, this chapter focuses on reasoning about a software solution for detecting control flow faults, which occur when a transient fault causes control to jump to an

---

*Material in this chapter is joint work with David Walker and has been previously published as [54] and [53].*

unexpected code address.  $TAL_{FT}$  avoids the issue by using a hardware comparison in the fetch stage and control-flow instructions to detect such faults.

Researchers [47, 60, 10] have developed techniques for detecting many, but not all, classes of control-flow errors entirely in software. Though these techniques are promising, they have not been proven sound and many theoretical questions remain. In particular, is it possible to characterize the effectiveness of these techniques *analytically* as opposed to empirically? In other words, can we prove that such techniques are sound with respect to an interesting and nontrivial, though incomplete, fault model? One of the key benefits of such an analysis is that it would guarantee that an important fragment of the problem has been thoroughly solved and thereby allow researchers to study auxiliary instrumentation techniques that address the remaining incompleteness. Perhaps more importantly, a formal fault model and proof of soundness would define an important hardware/software interface: The software has been proven to handle faults that lie within the model; future hardware designers need only provide mechanisms to catch those faults that lie outside the model. Such results would show how to shift a substantial portion of the control-flow checking burden from the hardware to software. This may lead to simpler hardware designs as well as the opportunity to trade performance for reliability at *compile time* as opposed to *hardware design time*.

This chapter presents  $TAL_{CF}$ , a type system for reasoning about a software-only technique for detecting a specific class of control-flow faults. The technique instruments each basic block with a certain instructions that implement a fault tolerance protocol. From a technical perspective, the type system introduces a novel way of classifying the reliability properties of program values and entire machine states, generalizing the earlier “color systems” used by  $\lambda_{zap}$  [73] and  $TAL_{FT}$ . The type system is also of interest for the way it uses a collection of abstract types to track the state of the fault tolerance

protocol. The key technical challenge in reasoning about such a solution is the fact that after a control-flow fault has occurred, it is impossible to count on almost *any* standard program invariant. So, how can one carry out a proof of type preservation under such circumstances?

The rest of the chapter is organized as follows. First, Section 4.1 gives additional intuition about the problem and solution by explaining a simple assembly-language protocol for detecting control-flow errors. This protocol is a simplified version of the protocols used by Oh [47] and Reis [60]. Section 4.2 begins the more formal work by defining the syntax and operational semantics of an idealized assembly language that includes rules to model transient faults. Section 4.3 defines the type system that guarantees that assembly code follows the simple protocol outlined earlier in Section 4.1. This type system, particularly the special value and machine state typing rules, codifies the major invariants needed to prove the subsequent type safety and strong reliability properties. Section 4.4 sketches the major components of the type safety and fault-tolerance proofs. Section 4.5 shows that our typed assembly language is sufficiently expressive to translate basic “while” programs into well-typed, fault-tolerant code. Finally, Section 4.6 summarizes.

## 4.1 Informal Overview

When a transient fault causes the actual sequence of control flow blocks visited by a program to deviate from the expected sequence, we say a control-flow error has occurred. In our model, control-flow errors arise in three different ways: (1) there may be a fault to the target address of a jump instruction; (2) there may be a fault to the target address of a conditional jump instruction; or (3) there may be a fault to the boolean used to

decide whether to jump or fall through a conditional. Such faults may occur immediately prior to attempting the control-flow transfer or at any other time during the computation. However, whenever a control-flow operation is executed, we assume execution is either transferred to the beginning of some valid block, or to some invalid block or illegal instruction. In the latter case, we assume the hardware immediately catches an attempt to execute the illegal instruction. We do not consider the possibility that a fault causes a control flow transfer to a legal instruction in the middle of some valid block. We discuss this limitation in Section 4.6.

As in previous chapters, we adhere to the *Single Event Upset* model, which states that only one fault may occur during an execution. However, even though just one fault occurs, faulty values may be copied, propagated and used in any way an ordinary value may be used. Hence, a single fault can lead to arbitrarily many corrupted values if not caught soon after it occurs.

The goal of this work is to develop and prove correct a software protocol that guarantees such control-flow errors can never go undetected. The central challenge in this endeavor is to overcome the problem that *no single value can ever be trusted to be correct* — a transient fault may strike any value in any register. Consequently, as is usual in fault tolerance, the solution is to avoid relying on any single value by replicating the critical state and checking replicas against one another. In this case, the critical state is the value of the program counter. Checking the correctness of a control-flow transfer involves creating a replica of the intended control-flow destination and then checking the replica against the real program counter to detect any difference.

To be more specific, compiled code creates the replica prior to any control-flow transfer by moving the intended destination into a designated register. We refer to this register as the *intention register*  $ri$ . This intention register is part of the global “calling

convention” for fault-tolerant control flow transfers. We fix the register so that all jump targets know where to find the intended destination, even when there has been a control-flow fault.

As an example, to jump to address L2, one might use the following code sequence. In this code, we leave ellipsis in between instructions to emphasize our system allows flexible scheduling of instructions — ordinary instructions may be interleaved with the instructions used to guarantee fault tolerance.

```
L1: ...; movi ri, L2; ...; movi r2, L2; ...; jmp r2;
```

Because the intention register *ri* plays a special role in the protocol for detecting control-flow errors, we will need to type-check the move instruction that loads this register in a special way. To designate the move as special, we henceforth write it *intend L2* rather than *movi ri, L2* as in the following example code.

```
L1: ...; intend L2; ...; movi r2, L2; ...; jmp r2
```

If the intention register has been set properly prior to all jump instructions, the jump targets are able to catch control flow errors. To be specific, all jump targets should be instrumented with the following code.

```
Lk: movi r2, Lk; ...; sub r2, r2, ri; ...; brnz r2, lrecover; ...
```

Here, the current block address *Lk* is loaded into some register *r2*. That register is then compared with the contents of *ri* and if there is any difference, control is transferred to *lrecover*, an address containing recovery code. Once again, since the branch to the recovery code plays a special role in the fault-tolerance protocol, we give it the special syntax *recovernz r2*. Thus, our detection code will henceforth be written as follows.

```
L2: movi r2, L2; ...; sub r2, r2, ri; ...; recovernz r2; ...
```

As an example of how a transient fault might be caught using our protocol, suppose register  $r_2$  is corrupted just prior to attempting to execute the jump to  $L_2$  in block  $L_1$ . Upon arrival at some erroneous control flow block, say  $L_3$ , the intended destination  $L_2$  remains safely untouched in register  $r_i$ , though, unnervingly, all other program invariants may be disrupted. The target code compares the contents of  $r_i$  (*i.e.*,  $L_2$ ) with  $L_3$ , which it loaded into  $r_2$  after arriving at the current block. It detects a difference and jumps to the recovery code.

One must also consider what happens if faults strike at different times or in different places. For instance, the jump target might have been corrupted much earlier than we suggested above, perhaps just after being initially loaded into  $r_2$ , instead of just prior to the jump. Will that make a difference? In this case, no. Likewise,  $r_i$  might be corrupted, either before or after jumping. In this case, we reach the correct destination, but it appears as though there was a fault because  $r_i$  differs from the current block label (assuming the fault occurs prior to the subtraction). Unable to tell the difference between a fault in the intention register and a fault in the control-flow transfer itself, we jump to recovery code. A number of other scenarios must also be analyzed — in order to have confidence in the solution, one must do so in a principled, disciplined fashion.

It is important to observe that similar, but subtly different code sequences do not adequately protect against faults. In particular, optimizations like copy propagation, common subexpression elimination and some code motion transformations are not always semantics-preserving in the context of transient faults. For instance, the following simple change to the way block  $L_1$  was written above leads to a vulnerability.

```
L1: movi r2, L2(*); ...; movi ri, r2(**); ...; jmp r2
```

Here, a single transient fault to `r2` anywhere between execution of instructions `(*)` and `(**)` results in an uncaught control-flow fault as both the jump target and the intention register will simultaneously be incorrect.

Likewise, the code motion transformation illustrated below shifts the move from a target block into the jumping block and creates a vulnerability.

```
L1: movi r2, L2; intend L2; movi r3, L2; jmp r2;
Lk: sub r3, r3, ri(**); recovernz r3; ...
```

Above, a fault to `r2` causes a control-flow error, but testing `r3` against `ri` at instruction `(***)` will not help detect the fault. The conclusions to draw from these examples are that the correctness properties of this code are indeed subtle and that verifying fault tolerance properties *after* the compiler has completed its suite of performance optimizations may help detect errors in code generation.

**Conditional Branches.** The protocol for handling conditional branches is slightly more involved than the case for jumps, but follows a similar pattern. We begin by assuming that the condition for the jump is held in registers `r4` and `r4'`. These two registers must be *independent replicas* of one another. In other words, in the absence of faults, they should contain the same boolean value, and moreover, a fault to one should have no impact on the value of the other. Given this assumption (which will be verified by our type system), the following code sequence sets up a conditional branch, which may fall through to `L2` or may jump to `L3`. The code uses a conditional branch `brz r4, r3`, which jumps to `r3` if `r4` is zero and otherwise falls through to `L2`. It also uses a conditional move `cmovz`



$r4'$ ,  $ri$ ,  $r3'$ , which moves the contents of  $r3'$  into  $ri$  if  $r4'$  is zero, and otherwise does nothing. <sup>1</sup>

```
L1:  ...; movi r3, L3; movi r3', L3; ...; intend L2;
      cmovz r4', ri, r3'; brz r4, r3;
L2:  ...
L3:  ...
```

Again, to indicate the special role of  $ri$  and simplify the presentation, we will henceforth write the conditional move `cmovz r4', ri, r3'` as `intendz r4', r3'`. Intuitively, the `intend` instruction unconditionally sets the intention register, whereas the `intendz` instruction conditionally sets the intention register. The error-detection code in blocks labeled L2 and L3 is identical to the error-detection code discussed earlier for jumps, as it must be.

**Summary** With just a few, well-thought-through instructions, it is possible to create a redundant copy of the intended destination of any control flow transfer prior to initiating the transfer itself. Moreover, at any control-flow target, it is possible to use that redundant copy to check that control has actually arrived at the proper place. However, as our examples illustrated, it is also easy to make slight errors in the process. In addition, because transient faults can occur at so many different places in the protocol and influence so many different bits of state, one needs a proof to believe such a protocol will work. Hence, in the following sections, we make the machine's operational semantics and fault model precise and develop a sound type system strong enough to verify that the “good” instruction sequences we have discussed in this section are indeed fault tolerant.

---

<sup>1</sup>Many architectures including the IA-32 following the Pentium Pro, the Sparc V-9 and the IA-64 have conditional moves. If the architecture does not have a conditional move, a conditional branch and a move instruction can be used instead, but this branch will not be protected against faults.

<i>colors</i>	$c$	$::= G \mid B \mid O$
<i>integers</i>	$n$	$::= \dots \mid -1 \mid 0 \mid 1 \mid \dots$
<i>colored values</i>	$v$	$::= c \ n$
<i>code memory</i>	$C$	$::= \cdot \mid C, \ell \rightarrow b$
<i>registers</i>	$r$	$::= r_i \mid r_1 \mid \dots \mid r_n$
<i>register file</i>	$R$	$::= \cdot \mid R, r \rightarrow v$
<i>history</i>	$h$	$::= \ell_1, \dots, \ell_n$
<i>instructions</i>	$i$	$::= \text{movi } r_d \ v \mid \text{sub } r_d \ r_s \ r_s$ $\mid \text{intend } r_t \mid \text{intendz } r_z \ r_t$ $\mid \text{recovernz } r_z$
<i>blocks</i>	$b$	$::= i; b \mid \text{jmp } r_t \mid \text{brz } r_z \ r_t$
<i>state</i>	$\Sigma$	$::= (C, h, R, b)$
<i>final states</i>	$\mathcal{F}$	$::= \Sigma \mid \text{recover}(h) \mid \text{herror}(h)$

Figure 4.1: Machine State Syntax.

## 4.2 The Control-Flow Machine

Figure 4.1 summarizes the syntax of the assembly language and machine states. For clarity and elegance, we will work with a minimal assembly instruction set involving move (`movi`), subtraction (`sub`), jump (`jmp`) and conditional branch if zero (`brz`) instructions as well as the special macros `intend`, `intendz` and `recovernz`. Instruction operands include constant values  $v$  and registers  $r$ . Similar to TAL<sub>FT</sub>, we annotate every value with a *color*, although TAL<sub>CF</sub> has three colors: green  $G$ , blue  $B$  and orange  $O$ . Again, these colors have no operational significance, but will be used by the simulation relation. The only kind of value is an integer. In general, meta-variable  $n$  ranges over integers, but when we wish to emphasize that an integer will be used as an address, we use the meta-variable  $\ell$ .

Instructions are grouped together in code blocks  $b$ . These blocks are always terminated by either a jump or a conditional branch instruction. Code memory  $C$  is a partial map from addresses to valid code blocks  $b$ . Addresses are ordered, and we use the notation  $\ell + 1$  to refer to the address of the block following the block at  $\ell$ . If a block at  $\ell$  ends with a conditional branch, we assume  $\ell + 1$  inhabits the domain of  $C$  — in other words, conditional branches always have a block to fall through to.

As before, the register file  $R$  is a mapping from registers to the colored values they contain. The registers include the intention register  $r_i$  and a number of general-purpose registers  $r_1$  through  $r_n$ . We use the notation  $R(r)$  to denote the contents of  $r$  in  $R$ . We use the notation  $R[r \mapsto v]$  to denote a new register file  $R'$  created by updating  $R$  so it maps  $r$  to  $v$ . When we wish to refer to the unannotated integer  $n$  as opposed to the colored value  $c\ n$  in a register  $r$  in  $R$ , we use the notation  $R_{val}(r)$ . Similarly,  $R_{col}(r)$  refers to the color annotating the value in  $r$ .

An ordinary abstract machine state  $\Sigma$  is a tuple containing code  $C$ , history  $h$ , register file  $R$  and code block to be executed  $b$ . The history  $h$  is a sequence of labels. It records the code blocks visited during the current execution. As  $TAL_{CF}$  has no concept of memory, we will define “observable behavior” and the sequence of blocks visited. In addition to ordinary abstract machine states, there are two special “final states.” The state  $recover(h)$  represents a state in which a transient fault has occurred and has been caught. The labels in history  $h$  were visited during the execution. The state  $hwerror(h)$  represents a state in which a transient fault causes transition to an invalid address.

### 4.2.1 Dynamic Semantics

We model the dynamic semantics of the assembly language using a small-step operational semantics. In general, the single-step operational judgments have the form  $\Sigma \longrightarrow_k \mathcal{F}$

where  $k$ , which is either zero or one, records the number of faults that occur during the step.

**The Fault Model.** The most interesting rules in the system are the rules modeling faults. The primary rule (*zap-reg*) is familiar and again states that the value in any register may be corrupted arbitrarily, though its color tag (which has no operational significance) remains unchanged.

$$\frac{R(r) = c \ n}{(C, h, R, b) \longrightarrow_1 (C, h, R[r \mapsto c \ n'], b)} \text{ (zap-reg)}$$

The rule above may fire at any time. In particular, it may fire just prior to execution of a jump (*jmp*  $r_t$ ) or a branch (*brz*  $r_z \ r_t$ ), corrupting the jump target in register  $r_t$ . Such a fault models a control-flow error. Of course, it is equally possible that any other register is corrupted.

For uniformity in our fault model, we also consider errors in the execution of the *recovernz*  $r_z$  instruction. Recall, this instruction is merely a macro for the conditional branch *brz*  $r_z \ \ell_{recover}$ . However, since  $\ell_{recover}$  is a constant, it is unaffected by faults in registers modeled by the *zap-reg* rule (our other branching instructions take arguments in registers). To simulate a fault that causes control to jump somewhere other than the  $\ell_{recover}$  label when the  $r_z$  register contains a non-zero value, we add the following rules.

$$\frac{R_{val}(r_z) \neq 0}{(C, h, R, \text{recovernz } r_z; b) \longrightarrow_1 (C, h, R, C(\ell))} \text{ (zap-recovernz1)}$$

$$\frac{R_{val}(r_z) \neq 0}{(C, h, R, \text{recovernz } r_z; b) \longrightarrow_1 \text{hwerror}(h)} \text{ (zap-recovernz2)}$$

The *zap-recovernz1* rule expresses the possibility that a fault causes execution to jump to some random block labeled  $\ell$  rather than the recovery code block. The *zap-recovernz2* rule expresses the possibility that a fault causes control to jump to an illegal address. Attempted execution of code at this address results in immediate transition to the final state  $hwerror(h)$ , where  $h$  represents the sequence of blocks visited not including the illegal address.

**Other Operational Rules.** All other operational rules are presented in Figure 4.2. The majority of these rules are quite unsurprising. For instance, the *movi* rule implements the move by updating the register file. Notice that the index on the arrow is “0” indicating no fault occurs during this transition. Naturally, the *intend* rule is similar to *movi* as *intend* is just a macro for a move into  $r_i$ .

Skipping to the bottom of the figure, it is important to notice there are two rules for expressing the semantics of a `jmp  $r_t$`  instruction. The first rule *jmp* fires whenever  $r_t$  contains the address of a valid block. Of course, due to a fault earlier in execution, the address in  $r_t$  may not be the intended destination for this jump. In addition to transferring control to the new block, this instruction does some bookkeeping. In particular, it extends the current history with the destination address and it changes the color of  $r_t$  to be orange. The latter effect facilitates the proof of correctness and will be explained in detail in Section 4.3. The second rule *jmp-hw-error* fires whenever  $r_t$  does *not* contain the address of a valid block. In this case, there is an attempt to transfer control to an illegal address, which is caught by the hardware. The rules for conditional branches follow a similar pattern to those for the unconditional jumps.

$$\frac{}{(C, h, R, \text{movi } r_d \ v; b) \longrightarrow_0 (C, h, R[r_d \mapsto v], b)} \text{ (movi)}$$

$$\frac{v' = R_{\text{col}}(r_a) (R_{\text{val}}(r_a) - R_{\text{val}}(r_b))}{(C, h, R, \text{sub } r_d \ r_a \ r_b; b) \longrightarrow_0 (C, h, R[r_d \mapsto v'], b)} \text{ (sub)}$$

$$\frac{}{(C, h, R, \text{intend } r_t; b) \longrightarrow_0 (C, h, R[r_t \mapsto R(r_t)], b)} \text{ (intend)}$$

$$\frac{R_{\text{val}}(r_z) = 0}{(C, h, R, \text{intendz } r_z \ r_t; b) \longrightarrow_0 (C, h, R[r_t \mapsto R(r_t)], b)} \text{ (intendz-set)}$$

$$\frac{R_{\text{val}}(r_z) \neq 0}{(C, h, R, \text{intendz } r_z \ r_t; b) \longrightarrow_0 (C, h, R, b)} \text{ (intendz-unset)}$$

$$\frac{R_{\text{val}}(r_z) = 0}{(C, h, R, \text{recovernz } r_z; b) \longrightarrow_0 (C, h, R, b)} \text{ (recovernz-ok)}$$

$$\frac{R_{\text{val}}(r_z) \neq 0}{(C, h, R, \text{recovernz } r_z; b) \longrightarrow_0 \text{recover}(h)} \text{ (recovernz-halt)}$$

$$\frac{R_{\text{val}}(r_z) = 0 \quad R_{\text{val}}(r_t) \in \text{Dom}(C)}{(C, h, R, \text{brz } r_z \ r_t) \longrightarrow_0 (C, (h, R_{\text{val}}(r_t)), R[r_t \mapsto O R_{\text{val}}(r_t)], C(R_{\text{val}}(r_t)))} \text{ (brz-taken)}$$

$$\frac{R_{\text{val}}(r_z) \neq 0 \quad \ell+1 \in \text{Dom}(C)}{(C, h, R, \text{brz } r_z \ r_t) \longrightarrow_0 (C, (h, \ell+1), R[r_t \mapsto O R_{\text{val}}(r_t)], C(\ell+1))} \text{ (brz-untaken)}$$

$$\frac{R_{\text{val}}(r_z) = 0 \quad R_{\text{val}}(r_t) \notin \text{Dom}(C)}{(C, h, R, \text{brz } r_z \ r_t) \longrightarrow_0 \text{hwerror}(h)} \text{ (brz-hw-error)}$$

$$\frac{R_{\text{val}}(r_t) \in \text{Dom}(C)}{(C, h, R, \text{jmp } r_t) \longrightarrow_0 (C, (h, R_{\text{val}}(r_t)), R[r_t \mapsto O R_{\text{val}}(r_t)], C(R_{\text{val}}(r_t)))} \text{ (jmp)}$$

$$\frac{R_{\text{val}}(r_t) \notin \text{Dom}(C)}{(C, h, R, \text{jmp } r_t) \longrightarrow_0 \text{hwerror}(h)} \text{ (jmp-hw-error)}$$

Figure 4.2: Operational Semantics.

### 4.3 Typing

The design of the type system is based on three main concepts:

- Classifying the reliability properties of values.
- Using abstract types to make sure that the fault tolerance protocol proceeds in the correct order, with no steps omitted or inappropriate steps inserted.
- Equivalence checking to ensure that redundant values act as proper backups to the original.

The following paragraphs explain the main intuitions behind each concept. Later subsections will give precise details.

**Classifying the Reliability Properties of Values.** Because faults occur completely unpredictably and at run time, it is not possible for the type system to know which values have incurred faults or to track the propagation of presumed faulty values precisely. It is not possible to know exactly values may or may not be trusted. The type system will have to approximate these properties somehow. It does so by assigning each value to one of several compile-time “groups” and ensuring that each member of a group has related reliability properties. As a mnemonic, each group has an associated color  $c$ , which may be either *green*, *blue* or *orange*. This is a generalized version of the color scheme used by TAL<sub>FT</sub>.

As we saw in Section 4.1, the protocol for detecting faults in software involves keeping redundant copies of the values used in control flow transfers and using these to check for correct control flow. We will refer to the main computation as the *green* computation, and the redundant copies as the *blue* (or “backup”) computation. Most

values either belong to the green group or to the blue group. As in  $TAL_{FT}$ , these two groups have the property that they are *redundant* and *independent*. In other words, a fault in a green value can never percolate to a blue value and vice versa. Consequently, when corresponding green and blue values are compared at least one of them must be correct, even when a fault has occurred. This mutual independence property is ensured by a series of simple checks in the type system that guarantee that green values are not used to construct blue values and vice versa.

But what if a control-flow fault *has* occurred? In that case, almost all program invariants are invalidated, including any properties of either blue or green values. Fortunately, though, the defining characteristic of *orange* values is preservation of their properties in just this situation.

There are two general mechanisms by which one can guarantee orange values maintain their expected properties in the face of a control-flow fault. The first mechanism is to ensure that the orange value in question is not live across the control-flow transfer: If the value has been constructed in the current block and does not depend upon values in previous blocks, a control-flow error will not influence its properties. This first mechanism is used in the checking code at the beginning of each program block. In particular, the operation that moves a label into a register at the beginning of a block may label its results orange:

```
Lk: movi r2, Lk;           // r2 is orange
    ...;
    sub r2, r2, ri;
    ...;
    recovernz r2;
    ...
```



The second mechanism involves ensuring that every possible control-flow transfer maintains the invariant in question. If the invariant is true across *every* control-flow transfer, then it is true no matter where control winds up. This second mechanism is used to classify the contents of  $r_i$  as orange across every control-flow transfer. Just as the type system isolates green values from blue and blue from green, orange is also isolated from the other two. Again, the purpose is to avoid having a fault in one color influence the others.

Although values are classified using colors, we again use the idea of *zap tags* to classify machine states. Intuitively, each zap tag specifies which colors may no longer be trusted. For example, if zap tag  $Z$  is empty (written “.”), then there have been no faults during the computation, and all values, no matter what their color, satisfy the standard invariants associated with their compile-time type. On the other hand, if  $Z$  is a color  $c$ , then there has been a fault to a value colored  $c$  and, moreover, the corruption may have spread to any other value colored  $c$ . Consequently, values colored  $c$  will not necessarily satisfy any particular properties associated with their compile-time type.

The final zap tag  $CF$  classifies machine states after a control-flow error has occurred. In this case, control may have transferred somewhere totally unexpected, and so we know nothing about green *or* blue values. Fortunately, though, the properties of orange values remain valid.

Figure 4.3 summarizes the properties that hold under each zap tag while in block  $\ell$ . We say a value is *trusted* if it satisfies standard canonical forms properties (e.g., a value with code type is actually a pointer to valid code). We say a value is *untrusted* when we cannot guarantee standard canonical forms properties hold.

We say a zap tag  $Z$  is a subtype of another  $Z'$ , written  $Z \leq Z'$ , when the values in machine states classified by  $Z$  are more trusted than the values in machine states classified

Zap Tag	$G$ values	$B$ values	$O$ values	$\ell$ correct
.	trusted	trusted	trusted	yes
$G$	<i>untrusted</i>	trusted	trusted	yes
$B$	trusted	<i>untrusted</i>	trusted	yes
$O$	trusted	trusted	<i>untrusted</i>	yes
$CF$	<i>untrusted</i>	<i>untrusted</i>	trusted	<i>no</i>

Figure 4.3: Properties of colored values and zap tags.

by  $Z'$ . Hence the empty zap tag is a subtype of all other zap tags, and both  $B$  and  $G$  zap tags are a supertype of  $CF$ . We use this relationship in the Preservation Theorem. Well-typed states will remain well-typed, although the zap tag may escalate to a subtype. For example, a program may start out well-typed under the empty zap tag. After a fault affects a green value, the program is well-typed under the  $G$  zap tag. If a corrupt green value is eventually used in a control flow transfer, then the program becomes well-typed under the  $CF$  zap tag. Each time the zap tag changes, more values become untrusted.

**Typing Protocol Stages.** The instructions in each block can be thought of as being divided into three distinct stages – the *checking code*, the *block body*, and the *exit code*. Each of these stages has its own distinct invariants. The type of intention register  $r_i$  encodes the current stage and ensures that the stages occur in the correct order. It also guarantees no part of the protocol can be omitted or any inappropriate instruction added. These stages may be summarized as follows.

1. The checking code compares the intended target with the current location to determine if there has been a control flow fault. In this region,  $r_i$  must be colored orange and have basic type *check*.

```

L1: movi r1, L1;
    ...
    sub r1, r1, ri;
    ...
    recovernz r1;
    .
    .
    intend L2;
    ...
    movi r2, L2;
    ...
    jmp r2;

```

} *checking code*  
 } *block body*  
 } *exit code*

Figure 4.4: Example: Protocol Stages.

2. In the block body, we already know the control flow correctly transferred to this block. At the end of this sequence, there is some green register that holds the target label for the next control flow transfer and some blue register that holds the duplicate copy of this label. In the absence of faults, these two values are equal. In this region,  $r_i$  must have basic type *ok*.
3. The exit code sequence sets the intended target and transfers control to the new block. In the exit code sequence,  $r_i$  is colored blue and has type *go* when an intention has been set, and type *goz* when a conditional intention has been set. As we saw in Section 4.2.1,  $r_i$  is recolored orange during the execution of the control flow transfer.

For example, consider the example code sequences from Section 4.1 shown in Figure 4.4. On entry, each block first checks that control has reached this block correctly, and sets its intention before transferring control to another block.

## Static Expressions

<i>exp kinds</i>	$\kappa ::= \kappa_{int} \mid \kappa_{hist}$
<i>exp contexts</i>	$\Delta ::= \cdot \mid \Delta, x : \kappa$
<i>exps</i>	$E ::= x \mid n \mid E - E \mid E?E : E$
<i>substitutions</i>	$S ::= \cdot \mid S, E/x$

## Types

<i>stage description</i>	$\rho ::= check \mid ok \mid go \mid goz$
<i>basic types</i>	$\tau ::= int \mid \rho \mid \forall[\Delta](\Gamma, \sigma)$
<i>value types</i>	$t ::= \langle c, \tau, E \rangle$
<i>type option</i>	$\tau_{opt} ::= \tau \mid undef$

## Context Typing

<i>heap typing</i>	$\Psi ::= \cdot \mid \Psi, \ell \rightarrow \tau$
<i>reg file types</i>	$\Gamma ::= \cdot \mid \Gamma, r \rightarrow t$
<i>history typing</i>	$\sigma ::= \varepsilon \mid x \mid \sigma \circ E$
<i>zap tags</i>	$Z ::= \cdot \mid c \mid CF$

Figure 4.5: Typing Syntax.

**Testing Value Equivalence.** There are many places in the fault tolerance protocol where we require a blue value to be an independent and redundant copy of a green value. To ensure that blue and green values are equal in the absence of faults, we use the same technique using static expressions as in previous chapters, though the language of static expressions is different.

Now that we have given the intuition behind the type system design, we will move on to the technical details. The syntax for the type system is presented in Figure 4.5, and the next few subsections explain the elements and the corresponding judgments in detail.

### 4.3.1 Value Typing

The type of a value is a triple  $\langle c, \tau, E \rangle$ . The color  $c$  is assigned according to the intuitions expressed in the previous subsection. A basic type  $\tau$  is either an integer, a code type, or a special type that indicates the state of the fault tolerance protocol.

The third component  $E$  is a static expression that describes the value in more detail. These expressions are used to require that blue and green computations compute identical results in the absence of faults. In  $TAL_{CF}$ , expressions include variables  $x$ , integers  $n$ , subtraction  $E_1 - E_2$  and conditional expressions  $E_1 ? E_2 : E_3$  which equal  $E_2$  when  $E_1$  is non-zero and  $E_3$  when  $E_1$  is zero.

Expression judgments are shown in Figures 4.6 and 4.7. The kinding judgment  $\Delta \vdash E : \kappa$  holds when all the free variables in  $E$  are contained in  $\Delta$ . Expression  $E$  has kind  $\kappa_{int}$  when it describes an integer and kind  $\kappa_{hist}$  when it describes a history typing. Expression variables  $x$  are the only expressions that can have type  $\kappa_{hist}$ . Judgments  $\Delta \vdash \sigma$  wf,  $\Delta \vdash \Gamma$  wf, and  $\Delta \vdash \tau$  wf hold when all expressions used in these constructs are well-kinded. The judgment  $\Delta \vdash S : \Delta'$  holds when  $S$  provides substitutions for all variables in  $\Delta'$ , and the substituted expressions are well-formed in  $\Delta$ .

The function  $\llbracket E \rrbracket$  supplies the denotation of the closed static expression  $E$  as an integer. The judgments  $\Delta \vdash E_1 = E_2$  and  $\Delta \vdash E_1 \neq E_2$  hold when the relation holds for all substitutions of the variables in  $\Delta$ .  $\Delta \vdash \sigma_1 = \sigma_2$  simply extends this relationship to all expressions in the two sequences.

**Value Typing Judgment.** The value typing judgment has the form  $\Delta; \Psi \vdash^Z v : t$  and is shown in Figure 4.8. The context  $\Delta$  contains free expression variables, and the heap type  $\Psi$  maps integer addresses to basic types. The zap tag  $Z$  characterizes the current state of the machine as explained earlier.  $Z$  is always the empty tag when a user checks a program

$\Delta \vdash E : \kappa$ 

$$\frac{x \in \text{Dom}(\Delta)}{\Delta \vdash x : \Delta(x)} \text{ (wf-var)} \quad \frac{\Delta \vdash E_1 : \kappa_{int} \quad \Delta \vdash E_2 : \kappa_{int}}{\Delta \vdash E_1 - E_2 : \kappa_{int}} \text{ (wf-sub)}$$

$$\frac{}{\Delta \vdash n : \kappa_{int}} \text{ (wf-int)} \quad \frac{\Delta \vdash E_1 : \kappa_{int} \quad \Delta \vdash E_2 : \kappa \quad \Delta \vdash E_3 : \kappa}{\Delta \vdash E_1 ? E_2 : E_3 : \kappa} \text{ (wf-ifexp)}$$

 $\Delta \vdash \sigma \text{ wf}$ 

$$\frac{}{\Delta \vdash \varepsilon \text{ wf}} \text{ (wf-}\sigma\text{-}\varepsilon) \quad \frac{\Delta \vdash x : \kappa_{hist}}{\Delta \vdash x \text{ wf}} \text{ (wf-}\sigma\text{-var)} \quad \frac{\Delta \vdash \sigma \text{ wf} \quad \Delta \vdash E : \kappa_{int}}{\Delta \vdash \sigma \circ E \text{ wf}} \text{ (wf-}\sigma)$$

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

$$\frac{\forall r. \Gamma(r) = \langle c, \tau, E \rangle \wedge \Delta \vdash \tau \text{ wf} \wedge \Delta \vdash E : \kappa_{int}}{\Delta \vdash \Gamma \text{ wf}} \text{ (wf-R)}$$

 $\Delta \vdash \tau \text{ wf}$ 

$$\frac{}{\Delta \vdash \text{int wf}} \text{ (wf-int)} \quad \frac{}{\Delta \vdash \rho \text{ wf}} \text{ (wf-}\rho)$$

$$\frac{(\Delta \cup \Delta') \vdash \Gamma' \text{ wf} \quad (\Delta \cup \Delta') \vdash \sigma' \text{ wf}}{\Delta \vdash \forall[\Delta'](\Gamma', \sigma') \text{ wf}} \text{ (wf-}\forall[\Delta'](\Gamma', \sigma'))$$

 $\Delta \vdash S : \Delta'$ 

$$\frac{}{\Delta \vdash \dots} \text{ (subst-emp-t)} \quad \frac{\Delta \vdash S' : \Delta'' \quad \Delta \vdash E : \kappa \quad x \notin (\Delta \cup \Delta'')}{\Delta \vdash S', E/x : (\Delta'', x : \kappa)} \text{ (subst-t)}$$

Figure 4.6: Static Expression Judgments, Part 1.

[[E]]

$$\begin{aligned}
[[n]] &= n \\
[[E_1 - E_2]] &= [[E_1]] - [[E_2]] \\
[[E_b ? E_t : E_f]] &= \text{if } [[E_b]] \text{ then } [[E_t]] \text{ else } [[E_f]]
\end{aligned}$$

$\Delta \vdash E = E$

$$\frac{\Delta \vdash E_1 : \kappa_{int} \quad \Delta \vdash E_2 : \kappa_{int} \quad \forall S. \cdot \vdash S : \Delta \implies [[S(E_1)]] = [[S(E_2)]]}{\Delta \vdash E_1 = E_2} \text{ (E-eq)}$$

$$\frac{\Delta \vdash E_1 : \kappa_{int} \quad \Delta \vdash E_2 : \kappa_{int} \quad \forall S. \cdot \vdash S : \Delta \implies [[S(E_1)]] \neq [[S(E_2)]]}{\Delta \vdash E_1 \neq E_2} \text{ (E-neq)}$$

$\Delta \vdash \sigma = \sigma$

$$\frac{}{\Delta \vdash \varepsilon = \varepsilon} \text{ (\varepsilon-eq)} \quad \frac{}{\Delta \vdash x = x} \text{ (\sigma-var)}$$

$$\frac{\Delta \vdash E_1 = E_2 \quad \Delta \vdash \sigma_1 = \sigma_2}{\Delta \vdash \sigma_1 \circ E_1 = \sigma_2 \circ E_2} \text{ (\sigma-eq)}$$

Figure 4.7: Static Expression Judgments, Part 2.

$$\boxed{\Psi \vdash n : \tau}$$

$$\frac{}{\Psi \vdash n : \text{int}} \text{ (int-t)} \quad \frac{}{\Psi \vdash n : \Psi(n)} \text{ (address-t)} \quad \frac{}{\Psi \vdash n : \rho} \text{ (\rho-t)}$$

$$\boxed{\Delta; \Psi \vdash^Z v : t}$$

$$\frac{\Psi \vdash n : \tau \quad \Delta \vdash E = n}{\Delta; \Psi \vdash^Z c n : \langle c, \tau, E \rangle} \text{ (val-t)}$$

$$\frac{\Delta \vdash E : \kappa_{\text{int}}}{\Delta; \Psi \vdash^c c n : \langle c, \tau, E \rangle} \text{ (val-zap-c-t)} \quad \frac{\Delta \vdash E : \kappa_{\text{int}} \quad c' = B \text{ or } c' = G}{\Delta; \Psi \vdash^{CF} c n : \langle c', \tau, E \rangle} \text{ (val-zap-CF-t)}$$

$$\boxed{\Delta \vdash t \leq t'}$$

$$\frac{\Delta \vdash E_1 = E_2}{\Delta \vdash \langle c, \tau, E_1 \rangle \leq \langle c, \tau, E_2 \rangle} \text{ (subtp-reflex)} \quad \frac{\Delta \vdash E_1 = E_2}{\Delta \vdash \langle c, \tau, E_1 \rangle \leq \langle c, \text{int}, E_2 \rangle} \text{ (subtp-int)}$$

$$\boxed{\Delta \vdash \Gamma \leq \Gamma'}$$

$$\frac{\forall r. \Gamma_1(r) \leq \Gamma_2(r)}{\Delta \vdash \Gamma_1 \leq \Gamma_2} \text{ (\Gamma-subtp)}$$

Figure 4.8: Value Typing Judgment and Subtyping Judgment.



at compile time. It only takes on other values at run time for the purposes of the proof of type preservation.

The main value typing judgment depends upon an auxiliary judgment with the form  $\Psi \vdash n : \tau$ . This auxiliary judgment allows integer  $n$  to be given either a basic *int* type, a stage description type  $\rho$ , or a code type  $\Psi(n)$ . If  $E$  is equal to  $n$  and  $\Psi \vdash n : \tau$ , then  $c n$  can always be given the type  $\langle c, \tau, E \rangle$ . However, if the zap tag  $Z$  is a color  $c$ , then all values  $c n$  can also be typed using any basic type and any well-formed expression — such a general rule reflects the fact that we can make no guarantees about such values. When the zap tag is *CF*, then *any* green and blue value can be given *any* type, including giving green values blue types and vice versa. In other words, as mentioned earlier, when there has been a control-flow fault, we can make no assumptions about either green or blue values.

**Value Subtyping.** There is also a subtyping relationship  $\Delta \vdash t \leq t'$  shown in Figure 4.8. This judgment allows type  $\langle c, \tau, E \rangle$  to be subtype of  $\langle c, int, E' \rangle$  whenever  $\Delta \vdash E = E'$ . Register File subtyping is a basic extension of value subtyping that requires every register in the first register file type to be a subtype of the corresponding register in the second. The subtyping judgment is used to type check control flow transfers.

### 4.3.2 Instruction Typing

Figure 4.9 presents the instruction typing judgment, which has the form  $\Delta; \Psi; \Gamma \vdash i : \Gamma'$ . As before,  $\Delta$  contains free expression variables and  $\Psi$  types heap addresses.  $\Gamma$  acts as the precondition for the instruction, mapping registers to their corresponding types prior to execution of the instruction.  $\Gamma'$  acts as the postcondition for the instruction, mapping registers to types guaranteed after execution of the instruction.

$$\boxed{\Delta; \Psi; \Gamma \vdash i : \Gamma'}$$

$$\frac{r_d \neq r_i}{\Delta; \Psi; \Gamma \vdash \text{movi } r_d \ c \ n : \Gamma[r_d \mapsto \langle c, \text{int}, n \rangle]} \text{ (movi-t)}$$

$$\frac{r_d \neq r_i \quad \Gamma(r_a) = \langle c, \text{int}, E_a \rangle \quad \Gamma(r_b) = \langle c, \text{int}, E_b \rangle}{\Delta; \Psi; \Gamma \vdash \text{sub } r_d \ r_a \ r_b : \Gamma[r_d \mapsto \langle c, \text{int}, E_a - E_b \rangle]} \text{ (sub-t)}$$

$$\frac{\Gamma(r_i) = \langle c_i, \text{ok}, E_i \rangle \quad \Gamma(r_t) = \langle B, \forall[\Delta_t](\Gamma_t, \sigma_t), E_t \rangle}{\Delta; \Psi; \Gamma \vdash \text{intend } r_t : \Gamma[r_i \mapsto \langle B, \text{go}, E_t \rangle]} \text{ (intend-t)}$$

$$\frac{\Gamma(r_i) = \langle B, \text{go}, E_i \rangle \quad \Gamma(r_t) = \langle B, \forall[\Delta_t](\Gamma_t, \sigma_t), E_t \rangle}{\Gamma(r_z) = \langle B, \text{int}, E_z \rangle \quad t' = \langle B, \text{goz}, E_z ? E_i : E_t \rangle} \text{ (intendz-t)}$$

$$\Delta; \Psi; \Gamma; \vdash \text{intendz } r_z \ r_t : \Gamma[r_i \mapsto t']$$

Figure 4.9: Instruction Typing Judgment.

The simplest instruction to type check is the `movi  $r_d \ c \ n$`  instruction. It merely updates the type of  $r_d$  to be  $\langle c, \text{int}, n \rangle$ . The subtraction instruction `sub  $r_d \ r_a \ r_b$`  requires that the values being subtracted are integers. Notice it also requires the integer arguments have the same color as the result – this restriction prevents faults in values with one color from influencing another. These two instructions place no restrictions on the type of  $r_i$ , so they can occur during any stage of a block.

The unconditional intention instruction `intend  $r_t$`  requires that  $r_i$  has basic type *ok*. This restriction guarantees any new intend will occur after the checking code has been completed. Intentions are part of the blue computation, so the register that is used to set the intention must contain a blue value with code type. The type of  $r_i$  is updated to reflect the new static expression and the new stage *go*.

The conditional intention instruction `intendz  $r_z \ r_t$`  is similar, although it must occur after an unconditional intention. In other words, to set the intention for a conditional

branch, first use `intend` to set  $r_i$  to contain the address of the fall through block, and then conditionally set it to contain the branch target. The resulting type of  $r_i$  has basic type  $goz$  and a conditional expression guarded by the expression  $E_z$  describing  $r_z$ . If  $E_z$  is nonzero, then  $r_i$  will be described by  $E_i$ , which describes the fall through branch. Otherwise, it is described by  $E_t$ , which describes the branch target.

Despite the fact that `recovernz` is syntactically an instruction, it is type-checked using the block typing judgment because it affects the set of free expression variables.

### 4.3.3 Block Typing

Figure 4.10 presents the block typing judgment, which has the form  $\Delta; \Psi; \Gamma; \sigma; E_i; \tau \text{ opt} \vdash b$ . In addition to  $\Delta$ ,  $\Psi$ , and  $\Gamma$ , the block typing judgment is parameterized by a sequence  $\sigma$ , an expression  $E_i$ , and a type option  $\tau \text{ opt}$ .

The sequence  $\sigma$  contains a list of expressions that describe the locations in the current history  $h$ . While typing a block at location  $\ell$ ,  $\sigma$  has the form  $x_h \circ \ell$  meaning that the program has already visited some unknown sequence of locations ( $x_h$ ) leading up to this point and that the label of the current block is  $\ell$ . The expression  $E_i$  describes the intended target when the transfer occurred to the current label  $\ell$ . If control flow correctly transferred to  $\ell$ , then  $\Delta \vdash E_i = \ell$ . The option type  $\tau \text{ opt}$  contains the type of the label  $\ell + 1$  if such a label exists. It is used when a branch falls through to the subsequent block to determine the type of that block.

The first rule, *sequence-t*, is used when the first instruction in a block is one of the basic instructions described previously. Descriptions of the other rules follow.

$$\boxed{\Delta; \Psi; \Gamma; \sigma; E_i; \tau \text{ opt} \vdash b}$$

$$\frac{\Delta; \Psi; \Gamma \vdash i : \Gamma' \quad \Delta; \Psi; \Gamma'; \sigma; E_i; \tau \text{ opt} \vdash b}{\Delta; \Psi; \Gamma; \sigma; E_i; \tau \text{ opt} \vdash i; b} \text{ (sequence-}t\text{)}$$

$$\frac{\begin{array}{l}
\Gamma(r_i) = \langle O, \text{check}, x_i \rangle \\
\Gamma(r_z) = \langle O, \text{int}, E_z \rangle \\
\Delta, x : \kappa_{\text{int}} \vdash E_z = E_\ell - x_i \\
\Delta \vdash \Gamma/r_i/r_z \text{ wf} \quad \Delta \vdash \sigma \text{ wf} \quad \Delta \vdash E_\ell : \kappa_{\text{int}} \\
\Gamma' = \Gamma[r_z \mapsto \langle O, \text{int}, 0 \rangle][r_i \mapsto \langle B, \text{ok}, E_\ell \rangle] \\
\Delta; \Psi; \Gamma'; \sigma \circ E_\ell; E_\ell; \tau \text{ opt} \vdash b
\end{array}}{(\Delta, x : \kappa_{\text{int}}); \Psi; \Gamma; \sigma \circ E_\ell; x_i; \tau \text{ opt} \vdash \text{recovernz } r_z; b} \text{ (recovernz-}t\text{)}$$

$$\frac{\begin{array}{l}
\Gamma(r_z) = \langle O, \text{int}, E_z \rangle \quad \cdot \vdash E_z = E_i - E_\ell \\
\Gamma(r_i) = \langle O, \text{check}, E_i \rangle \quad \cdot \vdash E_i = E_\ell \\
;\Psi; \Gamma[r_i \mapsto \langle O, \text{ok}, E_i \rangle]; \sigma \circ E_\ell; E_i; \tau \text{ opt} \vdash b
\end{array}}{;\Psi; \Gamma; \sigma \circ E_\ell; E_i; \tau \text{ opt} \vdash \text{recovernz } r_z; b} \text{ (recovernz-}eq\text{-}t\text{)}$$

$$\frac{\begin{array}{l}
\Gamma(r_z) = \langle O, \text{int}, E_z \rangle \quad \cdot \vdash E_z = E_i - E_\ell \\
\Gamma(r_i) = \langle O, \text{check}, E_i \rangle \quad \cdot \vdash E_i \neq E_\ell \\
;\Psi; \Gamma; \sigma \circ E_\ell; E_i; \tau \text{ opt} \vdash \text{recovernz } r_z; b
\end{array}}{\text{;}\Psi; \Gamma; \sigma \circ E_\ell; E_i; \tau \text{ opt} \vdash \text{recovernz } r_z; b} \text{ (recovernz-}neq\text{-}t\text{)}$$

$$\frac{\begin{array}{l}
\Gamma(r_i) = \langle B, \text{goz}, E'_z ? E'_f : E'_t \rangle \quad \Delta \vdash E'_f = E_\ell + 1 \\
\Gamma(r_z) = \langle G, \text{int}, E_z \rangle \quad \Delta \vdash E_z = E'_z \\
\Gamma(r_t) = \langle G, \forall[\Delta_t](\Gamma_t, \sigma_t), E_t \rangle \quad \Delta \vdash E_t = E'_t \\
\Delta \vdash \Gamma[r_i \mapsto \langle O, \text{check}, E'_t \rangle] \leq S_t(\Gamma_t) \quad \exists S_t . \Delta \vdash S_t : \Delta_t \\
\Delta \vdash \Gamma[r_i \mapsto \langle O, \text{check}, E'_f \rangle] \leq S_f(\Gamma_f) \quad \exists S_f . \Delta \vdash S_f : \Delta_f \\
\Delta \vdash \sigma \circ E_\ell \circ E'_t = S_t(\sigma_t) \quad \Delta \vdash \sigma \circ E_\ell \circ E'_f = S_f(\sigma_f)
\end{array}}{\Delta; \Psi; \Gamma; \sigma \circ E_\ell; E_i; \forall[\Delta_f](\Gamma_f, \sigma_f) \vdash \text{brz } r_z r_t} \text{ (brz-}t\text{)}$$

$$\frac{\begin{array}{l}
\Gamma(r_i) = \langle B, \text{go}, E'_t \rangle \quad \Delta \vdash E_t = E'_t \\
\Delta \vdash \Gamma[r_i \mapsto \langle O, \text{check}, E'_t \rangle] \leq S_t(\Gamma_t) \quad \Delta \vdash \sigma \circ E_\ell \circ E_t = S_t(\sigma_t)
\end{array}}{\Delta; \Psi; \Gamma; \sigma \circ E_\ell; E_i; t \vdash \text{jmp } r_t} \text{ (jmp-}t\text{)}$$

Figure 4.10: Block Typing Judgment.

**Recovery.** There are three distinct rules for checking  $recovernz\ r_z$ . All of them require the instruction to occur in the first stage of the block when  $r_i$  contains an orange value with basic type *check*. The operand register  $r_z$  compares this value to the current label.

The first rule *recovernz-t* applies when  $r_i$  is described by variable  $x_i$ . This is the rule used by a programmer to check correctness of their program at compile time. Control only proceeds past this point in the block if  $x_i$  is equal to the expression  $E_\ell$ , which describes the current location, so the remainder of the block is typed by substituting  $E_\ell$  for  $x_i$ . The types of  $r_i$  and  $r_z$  are updated to reflect the deletion of  $x_i$ . Judgment  $\Delta \vdash \Gamma / r_i / r_z$  wf and  $\Delta \vdash \sigma$  wf hold when all variables used in registers other than  $r_i$  and  $r_z$  as well as the expressions in  $\sigma$  are all contained in  $\Delta$ . Because none of these pieces of state contain  $x_i$ , they do not need to be modified.

The other two rules *recovernz-eq-t* and *recovernz-neq-t* are needed to carry out the proof of type preservation (particularly the substitution lemma), but would never be used to type check programs prior to execution. In these situations,  $x_i$  has already been replaced with a closed expression  $E_i$  that describes the intention register at block entry. Here, it is evident that either  $\cdot \vdash E_i = E_\ell$  or not, so there is one typing rule for each situation. The rule *recovernz-neq-t* does not place any requirements on the remainder of the block since control does not proceed past this point.

**Control Flow Transfers.** In order to verify unexpected transfers from the end of one block to the beginning of any other, code blocks must have the same basic precondition. To be specific, each block must expect that the intention register  $r_i$  contains an orange value with basic type *check* that is described by a variable  $x_i$ . This variable does not occur anywhere else in the function precondition. This condition entails every target block can accept any orange value in  $r_i$ .

The rule  $jmp-t$  requires that  $r_i$  has type  $\langle B, go, E'_i \rangle$  specifying that the intention must already have been set before the jump. Also, the current jump target has a code type and is described by an expression  $E_t$  that is equal to  $E'_t$ . This enforces that in the absence of faults, the duplicate target is equal to the target.

The target label precondition contains a set of expression variables  $\Delta_t$  and requires a register file described by  $\Gamma_t$  and a history described by  $\sigma_t$ . There is some substitution  $S_t$  for the variables in  $\Delta_t$  so that the current register file type and sequence are subtypes of those required by the target.

The  $jmp r_t$  and  $brz r_z r_t$  instructions recolor the blue intention register to be orange when control is transferred to a new block. At first, this seems to contradict the rule that faults to a value of one color should never corrupt values of other colors. However, because the target block does not place any restrictions on the expression describing  $r_i$ , the variable  $x_i$  that describes the value can be instantiated with the value itself. Because of this, a blue value that is not trusted can become a trusted orange value during a control flow transfer, continuing to leave only the blue values untrusted.

The rule  $brz-t$  is similar, but adds in the conditional register  $r_z$  and specifies both the fall through and the branch cases.

### 4.3.4 Machine State Typing

**Code Memory Typing.** The judgment  $\vdash C : \Psi$  describes the invariants for code memory. As described previously, all blocks must have the same basic precondition. The register  $r_i$  is described by the type  $\langle O, check, x_i \rangle$ . The other registers are colored either blue or green, and their static expressions do not contain the variable  $x_i$ . If a label  $\ell$  has type  $\forall[\Delta](\Gamma, x_h \circ \ell)$ , then code at that label must be well-typed given  $\Psi$ ,  $\Delta$ ,  $\Gamma$ ,  $x_h \circ \ell$ , the intention expression  $x_i$ , and the fall-through label type  $\Psi(\ell + 1)$ .

$$\boxed{\vdash C : \Psi}$$

$$\begin{array}{l}
\forall \ell \in \text{Dom}(C) \cup \text{Dom}(\Psi) . \\
\Psi(\ell) = \forall[\Delta](\Gamma, x_h \circ \ell) \\
\Delta = \Delta', x_i : \kappa_{int}, x_h : \kappa_{hist} \\
\Gamma = \Gamma', r_i \mapsto \langle O, check, x_i \rangle \\
\forall r' \in \text{Dom}(\Gamma') . \Gamma'(r') \neq \langle O, b', E' \rangle \\
\Delta' \vdash \Psi(\ell + 1) \text{ wf} \quad \Delta' \vdash \Gamma' \text{ wf} \\
\Delta; \Psi; \Gamma; x_h \circ \ell; x_i; \Psi(\ell + 1) \vdash C(\ell) \\
\hline
\vdash C : \Psi \quad (C-t)
\end{array}$$

$$\boxed{\Psi \vdash R : \Gamma}$$

$$\frac{\forall r. ; \Psi \vdash^Z R(r) : \Gamma(r)}{\Psi \vdash^Z R : \Gamma} \quad (R-t)$$

$$\boxed{\vdash h : \sigma}$$

$$\frac{}{\vdash () : \varepsilon} \quad (h\text{-empty-}t) \quad \frac{\vdash h : \sigma \quad \cdot \vdash E = n}{\vdash (h, n) : \sigma \circ E} \quad (h\text{-app-}t)$$

$$\boxed{\vdash^Z (C, h, R, b)}$$

$$\begin{array}{l}
\vdash C : \Psi \\
\Psi \vdash^Z R : \Gamma \\
\vdash (h, \ell) : \sigma \\
(Z = CF) ? (\cdot \vdash E_i \neq \ell) : (\cdot \vdash E_i = \ell) \\
\Gamma(r_i) \neq \langle O, check, E_i \rangle \implies \cdot \vdash E_i = \ell \\
; \Psi; \Gamma; \sigma; E_i; \Psi(\ell + 1) \vdash b \\
\hline
\vdash^Z (C, (h, \ell), R, b) \quad (\Sigma-t)
\end{array}$$

Figure 4.11: Machine State Typing.

**Register File Typing.** The judgment  $\Psi \vdash^Z R : \Gamma$  states that register file  $R$  has type  $\Gamma$  under zap tag  $Z$  given heap typing  $\Psi$ . It holds when each register in  $R$  has the corresponding type in  $\Gamma$  under  $Z$ . Again, values with colors that are affected by  $Z$  are not trusted to have their given types.

**History Typing.** A history  $h$  is described by sequence  $\sigma$  when each location is equal to the corresponding expression.

**Machine State Typing.** A machine state  $\Sigma$  is well-typed under zap tag  $Z$  when each of its elements is well-typed, and two additional invariants hold. (1) If  $Z$  is  $CF$  then the current location  $\ell$  is not equal to the intended location  $E_i$ . Otherwise, if  $Z$  is not  $CF$ , then these two are equal. (2) If the current block  $b$  has proceeded past the checking stage, then it must be the case that  $\ell$  is equal to  $E_i$ . These two invariants together imply it is not possible for code past the checking stage of a block to be well-typed under the  $CF$  zap tag. Consequently, a proof of type preservation will imply that any control-flow error will be caught in the checking stage of the next block.

## 4.4 Formal Results

We have proven a number of properties of our type system including variants of the standard Progress, Preservation and Type Safety theorems. The statement of the Progress Theorem is standard. The statement of the Preservation Lemma relies on the subtyping relationship between zap tags. As the program executes, it remains well-typed, although the zap tag used to type machine states may elevate to a supertype. In other words, the program may start out well-typed under the empty zap tag, become well-typed under a colored zap tag after a fault, and then become well typed under  $CF$  if a control-flow fault



occurs. Our most important result is a Fault Tolerance theorem, which we sketch briefly below. More detailed proof sketches appear in Appendix F.

In order to explain the theorem, we require a couple of additional concepts which should be familiar. First, we say a machine state  $\Sigma$  is well-formed (written  $\vdash^Z \Sigma$ ) when all code and state are well-typed relative to the zap tag  $Z$ . Second, we say a faulty machine state  $\Sigma_f$  simulates a fault-free state  $\Sigma$  under color  $c$  (written  $\Sigma_f \text{ sim}^c \Sigma$ ) whenever the two states are identical except for values colored  $c$ . In other words, values colored  $c$  may be completely different from one another, but otherwise the two states are identical.

The judgment  $\Sigma \xRightarrow[k]{h} \mathcal{F}$  states that machine state  $\Sigma$  executes through a sequence of blocks  $h$  to reach state  $\mathcal{F}$  while incurring  $k$  faulty transitions. So if  $\Sigma = (C, h_1, R, b)$ , then  $\mathcal{F}$  is either  $(C, (h_1, h), R', b')$ ,  $\text{hwerror}(h_1, h)$ , or  $\text{recover}(h_1, h)$ .

We say a program is fault-tolerant if any execution of the program with a single fault behaves in one of four possible ways with regards to the original, non-faulty computation: (1) The faulty computation visits the same sequence of blocks as the original, and the final faulty state simulates the original result state under some color  $c$ . (2) The faulty computation attempts to transfer control to an invalid address outside the domain of code memory and triggers a hardware fault. Prior to the occurrence of the hardware fault, the faulty computation visited the same blocks as the original computation. (3) A fault affecting the intention register or checking code cause the faulty computation to detect a fault in software and jump to recovery code. (4) The faulty computation veers off course to a block that does not match the corresponding block in the original computation. In this case, the checking code in the invalid block catches the error and transfers control to the recovery code.

**Theorem 9 (Fault Tolerance)**

If  $\vdash \Sigma$  and  $\Sigma \Longrightarrow_0^h \Sigma'$  then at least one of the following cases applies and all derivations  $\Sigma \Longrightarrow_1^{h_f} \mathcal{F}$  where  $\text{length}(h_f) \leq \text{length}(h)$  fit one of these cases:

1.  $\Sigma \Longrightarrow_1^h \Sigma'_f$  and  $\exists c . \Sigma'_f \text{ sim}^c \Sigma'$
2.  $\Sigma \Longrightarrow_1^{h_f} \text{hwerror}(h', h_f)$  and  $h_f$  is a prefix of  $h$
3.  $\Sigma \Longrightarrow_1^{h_f} \text{recover}(h', h_f)$  and  $h_f$  is a prefix of  $h$
4.  $\Sigma \Longrightarrow_1^{h_f} \text{recover}(h', h_f)$  and  $h_f = (h_1, l')$  and  $h = (h_1, l, h_2)$

## 4.5 Translation

In order to show that TAL<sub>CF</sub> is sufficiently expressive to be of interest, we define a simple language of while loops and show how to compile statements in this language into well-typed TAL<sub>CF</sub> programs. We only sketch the results here, but more details are available in Appendix G.

The while loop language statements consist of simple assignment, subtraction, if statements, while loops, and sequences of statements. As all the variables in this language contain integers, the well-formedness judgment  $V \vdash s$  simply enforces that all variables  $v$  in  $s$  exist in the variable context  $V$ .

$$\begin{aligned}
 s \quad ::= \quad & v := n \mid v_d := v_a - v_b \\
 & \mid \text{if } 0 \ v_z \text{ then } s_1 \text{ else } s_2 \mid \text{while } v_z \neq 0 \text{ do } s \\
 & \mid s_1; s_2
 \end{aligned}$$

To translate a statement  $s$  as a stand-alone program, it is translated with 1 as the starting label. The result of the translation is a code memory  $C$ , the instructions  $\vec{i}$  from the body of the final block of the program, and the label  $\ell$  of that last block. Because there is no halt instruction in  $TAL_{CF}$ , code is added to the last block in the translation to create an infinite loop at label  $\ell_{halt}$ . The function  $InitRegFile(V)$  creates an initial register file that maps each register used to translate  $V$  to 0. The assembly language program corresponding to  $s$  is the  $TAL_{CF}$  state consisting of the generated code memory, a history with only the first location, an initial register file, and code to jump to the first label in code memory. If the original statement is well-formed, then the translation is well-typed.

**Theorem 10 (Translation)**

If  $\llbracket V \vdash s \rrbracket (\cdot, \cdot, \cdot, 1) = (\cdot, C, \vec{i}, \ell)$  then

$\vdash (C', 0, InitRegFile(V), intendjmp\ 1)$

where  $C' = [\ell \mapsto \text{check } \ell; \vec{i}; intendjmp\ \ell_{halt}] [\ell_{halt} \mapsto \text{check } \ell_{halt}; intendjmp\ \ell_{halt}]$

## 4.6 Summary and Future Work

This chapter presents  $TAL_{CF}$ , a typed assembly language design for reasoning about software-only transient fault solutions. Well-typed  $TAL_{CF}$  programs will always detect a single transient fault that causes control to incorrectly transfer between basic blocks before control has reached more than one incorrect block. A translation from a simple language of while programs into  $TAL_{CF}$  gives evidence that it is sufficiently expressive to serve as a target of compilation.

We acknowledge that the fault model used in this chapter is simplistic. By assuming hardware support to catch control transfers into the middle of blocks, we avoid dealing

with many interesting and likely situations. This assumption is required because stating intentions involves resetting  $r_i$ , so an incorrect transfer into a block before the `intend  $r_i$`  instruction may not be caught.

A sequence of existing work discussed in Section 5.2.2 handles increasing classes of erroneous transfers. By ensuring that the intention register is a function of the entire control-flow path, not just the current block, they can detect most jumps into the middle of blocks. The classification scheme of values and reliability properties from this chapter does not transfer directly to these more complex solutions, but we believe we can develop a similar classification to capture the necessary invariants. In doing so, the Fault Tolerance Theorem becomes more difficult to state and prove due to the increase of possible scenarios a single fault may cause. (For example, a fault may cause control to transfer from the middle of one block to the middle of a second block. This second block may transfer control to a third block before the error is finally detected.) In essence, the current work and proof strategy are an important building block for reasoning about more complex solutions.

# Chapter 5

## Related Work and Conclusion

This dissertation applies techniques for reasoning about low-level software to the problem of verifying fault tolerance techniques. As such, the related work draws from diverse topics in programming languages and computer architecture.

### 5.1 PCC and TAL

Proof-carrying Code (PCC) [44, 43, 3] is a method for guaranteeing properties of low-level code. In PCC, the compiler generates a safety proof along with the code binary, and the end user can automatically verify the safety proof before executing the binary.

The most popular way of implementing PCC is using a *type-preserving* compiler. Like ordinary compilers, type-preserving compilers are organized as a series of translations between intermediate languages. Unlike in ordinary compilers, each of these intermediate languages has a corresponding type system and may be type checked. In the final compilation phase, a fully type-preserving compiler will output an assembly or native code binary as well as sufficient type information for a type checker to reconstruct

a typing derivation for the assembly code. The first such language was called TAL (Typed Assembly Language) [40]. Over the years, many variants have been developed [66, 62, 17, 16, 14].

Working at such a low level leads to a number of technical challenges not seen in high-level type systems. For example, creating a new object actually consists of allocating the memory and then initializing the contents. TAL [40] handled this by using *initialization flags* to track which locations had been initialized. In addition, representing class and object encodings requires new technology [15, 13]. The stack is a concept that is not part of most high-level languages, and it turns out to be quite tricky. Unlike memory allocated in the heap, the type of the value stored in a given stack slot changes during the execution of the program. As the stack grows and shrinks, a location may hold a code pointer in one stack frame and an integer in some later stack frame. Assembly-level type systems need to track these changing types carefully to avoid unsafe behavior. Things get even more complicated when allocating program data on the stack because stack locations may be aliased. At this point, it becomes very difficult to develop a type system that remains sound and yet is expressive enough to handle advanced stack allocation optimizations with. A sequence of existing work [38, 30, 50] has investigated this issue. Other current areas of interest include handling interrupts [22], interfacing with the garbage collector [27, 33], and dealing with concurrency [21].

Recently, the Bartok Compiler [14] was modified to be fully type-preserving. Bartok is an optimizing, object-oriented compiler of approximately 200,000 lines. Not only does the type-preserving version generate well-typed assembly code and support almost all the original compiler's optimizations, the additional compilation overhead is only 42%<sup>1</sup> and

---

<sup>1</sup>Personal communication with J. Chen. August 2008. Compilation overhead published in June 2008 [14] was 83% but has since been reduced by performance tuning.

the generated code is only 2.3% slower. In other words, type-preserving compilers are reaching the point of being practical for use in real-world situations.

However, one assumption made by all of these languages is that the hardware will faithfully execute programs. No guarantees are provided about the behavior of programs in the presence of transient faults.

## 5.2 Research in Transient Fault Tolerance

Researchers in the computer architecture and compiler communities have been working on solutions to transient faults for a long time. This section is not meant to be a full summary of the research in this area. Instead, it will provide a high-level overview of some of the differing types of solutions and focus on those techniques most similar to the ideas used in this dissertation.

### 5.2.1 Hardware-Based Solutions

Hardware-based techniques add hardware to the processor to detect transient faults. These include fine-grained techniques such as error-correcting codes and parity bits, as well as more course-grained solutions that duplicate entire structures within the processor.

For example, the 777 primary flight computer triplicates all hardware resources [75]. The Compaq Nonstop Himalaya [74] has two identical processors that run in lockstep executing their own copies of the same program. The external pins of the processors are compared on every cycle to detect faults, but recovery is left to the software. The IBM S/390 [65] duplicates units within the processor and compares their outputs on each cycle, invoking hardware recovery if necessary. It also uses parity and ECC to protect the L1 and L2 caches.

One of the main drawbacks of executing in lockstep is that it may reduce performance because processor resources are partitioned statically. Simultaneous Multithreading (SMT) [70] is a technique that allows two independent threads to run at the same time while issuing instructions to different functional units. Researchers exploited SMT to develop fault detection techniques that allow two copies of the same program to run simultaneously while still allowing dynamic scheduling of the hardware resources [57]. Others have extended this idea to handle recovery [71]. Similarly, two copies of a program can be run independently on separate cores of a processor [41], and also deal with recovery [23].

Watchdog coprocessors [35] involve a second simple processor which monitors the inputs and outputs of the main processor in order to detect faults. Despite requiring less hardware than full replication, this technique can detect a large number of faults by looking only at control flow and memory access behavior.

## 5.2.2 Software-Based Solutions

Just as the hardware-based solutions work by adding hardware redundancy, the software-based solutions add redundant instructions. The techniques can be divided into two main camps: protecting data and protecting control flow.

### Protecting Data

An very general method for protecting data is to duplicate the entire computation and insert extra instructions to compare the two copies and check for consistency [48]. SWIFT [60] is one solution that uses this technique. Comparisons are inserted before each store instruction to ensure that the two copies of the address and the two copies of the value agree. Once the comparison succeeds, a single store instruction commits the change to



memory. However, there is a *window of vulnerability* between the comparison and single store instruction where faults may go undetected and cause memory corruption. It appears impossible to address this vulnerability without relying on additional, nonstandard hardware support.

### Protecting Control Flow

One particular challenge for software-based solutions involves the detection of control-flow faults. Our language  $TAL_{CF}$  in Chapter 4 keeps an *approximate program counter* and use an *intentions register* across control flow transfers to ensure that the location reached was the same as the intended location. Of course, as we discussed, this solution has the immediate drawback of being unable to detect faults into or out of the middle of code blocks.

A sequence of existing work [48, 60, 10] uses a more sophisticated version of this idea to handle increasing classes of erroneous transfers. By ensuring that the intentions register is a function of the entire control-flow path, not just the current block, they can detect most jumps into the middle of blocks.

CFCSS [48] uses an approximate program counter and assigns each block a static signature. After a control flow transfer, the approximate pc is updated to contain the xor of its old value and a precalculated value containing the xor of the current signature and the predecessors' signature. (A correct transfer from block  $A$  to block  $B$  will result in a new approximate program counter of  $A \otimes (A \otimes B)$ , which is equal to  $B$ .) Because the predecessor is clearly defined, there is no need for this technique to duplicate both computations. However, since the true block and the false block of a branch have the same predecessor block, this technique cannot detect a fault that causes control to take the incorrect arm of a branch. The authors present experimental results for CFCSS show

that only 3.1% of the injected faults were undetected and resulted in incorrect outputs. The authors include high-level arguments about the ability to catch certain classes of faults but do not clearly specify what classes of faults were resulting in that undetected 3.1%.

SWIFT [60] uses a slightly different technique that adds a *transition register* to specify the intended transition. Before each control-flow transfer, the current block and the intended target block are xored together and the result is put in a designated transition register. At the beginning of each block, the transition register is xored with the approximate program counter to give the new approximation. Again, without faults, this should result in  $A \otimes (A \otimes B) = B$ . However, since  $A \otimes B$  is computed at run time, this is more precise than the static method used by CFCSS. Because the intended targets cannot be determined statically, SWIFT must duplicate all the values used in control flow, but this is already being done to protect data. Again, the authors present experimental results show an impressive reduction in uncaught errors, and though the work includes itemized lists of classes of faults that are believed to be caught and not caught, there is no attempt to prove their fault tolerance scheme is correct.

One class of faults that SWIFT cannot always detect is faults from the end of a block to an instruction within the body of that same block. Borin et al. [10] introduce two additional techniques to handle this. The first technique zeros out the approximate program counter within the middle of block to detect situations where the program has incorrectly entered the middle of the block. The second technique, RCF (Region Based Control Flow) Checking, protects the branches that must be used to set the transition register.  $TAL_{CF}$  assumed the existence of a conditional move instruction to avoid unprotected branches for setting the intentions register. The additional branch is given its own signature, allowing faults there to be detected. The authors give a clear diagram-based

explanation of the types of faults handled by each solution. In addition, they give a strong argument that the high-level algorithm they provide is correct. However, that does not guarantee that the generated code correctly implements the algorithm.

Despite the improvements made in this area, there are still situations (such as jumping back two instructions within a block) that cannot be handled.

### **Dealing with Performance**

Even though many software-based solutions more than double the number of instructions, the performance overhead is generally around 40% [60, 10]. Superscaler processors are rarely able to make full use of their resources, so by adding an additional independent computation, the processor is able to achieve higher throughput.

Even so, 40% overhead is still substantial. There may be times when this overhead is acceptable, and times when it is not. But luckily, being software-based, these kinds of solutions can be applied *selectively*. Reis [58] shows that often protecting only a small number of functions can obtain nearly as much benefit as protecting the entire program, while incurring a much smaller overhead.

### **Adding Recovery**

The software techniques discussed so far have focused on error *detection*, but being able to *recover* from errors is also extremely important.

Two copies of the program are enough to detect errors, but three copies can be used to detect *and* recover by using majority voting. SWIFT-R [59] is a version of SWIFT extended to triple redundancy. Though it can recover from faults, the performance overhead is higher, and there are still windows of vulnerability. The authors also introduce a different recovery technique that only requires two copies of a computation. The difference is that

in the second copy, all values are multiplied by three. Single-bit errors can be detected by checking that the original value multiplied by three is the same as the duplicate. If an error is detected, the program can determine which one is corrupt by using the modulo operation on the duplicate value. Because a single-bit error will manifest as a difference of  $2^n$  for some  $n$ , the duplicate is only evenly divisible by three when it has not been corrupted.

### 5.2.3 Hybrid Solutions

There are clear trade-offs between hardware-based solutions and software-based solutions. *Hybrid* solutions use additional hardware and additional software to try to exploit the best of each type.

CRAFT [61], the system  $TAL_{FT}$  is modeled after, is a hybrid version of SWIFT. It duplicates the computation in software, but uses an additional hardware buffer to close the window of vulnerability around store instructions.

## 5.3 Formal Reasoning about Faults

There are a number of efforts that look at formalizing algorithms for fault tolerance, both for systems affected by transient faults and concurrent systems. A couple of the software-based systems for control flow that we have already discussed [10, 48] include formal reasoning about the high-level algorithms involved.

Di Vito and Bulter verify the system design for a fault-tolerant system for flight control [72] that votes using results from replicated processors. The system is specified at a high-level using an abstract model of real-time computation, and the proofs are completed using algebraic techniques and verified by an automated proof assistant.

Lamport and Merz [32] give a proof of correctness for a solution to the Byzantine generals problem in which the states and behaviors of the concurrent system are specified using logic, and the desired properties are checked using automated tools. However, as they themselves admit, this is very far from being able to reason at the level of executable code.

Argus [36] is a hardware-based solution with the goal of providing low-cost, low-power reliability using simple cores. It dynamically checks four invariants about control flow, computations, dataflow, and memory correctness. The authors model the system as an abstract von Neumann processor, similar to the operational semantics used in this dissertation. They prove that all executions that satisfy the four invariants will output the same values as a correct execution by using induction on the length of the execution. However, the actual algorithms for checking the invariants require hardware support and are specified at a high level.

Walker et al. [73] were the first to use a type system to reason about fault-tolerant code. They defined a typed  $\lambda$ -calculus called  $\lambda_{\text{zap}}$  in which computation are triplicated and compared at control flow transfers and before output instructions using a special atomic *vote* operation. Later work by Elsmann [20] showed how to break the *vote* operation down into a sequence of conditional statements while still being able to prove correctness. There are two fundamental differences between these results and the research presented here. First, by working at such a high level of abstraction, they are able to avoid dealing with the low-level hardware structures. Second, their type systems are much weaker than those we have presented in previous chapters. For example, types in  $\lambda_{\text{zap}}$  consist of a color and a base type, but there are no static expressions. This means that all the formal results depend on the correctness of the translation into  $\lambda_{\text{zap}}$  which adds the duplication. As we have discussed, trusting the compiler can be dangerous.

Closely related to our work on  $TAL_{CF}$  is work by Abadi et al. [1, 2] on secure control flow. Their system CFI (Control Flow Integrity) prevents attackers from manipulating the control flow of a program. The goal is to guarantee that an executing program will always follow a path through a statically determined control flow graph. However, an essential difference is the fault model. CFI assumes that an attacker can make arbitrary changes to data memory and most registers, but that three designated registers are protected. So while CFI must worry about arbitrary amounts of data corruption but can rely on three values to be correct,  $TAL_{CF}$  only has to worry about a single fault, but can never trust a single value. In addition, the formal results for CFI state that the attacked program will follow *some* path through the existing control graph, while  $TAL_{CF}$  guarantees that a faulty program will follow the *same* path as the original program.

## 5.4 Concluding Remarks

Transient faults are caused when energetic particles strike a processor and deposit charge. They may corrupt executing programs and have caused crashes at major companies. Processor technology trends, including decreasing feature size, decreasing voltages, and increasing clock rates, will result in future processors that are even more susceptible. Existing solutions that involve additional software often have a simple intuition but are implemented as just one of the many phases in a large, optimizing compiler. At the same time, recent work has shown that full-scale, type-preserving compilers are a practical way for reasoning about the behavior of low-level code.

This dissertation develops the first set of techniques for verifying fault-tolerant executable code and gives solid evidence that typed assembly languages are a promising direction for continued research. Specifically, we introduce three new typed assembly

languages.  $TAL_{FT}$  is used to reason about a hybrid fault detection scheme. Well-typed programs are guaranteed to detect a single fault before that fault causes a change in observable behavior.  $ETAL_{FT}$ , an extended version of  $TAL_{FT}$ , is expressive enough to be used as the target language in a compiler and also allows us to investigate the interactions between fault tolerance and other current research areas in typed assembly languages. In a slightly different direction,  $TAL_{CF}$  takes a first step towards formalizing the guarantees that can be provided by a purely software-based solution.

The general approach used by all three languages includes some modifications to the standard typed assembly language methodology and can be thought of in five stages. First, in order to reason about machine execution, we define an operational semantics that includes specific rules to model faults. One example is the rule *reg-zap* that can occur at any point during execution and corrupt the contents of a random register in the register file. Second, it is necessary to formally define what “fault tolerance” means in each specific situation. For  $TAL_{FT}$  and  $ETAL_{FT}$ , we assume that memory is externally visible and so faults may never corrupt memory writes. Since  $TAL_{FT}$  focuses only on control flow, it only allows faulty programs to visit one incorrect block before detecting the error. Next, the type system is designed to combine standard type safety invariants with invariants about the fault tolerance solution, particularly specifying the relationship between different copies of a computation. Once we have these definitions, we show that the type system is sound – well-typed programs will always be fault-tolerant when executed. The main theorem for each language relates an execution and a faulty version of the same execution and shows that they will be indistinguishable to an observer. The final step is to show a language can be used as the target language of a compiler.

In addition to being the first to verify low-level code, these languages make three important contributions. The type systems combine basic typing information, computation

colors, and a language of static expressions to enforce that the two different computations perform equivalent tasks. This allows us to verify low-level code without relying on how it was created. By capturing different fault models, we show that our methodology is generally applicable.  $TAL_{FT}$  and  $ETAL_{FT}$  focus on a hybrid solution, where additional hardware is used to address the vulnerability between comparing two values and acting on the result of the comparison.  $TAL_{FT}$  is purely software-based, which leads to a number of challenges, including the need to reason about situations where control ends up in an unexpected location. Finally, the translations into typed assembly, particular that of  $ETAL_{FT}$ , show that these languages can be used as the final result of a realistic, optimizing compiler.



# Appendix A

## TAL<sub>FT</sub> Proof Details

This appendix gives expanded details on the formal results in Section 2.3, including sketches of the corresponding lemmas and proofs. Working notes for the complete proofs appear in the companion technical report [49].

Section A.1 discusses useful lemmas used throughout the proofs. Section A.2 proves the standard notions of type safety. Section A.3 defines a multistep operation semantics and some associated lemmas, including the No False Positives Corollary. Section A.4 contains the Fault Tolerance Theorem and its associated lemmas. Each lemma and theorem is preceded by a brief English explanation of its role in the proof, and followed by details on how the proof was constructed.

### A.1 Lemmas

The proofs of the theorems in the remainder of this section rely on the lemmas discussed below.

### A.1.1 Properties of Static Expressions

When an expression is closed, applying substitutions to that expression results in a syntactically equivalent expression.

#### Lemma 11 (Substituting Closed Expressions)

1. If  $\cdot \vdash E : \kappa$  then  $\forall S. S(E) = E$

*Proof* By induction on the structure of  $\cdot \vdash E : \kappa$ . ■

Even though  $\llbracket E \rrbracket$  is a partial function, it is always defined over well-kinded closed expressions.

#### Lemma 12 (Expression Denotation)

1. If  $\cdot \vdash E : \kappa_{int}$  then  $\exists n. \llbracket E \rrbracket = n$ .
2. If  $\cdot \vdash E : \kappa_{mem}$  then  $\exists M. \llbracket E \rrbracket = M$ .

*Proof* By induction on the structure of  $\cdot \vdash E : \kappa$ . ■

The equality of expressions is transitive.

#### Lemma 13 (Expression Equality Transitivity)

If  $\Delta \vdash E_1 = E_2$  and  $\Delta \vdash E_2 = E_3$  then  $\Delta \vdash E_1 = E_3$ .

*Proof* By inspection of the definition of  $\Delta \vdash E_1 = E_2$ . ■

Substituting an expression of kind  $\kappa$  for a free variable of type  $\kappa$  preserves typing.

#### Lemma 14 (Substitution Lemma)

1. If  $\Delta, x : \kappa \vdash E' : \kappa'$  and  $\Delta \vdash E : \kappa$  then  $\Delta \vdash E'[E/x] : \kappa'$ .
2. If  $\Delta, x : \kappa \vdash E_1 = E_2$  and  $\Delta \vdash E : \kappa$  then  $\Delta \vdash E_1[E/x] = E_2[E/x]$ .

3. If  $\Delta, x : \kappa \vdash E_1 \neq E_2$  and  $\Delta \vdash E : \kappa$  then  $\Delta \vdash E_1[E/x] \neq E_2[E/x]$ .
4. If  $\Psi; \Delta, x : \kappa \vdash^Z v : t$  and  $\Delta \vdash E : \kappa$  then  $\Psi; \Delta \vdash^Z v : t[E/x]$ .
5. If  $\Psi; (\Delta, x : \kappa; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash^Z ir \Rightarrow RT$  and  $\Delta \vdash E : \kappa$   
then  $\Psi; (\Delta; \Gamma[E/x]; \overline{(E_d, E_s)}[E/x]; E_m[E/x]) \vdash^Z ir \Rightarrow RT[E/x]$ .
6. If  $\cdot \vdash S : \Delta$  and  $\Psi; \Delta \vdash^Z v : t$  then  $\Psi; \cdot \vdash^Z v : S(t)$ .
7. If  $\cdot \vdash S : \Delta$  and  $\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \vdash^Z ir \Rightarrow (\Delta; \Gamma'; \overline{(E'_d, E'_s)}; E'_m)$   
then  $\Psi; (\cdot; S(\Gamma); S(\overline{(E_d, E_s)}); S(E_m)) \vdash^Z ir \Rightarrow (\cdot; S(\Gamma'); S(\overline{(E'_d, E'_s)}); S(E'_m))$ .

*Proof* Parts 1, 4, and 5 – by induction on the respective typing derivation. Parts 2 and 3 – by inspection of the equality judgment definition. Parts 6 and 7 – by induction on the size of  $\Delta$ , using parts 4 and 5 respectively. ■

## A.1.2 Properties of Well-Typed Values

Well-typed values are described by static expressions denoting integers and not memories.

### Lemma 15 (Value Kinding)

1. If  $\Psi; \Delta \vdash^Z v : \langle c, b, E \rangle$  then  $\Delta \vdash E : \kappa_{int}$ .
2. If  $\Psi; \Delta \vdash^Z v : (E' = 0) \Rightarrow \langle c, b, E \rangle$  then  $\Delta \vdash E : \kappa_{int}$ .

*Proof* By case analysis on the value typing judgment. ■

The type of a value gives us information about the shape of the value.

### Lemma 16 (Canonical Forms)

If  $\Psi; \cdot \vdash^Z c n : t$  and  $Dom(\Psi) = Dom(C) \cup Dom(M)$  and  $\Psi \vdash M : E_m$  and  $\Psi \vdash C$ , then

1. If  $t = \langle c, b, E \rangle$  or  $t = (E' = 0) \Rightarrow \langle c, b, E \rangle$  and  $c = Z$  then no particular properties of  $n$  are known.
2. If  $t = \langle c, \text{int}, E \rangle$  and  $c \neq Z$  then  $\cdot \vdash E = n$ .
3. If  $t = \langle c, \Theta \rightarrow \text{void}, E \rangle$  and  $c \neq Z$  then  $\Psi(n) = \Theta \rightarrow \text{void}$  and  $n \in \text{Dom}(C)$  and  $\cdot \vdash E = n$  and  $n \neq 0$ .
4. If  $t = \langle c, b \text{ ref}, E \rangle$  and  $c \neq Z$  then  $\Psi(n) = b \text{ ref}$  and  $n \in \text{Dom}(M)$  and  $\cdot \vdash E = n$ .
5. If  $t = (E' = 0) \Rightarrow t'$  and  $c \neq Z$  and  $\cdot \vdash E' = 0$  then  $n \neq 0$ .
6. If  $t = (E' = 0) \Rightarrow t'$  and  $c \neq Z$  and  $\cdot \vdash E' \neq 0$  then  $n = 0$ .

*Proof* By induction on the structure of  $\Psi; \cdot \vdash^Z c n : t$ . ■

If a value has a type, and this type has a supertype, then the value also has the supertype.

**Lemma 17 (Subtyping)**

If  $\Psi; \Delta \vdash^Z v : t$  and  $\Delta \vdash t \leq t'$  then  $\Psi; \Delta \vdash^Z v : t'$ .

*Proof* By induction on the derivation of  $\Psi; \Delta \vdash^Z v : t$ . ■

If a value is well-typed under the empty zap tag, then that value is well-typed under all colors.

**Lemma 18 (Color Weakening)**

If  $\Psi; \cdot \vdash v : t$  then  $\forall c. \Psi; \cdot \vdash^c v : t$ .

*Proof* By induction on the value typing judgment. ■

Color weakening extends to register files.

**Lemma 19 (Register File Color Weakening)**

If  $\Psi \vdash R : \Gamma$  then  $\forall c. \Psi \vdash^c R : \Gamma$ .

*Proof* By inversion of the register file typing rule  $R-t$  and the Color Weakening Lemma. ■

### A.1.3 Properties of Well-typed Memories

Well-typed programs with no faults only load values from valid locations.

**Lemma 20 (Well-typed Domain)**

If  $\vdash (R, C, M, Q, ld_G r_d, r_s)$  then  $R_{val}(r_d) \in Dom(M)$ .

*Proof* By inversion of  $\vdash (R, C, M, Q, ir)$ , inversion of the  $ld_G-t$  typing rule, and the Canonical Forms lemma. ■

When looking up a value in a memory with an update, if the updated location is not the requested location, then this is equivalent to looking up the location in the memory without the update.

**Lemma 21 (Irrelevant Update)**

If  $E = sel (upd E_m E_d E_s) E_n$  and  $\cdot \vdash E_d \neq E_n$  then  $E = sel E_m E_n$ .

*Proof* By inspection of the denotation of  $sel$  and  $upd$ . ■

### A.1.4 Properties of Well-typed Queues

The length of the queue is the same as the length of the sequence that describes it. When the zap tag is not green, then each item in the queue is described by the corresponding expression. In addition, the first element of each pair has type  $b \text{ ref}$  and its value has type  $b$ .

#### Lemma 22 (Queue)

1. If  $\Psi \vdash^Z Q : \overline{(E_d, E_s)}$  then  $\text{length}(Q) = \text{length}(\overline{(E_d, E_s)})$ .
2. If  $\Psi \vdash^Z \overline{(n_1, n_2)} : \overline{(E_d, E_s)}$  and  $Z \neq G$  then for all  $k$  from 1 to  $\text{length}(\overline{(n_1, n_2)})$ ,  $\cdot \vdash E_{dk} = n_{1k}$  and  $\cdot \vdash E_{sk} = n_{2k}$  and there is some base type  $b$  such that  $\Psi \vdash n_{1k} : b \text{ ref}$  and  $\Psi \vdash n_{2k} : b$ .

*Proof* By induction on the structure of  $\Psi \vdash^Z Q : \overline{(E_d, E_s)}$ . ■

If the *find* function returns no match on an address  $n_1$ , and the queue is described by the sequence of address-value pairs  $\overline{(E_d, E_s)}$ , then none of the address expressions are equal to  $n_1$ .

#### Lemma 23 (Find)

If  $\text{find}(Q, n_1) = ()$  and  $\Psi \vdash Q : \overline{(E_d, E_s)}$  then for  $k$  from 1 to  $\text{length}(Q)$ ,  $\cdot \vdash E_{dk} \neq n_1$ .

*Proof* By definition of the *find* function and the queue typing judgment. ■

If a queue is well-typed under the empty zap tag, then it is also well-typed under any colored zap tag.

#### Lemma 24 (Queue Color Weakening)

If  $\Psi \vdash Q : \overline{(E_d, E_s)}$  then  $\forall c. \Psi \vdash^c Q : \overline{(E_d, E_s)}$ .

*Proof* By induction on the queue typing judgment. ■

## A.2 Type Safety

Progress states that well-typed states can take a step. In particular, a machine state that is well-typed under the empty zap tag can take a non-faulty step to another ordinary, non-faulty machine state. A machine state that is well-typed under a zap tag of color  $c$  can take a step, but the result of that step may either be another ordinary machine state or the *fault* state.

### Theorem 25 (Progress)

1. If  $\vdash \Sigma$  then  $\Sigma \longrightarrow_0^s \Sigma'$  and  $\Sigma' \neq \text{fault}$ .
2. If  $\vdash^c \Sigma$  then  $\Sigma \longrightarrow_0^s \Sigma$ .

*Proof* By case analysis on the instruction  $ir$  in state  $\Sigma$ .

Part 1 uses inversion on the typing rules to determine that the preconditions hold for the appropriate operational rule. For example, Figure A.1 shows the case for  $st_B$ . By inverting various typing rules, we gather enough information to conclude that the last pair in the queue is equal to the contents of the two registers. Notice that step 9 inverts the *val-t* rule. This can only be done because we know that  $Z = \cdot$ , and so the *val-zap-t* rule cannot apply instead. (And similarly, step 12 inverts rule *Q-t*.) Once we have these equalities, we can apply the operational rule  $st_B\text{-mem}$ . Other cases are similar, although the cases for  $ld_G$  and  $bz_B$  subdivide further based on the result of the *find* function and the value in the branch register.

Part 2 is simpler than part 1. Instead of using typing rules to gather as much information, we just further subdivided based on properties needed to apply the rules.

These subcases take either the normal operation rule, or the corresponding failure rule as necessary. Some inversion may be done on the typing rules to show that the case does not get stuck. Again, Figure A.1 shows the case for  $st_B$ . From the typing information, we know that the Q is not empty. Then either the last pair in the queue is equal to the registers (rule  $st_B$  applies) or it is not (rule  $st_B$ -fail applies). ■

According to Preservation, if a machine state is well-typed under a zap tag  $Z$ , and it takes a non-faulty step to another machine state, then that resulting state will also be well-typed under  $Z$ . Additionally, if a state is well-typed under the empty zap tag, and it takes a faulty step, then there is some color  $c$  such that the resulting state is well-typed under  $c$ .

**Theorem 26 (Preservation)**

1. If  $\vdash^Z \Sigma$  and  $\Sigma \longrightarrow_0^s \Sigma'$  and  $\Sigma' \neq \text{fault}$  then  $\vdash^Z \Sigma'$ .
2. If  $\vdash \Sigma$  and  $\Sigma \longrightarrow_1^s \Sigma'$  then  $\exists c. \vdash^c \Sigma'$ .

*Proof* By case analysis of the structure of the derivation  $\Sigma \longrightarrow_k^s \Sigma'$ .

Part 1 only applies to cases where  $\Sigma'$  is not *fault*. We use inversion to take apart the judgment for  $\vdash^Z \Sigma$ , modify it as necessary, and build the judgment  $\vdash^Z \Sigma'$ . Each case is subdivided based on the actual value of the zap tag. In subcases where the zap tag is not the same as the color of a value, we can invert rule  $val$ - $t$  on the typing of a value and use that information to construct a typing derivation for  $\Sigma'$ . In cases where the zap tag has the same color as the value, modified values in  $\Sigma'$  can be trivially typed using rules  $val$ -zap- $t$  and  $Q$ -zap- $t$ . Figures A.2 and A.3 show an example case for rule  $st_B$ - $mem$ .

Part 2 chooses  $c$  to be the color of the zapped value. It uses rules  $val$ -zap- $t$ ,  $val$ -zap- $cond$  or  $Q$ -zap- $t$  to show that the zapped value can be typed under zap tag  $c$ . The remainder of the state can be typed as before using the Color Weakening Lemmas 18, 19,



**Part 1 Example Case:  $st_B$** 

- a1.  $\vdash (R, C, M, Q, st_B r_d, r_s)$  [ assumption ]
1.  $\Psi \vdash C$  [ Inversion of  $\Sigma$ -t, a1 ]  
 2.  $\forall c. C(R_{val}(pc_c)) = st_B r_d, r_s$  [ Inversion of  $\Sigma$ -t, a1 ]  
 3.  $\Psi; (\Delta; \Gamma; (\overline{E_d, E_s}), (E'_d, E'_s); E_m) \vdash st_B r_d r_s$  [ Inversion of C-t, 1, 2, inspection of  $st_B$ -t ]  
 $\Rightarrow (\Delta; \Gamma++; (\overline{E_d, E_s}); upd E_m E'_d E'_s)$   
 4.  $\exists S. \cdot \vdash S : \Delta$  [ Inversion of  $\Sigma$ -t, a1 ]  
 5.  $\Psi; (\cdot; S(\Gamma); S(\overline{E_d, E_s}), (E'_d, E'_s); S(E_m)) \vdash st_B r_d r_s$  [ Lemma 14 (Substitution), 4, 3 ]  
 $\Rightarrow (\cdot; S(S(\Gamma)++); S(\overline{E_d, E_s}); upd S(E_m E'_d E'_s))$
6.  $\Psi \vdash R : S(\Gamma)$  [ Inversion of  $\Sigma$ -t, a1 ]  
 7.  $S(\Gamma)(r_d) = \langle B, b \text{ ref}, E''_d \rangle$  [ Inversion of  $st_B$ -t, 5 ]  
 $S(\Gamma)(r_s) = \langle B, b, E''_s \rangle$   
 8.  $\Psi; \cdot \vdash R(r_d) : \langle B, b \text{ ref}, E''_d \rangle$  [ Inversion of R-t, 6, 7 ]  
 $\Psi; \cdot \vdash R(r_s) : \langle B, b, E''_s \rangle$   
 9.  $\cdot \vdash R_{val}(r_d) = E''_d$  [ Inversion of val-t, 8 ]  
 $\cdot \vdash R_{val}(r_s) = E''_s$
10.  $\cdot \vdash S(E'_d) = E''_d$  [ Inversion of  $st_B$ -t, 5 ]  
 $\cdot \vdash S(E'_s) = E''_s$
11.  $\Psi \vdash Q : S(\overline{E_d, E_s}), (E'_d, E'_s)$  [ Inversion of  $\Sigma$ -t, a1 ]  
 12.  $Q = (\overline{(n, n')}, (n_l, n'_l))$  [ Lemma 22 (Queue), 11 ]  
 $\cdot \vdash S(E'_d) = n_l$  and  $\cdot \vdash S(E'_s) = n'_l$
13.  $R_{val}(r_d) = n_l$  and  $R_{val}(r_s) = n'_l$  [ Lemma 13 (Exp Eq Transitivity), 9, 10, 12 ]  
 14.  $(R, C, M, (\overline{(n, n')}, (n_l, n'_l)), st_B r_d, r_s) \xrightarrow{0}_{(n_l, n'_l)} (R++, C, M[n_l \mapsto n'_l], \overline{(n, n')}, \cdot)$  [  $st_B$ -mem, 14 ]
- Case complete.

**Part 2 Example Case:  $st_B$** 

- a1.  $\vdash^c (R, C, M, Q, st_B r_d, r_s)$  [ assumption ]
1.  $\Psi \vdash^c Q : S(\overline{E_d, E_s}), (E'_d, E'_s)$  [ Inversion of  $\Sigma$ -t, a1 ]  
 2.  $Q = (\overline{(n, n')}, (n_l, n'_l))$  [ Lemma 22 (Queue), 1 ]
- Case on whether  $R_{val}(r_d) \stackrel{?}{=} n_l$  and  $R_{val}(r_s) \stackrel{?}{=} n'_l$
- subcase a:**  $r_d$  and  $r_s$  are the same as the last pair in the queue  
 a4a.  $R_{val}(r_d) = n_l$  and  $R_{val}(r_s) = n'_l$  [ subcase assumption ]  
 3a.  $(R, C, M, Q, st_B r_d, r_s) \xrightarrow{0}_{(n_l, n'_l)} (R++, C, M[n_l \mapsto n'_l], \overline{(n, n')}, \cdot)$  [  $st_B$ -mem, a4a ]  
 Subcase complete.
- subcase b:** either  $r_s$  or  $r_d$  or the last pair in the queue has been corrupted  
 a4b.  $R_{val}(r_d) \neq n_l$  or  $R_{val}(r_s) \neq n'_l$  [ subcase assumption ]  
 3b.  $(R, C, M, Q, st_B r_d, r_s) \xrightarrow{0} \text{fault}$  [  $st_B$ -mem-fail, a4b ]  
 Subcase complete.
- Case complete.

Figure A.1: Example Cases of Theorem 25 (Progress).

and 24. Figure A.4 shows an example case for rule *reg-zap*. ■

The Progress and Preservation Theorems define the usual notion of type safety.

### A.3 Multistep Transitions

In order to prove properties of our type system, we extend our single-step transition  $\Sigma_1 \xrightarrow{k}^s \Sigma_2$  from Section 2.1 to a sequence of  $n$  transitions containing exactly  $k$  faults  $\Sigma_1 \xrightarrow{n}^s \Sigma_2$ , where  $n$  is greater than or equal to zero, and  $k$  is still either 0 or 1.

$$\frac{}{\Sigma \xrightarrow{0}^s \Sigma} \text{ (multi-base)} \quad \frac{\Sigma \xrightarrow{k_1}^{s_1} \Sigma'' \quad \Sigma'' \xrightarrow{k_2}^{(n-1)s_2} \Sigma'}{\Sigma \xrightarrow{n}^{(s_1, s_2)} \Sigma'} \text{ (multi-compose)}$$

#### A.3.1 No False Positives

By combining part one of Progress with part one of Preservation, we get the following important corollary: The hardware never claims to have detected a fault when no fault has occurred during execution of a well-typed program.

##### Corollary 27 (No False Positives)

If  $\vdash \Sigma$  then  $\forall n. \Sigma \xrightarrow{n}^s \Sigma'$  and  $\vdash \Sigma'$ .

*Proof* By Progress Part 1, Preservation Part 1, and induction on the derivation of  $\Sigma \xrightarrow{n}^s \Sigma'$ . ■

**Example Case:  $st_B$ -mem**

$$\frac{R_{val}(r_d) = n_l \quad R_{val}(r_s) = n'_l}{(R, C, M, (\overline{(n, n')}, (n_l, n'_l)), st_B \ r_d, r_s) \longrightarrow_0^{(n_l, n'_l)} (R++, C, M[n_l \mapsto n'_l], \overline{(n, n')}, \cdot)} \quad (st_B\text{-mem})$$

- |     |   |  |
|-----|---|--|
| a1. | $\vdash^Z (R, C, M, Q, st_B \ r_d, r_s)$  | [ assumption ]                         |
| p1. | $R_{val}(r_d) = n_l$  | [ premise ]                            |
| p2. | $R_{val}(r_s) = n'_l$   | [ premise ]                            |
| 1.  | $Dom(\Psi) = Dom(C) \cup Dom(M)$  | [ Inversion of $\Sigma$ -t, a1 ]       |
| 2.  | $Z \neq G \implies Dom(\overline{(n, n')}, (n_l, n'_l)) \subseteq Dom(M)$   | [ Inversion of $\Sigma$ -t, a1 ]       |
| 3.  | $\Psi \vdash C$   | [ Inversion of $\Sigma$ -t, a1 ]       |
| 4.  | $\forall c \neq Z. C(R_{val}(pc_c)) = st_B \ r_d, r_s$  | [ Inversion of $\Sigma$ -t, a1 ]       |
| 5.  | $\forall c \neq Z. \Psi(R_{val}(pc_c)) = (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \rightarrow void$   | [ Inversion of $\Sigma$ -t, a1 ]       |
| 6.  | $\exists S. \cdot \vdash S : \Delta$  | [ Inversion of $\Sigma$ -t, a1 ]       |
| 7.  | $\Psi \vdash M : S(E_m)$  | [ Inversion of $\Sigma$ -t, a1 ]       |
| 8.  | $\Psi \vdash^Z (\overline{(n, n')}, (n_l, n'_l)) : S(\overline{(E_d, E_s)}, (E'_d, E'_s))$  | [ Inversion of $\Sigma$ -t, a1 ]       |
| 9.  | $\Psi \vdash^Z R : S(\Gamma)$   | [ Inversion of $\Sigma$ -t, a1 ]       |
| 10. | $\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}, (E'_d, E'_s); E_m) \vdash st_B \ r_d \ r_s \Rightarrow \Theta'$  | [ Inversion of C-t, 3, 4 ]             |
| 11. | $\Theta' = (\Delta; \Gamma++; \overline{(E_d, E_s)}; upd \ E_m \ E'_d \ E'_s)$  | [ 10, inspection of $st_B$ -t ]        |
| 12. | $\forall c \neq Z. C(R_{val}(pc_c) + 1) = \Theta' \mapsto void$   | [ Inversion of C-t, 3, 4, 10, 11 ]     |
| 5'. | $\forall c \neq Z. C(R_{++val}(pc_c)) = \Theta' \mapsto void$   | [ 12, def of R++ ]                     |
| 9'. | $\Psi \vdash^Z R++ : S(\Gamma++)$   | [ 9, def of R++ and $\Gamma++$ ]       |
| 13. | $\Psi; (\cdot; S(\Gamma); S(\overline{(E_d, E_s)}, (E'_d, E'_s)); S(E_m)) \vdash st_B \ r_d \ r_s$<br>$\Rightarrow (\cdot; S(S(\Gamma)++; S(\overline{(E_d, E_s)}); upd \ S(E_m) \ E'_d \ E'_s))$ | [ Lemma 14 (Substitution), 10, 11, 6 ] |
| 14. | $S(\Gamma)(r_d) = \langle B, b \ ref, E''_d \rangle$  | [ Inversion of $st_B$ -t, 5 ]          |
| 15. | $S(\Gamma)(r_s) = \langle B, b, E''_s \rangle$  | [ Inversion of $st_B$ -t, 5 ]          |
| 16. | $\cdot \vdash S(E'_d) = E''_d$  | [ Inversion of $st_B$ -t, 5 ]          |
| 17. | $\cdot \vdash S(E'_s) = E''_s$  | [ Inversion of $st_B$ -t, 5 ]          |

Case on Z

**Subcase a1:  $Z = G$** 

- |      |   |  |
|------|---|--|
| a2a. | $Z = G$   | [ subcase assumption ]                                   |
| 18a. | $\Psi \vdash^Z \overline{(n, n')} : S(\overline{(E_d, E_s)})$ | [ Repeated Inversion of $Q$ -zap-t, a2a, 8, $Q$ -zap-t ] |

**Subcase b1:  $Z = \cdot$  or  $Z = B$** 

- |      |   |  |
|------|---|--|
| a2b. | $Z = \cdot$ or $Z = B$  | [ subcase assumption ]                           |
| 18b. | $\Psi \vdash^Z \overline{(n, n')} : S(\overline{(E_d, E_s)})$ | [ Repeated Inversion of $Q$ -t, a2a, 8, $Q$ -t ] |

Merge subcases a1 and b1:

- |     |   |             |
|-----|---|-------------|
| 8'. | $\Psi \vdash^Z \overline{(n, n')} : S(\overline{(E_d, E_s)})$ | [ 18a/18b ] |
| 2'. | $Z \neq G \implies Dom(\overline{(n, n')}) \subseteq Dom(M)$  | [ 2 ]       |

Continued in Figure A.3...

Figure A.2: Example Case from Theorem 26 (Preservation) Part 1.

**Example Case: st<sub>B</sub>-mem (...continued from Figure A.2)**Case on  $Z$ **Subcase a2:**  $Z = B$  (the queue is correct)a3a.  $Z = B$ 

[ subcase assumption ]

19a.  $\cdot \vdash S(E'_d) = n_l$ [ Inversion of  $Q$ -t, a3a, 8 ]20a.  $\Psi \vdash n_l : b \text{ ref}$ [ Inversion of  $Q$ -t, a3a, 8 ]21a.  $\Psi; \cdot \vdash^B n'_l : \langle B, b \text{ ref}, E''_d \rangle$ [  $val$ -t, 19a, 20a ]22a.  $n_l \in \text{Dom}(M)$ 

[ Lemma 16 (Canonical Forms), 7, 3, 21a ]

23a.  $\cdot \vdash S(E'_s) = n'_l$ [ Inversion of  $Q$ -t, a3a, 8 ]24a.  $\Psi \vdash n'_l : b$ [ Inversion of  $Q$ -t, a3a, 8 ]25a.  $\llbracket S(E'_d) \rrbracket = n_l$  and  $\llbracket S(E'_s) \rrbracket = n'_l$ [ Inversion of  $E$ -eq, 19a, 23a ]**Subcase b2:**  $Z = \cdot$  or  $Z = G$  ( $r_d$  and  $r_s$  are correct)a3b.  $Z = \cdot$  or  $Z = G$ 

[ subcase assumption ]

19b.  $\Psi; \cdot \vdash^Z R_{val}(r_d) : \langle B, b \text{ ref}, E''_d \rangle$ [ Inversion of  $R$ -t, 9, 14 ]20b.  $\Psi; \cdot \vdash^Z n_l : \langle B, b \text{ ref}, S(E'_d) \rangle$ 

[ 19b, p1, Lemma 13 (Exp Eq Transitivity), 16 ]

21b.  $n_l \in \text{Dom}(M)$ 

[ Lemma 16 (Canonical Forms), 7, 3, 20b ]

22b.  $\Psi \vdash n_l : b \text{ ref}$ [ Inversion of  $val$ -t, a3b, 20b ]23b.  $\Psi; \cdot \vdash^Z R_{val}(r_s) : \langle B, b, E''_s \rangle$ [ Inversion of  $R$ -t, 9, 15 ]24b.  $\Psi; \cdot \vdash^Z n'_l : \langle B, b, S(E'_s) \rangle$ 

[ 23b, p2, Lemma 13 (Exp Eq Transitivity), 17 ]

25b.  $\Psi \vdash n'_l : b$ [ Inversion of  $val$ -t, a3b, 24b ]26b.  $\cdot \vdash S(E'_d) = n_l$  and  $\cdot \vdash S(E'_s) = n'_l$ [ Inversion of  $val$ -t, a3b, 20b, 24b ]27b.  $\llbracket S(E'_d) \rrbracket = n_l$  and  $\llbracket S(E'_s) \rrbracket = n'_l$ [ Inversion of  $E$ -eq, 26b ]

Merge subcases a2 and b2:

30.  $\forall \ell \in \text{Dom}(M). \Psi \vdash \ell : b \text{ ref} \wedge \Psi \vdash M(\ell) : b$ [ Inversion of  $M$ -t, 7 ]31.  $\forall \ell \in \text{Dom}(M[n_l \mapsto n'_l]). \Psi \vdash \ell : b \text{ ref} \wedge \Psi \vdash M(\ell) : b$ 

[ 30, 20a/22b, 24a/25b ]

32.  $\llbracket S(E_m) \rrbracket = M$ [ Inversion of  $M$ -t, 7 ]33.  $\llbracket \text{upd } S(E_m) S(E'_d) S(E'_s) \rrbracket = \llbracket S(E_m) \rrbracket [ \llbracket S(E'_d) \rrbracket \mapsto \llbracket S(E'_s) \rrbracket ]$ [ def of  $\llbracket E \rrbracket$  ]34.  $\llbracket \text{upd } S(E_m) S(E'_d) S(E'_s) \rrbracket = M[n_l \mapsto n'_l]$ 

[ 33, 32, 25a/27b ]

35.  $\cdot \vdash S(E_m) : \kappa_{mem}$ [ Inversion of  $M$ -t, 7 ]36.  $\cdot \vdash S(E'_d) : \kappa_{int}$  and  $\cdot \vdash S(E'_s) : \kappa_{int}$ [ Lemma 15 (Value Kinding), Inversion of  $R$ -t, 9, 14, 15 ]37.  $\cdot \vdash \text{upd } S(E_m) S(E'_d) S(E'_s) : \kappa_{mem}$ [  $E$ -upd-t, 35, 36 ]7'.  $\Psi \vdash M[n_l \mapsto n'_l] : S(E_m E'_d E'_s)$ [  $M$ -t, 31, 34, 37 ]1'.  $\text{Dom}(\Psi) = \text{Dom}(C) \cup \text{Dom}(M[n_l \mapsto n'_l])$ 

[ 1, 22a/21b ]

38.  $\vdash^Z (R++, C, M[n_l \mapsto n'_l], \overline{(n, n')}, \cdot)$ [  $\Sigma$ -t, 1', 2', 3,  $ir = \cdot, 5', 6', 7', 8', 9'$  ]

Figure A.3: Example Case from Theorem 26 (Preservation) Part 1.

**Example Case: reg-zap**

$$\frac{R(a) = c \ n}{(R, C, M, Q, ir) \longrightarrow_1 (R[a \mapsto c \ n'], C, M, Q, ir)} \text{ (reg-zap)}$$

- |     |  |                                  |
|-----|--|----------------------------------|
| a1. | $\vdash (R, C, M, Q, ir)$  | [ assumption ]                   |
| 1.  | $Dom(\Psi) = Dom(C) \cup Dom(M)$   | [ Inversion of $\Sigma$ -t, a1 ] |
| 2.  | $Dom(Q) \subseteq Dom(M)$  | [ Inversion of $\Sigma$ -t, a1 ] |
| 3.  | $\Psi \vdash C$  | [ Inversion of $\Sigma$ -t, a1 ] |
| 4.  | $\forall c. ir \neq \cdot \implies C(R_{val}(pc_c)) = ir$  | [ Inversion of $\Sigma$ -t, a1 ] |
| 5.  | $\forall c. \Psi(R_{val}(pc_c)) = (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \rightarrow void$ | [ Inversion of $\Sigma$ -t, a1 ] |
| 6.  | $\exists S. \cdot \vdash S : \Delta$   | [ Inversion of $\Sigma$ -t, a1 ] |
| 7.  | $\Psi \vdash M : S(E_m)$   | [ Inversion of $\Sigma$ -t, a1 ] |
| 8.  | $\Psi \vdash Q : S(\overline{E_d, E_s})$   | [ Inversion of $\Sigma$ -t, a1 ] |
| 9.  | $\Psi \vdash R : S(\Gamma)$  | [ Inversion of $\Sigma$ -t, a1 ] |

Case on the shape of  $S(\Gamma)(a)$

**Subcase a:**  $S(\Gamma)(a)$  is a triple  $\langle c, b, E \rangle$

- |      |  |   |
|------|--|---|
| a2a. | Let $\langle c, b, E \rangle = S(\Gamma)(a)$                 | [ subcase assumption ]                                    |
| 10a. | $\cdot \vdash E : \kappa_{int}$                              | [ Lemma 15 (Value Kinding), Inversion of $R$ -t, 9, a2a ] |
| 11a. | $\Psi; \cdot \vdash^c c \ n' : \langle c, b, E \rangle$      | [ val-zap-t, 10a ]  |
| 12a. | $\Psi; \cdot \vdash^c R[a \mapsto c \ n'](a) : S(\Gamma)(a)$ | [ 11a, a2a, def of R [ ] ]                                |

**Subcase b:**  $S(\Gamma)(a)$  is a conditional type  $E' = 0 \Rightarrow \langle c, b, E \rangle$

- |      |  |   |
|------|--|---|
| a2b. | Let $(E' = 0 \Rightarrow \langle c, b, E \rangle) = S(\Gamma)(a)$          | [ subcase assumption ]                    |
| 10b. | $\cdot \vdash E : \kappa_{int}$  | [ Lemma 15, Inversion of $R$ -t, 9, a2b ] |
| 11b. | $\Psi; \cdot \vdash^c c \ n' : E' = 0 \Rightarrow \langle c, b, E \rangle$ | [ val-zap-cond, 10b ]                     |
| 12b. | $\Psi; \cdot \vdash^c R[a \mapsto c \ n'](a) : S(\Gamma)(a)$               | [ 11b, a2b, def of R [ ] ]                |

Merge subcases a and b:

- |     |   |   |
|-----|---|---|
| 9'. | $\Psi \vdash^c R[a \mapsto c \ n'] : S(\Gamma)$   | [ $R$ -t, 9, Lemma 19 (Reg File Color Weakening), 12a/12b ] |
| 8'. | $\Psi \vdash^c Q : S(\overline{E_d, E_s})$  | [ Lemma 24 (Queue Color Weakening), 8 ]                     |
| 2'. | $Z \neq G \implies Dom(Q) \subseteq Dom(M)$   | [ 2 ]   |
| 4'. | $\forall c \neq Z. ir \neq \cdot \implies C(R_{val}(pc_c)) = ir$  | [ 4 ]   |
| 5'. | $\forall c \neq Z. \Psi(R_{val}(pc_c)) = (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m) \rightarrow void$ | [ 5 ]   |
| 13. | $\vdash^c (R[a \mapsto c \ n'], C, M, Q, ir)$   | [ $\Sigma$ -t, 1, 2', 3, 4', 5', 6, 7, 8', 9' ]             |

Case complete.

Figure A.4: Example Case of Theorem 26 (Preservation) Part 2.

### A.3.2 Multistep Split and Combine

The following two lemmas about the multistep relation let us take apart and put together different sequences of steps.

If a machine state evaluates in a sequence of steps with no faults to a final state, then this computation can be divided into a sequence of non-faulty steps reaching an intermediate state, and a sequence of non-faulty steps from this intermediate state to the final state.

**Lemma 28 (Multistep Split)**

If  $\Sigma \xrightarrow[n]{}_0^s \Sigma'$  then there exists  $n_1, n_2, \Sigma'', s_1$ , and  $s_2$  such that  $n = n_1 + n_2$  and  $s = (s_1, s_2)$  and  $\Sigma \xrightarrow[n_1]{}_0^{s_1} \Sigma''$  and  $\Sigma'' \xrightarrow[n_2]{}_0^{s_2} \Sigma'$ .

*Proof* By induction on the structure of  $\Sigma \xrightarrow[n]{}_0^s \Sigma'$ . ■

If a machine state evaluates in a sequence of  $n_1$  non-faulty steps to another state, that state faults to a third state, and the third state evaluates in  $n_2$  non-faulty steps to a final state, then the original state can reach the faulty state in a sequence of  $n_1 + 1 + n_2$  steps including exactly one fault step.

**Lemma 29 (Multistep Combine)**

If  $\Sigma \xrightarrow[n_1]{}_0^{s_1} \Sigma'$  and  $\Sigma' \xrightarrow{1} \Sigma'_f$  and  $\Sigma'_f \xrightarrow[n_2]{}_0^{s_2} \Sigma''$  then  $\Sigma \xrightarrow[n']{}_1^{(s_1, s_2)} \Sigma''$  where  $n' = n_1 + 1 + n_2$ .

*Proof* By induction on the structure of  $\Sigma \xrightarrow[n_1]{}_0^{s_1} \Sigma'$ . ■

## A.4 Fault Tolerance

A program is fault tolerant when all the faulty executions of that program *simulate* fault-free executions of the program. More precisely, the sequence of outputs from the faulty

$$\begin{array}{c}
\boxed{v_1 \text{ sim}^Z v_2} \\
\frac{}{C \ n \ \text{sim}^Z \ C \ n} \text{ (sim-val)} \quad \frac{}{C \ n \ \text{sim}^C \ C \ n'} \text{ (sim-val-zap)} \\
\boxed{R \ \text{sim}^Z \ R'} \\
\frac{\forall a. R(a) \ \text{sim}^Z \ R'(a)}{R \ \text{sim}^Z \ R'} \text{ (sim-R)} \\
\boxed{Q \ \text{sim}^Z \ Q'} \\
\frac{}{\cdot \ \text{sim}^Z \ \cdot} \text{ (sim-Q-empty)} \\
\frac{G \ n_1 \ \text{sim}^Z \ G \ n'_1 \quad G \ n_2 \ \text{sim}^Z \ G \ n'_2 \quad Q \ \text{sim}^Z \ Q'}{((n_1, n_2), Q) \ \text{sim}^Z \ ((n'_1, n'_2), Q')} \text{ (sim-Q)} \\
\boxed{\Sigma_1 \ \text{sim}^Z \ \Sigma_2} \\
\frac{R \ \text{sim}^Z \ R' \quad Q \ \text{sim}^Z \ Q'}{(R, C, M, Q, ir) \ \text{sim}^Z \ (R', C, M, Q', ir)} \text{ (sim-}\Sigma\text{)}
\end{array}$$

Figure A.5: Similarity of Machine States.

executions are required either to be identical to the fault-free execution or, in the case the hardware detects the fault, a prefix of the fault-free execution.

### A.4.1 Simulation Relation

In order to reason about pairs of faulty and fault-free executions, we define similarity relations between values, register files, queues and machine states. Each of these relations is defined relative to the zap tag  $Z$ . Intuitively, if  $Z$  is empty, the related objects must be identical. If  $Z$  is a color  $c$ , the objects must be identical modulo values colored  $c$ . In the latter case, values colored  $c$  may be corrupted, and there is no hope they satisfy any particular relation. The formal definitions of these relations are shown in Figure A.5.

### A.4.2 Singlestep Fault Detection

We begin by defining fault detection for a single step in the execution of a program. It essentially says that if we have two similar computations, one with a fault and one without, then either the faulty computation can take a step indistinguishable from that of the non-faulty version, or the faulty computation reaches the fault state.

**Lemma 30 (Singlestep Fault Detection)**

If  $\vdash \Sigma$  and  $\Sigma \text{ sim}^c \Sigma_f$  and  $\Sigma \longrightarrow_0^s \Sigma'$  then  $\Sigma_f \longrightarrow_0^{s_f} \Sigma'_f$  and either

1.  $\Sigma' \text{ sim}^c \Sigma'_f$  and  $s = s_f$ , or
2.  $\Sigma'_f = \text{fault}$  and  $s_f = ()$ .

*Proof* By case analysis of  $\Sigma \longrightarrow_0^s \Sigma'$ . Each case is handled in one of three ways.

**Failure Cases.** Rules where  $\Sigma$  steps to *fault* are not applicable because  $\Sigma$  is well typed under the empty zap tag, and so Progress 1 tells us that  $\Sigma'$  is not *fault*.

**Random Cases.** Rules *ld<sub>B</sub>-rand* and *ld<sub>G</sub>-rand* only apply when loading from an invalid address that is not in the domain of memory. According to Lemma 20 (Well-typed Domain), states that are well-typed under the empty zap tag only load from valid addresses, so these cases can be ruled out.

**Standard Cases.** The remaining rules are all handled in approximately the same way.

Each rule has a handful of premises. These relationships may or may not hold in the faulty computation. It will depend on what exactly has been corrupted. The proof divides into subcases based on these relationships. Some subcases may further divide based on whether  $c$  is  $G$  or  $B$ .



Figure A.6 shows the case for  $st_B\text{-mem}$ . There are two premises relating the two registers to the last pair in the queue. If one of these premises does not hold in the faulty computation, then it will step to  $fault$  using the  $st_B\text{-fail}$  rule with no output.

However, if both equalities hold then the proof further subdivides based on the color  $c$ . If  $c$  is  $B$ , then we know that the faulty and non-faulty queues are equal because they contain green values and simulate each other under color  $B$ . From there, we can determine that the faulty computation also steps using  $st_B$ , the resulting state is similar to  $\Sigma'$ , and the outputs are equal.

If  $c$  is  $G$ , then we use the fact that  $\Sigma$  is well-typed to show that the registers  $r_d$  and  $r_s$  are colored blue, and so both the faulty and non-faulty registers are correct and equal to each other. And since the registers are assumed to be equal to the ends of the queues, we also know that the ends of the queues are equal. From there, we continue as in the previous subcase.

■

### A.4.3 Multistep Fault Detection

The Multistep Fault Detection Lemma extends the Singlestep Fault Detection Lemma for  $n$  steps. If a fault has occurred and the non-faulty computation takes  $n$  steps to a state  $\Sigma'$ , then the faulty computation with either take  $n$  steps to a state that simulates  $\Sigma'$ , or it will terminate in the fault state  $fault$  during this time.

#### Lemma 31 (Multistep Fault Detection)

If  $\vdash \Sigma$  and  $\Sigma \text{ sim}^c \Sigma_f$  and  $\Sigma \xrightarrow{n}^s \Sigma'$  then either

1.  $\Sigma_f \xrightarrow{n}^{s_f} \Sigma'_f$  and  $\Sigma' \text{ sim}^c \Sigma'_f$  and  $s_f = s$ , or

Example Case:  $st_B$ -mem

$$\frac{R_{val}(r_d) = n_l \quad R_{val}(r_s) = n'_f}{(R, C, M, ((\overline{n, n'}), (n_l, n'_f)), st_B \ r_d, r_s) \xrightarrow{0}^{(n_l, n'_f)} (R^{++}, C, M[n_l \mapsto n'_f], (\overline{n, n'}), \cdot)} \quad (st_B\text{-mem})$$

- |     |   |   |
|-----|---|---|
| a1. | $\vdash (R, C, M, ((\overline{n, n'}), (n_l, n'_f)), st_B \ r_d, r_s)$  | [ assumption ]                              |
| a2. | $(R, C, M, ((\overline{n, n'}), (n_l, n'_f)), st_B \ r_d, r_s) \text{ sim}^c \Sigma_f$  | [ assumption ]                              |
| a3. | $(R, C, M, ((\overline{n, n'}), (n_l, n'_f)), st_B \ r_d, r_s) \xrightarrow{0}^{(n_l, n'_f)} (R^{++}, C, M[n_l \mapsto n'_f], (\overline{n, n'}), \cdot)$ | [ assumption ]                              |
| p1. | $R_{val}(r_d) = n_l$  | [ premise ]                                 |
| p2. | $R_{val}(r_s) = n'_f$   | [ premise ]                                 |
| 1.  | $\Sigma_f = (R_f, C, M, Q_f, \text{tr})$  | [ a2, definition of $sim\text{-}\Sigma$ ]   |
| 2.  | $R \text{ sim}^c R_f$   | [ a2, 1, inversion of $sim\text{-}\Sigma$ ] |
| 3.  | $((\overline{n, n'}), (n_l, n'_f)) \text{ sim}^c Q_f$   | [ a2, 1, inversion of $sim\text{-}\Sigma$ ] |
| 4.  | $Q_f = (n_f, n'_f), (n_{lf}, n'_{lf})$  | [ 3, definition of $sim\text{-}Q$ ]         |

Case on whether  $R_{f\text{val}}(r_d) \stackrel{?}{=} n_{lf}$  and  $R_{f\text{val}}(r_s) \stackrel{?}{=} n'_{lf}$  and the color  $c$

**subcase a:** either  $r_s$  or  $r_d$  or the last pair in the queue has been corrupted

- |      |   |                          |
|------|---|--------------------------|
| a4a. | $R_{f\text{val}}(r_d) \neq n_{lf}$ or $R_{f\text{val}}(r_s) \neq n'_{lf}$ | [ subcase assumption ]   |
| 6a.  | $(R_f, C, M, Q_f, st_B \ r_d, r_s) \xrightarrow{0} \text{fault}$          | [ $st_B$ -fail, 4, a4a ] |
| 7a.  | $\Sigma_f \xrightarrow{0}^{sf} \text{fault}$ and $s_f = ()$               | [ 6a ]                   |
- subcase complete.

**subcase b:** blue values are corrupted, but not  $r_s$  or  $r_d$

- |      |  |   |
|------|--|---|
| a4b. | $R_{f\text{val}}(r_d) = n_{lf}$ and $R_{f\text{val}}(r_s) = n'_{lf}$   | [ subcase assumption ]  |
| a5b. | $c = B$  | [ subcase assumption ]  |
| 6b.  | $G \ n_l \ \text{sim}^B \ G \ n_{lf}$  | [ 3, inversion of $sim\text{-}Q$ ]                                      |
| 7b.  | $n_l = n_{lf}$   | [ 6b, inversion of $sim\text{-}val$ , a5b ]                             |
| 8b.  | $G \ n'_f \ \text{sim}^B \ G \ n'_{lf}$  | [ 3, inversion of $sim\text{-}Q$ ]                                      |
| 9b.  | $n'_f = n'_{lf}$   | [ 8b, inversion of $sim\text{-}val$ , a5b ]                             |
| 10b. | $M[n_l \mapsto n'_f] = M[n_{lf} \mapsto n'_{lf}]$  | [ 7b, 9b, a5b ]   |
| 11b. | $R^{++} \ \text{sim}^B \ R_f^{++}$   | [ 2, definition of $R^{++}$ , $sim\text{-}val$ , $sim\text{-}R$ , a5b ] |
| 12b. | $\frac{(n, n') \ \text{sim}^B \ (n_f, n'_f)}{(n, n') \ \text{sim}^B \ (n_f, n'_f)}$  | [ 3, 4, inversion of $sim\text{-}Q$ , a5b ]                             |
| 13b. | $(R^{++}, C, M[n_l \mapsto n'_f], (\overline{n, n'}), \cdot) \ \text{sim}^B \ (R_f^{++}, C, M[n_{lf} \mapsto n'_{lf}], (\overline{n_f, n'_f}), \cdot)$ | [ 10b, 11b, 12b ]   |
| 14b. | $\Sigma_f \xrightarrow{0}^{(n_{lf}, n'_{lf})} (R_f^{++}, C, M[n_{lf} \mapsto n'_{lf}], (\overline{n_f, n'_f}), \cdot)$                                 | [ $st_B$ -mem, a4b, 1, 4 ]  |
| 15b. | $(n_{lf}, n'_{lf}) = (n_l, n'_f)$  | [ 7b, 9b ]  |
| 16b. | $\Sigma_f \xrightarrow{0}^{sf} \Sigma'_f$ and $\Sigma' \ \text{sim}^c \Sigma'_f$ and $s_f = s$   | [ 14b, 13b, 15b ]   |
- subcase complete.

**subcase c:** green values are corrupted, but not the last pair in the queue

- |      |  |   |
|------|--|---|
| a4c. | $R_{f\text{val}}(r_d) = n_{lf}$ and $R_{f\text{val}}(r_s) = n'_{lf}$   | [ subcase assumption ]  |
| a5c. | $c = G$  | [ subcase assumption ]  |
| 6c.  | $R_{col}(r_d) = R_{col}(r_s) = B$  | [ a1, inversion of $\Sigma\text{-}t$ , inversion of $st_B\text{-}t$ ]   |
| 7c.  | $R(r_d) \ \text{sim}^G \ R_f(r_d)$   | [ 2, inversion of $sim\text{-}R$ , a5c ]                                |
| 8c.  | $R(r_s) \ \text{sim}^G \ R_f(r_s)$   | [ 2, inversion of $sim\text{-}R$ , a5c ]                                |
| 9c.  | $R_{val}(r_d) = R_{f\text{val}}(r_d)$  | [ 7c, 6c, inversion of $sim\text{-}val$ ]                               |
| 10c. | $R_{val}(r_s) = R_{f\text{val}}(r_s)$  | [ 8c, 6c, inversion of $sim\text{-}val$ ]                               |
| 11c. | $n_l = n_{lf}$   | [ 9c, a4c, p1 ]   |
| 12c. | $n'_f = n'_{lf}$   | [ 10c, a4c, p2 ]  |
| 13c. | $M[n_l \mapsto n'_f] = M[n_{lf} \mapsto n'_{lf}]$  | [ 11c, 12c ]  |
| 14c. | $R^{++} \ \text{sim}^G \ R_f^{++}$   | [ 2, definition of $R^{++}$ , $sim\text{-}val$ , $sim\text{-}R$ , a5c ] |
| 15c. | $\frac{(n, n') \ \text{sim}^G \ (n_f, n'_f)}{(n, n') \ \text{sim}^G \ (n_f, n'_f)}$  | [ 3, 4, inversion of $sim\text{-}Q$ , a5c ]                             |
| 16c. | $(R^{++}, C, M[n_l \mapsto n'_f], (\overline{n, n'}), \cdot) \ \text{sim}^G \ (R_f^{++}, C, M[n_{lf} \mapsto n'_{lf}], (\overline{n_f, n'_f}), \cdot)$ | [ 13c, 14c, 15c ]   |
| 17c. | $\Sigma_f \xrightarrow{0}^{(n_{lf}, n'_{lf})} (R_f^{++}, C, M[n_{lf} \mapsto n'_{lf}], (\overline{n_f, n'_f}), \cdot)$                                 | [ $st_B$ -mem, a4c ]  |
| 18c. | $(n_{lf}, n'_{lf}) = (n_l, n'_f)$  | [ 11c, 12c ]  |
| 19c. | $\Sigma_f \xrightarrow{0}^{sf} \Sigma'_f$ and $\Sigma' \ \text{sim}^c \Sigma'_f$ and $s_f = s$   | [ 17c, 16c, 18c ]   |
- subcase complete.

case complete.

Figure A.6: Example Case of Lemma 30.

2. Exists  $m \leq n$ .  $\Sigma_f \xrightarrow{m}_0^{s_f} \text{fault}$  and  $s_f$  is a prefix of  $s$ .

*Proof* By induction on the structure of  $\Sigma \xrightarrow{n}_0^s \Sigma'$ . The base case for *multi-base* is immediate.

In the case for *multi-compose*, we know from the premises that  $\Sigma$  takes a single step to some  $\Sigma''$ . Using this and the Singlestep Fault Detection Lemma, we know that  $\Sigma_f$  either takes a single step with no output to *fault* or takes a single step to some state  $\Sigma''_f$  that simulates  $\Sigma''$  while generating equal output.

In the former case, we can immediately prove the second possibility with  $m = 1$  and  $()$  as a prefix of  $s$ .

In the latter case, we call the Induction Hypothesis, which tells us that either  $\Sigma''_f$  takes  $n - 1$  steps to a state that simulates  $\Sigma'$  generating equal output, or it reaches *fault* in no more than  $n - 1$  steps with a prefix of the output. In the first case, where *fault* is never reached, we use *multi-compose* to conclude that  $\Sigma_f$  takes  $n$  steps to  $\Sigma'_f$  and the total output is equal. In the second case, where a *fault* is reached later in execution, we use *multi-compose* to show that  $\Sigma_f$  reaches *fault* in no more than  $n$  steps, and its output is a prefix of the non-faulty output. ■

#### A.4.4 Fault Similarity

The Fault Similarity Lemma states that if a non-faulty machine state takes a single faulty step, then the resulting machine state is similar to the original state under some color  $c$ .

##### Lemma 32 (Fault Similarity)

If  $\Sigma \longrightarrow_1 \Sigma_f$ , then  $\exists c. \Sigma \text{ sim}^c \Sigma'$ .

*Proof* By case analysis on the definition of  $\Sigma \longrightarrow_1 \Sigma'$ . In each case,  $c$  is assigned the color of the value that is zapped. The zapped value is similar to the original value by *sim-val-zap*. The remainder of the state is equal, and is similar using *sim-val*. ■

### A.4.5 Fault Tolerance Theorem

By using the three previous lemmas, we can state and prove the fault tolerance theorem for well-typed programs. Assume that machine state  $\Sigma$  is well-typed under the empty zap tag, and non-faulty execution of  $\Sigma$  for  $n$  steps results in a state  $\Sigma'$  and outputs a sequence of value-address pairs  $s$ . If somewhere during that execution a single fault is encountered, the faulty execution will either run for  $n + 1$  steps or terminate in the fault state during that time. If the faulty execution takes  $n + 1$  steps and reaches the non-faulty state  $\Sigma'_f$ , then  $\Sigma'$  simulates  $\Sigma'_f$  and the sequence of output pairs is identical the original execution. Alternatively, if the faulty execution reaches the fault state then the output pairs will be a prefix of the non-faulty output pairs.

#### Theorem 33 (Fault Tolerance)

If  $\vdash \Sigma$  and  $\Sigma \xrightarrow{0}^s \Sigma'$  then either  $\Sigma \xrightarrow{1}^{s'} \Sigma'_f$  or  $\exists m \leq (n+1) . \Sigma \xrightarrow{1}^m \text{fault}$ , and

1. For all derivations  $\Sigma \xrightarrow{1}^{s'} \Sigma'_f$  where  $\Sigma'_f \neq \text{fault}$ .  $s' = s$  and  $\exists c. \Sigma' \text{ sim}^c \Sigma'_f$ .
2. For all derivations  $\Sigma \xrightarrow{1}^m \text{fault}$  where  $m \leq (n+1)$ .  $s'$  is a prefix of  $s$ .

*Proof* By case analysis on the definition of  $\Sigma \xrightarrow{0}^s \Sigma'$ .

In the base case for *multi-base*, we can easily prove result 1 by having  $\Sigma$  take a faulty step to  $\Sigma_f$  and using the Fault Similarity Lemma to show that  $\exists c. \Sigma \text{ sim}^c \Sigma_f$ .

In the recursive case *multi-compose*, we use the Multistep Split Lemma to divide the computation into two pieces with an intermediate state  $\Sigma''$ .  $\Sigma''$  can take a faulty step

to  $\Sigma_f$ , and we can again use the Fault Similarity Lemma to show that they are similar. We then call the Multistep Fault Detection Lemma with the execution from  $\Sigma''$  to  $\Sigma'$  and the similarity between  $\Sigma''$  and  $\Sigma_f$ . This tells us that  $\Sigma_f$  can either step to a state  $\Sigma'_f$  that is similar to  $\Sigma'$  or reaches a fault before that point. Finally, we use the Multistep Combine Lemma to combine the first part of the non faulty execution, the fault step, and the resulting execution from the Multistep Fault Detection Lemma, to show that either  $\Sigma \xrightarrow{(n+1)}_1^{s'} \Sigma'_f$  or  $\Sigma \xrightarrow{m}_1^{s'} \text{fault}$ . The length of the faulty computation is at most  $n + 1$  because of the addition of the single fault step. Since the Multistep Split Lemma chooses some unspecified division of the original computation, the result hold regardless of exactly where the fault is injected. ■

# Appendix B

## MiniC Typing Rules

This appendix contains the typing rules for the MiniC Language defined in Section 3.1.

<i>type</i>	$\tau$	::=	$int \mid \tau \text{ ref} \mid X \rightarrow \tau$
<i>variable context</i>	$X$	::=	$\cdot \mid x : \tau, X$
<i>value</i>	$v$	::=	$n \mid x \mid \text{ref } v$
<i>value list</i>	$vs$	::=	$\cdot \mid v, vs$
<i>binaryop</i>	$op$	::=	$+ \mid - \mid *$
<i>statement</i>	$s$	::=	$x = v \mid x = v \text{ op } v$ $\mid x = !v \mid v := v$ $\mid x = f(vs)$ $\mid \text{if } v \text{ then } s \text{ else } s \mid \text{while } v \text{ do } s$ $\mid s; s$
<i>function declarations</i>	$fds$	::=	$\cdot \mid \tau f(X) \{lds; s; \text{return } v\} fds$
<i>local declarations</i>	$lds$	::=	$\cdot \mid \tau x = v; lds$
<i>program</i>	$p$	::=	$fds; lds; s; \text{return } v;$

$X \vdash v : \tau$

$$\frac{}{X \vdash n : int} \text{ (n-t)} \quad \frac{}{X \vdash x : X(x)} \text{ (var-t)} \quad \frac{}{X \vdash \text{ref } v : \tau \text{ ref}} \text{ (ref-t)}$$

$$\boxed{X \vdash vs : X}$$

$$\frac{}{X \vdash \dots} \text{ (vs}\dots\text{t)} \quad \frac{X \vdash v : \tau \quad X \vdash vs : X}{X \vdash (v, vs) : (x : \tau, X)} \text{ (vs-t)}$$

$$\boxed{F; A; L \vdash s \text{ wf}}$$

$$\frac{L \vdash x : \tau \quad (A \cup L) \vdash v : \tau}{F; A; L \vdash x = v \text{ wf}} \text{ (s-assign-wf)}$$

$$\frac{L \vdash x : \text{int} \quad (A \cup L) \vdash v_1 : \text{int} \quad (A \cup L) \vdash v_2 : \text{int}}{F; A; L \vdash x = v_1 \text{ op } v_2 \text{ wf}} \text{ (s-op-wf)}$$

$$\frac{L \vdash x : \tau \quad (A \cup L) \vdash v : \tau \text{ ref}}{F; A; L \vdash x = !v \text{ wf}} \text{ (s-deref-wf)}$$

$$\frac{(A \cup L) \vdash v_1 : \tau \text{ ref} \quad (A \cup L) \vdash v_2 : \tau}{F; A; L \vdash v_1 := v_2 \text{ wf}} \text{ (s-update-wf)}$$

$$\frac{(A \cup L) \vdash v : \text{int} \quad F; A; L \vdash s_1 \text{ wf} \quad F; A; L \vdash s_2 \text{ wf}}{F; A; L \vdash \text{if } v \text{ then } s_1 \text{ else } s_2 \text{ wf}} \text{ (s-if-wf)}$$

$$\frac{(A \cup L) \vdash v : \text{int} \quad F; A; L \vdash s \text{ wf}}{F; A; L \vdash \text{while } v \text{ do } s \text{ wf}} \text{ (s-while-wf)}$$

$$\frac{L \vdash x : \tau \quad F \vdash f : A' \rightarrow \tau \quad (A \cup L) \vdash vs : A'}{F; A; L \vdash x = f(vs) \text{ wf}} \text{ (s-call-wf)}$$

$$\frac{F; A; L \vdash s_1 \text{ wf} \quad F; A; L \vdash s_2 \text{ wf}}{F; A; L \vdash s_1; s_2 \text{ wf}} \text{ (s-seq-wf)}$$

$$\boxed{F; A; L \vdash lds : L'}$$

$$\frac{}{F; A; L \vdash \cdot : L} \text{ (lds---t)}$$

$$\frac{x \notin \text{Dom}(F \cup A \cup L) \quad (A \cup L) \vdash v : \tau \quad F; A; L[x : \tau] \vdash lds : L'}{F; A; L \vdash \tau x = v; lds : L'} \text{ (lds-t)}$$

$$\boxed{F \vdash fds : F'}$$

$$\frac{}{F \vdash \cdot : F} \text{ (fds---t)}$$

$$\frac{f \notin \text{Dom}(F) \quad F[f : (A \rightarrow \tau)]; A; \cdot \vdash lds : L \quad F[f : (A \rightarrow \tau)]; A; L \vdash s \text{ wf} \quad (A \cup L) \vdash v : \tau \quad F[f : (A \rightarrow \tau)] \vdash fds : F'}{F \vdash \tau f(A) \{lds; s; \text{return } v\} fds : F'} \text{ (fds-t)}$$

$$\boxed{\vdash p \text{ wf}}$$

$$\frac{\cdot \vdash fds : F \quad F; \cdot \vdash lds : L \quad F; \cdot; L \vdash s \text{ wf} \quad L \vdash v : \text{int}}{\vdash fds; lds; s; \text{return } v \text{ wf}} \text{ (p-wf)}$$



# Appendix C

## Complete Rules for ETAL<sub>FT</sub>

This appendix contains the complete dynamic and static semantics for ETAL<sub>FT</sub>. Section 3.2 discusses the main differences between ETAL<sub>FT</sub> and TAL<sub>FT</sub>.

### C.1 Syntax of Machine States

<i>colors</i>	$c$	$::= G \mid B$
<i>colored values</i>	$v$	$::= n$
<i>registers</i>	$r$	$::= r_n$
<i>general regs</i>	$a$	$::= r \mid gd \mid pc_c \mid sp_c$
<i>register file</i>	$R$	$::= \cdot \mid R, a \rightarrow n$
<i>code memory</i>	$C$	$::= \cdot \mid C, n \rightarrow i$
<i>value memory</i>	$M$	$::= \cdot \mid M, n \rightarrow n$
<i>store queue</i>	$Q$	$::= \overline{(n, n)}$
<i>ALU ops</i>	$op$	$::= add \mid sub \mid mul$
<i>instructions</i>	$i$	$::= op\ r_d, r_s, r_t \mid op\ r_d, r_s, n \mid mov\ r_d, n \mid mov\ r_d, r_s$ $\mid ld_c\ r_d, r_s \mid sld_c\ r_d\ n \mid st_c\ r_d, r_s \mid sst\ n\ r_v$ $\mid bz_c\ r_z, r_d \mid jmp_c\ r_d$ $\mid malloc[b]\ r_g, r_b \mid salloc\ n \mid sfree\ n$
<i>inst register</i>	$ir$	$::= i \mid \cdot$
<i>state</i>	$\Sigma$	$::= (R, C, M, Q, ir) \mid fault$

## C.2 Dynamic Semantics

### C.2.1 Fault Rules

$$\boxed{\Sigma \longrightarrow_1^s \Sigma'}$$

$$\frac{R(a) = n}{(R, C, M, Q, ir) \longrightarrow_1 (R[a \mapsto n'], C, M, Q, ir)} \text{ (reg-zap)}$$

$$\frac{Q_1 = \overline{(n_1, n'_1)}, (m_1, m'), \overline{(n_2, n'_2)}}{Q_2 = \overline{(n_1, n'_1)}, (m_2, m'), \overline{(n_2, n'_2)}} \text{ (Q}_1\text{-zap)}$$

$$\frac{}{(R, C, M, Q_1, ir) \longrightarrow_1 (R, C, M, Q_2, ir)}$$

$$\frac{Q_1 = \overline{(n_1, n'_1)}, (m, m'), \overline{(n_2, n'_2)}}{Q_2 = \overline{(n_1, n'_1)}, (m, m'_2), \overline{(n_2, n'_2)}} \text{ (Q}_2\text{-zap)}$$

$$\frac{}{(R, C, M, Q_1, ir) \longrightarrow_1 (R, C, M, Q_2, ir)}$$

### C.2.2 Normal Execution Rules

$$\boxed{\Sigma \longrightarrow_0^s \Sigma'}$$

$$\frac{R(pc_G) = R(pc_B) \quad R(pc_G) \in \text{Dom}(C)}{(R, C, M, Q, \cdot) \longrightarrow_0 (R, C, M, Q, C(R(pc_G)))} \text{ (fetch)}$$

$$\frac{R(pc_G) \neq R(pc_B)}{(R, C, M, Q, \cdot) \longrightarrow_0 \text{fault}} \text{ (fetch-fail)}$$

$$\frac{R' = R++[r_d \mapsto R(r_s) \text{ op } R(r_t)]}{(R, C, M, Q, \text{op } r_d, r_s, r_t) \longrightarrow_0 (R', C, M, Q, \cdot)} \text{ (op2r)}$$

$$\frac{R' = R++[r_d \mapsto R(r_s) \text{ op } n]}{(R, C, M, Q, \text{op } r_d, r_s, n) \longrightarrow_0 (R', C, M, Q, \cdot)} \text{ (op1r)}$$

$$\frac{R' = R++[r_d \mapsto n]}{(R, C, M, Q, \text{mov } r_d \text{ } n) \longrightarrow_0 (R', C, M', Q, \cdot)} \text{ (mov-n)}$$

$$\frac{R' = R++[r_d \mapsto R(r_s)]}{(R, C, M, Q, \text{mov } r_d \ r_s) \longrightarrow_0 (R', C, M', Q, \cdot)} \text{ (mov-reg)}$$

$$\frac{\begin{array}{l} n = \max(\text{Dom}(M)) + 1 \\ R' = R++[r_g \mapsto n][r_b \mapsto n] \end{array}}{(R, C, M, Q, \text{malloc}[b] \ r_g \ r_b) \longrightarrow_0 (R', C, (M, n \mapsto 0), Q, \cdot)} \text{ (malloc)}$$

$$\frac{\begin{array}{l} R' = R++[sp_G \mapsto R(sp_G) - n][sp_B \mapsto R(sp_B) - n] \\ m = \min(\text{Dom}(M)) \\ M' = (M, m - 1 \mapsto 0, \dots, m - n \mapsto 0) \end{array}}{(R, C, M, Q, \text{salloc } n) \longrightarrow_0 (R', C, M', Q, \cdot)} \text{ (salloc)}$$

$$\frac{\begin{array}{l} R' = R++[sp_G \mapsto R(sp_G) + n][sp_B \mapsto R(sp_B) + n] \\ m = \min(\text{Dom}(M)) \\ M = M', m \mapsto v_m, \dots, (m + n - 1) \mapsto v_1 \end{array}}{(R, C, M, Q, \text{sfree } n) \longrightarrow_0 (R', C, M', Q, \cdot)} \text{ (sfree)}$$

$$\frac{Q' = ((R(r_d), R(r_s)), Q)}{(R, C, M, Q, \text{st}_G \ r_d, r_s) \longrightarrow_0 (R++, C, M, Q', \cdot)} \text{ (st}_G\text{-queue)}$$

$$\frac{R(rd) = n_l \quad R(rs) = n'_l}{(R, C, M, (\overline{(n, n')}, (n_l, n'_l)), \text{st}_B \ r_d, r_s) \longrightarrow_0^{(n_l, n'_l)} (R++, C, M[n_l \mapsto n'_l], \overline{(n, n')}, \cdot)} \text{ (st}_B\text{-mem)}$$

$$\frac{Q = (\overline{(n, n')}, (n_l, n'_l)) \quad R(rd) \neq n_l \text{ or } R(rs) \neq n'_l}{(R, C, M, Q, \text{st}_B \ r_d, r_s) \longrightarrow_0 \text{ fault}} \text{ (st}_B\text{-mem-fail)}$$

$$\frac{R(sp_B) = R(sp_G) \quad R(sp_G) + n \in \text{Dom}(M)}{(R, C, M, Q, \text{sst } n, r_v) \longrightarrow_0^{(R(sp_G)+n, R(r_v))} (R++, C, M[R(sp_G) + n \mapsto R(r_v)], Q, \cdot)} \text{ (sst)}$$

$$\frac{R(sp_G) \neq R(sp_B) \text{ or } R(sp_B) + n \notin \text{Dom}(M)}{(R, C, M, Q, \text{sst } n, r_v) \longrightarrow_0 \text{ fault}} \text{ (sst-fail)}$$

$$\frac{\begin{array}{l} \text{find}(Q, R(r_s)) = (R(r_s), n) \\ R' = R++[r_d \mapsto n] \end{array}}{(R, C, M, Q, ld_G \ r_d, r_s) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (ld_G\text{-queue})$$

$$\frac{\begin{array}{l} \text{find}(Q, R(r_s)) = () \\ R(r_s) \in \text{Dom}(M) \\ R' = R++[r_d \mapsto M(R(r_s))] \end{array}}{(R, C, M, Q, ld_G \ r_d, r_s) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (ld_G\text{-mem})$$

$$\frac{\begin{array}{l} R(r_s) \in \text{Dom}(M) \\ R' = R++[r_d \mapsto M(R(r_s))] \end{array}}{(R, C, M, Q, ld_B \ r_d, r_s) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (ld_B\text{-mem})$$

$$\frac{\begin{array}{l} \text{find}(Q, R(r_s)) = () \\ R(r_s) \notin \text{Dom}(M) \end{array}}{(R, C, M, Q, ld_G \ r_d, r_s) \longrightarrow_0 \text{fault}} \quad (ld_G\text{-fail})$$

$$\frac{R(r_s) \notin \text{Dom}(M)}{(R, C, M, Q, ld_B \ r_d, r_s) \longrightarrow_0 \text{fault}} \quad (ld_B\text{-fail})$$

$$\frac{\begin{array}{l} \text{find}(Q, R(r_s)) = () \\ R(r_s) \notin \text{Dom}(M) \\ R' = R++[r_d \mapsto n] \end{array}}{(R, C, M, Q, ld_G \ r_d, r_s) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (ld_G\text{-rand})$$

$$\frac{\begin{array}{l} R(r_s) \notin \text{Dom}(M) \\ R' = R++[r_d \mapsto n] \end{array}}{(R, C, M, Q, ld_B \ r_d, r_s) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (ld_B\text{-rand})$$

$$\frac{\begin{array}{l} R(sp_c) + n \in \text{Dom}(M) \\ R' = R++[r_d \mapsto M(R(sp_c) + n)] \end{array}}{(R, C, M, Q, sld_c \ r_d, n) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (sld_c)$$

$$\frac{R(sp_c) + n \notin \text{Dom}(M)}{(R, C, M, Q, sld_c \ r_d, n) \longrightarrow_0 \text{fault}} \quad (sld_c\text{-fail})$$

$$\frac{R(gd) = 0 \quad R' = R_{++}[gd \mapsto R(r_d)]}{(R, C, M, Q, jmp_G r_d) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (jmp_G)$$

$$\frac{R(gd) \neq 0}{(R, C, M, Q, jmp_G r_d) \longrightarrow_0 \text{fault}} \quad (jmp_G\text{-fail})$$

$$\frac{R(gd) \neq 0 \quad R(r_d) = R(gd) \quad R' = R[pc_G \mapsto R(gd)][pc_B \mapsto R(r_d)][d \mapsto 0]}{(R, C, M, Q, jmp_B r_d) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (jmp_B)$$

$$\frac{R(r_d) \neq R(gd) \text{ or } R(gd) = 0}{(R, C, M, Q, jmp_B r_d) \longrightarrow_0 \text{fault}} \quad (jmp_B\text{-fail})$$

$$\frac{R(gd) = 0 \quad R(r_z) \neq 0}{(R, C, M, Q, bz_c r_z, r_d) \longrightarrow_0 (R_{++}, C, M, Q, \cdot)} \quad (bz\text{-untaken})$$

$$\frac{R(r_z) \neq 0 \quad R(gd) \neq 0}{(R, C, M, Q, bz_c r_z, r_d) \longrightarrow_0 \text{fault}} \quad (bz\text{-untaken-fail})$$

$$\frac{R(gd) = 0 \quad R(r_z) = 0 \quad R' = R_{++}[gd \mapsto R(r_d)]}{(R, C, M, Q, bz_G r_z, r_d) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (bz_G\text{-taken})$$

$$\frac{R(r_z) = 0 \quad R(gd) \neq 0}{(R, C, M, Q, bz_G r_z, r_d) \longrightarrow_0 \text{fault}} \quad (bz_G\text{-taken-fail})$$

$$\frac{R(gd) \neq 0 \quad R(r_z) = 0 \quad R(r_d) = R(gd) \quad R' = R[pc_G \mapsto R(gd)][pc_B \mapsto R(r_d)][gd \mapsto 0]}{(R, C, M, Q, bz_B r_z, r_d) \longrightarrow_0 (R', C, M, Q, \cdot)} \quad (bz_B\text{-taken})$$

$$\frac{R(r_z) = 0 \quad R(r_d) \neq R(gd) \text{ or } R(gd) = 0}{(R, C, M, Q, bz_B r_z, r_d) \longrightarrow_0 \text{fault}} \quad (bz_B\text{-taken-fail})$$

## C.3 Static Semantics

### C.3.1 Syntax

Static Expressions

<i>exp kinds</i>	$\kappa ::= \kappa_{int} \mid \kappa_{mem} \mid \kappa_{\sigma}$
<i>exp contexts</i>	$\Delta ::= \cdot \mid \Delta, x : \kappa$
<i>exps</i>	$E ::= x \mid n \mid E \text{ op } E \mid \text{sel } E_m E_n$ $\quad \mid \text{emp} \mid \text{upd } E_m E_{n_1} E_{n_2}$
<i>substitutions</i>	$S ::= \cdot \mid S, E/x$

Types

<i>zap tags</i>	$Z ::= \cdot \mid c$
<i>initialization flags</i>	$\wp ::= 1 \mid \frac{1}{2} \mid 0$
<i>base types</i>	$b ::= \text{int} \mid \Theta \rightarrow \text{void} \mid b \text{ ref}^{\wp} \mid \text{sptr}$
<i>reg types</i>	$t ::= \langle c, b, E \rangle \mid E' = 0 \Rightarrow \langle c, b, E \rangle \mid ns$
<i>reg file types</i>	$\Gamma ::= \cdot \mid \Gamma, a \rightarrow t$
<i>unlabeled stack</i>	$\sigma ::= \text{sbase} \mid \rho \mid t :: \zeta$
<i>labeled stack</i>	$\zeta ::= E : \sigma$
<i>result types</i>	$RT ::= \Theta \mid \text{void}$

Contexts

<i>heap typing</i>	$\Psi ::= \cdot \mid \Psi, n : b$
<i>static context</i>	$\Theta ::= \Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta$

### C.3.2 Properties of Static Expressions

$\Delta \vdash E : \kappa$

$$\frac{x \in \text{Dom}(\Delta)}{\Delta \vdash x : \Delta(x)} \text{ (E-var-t)}$$

$$\frac{}{\Delta \vdash n : \kappa_{int}} \text{ (E-int-t)}$$

$$\frac{\Delta \vdash E_1 : \kappa_{int} \quad \Delta \vdash E_2 : \kappa_{int}}{\Delta \vdash E_1 \text{ op } E_2 : \kappa_{int}} \text{ (E-op-t)}$$

$$\frac{\Delta \vdash E_m : \kappa_{mem} \quad \Delta \vdash E_n : \kappa_{int}}{\Delta \vdash sel E_m E_n : \kappa_{int}} \quad (E\text{-sel-}t)$$

$$\frac{\Delta \vdash E_m : \kappa_{mem} \quad \Delta \vdash E_{n_1} : \kappa_{int} \quad \Delta \vdash E_{n_2} : \kappa_{int}}{\Delta \vdash upd E_m E_{n_1} E_{n_2} : \kappa_{mem}} \quad (E\text{-upd-}t)$$

$$\frac{}{\Delta \vdash emp : \kappa_{mem}} \quad (E\text{-emp-}t)$$

$$\boxed{\Delta \vdash S : \Delta'}$$

$$\frac{}{\Delta \vdash \dots} \quad (sub\text{-emp-}t)$$

$$\frac{\Delta \vdash S : \Delta' \quad \Delta \vdash E : \kappa \quad x \notin Dom(\Delta) \cup Dom(\Delta')}{\Delta \vdash S, E/x : \Delta', x : \kappa} \quad (sub\text{-}t)$$

$$\boxed{\llbracket E \rrbracket}$$

$$\begin{aligned} \llbracket n \rrbracket &= n \\ \llbracket emp \rrbracket &= \cdot \\ \llbracket E_1 op E_2 \rrbracket &= \llbracket E_1 \rrbracket op \llbracket E_2 \rrbracket \\ \llbracket sel E_m E_n \rrbracket &= \llbracket E_m \rrbracket (\llbracket E_n \rrbracket) \\ \llbracket upd E_m E_1 E_2 \rrbracket &= \llbracket E_m \rrbracket [ \llbracket E_1 \rrbracket \mapsto \llbracket E_2 \rrbracket ] \end{aligned}$$

$$\boxed{\Delta \vdash E_1 op E_2}$$

$$\frac{\Delta \vdash E_1 : \kappa_{int} \quad \Delta \vdash E_2 : \kappa_{int} \quad \forall S. \cdot \vdash S : \Delta \implies \llbracket S(E_1) \rrbracket = \llbracket S(E_2) \rrbracket}{\Delta \vdash E_1 = E_2} \quad (E\text{-eq})$$

$$\frac{\Delta \vdash E_1 : \kappa_{int} \quad \Delta \vdash E_2 : \kappa_{int} \quad \forall S. \cdot \vdash S : \Delta \implies \llbracket S(E_1) \rrbracket \neq \llbracket S(E_2) \rrbracket}{\Delta \vdash E_1 \neq E_2} \text{ (E-ineq)}$$

$$\frac{\Delta \vdash E_1 : \kappa_{mem} \quad \Delta \vdash E_2 : \kappa_{mem} \quad \forall \ell \in \text{Dom}(\llbracket S(E_1) \rrbracket) \cup \text{Dom}(\llbracket S(E_2) \rrbracket). \llbracket S(E_1) \rrbracket(\ell) = \llbracket S(E_2) \rrbracket(\ell)}{\Delta \vdash E_1 = E_2} \text{ (E-mem-eq)}$$

### C.3.3 Value Typing

$$\boxed{\Psi \vdash n : b}$$

$$\overline{\Psi \vdash n : int} \text{ (int-t)} \quad \overline{\Psi \vdash n : \Psi(n)} \text{ (addr-heap-t)}$$

$$\overline{\Psi \vdash n : sptr} \text{ (addr-stack-t)} \quad \frac{\Psi \vdash n : b \text{ ref}^\varphi \quad \varphi \leq \varphi'}{\Psi \vdash n : b \text{ ref}^{\varphi'}} \text{ (addr-subtp-t)}$$

$$\boxed{\Psi; \Delta \vdash^Z n : t}$$

$$\frac{\Psi \vdash n : b \quad \Delta \vdash E = n}{\Psi; \Delta \vdash^Z n : \langle c, b, E \rangle} \text{ (val-t)}$$

$$\frac{\Psi; \Delta \vdash^Z n : t' \quad \Delta \vdash t' \leq t}{\Psi; \Delta \vdash^Z n : t} \text{ (val-subtp-t)}$$

$$\frac{n \neq 0 \quad \Psi; \Delta \vdash^Z n : \langle c, b, E \rangle \quad \Delta \vdash E' = 0}{\Psi; \Delta \vdash^Z n : E' = 0 \Rightarrow \langle c, b, E \rangle} \text{ (cond-t)}$$

$$\frac{\Delta \vdash E' \neq 0}{\Psi; \Delta \vdash^Z 0 : E' = 0 \Rightarrow \langle c, b, E \rangle} \text{ (cond-n0-t)}$$



$$\frac{\Delta \vdash E : \kappa_{int}}{\Psi; \Delta \vdash^c n : \langle c, b, E \rangle} \text{ (val-zap-t)}$$

$$\frac{\Delta \vdash E' : \kappa_{int} \quad \Delta \vdash E : \kappa_{int}}{\Psi; \Delta \vdash^c n : E' = 0 \Rightarrow \langle c, b, E \rangle} \text{ (val-zap-cond-t)}$$

$$\frac{}{\Psi; \Delta \vdash^Z n : ns} \text{ (ns-t)}$$

### C.3.4 Subtyping

$$\boxed{\varphi \uparrow}$$

$$\begin{aligned} 0 \uparrow &= \frac{1}{2} \\ \frac{1}{2} \uparrow &= \frac{1}{2} \\ 1 \uparrow &= 1 \end{aligned}$$

$$\boxed{\varphi \leq \varphi'}$$

$$1 \leq \frac{1}{2} \leq 0 \quad \varphi \leq \varphi$$

$$\boxed{b \leq b'}$$

$$\frac{}{b \leq b} \text{ (subtp-b-reflex)} \quad \frac{\varphi \leq \varphi'}{b \text{ ref}^\varphi \leq b \text{ ref}^{\varphi'}} \text{ (subtp-b-ref)} \quad \frac{}{b \leq int} \text{ (subtp-b-int)}$$

$$\boxed{\Delta \vdash t \leq t'}$$

$$\frac{\Delta \vdash E_1 = E_2 \quad b_1 \leq b_2}{\Delta \vdash \langle c, b_1, E_1 \rangle \leq \langle c, b_2, E_2 \rangle} \text{ (subtp-t-triple)}$$

$$\frac{\Delta \vdash t \leq t' \quad \Delta \vdash E = E'}{\Delta \vdash (E = 0 \Rightarrow t) \leq (E' = 0 \Rightarrow t')} \text{ (subtp-t-cond)}$$

$$\frac{}{\Delta \vdash t \leq ns} \text{ (subtp-t-ns)}$$

$$\boxed{\Delta \vdash \Gamma_1 \leq \Gamma_2}$$

$$\frac{\forall r \in \text{Dom}(\Gamma_2). \Gamma_1(r) \leq \Gamma_2(r)}{\Delta \vdash \Gamma_1 \leq \Gamma_2} \text{ (reg-file-comp)}$$

$$\boxed{\Delta \vdash \zeta \leq \zeta'}$$

$$\frac{\Delta \vdash E = E'}{\Delta \vdash E : sbase \leq E' : sbase} \text{ (subtp-}\zeta\text{-base)}$$

$$\frac{\Delta \vdash E = E'}{\Delta \vdash E : \rho \leq E' : \rho} \text{ (subtp-}\zeta\text{-var)}$$

$$\frac{\Delta \vdash E = E' \quad \Delta \vdash t \leq t' \quad \Delta \vdash \zeta \leq \zeta'}{\Delta \vdash E : (t :: \zeta) \leq E' : (t' :: \zeta')} \text{ (subtp-}\zeta\text{-cons)}$$

### C.3.5 Stack Typing Judgments

$$\boxed{\Delta \vdash \zeta \text{ wf}}$$

$$\frac{\Delta \vdash E : \kappa_{int}}{\Delta \vdash E : sbase \text{ wf}} \text{ (}\zeta\text{-wf-base)} \quad \frac{\Delta(\rho) = \kappa_\sigma \quad \Delta \vdash E : \kappa_{int}}{\Delta \vdash E : \rho \text{ wf}} \text{ (}\zeta\text{-wf-var)}$$

$$\frac{\Delta \vdash E + 1 = E' \quad \Delta \vdash (E' : \sigma') \text{ wf}}{\Delta \vdash E : (t :: (E' : \sigma')) \text{ wf}} \text{ (}\zeta\text{-wf-cons)}$$

$$\boxed{\Delta; \zeta \vdash E : t}$$

$$\frac{\Delta \vdash E_s = E}{\Delta; E_s : (t :: \zeta) \vdash E : t} (\zeta\text{-lookup-top}) \quad \frac{\Delta \vdash E_s \neq E \quad \Delta; \zeta' \vdash E : t}{\Delta; E_s : (t :: \zeta') \vdash E : t} (\zeta\text{-lookup-tail})$$

$$\boxed{\Delta \vdash \zeta[E \mapsto t] = \zeta'}$$

$$\frac{\Delta \vdash E_s = E}{\Delta \vdash (E_s : (t_s :: \zeta))[E \mapsto t] = E_s : (t :: \zeta)} (\zeta\text{-update-top})$$

$$\frac{\Delta \vdash E_s \neq E \quad \Delta \vdash \zeta[E \mapsto t] = \zeta'}{\Delta \vdash (E_s : (t_s :: \zeta))[E \mapsto t] = E_s : (t_s :: \zeta')} (\zeta\text{-update-tail})$$

### C.3.6 Instruction Typing Rules

$$\boxed{\Psi; \Theta \vdash ir \Rightarrow RT}$$

$$\frac{}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \cdot \Rightarrow (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta)} (\cdot\text{-}t)$$

$$\frac{\Gamma(r_s) = \langle c, \text{int}, E'_s \rangle \quad \Gamma(r_t) = \langle c, \text{int}, E'_t \rangle \quad \Gamma' = \Gamma_{++}[r_d \mapsto \langle c, \text{int}, E'_s \text{ op } E'_t \rangle]}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{op } r_d, r_s, r_t \Rightarrow (\Delta; \Gamma'; \overline{(E_d, E_s)}; E_m; \zeta)} (\text{op}2r\text{-}t)$$

$$\frac{\Gamma(r_s) = \langle c, \text{int}, E'_s \rangle \quad \Gamma' = \Gamma_{++}[r_d \mapsto \langle c, \text{int}, E'_s \text{ op } n \rangle]}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{op } r_d, r_s, n \Rightarrow (\Delta; \Gamma'; \overline{(E_d, E_s)}; E_m; \zeta)} (\text{op}1r\text{-}t)$$

$$\frac{\Psi; \Delta \vdash n : t}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{mov } r_d, n \Rightarrow (\Delta; \Gamma_{++}[r_d \mapsto t]; \overline{(E_d, E_s)}; E_m; \zeta)} (\text{mov-}n\text{-}t)$$

$$\frac{\Gamma(r_s) = t}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{mov } r_d, r_s \Rightarrow (\Delta; \Gamma_{++}[r_d \mapsto t]; \overline{(E_d, E_s)}; E_m; \zeta)} \text{ (mov-reg-t)}$$

$$\frac{\begin{array}{c} x \notin \Delta \\ \Gamma' = \Gamma_{++}[r_g \mapsto \langle G, b \text{ ref}^0, x \rangle][r_b \mapsto \langle B, b \text{ ref}^0, x \rangle] \\ E'_m = \text{upd } E_m \ x \ 0 \end{array}}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{malloc}[b] \ r_g \ r_b \Rightarrow (\Delta, x : \kappa_{\text{int}}; \Gamma'; \overline{(E_d, E_s)}; E'_m; \zeta)} \text{ (malloc-t)}$$

$$\frac{\begin{array}{c} \Gamma(sp_G) = \langle G, \text{sptr}, E_g \rangle \quad \Gamma(sp_B) = \langle B, \text{sptr}, E_b \rangle \\ \Delta \vdash E_g = E_b \quad \Delta \vdash E_g = E_t \\ \Gamma' = \Gamma_{++}[sp_G \mapsto \langle G, \text{sptr}, (E_g - n) \rangle][sp_B \mapsto \langle B, \text{sptr}, (E_b - n) \rangle] \\ \zeta' = (E_t - n) : ns :: (E_t - (n+1)) : ns :: \dots :: E_t : \sigma \end{array}}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; (E_t : \sigma)) \vdash \text{salloc } n \Rightarrow (\Delta; \Gamma'; \overline{(E_d, E_s)}; E_m; \zeta')} \text{ (salloc-t)}$$

$$\frac{\begin{array}{c} \Gamma(sp_G) = \langle G, \text{sptr}, E_g \rangle \quad \Gamma(sp_B) = \langle B, \text{sptr}, E_b \rangle \\ \Delta \vdash E_g = E_b \quad \Delta \vdash E_g = E_t \\ \zeta = E_t : t :: \dots :: E_f : \sigma \quad \Delta \vdash E_f = E_g + n \\ \Gamma' = \Gamma_{++}[sp_G \mapsto \langle G, \text{sptr}, (E_g + n) \rangle][sp_B \mapsto \langle B, \text{sptr}, (E_b + n) \rangle] \end{array}}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{sfree } n \Rightarrow (\Delta; \Gamma'; \overline{(E_d, E_s)}; E_m; E_f : \sigma)} \text{ (sfree-t)}$$

$$\frac{\Delta \vdash \Gamma(r_s) \leq \langle G, b \text{ ref}^{\frac{1}{2}}, E'_s \rangle \quad E = \text{sel } (\overline{\text{upd}} \ E_m \ \overline{(E_d, E_s)}) \ E'_s}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{ld}_G \ r_d \ r_s \Rightarrow (\Delta; \Gamma_{++}[r_d \mapsto \langle G, b, E \rangle]; \overline{(E_d, E_s)}; E_m; \zeta)} \text{ (ld}_G\text{-t)}$$

$$\frac{\Gamma(r_s) = \langle B, b \text{ ref}^1, E'_s \rangle \quad E = \text{sel } E_m \ E'_s}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{ld}_B \ r_d \ r_s \Rightarrow (\Delta; \Gamma_{++}[r_d \mapsto \langle B, b, E \rangle]; \overline{(E_d, E_s)}; E_m; \zeta)} \text{ (ld}_B\text{-t)}$$

$$\frac{\Gamma(sp_c) = \langle c, \text{sptr}, E_c \rangle \quad \Delta \vdash E_c + n = E_n \quad \Delta; \zeta \vdash E_n : \langle c, b, E \rangle}{\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{sld}_c \ r_d \ n \Rightarrow (\Delta; \Gamma_{++}[r_d \mapsto \langle c, b, E \rangle]; \overline{(E_d, E_s)}; E_m; \zeta)} \text{ (sld}_c\text{-t)}$$

$$\begin{array}{c}
\Gamma(r_d) = \langle G, b \text{ ref}^\Phi, E'_d \rangle \quad \Gamma(r_s) = \langle G, b, E'_s \rangle \\
\Gamma' = \Gamma_{++} \text{ except } \forall r \text{ where } \Gamma(r) = \langle c_r, b \text{ ref}^\Phi, E_r \rangle \text{ and } \Delta \vdash E_r = E'_d . \\
\Gamma'(r) = \langle c_r, b \text{ ref}^{\Phi^\uparrow}, E_r \rangle \\
\hline
\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{st}_G r_d r_s \Rightarrow (\Delta; \Gamma'; (E'_d, E'_s), \overline{(E_d, E_s)}; E_m; \zeta) \quad (\text{st}_G\text{-t})
\end{array}$$

$$\begin{array}{c}
\Delta \vdash \Gamma(r_d) \leq \langle B, b \text{ ref}^{\frac{1}{2}}, E''_d \rangle \quad \Gamma(r_s) = \langle B, b, E''_s \rangle \\
\Delta \vdash E'_s = E''_s \quad \Delta \vdash E'_d = E''_d \\
\Gamma' = \Gamma_{++} \text{ except } \forall r \text{ where } \Gamma(r) = \langle c_r, b \text{ ref}^{\frac{1}{2}}, E_r \rangle \text{ and } \Delta \vdash E_r = E'_d . \\
\Gamma'(r) = \langle c_r, b \text{ ref}^1, E_r \rangle \\
\hline
\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}, (E'_d, E'_s); E_m; \zeta) \vdash \text{st}_B r_d r_s \Rightarrow (\Delta; \Gamma'; \overline{(E_d, E_s)}; \text{upd } E_m E'_d E'_s; \zeta) \quad (\text{st}_B\text{-t})
\end{array}$$

$$\begin{array}{c}
\Gamma(\text{sp}_G) = \langle G, \text{sptr}, E_g \rangle \quad \Gamma(\text{sp}_B) = \langle B, \text{sptr}, E_b \rangle \quad \Delta \vdash E_g = E_b \\
\Delta \vdash E_g + n = E_n \quad \Gamma(r_v) = \langle c, b, E_v \rangle \quad \Delta \vdash \zeta[E_n \mapsto \langle c, b, E_v \rangle] = \zeta' \\
\hline
\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{sst } n r_v \Rightarrow (\Delta; \Gamma_{++}; \overline{(E_d, E_s)}; E_m; \zeta') \quad (\text{sst-t})
\end{array}$$

$$\begin{array}{c}
\Gamma(\text{gd}) = \langle G, \text{int}, 0 \rangle \quad \Gamma(r_z) = \langle G, \text{int}, E_z \rangle \\
\Theta = (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m) \quad \Gamma(r_d) = \langle G, \Theta \rightarrow \text{void}, E'_d \rangle \\
\Gamma'(\text{gd}) = \langle G, \text{int}, 0 \rangle \quad \Gamma'' = \Gamma_{++}[\text{gd} \mapsto E_z = 0 \Rightarrow \langle G, \Theta \rightarrow \text{void}, E'_d \rangle] \\
\hline
\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{bz}_G r_z r_d \Rightarrow (\Delta; \Gamma''; \overline{(E_d, E_s)}; E_m; \zeta) \quad (\text{bz}_G\text{-t})
\end{array}$$

$$\begin{array}{c}
\Gamma(r_d) = \langle G, \Theta \rightarrow \text{void}, E_{rd'} \rangle \quad \Theta = (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m) \\
\Gamma(\text{gd}) = \langle G, \text{int}, 0 \rangle \quad \Gamma'(\text{gd}) = \langle G, \text{int}, 0 \rangle \\
\Gamma'' = \Gamma_{++}[\text{gd} \mapsto \langle G, \Theta \rightarrow \text{void}, E_{rd'} \rangle] \\
\hline
\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash \text{jmp}_G r_d \Rightarrow (\Delta; \Gamma''; \overline{(E_d, E_s)}; E_m; \zeta) \quad (\text{jmp}_G\text{-t})
\end{array}$$

$$\begin{array}{c}
\Gamma(r_z) = \langle B, \text{int}, E_z \rangle \\
\Gamma(r_d) = \langle B, (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m; \zeta') \rightarrow \text{void}, E_r \rangle \\
\Gamma(gd) = E'_z = 0 \Rightarrow \langle G, (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m; \zeta') \rightarrow \text{void}, E'_r \rangle \\
\Delta \vdash E_z = E'_z \\
\Delta \vdash E_r = E'_r \\
\exists S. \Delta \vdash S : \Delta' \\
S(\Gamma')(gd) = \langle G, \text{int}, 0 \rangle \\
S(\Gamma')(pc_G) = \langle G, \text{int}, E'_r \rangle \\
S(\Gamma')(pc_B) = \langle B, \text{int}, E_r \rangle \\
\Delta \vdash \Gamma \leq S(\Gamma') \\
\Delta \vdash \overline{(E_d, E_s)} = S(\overline{(E'_d, E'_s)}) \\
\Delta \vdash E_m = S(E'_m) \\
\Delta \vdash \zeta \leq S(\zeta') \\
\hline
\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash bz_B r_z r_d \Rightarrow \\
(\Delta; \Gamma++; \overline{(E_d, E_s)}; E_m; \zeta) \quad (bz_B-t)
\end{array}$$

$$\begin{array}{c}
\Gamma(gd) = \langle G, (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m; \zeta') \rightarrow \text{void}, E'_r \rangle \\
\Gamma(r_d) = \langle B, (\Delta'; \Gamma'; \overline{(E'_d, E'_s)}; E'_m; \zeta') \rightarrow \text{void}, E_r \rangle \\
\Delta \vdash E_r = E'_r \\
\exists S. \Delta \vdash S : \Delta' \\
S(\Gamma')(gd) = \langle G, \text{int}, 0 \rangle \\
S(\Gamma')(pc_G) = \langle G, \text{int}, E'_r \rangle \\
S(\Gamma')(pc_B) = \langle B, \text{int}, E_r \rangle \\
\Delta \vdash \Gamma \leq S(\Gamma') \\
\Delta \vdash \overline{(E_d, E_s)} = S(\overline{(E'_d, E'_s)}) \\
\Delta \vdash E_m = S(E'_m) \\
\Delta \vdash \zeta \leq S(\zeta') \\
\hline
\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash jmp_B r_d \Rightarrow \text{void} \quad (jmp_B-t)
\end{array}$$

### C.3.7 Machine State Typing

$$\boxed{\Psi \vdash^Z R : \Gamma}$$

$$\frac{\begin{array}{l} \forall a \in \text{Dom}(\Gamma). \Psi; \cdot \vdash^Z R(a) : \Gamma(a) \\ \cdot \vdash \Gamma(pc_G) \leq \langle G, \text{int}, E_G \rangle \\ \cdot \vdash \Gamma(pc_B) \leq \langle B, \text{int}, E_B \rangle \\ \cdot \vdash E_G = E_B \end{array}}{\Psi \vdash^Z R : \Gamma} \text{ (reg-file-t)}$$

$$\boxed{\Psi \vdash C}$$

$$\frac{\begin{array}{l} 0 \notin \text{Dom}(C) \\ \forall n \in \text{Dom}(C). \Psi(n) = \Theta \rightarrow \text{void} \wedge \Psi; \Theta \vdash C(n) \Rightarrow RT \wedge \\ (RT = \Theta' \text{ implies } \Psi(n+1) = \Theta' \rightarrow \text{void}) \end{array}}{\Psi \vdash C} \text{ (C-t)}$$

$$\boxed{M = M_s \#^{\mathcal{L}} M_m}$$

$$\frac{\begin{array}{l} \text{Dom}(M) = \text{Dom}(M_1) \cup \text{Dom}(M_2) \\ \text{Dom}(M_1) \cap \text{Dom}(M_2) = \emptyset \\ \forall \ell_1 \in \text{Dom}(M_1). \forall \ell_{\mathcal{L}} \in \mathcal{L}. \forall \ell_2 \in \text{Dom}(M_2). \ell_1 < \ell_{\mathcal{L}} < \ell_2 \end{array}}{M = M_1 \#^{\mathcal{L}} M_2} \text{ (#-def)}$$

$$\boxed{\Psi; M; Q \vdash^Z \ell : b \text{ ref}^\Phi}$$

$$\frac{\Psi(\ell) = b \text{ ref}^1 \quad \Psi \vdash M(\ell) : b}{\Psi; M; Q \vdash^Z \ell : b \text{ ref}^1} \text{ (init-t)}$$

$$\frac{\Psi(\ell) = b \text{ ref}^0}{\Psi; M; Q \vdash^Z \ell : b \text{ ref}^0} \text{ (uninit-t)}$$

$$\frac{\Psi(\ell) = b \text{ ref}^{\frac{1}{2}} \quad Z \neq G \implies \exists n. (\ell, n) \in Q}{\Psi; M; Q \vdash^Z \ell : b \text{ ref}^{\frac{1}{2}}} \text{ (halfinit-t)}$$

$$\boxed{\Psi \vdash^Z Q : (\overline{E_d}, \overline{E_s})}$$

$$\frac{}{\Psi \vdash^Z () : ()} \text{ (Q-emp-t)} \quad \frac{\Psi \vdash^G \overline{(n'_1, n'_2)} : (\overline{E'_d}, \overline{E'_s}) \quad \cdot \vdash E_d : \kappa_{int} \quad \cdot \vdash E_s : \kappa_{int}}{\Psi \vdash^G (n_1, n_2), \overline{(n'_1, n'_2)} : (E_d, E_s), (\overline{E'_d}, \overline{E'_s})} \text{ (Q-zap-t)}$$

$$\frac{Z \neq G \quad \Psi \vdash^Z \overline{(n'_1, n'_2)} : (\overline{E'_d}, \overline{E'_s}) \quad \cdot \vdash E_d = n_1 \quad \cdot \vdash E_s = n_2 \quad \Psi \vdash n_1 : b \text{ ref}^\phi \quad \phi \leq \frac{1}{2} \quad \Psi \vdash n_2 : b}{\Psi \vdash^Z (n_1, n_2), \overline{(n'_1, n'_2)} : (E_d, E_s), (\overline{E'_d}, \overline{E'_s})} \text{ (Q-t)}$$

$$\boxed{\Psi \vdash^Z (M, Q) : (E_m, \overline{(E_d, E_s)})}$$

$$\frac{\forall \ell \in \text{Dom}(M). \exists \phi. \Psi; M; Q \vdash^Z \ell : b \text{ ref}^\phi \quad \llbracket E_m \rrbracket = M \quad \Psi \vdash^Z Q : (\overline{E_d}, \overline{E_s})}{\Psi \vdash^Z (M, Q) : (E_m, \overline{(E'_d, E'_s)})} \text{ (heap-t)}$$

$$\boxed{\Psi \vdash^Z M : \zeta}$$

$$\frac{\cdot \vdash E = \ell \quad \text{Dom}(M) = \{\ell\}}{\Psi \vdash^Z M : (E : \text{sbase})} \text{ (\zeta-t-base)}$$

$$\frac{\cdot \vdash (E : t :: \zeta') \text{ wf} \quad \cdot \vdash E = \ell \quad M = \{\ell \rightarrow n\} \# M' \quad \Psi; \cdot \vdash^Z n : t \quad \Psi \vdash^Z M' : \zeta'}{\Psi \vdash^Z M : (E : t :: \zeta')} \text{ (\zeta-t-cons)}$$



$$\boxed{\vdash^Z (R, C, M, Q, ir)} \\
\boxed{\vdash^Z (R, C, M, Q, ir) : \Psi, \Gamma, \zeta} \\
\begin{array}{l}
Dom(\Psi) = Dom(C) \cup Dom(M_m) \\
M = M_s \#^{Dom(C)} M_m \\
\Psi \vdash C \\
\forall c \neq Z. ir \neq \cdot \implies C(R(pc_c)) = ir \\
\forall c \neq Z. \Psi(R(pc_c)) = (\Delta; \Gamma; (\overline{E_d, E_s}); E_m; \zeta) \rightarrow void \\
\exists S. \cdot \vdash S : \Delta \\
\Psi \vdash^Z M_s : S(\zeta) \\
\Psi \vdash^Z (M_m, Q) : (S(E_m), S(\overline{(E_d, E_s)})) \\
\Psi \vdash^Z R : S(\Gamma) \\
K = extractK(R, \Gamma), extractK(M_h), extractK(M_s, \zeta)
\end{array} \\
\hline
\vdash^Z (R, C, M, Q, ir) : K \quad (\Sigma-t)$$

## C.4 Simulation of Machine States

### C.4.1 Color Extraction

$$\begin{array}{ll}
extended\ color\ k & ::= c \mid GB \mid none \\
coloring\ K & ::= \cdot \mid a \mapsto k \mid \ell \mapsto k
\end{array}$$

$$\begin{array}{ll}
extractK_t(\langle c, b, E \rangle) & = c \\
extractK_t(E_z = 0 \Rightarrow \langle c, b, E \rangle) & = c \\
extractK_t(ns) & = none \\
\\ 
extractK(R, \Gamma) & = \forall a \in R. a \in Dom(\Gamma) ? \\
& \quad a \mapsto extractK_t(\Gamma(a)) \quad : \quad a \mapsto none \\
\\ 
extractK(M_s, \zeta) & = \forall \ell \in Dom(M_s). \cdot; \zeta \vdash \ell : t \Rightarrow \ell \mapsto extractK_t(t) \\
\\ 
extractK(M_h) & = \forall \ell \in Dom(M_h). \ell \mapsto GB
\end{array}$$

### C.4.2 Simulation

$$\boxed{k n_1 \text{ sim}^Z k n_2}$$

$$\frac{}{k n \text{ sim}^Z k n} \text{ (sim-val)} \quad \frac{}{c n \text{ sim}^c c n'} \text{ (sim-val-zap)}$$

$$\frac{}{\text{none } n \text{ sim}^Z \text{ none } n'} \text{ (sim-val-no-color)}$$

$$\boxed{K \vdash R \text{ sim}^Z R'}$$

$$\frac{\forall a. K(a) R(a) \text{ sim}^Z K(a) R'(a)}{K \vdash R \text{ sim}^Z R'} \text{ (sim-R)}$$

$$\boxed{K \vdash M \text{ sim}^Z M'}$$

$$\frac{\text{Dom}(M) = \text{Dom}(M') \quad \forall \ell \in \text{Dom}(M). K(\ell) M(\ell) \text{ sim}^Z K(\ell) M'(\ell)}{K \vdash M \text{ sim}^Z M'} \text{ (sim-M)}$$

$$\boxed{Q \text{ sim}^Z Q'}$$

$$\frac{}{\cdot \text{ sim}^Z \cdot} \text{ (sim-Q-empty)}$$

$$\frac{G n_1 \text{ sim}^Z G n'_1 \quad G n_2 \text{ sim}^Z G n'_2 \quad Q \text{ sim}^Z Q'}{((n_1, n_2), Q) \text{ sim}^Z ((n'_1, n'_2), Q')} \text{ (sim-Q)}$$

$$\boxed{\Sigma_1 \text{ sim}^Z \Sigma_2}$$

$$\begin{array}{c}
\vdash (R, C, M, Q, ir) : K \\
\vdash^Z (R', C, M', Q', ir) : K \\
K \vdash R \text{ sim}^Z R' \\
K \vdash M \text{ sim}^Z M' \\
Q \text{ sim}^Z Q' \\
\hline
(R, C, M, Q, ir) \text{ sim}^Z (R', C, M', Q', ir) \quad (\text{sim-}\Sigma)
\end{array}$$

# Appendix D

## ETAL<sub>FT</sub> Formal Results

This appendix gives expanded details on the formal results for ETAL<sub>FT</sub> that are summarized in Section 3.3. All lemmas and theorems are included, along with proof sketches. The complete proofs in note form appear in the companion technical report [49].

### D.1 Modified Lemmas

Many of the lemmas used in Appendix A to prove properties of TAL<sub>FT</sub> must be modified for use with ETAL<sub>FT</sub>. A selection of the modified lemmas is provided below to illustrate the types of changes that are necessary.

Substituting an expression of kind  $\kappa$  for a free variable of type  $\kappa$  preserves typing. (*Modifications: Parts 5 and 7 contain the stack type and also to substitute in to the instruction as well (since malloc contains a base type). Parts 6 and 7 hold for incomplete substitutions as  $\Delta$  is extended while type checking malloc.*)

**Lemma 34 (Substitution Lemma)**

1. If  $\Delta, x : \kappa \vdash E' : \kappa'$  and  $\Delta \vdash E : \kappa$  then  $\Delta \vdash E'[E/x] : \kappa'$ .
2. If  $\Delta, x : \kappa \vdash E_1 = E_2$  and  $\Delta \vdash E : \kappa$  then  $\Delta \vdash E_1[E/x] = E_2[E/x]$ .
3. If  $\Delta, x : \kappa \vdash E_1 \neq E_2$  and  $\Delta \vdash E : \kappa$  then  $\Delta \vdash E_1[E/x] \neq E_2[E/x]$ .
4. If  $\Psi; \Delta, x : \kappa \vdash^Z v : t$  and  $\Delta \vdash E : \kappa$  then  $\Psi; \Delta \vdash^Z v : t[E/x]$ .
5. If  $\Psi; (\Delta, x : \kappa; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash^Z ir \Rightarrow RT$  and  $\Delta \vdash E : \kappa$   
then  $\Psi; (\Delta; \Gamma[E/x]; \overline{(E_d, E_s)}[E/x]; E_m[E/x]; \zeta[E/x]) \vdash^Z ir[E/x] \Rightarrow RT[E/x]$ .
6. If  $\Delta' \vdash S : \Delta$  and  $\Psi; \Delta \vdash^Z v : t$  then  $\Psi; \Delta' \vdash^Z v : S(t)$ .
7. If  $\Delta' \vdash S : \Delta$  and  $\Psi; (\Delta; \Gamma; \overline{(E_d, E_s)}; E_m; \zeta) \vdash^Z ir \Rightarrow (\Delta; \Gamma'; \overline{(E'_d, E'_s)}; E'_m; \zeta')$   
then  $\Psi; (\Delta'; S(\Gamma); S(\overline{(E_d, E_s)}); S(E_m)) \vdash^Z S(ir) \Rightarrow (\cdot; S(\Gamma'); S(\overline{(E'_d, E'_s)}); S(E'_m); S(\zeta'))$ .

*Proof* Parts 1, 4, and 5 – by induction on the respective typing derivation. Parts 2 and 3 – by inspection of the equality judgment definition. Parts 6 and 7 – by induction on the size of  $\Delta$ , using parts 4 and 5 respectively. ■

The type of a value gives us information about the shape of the value. (*Modifications: The value's color tag is removed. The new memory typing judgment including the queue is used instead of the old judgment. Case 4 handles the initialization flag. Cases 7 and 8 are new. Note that the existence of a stack pointer tells nothing about the contents of the location it refers to.*)

**Lemma 35 (Canonical Forms)**

If  $\Psi; \cdot \vdash^Z n : t$  and  $Dom(\Psi) = Dom(C) \cup Dom(M)$  and  $\Psi \vdash^Z (M, Q) : (E_m, \overline{(E_d, E_s)})$  and  $\Psi \vdash C$ , then

1. If  $t = \langle c, b, E \rangle$  or  $t = (E' = 0) \Rightarrow \langle c, b, E \rangle$  and  $c = Z$  then no particular properties of  $n$  are known.
2. If  $t = \langle c, \text{int}, E \rangle$  and  $c \neq Z$  then  $\cdot \vdash E = n$ .
3. If  $t = \langle c, \Theta \rightarrow \text{void}, E \rangle$  and  $c \neq Z$  then  $\Psi(n) = \Theta \rightarrow \text{void}$  and  $n \in \text{Dom}(C)$  and  $\cdot \vdash E = n$  and  $n \neq 0$ .
4. If  $t = \langle c, b \text{ ref}^\varphi, E \rangle$  and  $c \neq Z$  then
  - (a)  $\Psi(n) = b \text{ ref}^{\varphi'}$  and  $\varphi' \leq \varphi$ , and
  - (b)  $n \in \text{Dom}(M)$ , and
  - (c)  $\cdot \vdash E = n$ , and
  - (d)  $\varphi' = \text{ref}^{\frac{1}{2}} \wedge Z \neq G \implies (\exists (n, v) \in Q) \wedge (\forall (n, v) \in Q. \Psi \vdash v : b)$
  - (e)  $\varphi' = \text{ref}^1 \implies \Psi \vdash M(\ell) : b \wedge (Z \neq G \implies \forall (n, v) \in Q. \Psi \vdash v : b)$
5. If  $t = (E' = 0) \Rightarrow t'$  and  $c \neq Z$  and  $\cdot \vdash E' = 0$  then  $n \neq 0$ .
6. If  $t = (E' = 0) \Rightarrow t'$  and  $c \neq Z$  and  $\cdot \vdash E' \neq 0$  then  $n = 0$ .
7. If  $t = ns$  then no particular properties of  $n$  are known.
8. If  $t = \langle c, \text{sptr}, E \rangle$  and  $c \neq Z$  then  $\cdot \vdash E = n$

*Proof* By induction on the structure of  $\Psi; \cdot \vdash^Z n : t$ . ■

If a value has a type, and this type has a supertype, then the value also has the supertype.

(Modifications: Part 3 is new.)

### Lemma 36 (Subtyping)

1. If  $\Psi; \Delta \vdash^Z v : t$  and  $\Delta \vdash t \leq t'$  then  $\Psi; \Delta \vdash^Z v : t'$ .

2. If  $\Psi \vdash^Z R : \Gamma$  and  $\Delta \vdash \Gamma \leq \Gamma'$  then  $\Psi \vdash^Z R : \Gamma'$ .

3. If  $\Psi \vdash^Z M : \zeta$  and  $\Delta \vdash \zeta \leq \zeta'$  then  $\Psi \vdash^Z M : \zeta'$ .

*Proof* Part 1 is by induction on the derivation of  $\Psi; \Delta \vdash^Z v : t$ . Parts 2 and 3 are by inversion and reconstruction of the appropriate typing rules and using Part 1. ■

If a judgment holds under the empty zap tag, then that judgment holds for a colored zap tag as well. (*Modifications: Part 3 is new.*)

### Lemma 37 (Color Weakening)

1. If  $\Psi; \cdot \vdash v : t$  then  $\forall c. \Psi; \cdot \vdash^c v : t$ .

2. If  $\Psi \vdash R : \Gamma$  then  $\forall c. \Psi \vdash^c R : \Gamma$ .

3. If  $\Psi \vdash M : \zeta$  then  $\forall c. \Psi \vdash^c M : \zeta$ .

*Proof* Part 1 by induction on the value typing judgment. The remaining parts are by inversion and reconstruction of the judgment using Part 1 as necessary. ■

## D.2 New Lemmas For Stacks

If an expression in a stack type is referred to explicitly and memory  $M$  has that stack type, then there is some location  $\ell$  equal to that expression that exists in the domain of  $M$ .

### Lemma 38 (Valid Stack Location)

1. If  $\Psi \vdash^Z M : \zeta$  and  $\cdot; \zeta \vdash E' : t$  then  $\exists \ell. \cdot \vdash E' = \ell$  and  $\ell \in \text{Dom}(M)$ .

2. If  $\Psi \vdash^Z M : \zeta$  and  $\cdot \vdash \zeta[E' \mapsto t] = \zeta'$  then  $\exists \ell. \cdot \vdash E' = \ell$  and  $\ell \in \text{Dom}(M)$ .

3. If  $\Psi \vdash^Z M : \zeta$  and  $\zeta = E_1 : t_1 :: \dots :: E_n : t_n :: \zeta'$  then  $\exists \ell. \cdot \vdash E_i = \ell_i$  and  $\ell_i \in \text{Dom}(M)$ .

*Proof* By repeated inversion of  $\zeta$ -*t-cons* and  $\zeta$ -*t-base*, all expressions  $E$  on the spine of  $\zeta$  are equal to a location  $\ell$  in the domain of  $M$ . By  $\zeta$ -*update* and  $\zeta$ -*lookup*,  $E'$  is equal to an expression on the spine of  $\zeta$ . ■

In addition, if an expression is looked up on the stack and is equal to a location, then the contents of that location have the type given by the stack type.

**Lemma 39 (Stack Lookup)**

If  $\Psi \vdash^Z M : \zeta$  and  $\cdot; \zeta \vdash E : t$  and  $\cdot \vdash E = \ell$  then  $\Psi; \cdot \vdash^Z M(\ell) : t$ .

*Proof* By induction on the structure of  $\cdot; \zeta \vdash E : t$ . ■

If a memory has a stack type and we update that stack type, then the memory updated with a value of the same type has the new stack type.

**Lemma 40 (Stack Update)**

If  $\Psi \vdash^Z M : \zeta$  and  $\cdot \vdash \zeta[E' \mapsto t] = \zeta'$  and  $\cdot \vdash E = \ell$  and  $\Psi; \cdot \vdash^Z n : t$  then  $\Psi \vdash^Z M[\ell \mapsto n] : \zeta'$ .

*Proof* By induction on the structure of  $\cdot \vdash \zeta[E' \mapsto t] = \zeta'$ . ■

### D.3 New Lemmas For Dynamic Memory Allocation

Judgments that hold with  $\Psi$  continue to hold after adding extra information to  $\Psi$ .

**Lemma 41 (Heap Extension Lemma)**

1. If  $\Psi; \cdot \vdash^Z v : t$  and  $n \notin \text{Dom}(\Psi)$  then  $\Psi, n \mapsto b; \cdot \vdash^Z v : t$
2. If  $\Psi \vdash^Z M : \zeta$  and  $n \notin \text{Dom}(\Psi)$  then  $\Psi, n \mapsto b \vdash^Z M : \zeta$
3. If  $\Psi; M; Q \vdash^Z \ell : b \text{ ref}^\Phi$  and  $n \notin \text{Dom}(\Psi)$  then  $\Psi, n \mapsto b; M; Q \vdash^Z \ell : b \text{ ref}^\Phi$



4. If  $\Psi; \Theta \vdash ir \Rightarrow RT$  and  $n \notin \text{Dom}(\Psi)$  then  $\Psi, n \mapsto b; \Theta \vdash ir \Rightarrow RT$

*Proof* By induction on the appropriate derivations. ■

Judgments that hold with  $\Psi$  continue to hold after updating an entry in  $\Psi$  to contain a subtype of the original value.

**Lemma 42 ( $\Psi$  Subtyping Lemma)**

1. If  $\Psi; \cdot \vdash^Z v : t$  and  $b' \leq \Psi(n)$  then  $\Psi[n \mapsto b']; \cdot \vdash^Z v : t$
2. If  $\Psi; M; Q \vdash^Z \ell : b \text{ ref}^\Phi$  and  $b' \leq \Psi(n)$  then  $\Psi[n \mapsto b']; M; Q \vdash^Z \ell : b \text{ ref}^\Phi$
3. If  $\Psi; \Theta \vdash ir \Rightarrow RT$  and  $b' \leq \Psi(n)$  then  $\Psi[n \mapsto b']; \Theta \vdash ir \Rightarrow RT$

*Proof* By case analysis of the appropriate derivations. Insert *addr-subtp-t* where the derivation uses  $\Psi(n)$ . ■

Judgments that hold with a substitution  $S$  continue to hold after adding extra information to  $S$ .

**Lemma 43 (Substitution Extension Lemma)**

1. If  $\Psi; \cdot \vdash^Z v : S(t)$  and  $\cdot \vdash S : \Delta$  and  $x \notin \Delta$  then  $\Psi; \cdot \vdash^Z v : (S, E/x)(t)$
2. If  $\Psi \vdash^Z M : S(\zeta)$  and  $\cdot \vdash S : \Delta$  and  $x \notin \Delta$  then  $\Psi \vdash^Z M : (S, E/x)(\zeta)$
3. If  $\Psi \vdash^Z Q : S(\overline{(E_d, E_s)})$  and  $\cdot \vdash S : \Delta$  and  $x \notin \Delta$  then  $\Psi \vdash^Z Q : (S, E/x)(\overline{(E_d, E_s)})$
4. If  $\Psi \vdash^Z R : S(\Gamma)$  and  $\cdot \vdash S : \Delta$  and  $x \notin \Delta$  then  $\Psi \vdash^Z R : (S, E/x)(\Gamma)$
5. If  $M = \llbracket S(E_m) \rrbracket$  and  $\cdot \vdash S : \Delta$  and  $x \notin \Delta$  then  $M = \llbracket (S, E/x)(E_m) \rrbracket$

*Proof* By induction on the appropriate derivations. ■

## D.4 Type Safety

The statements of Progress and Preservation do not change, but the proofs are modified significantly.

### Theorem 44 (Progress)

1. If  $\vdash \Sigma$  then  $\Sigma \longrightarrow_0^s \Sigma'$  and  $\Sigma' \neq \text{fault}$ .
2. If  $\vdash^c \Sigma$  then  $\Sigma \longrightarrow_0^s \Sigma$ .

*Proof* By case analysis on the instruction  $ir$  in state  $\Sigma$ . (*Modifications: The existing cases require minor modifications to remove the color tags and use the modified typing judgments. Cases are added for the new instructions.*) ■

### Theorem 45 (Preservation)

1. If  $\vdash^Z \Sigma$  and  $\Sigma \longrightarrow_0^s \Sigma'$  and  $\Sigma' \neq \text{fault}$  then  $\vdash^Z \Sigma'$ .
2. If  $\vdash \Sigma$  and  $\Sigma \longrightarrow_1^s \Sigma'$  then  $\exists c. \vdash^c \Sigma'$ .

*Proof* By case analysis of the structure of the derivation  $\Sigma \longrightarrow_k^s \Sigma'$ . (*Part 1 Modifications: The new structure of the top level typing rules requires more significant changes to the existing cases. The new cases for stack load and store use Lemmas 38 (Valid Stack Location), 39 (Stack Update), and 40 (Stack Lookup). The rules for stack allocation and deallocation require proof that the affected locations are part of the stack and not the heap. The rule for memory allocation adds new information to  $\Psi$ , and heap stores modify entries in  $\Psi$  to contain subtypes. These rules use Lemmas 41 (Heap Extension), 42 ( $\Psi$  Subtyping), and 43 (Substitution Extension). Part 2 Modifications: The main change is handling faults to registers with no corresponding colors information. In this case, the resulting state can be typed using any zap tag.*) ■

## D.5 Fault Tolerance

The main change affecting the fault tolerance results comes from the reorganization of the simulation relation explained in Section 3.3.2. This change affects the statements of the lemmas and theorems since the meaning of the  $sim^Z$  relation now includes typing information, but actually has little significant impact on the structure of the proofs.

The separation of the value of memory into a heap and a stack changes the definition of "observable" memory. The new relationships  $s' \stackrel{\ell}{\simeq} s$  and  $s' \stackrel{\ell}{\preceq} s$  are used to relate two output sequences of address-value pairs. Because the faulty computation may store faulty values into the stack portion of memory, we can no longer use simple equality to compare address-value pairs. The judgment  $\stackrel{\ell}{\simeq}$  say that the addresses in the two sequences are equal, and for all addresses greater than  $\ell$ , the values are also equal. The judgment  $s' \stackrel{\ell}{\preceq} s$  is similar, but only requires that the locations in the first sequence are a subsequence of those in the second sequence. In other words, when given  $\max(Dom(C))$  as  $\ell$ , these judgments check that the stores committed to the heap are identical and the stores committed to the stack are to the same locations, though the values may differ.

The lemmas and theorems are restated with these modifications.

### Lemma 46 (Singlestep Fault Detection)

If  $\Sigma \text{ sim}^c \Sigma_f$  and  $\Sigma \xrightarrow{s}_0 \Sigma'$  then  $\Sigma_f \xrightarrow{s_f}_0 \Sigma'_f$  and either

1.  $\Sigma'_f \text{ sim}^c \Sigma_f$  and  $s' \stackrel{\max(Dom(\Sigma.C))}{\simeq} s$ , or
2.  $\Sigma'_f = \text{fault}$  and  $s_f = ()$ .

*Proof* By case analysis of  $\Sigma \xrightarrow{s}_0 \Sigma'$ . (Modifications: Other than the slight restructurings to access the color information, the existing cases are essentially unmodified. New cases are added for each of the new singlestep rules.) ■

**Lemma 47 (Multistep Fault Detection)**

If  $\Sigma \text{ sim}^c \Sigma_f$  and  $\Sigma \xrightarrow{n}_0^s \Sigma'$  then either

1.  $\Sigma_f \xrightarrow{n}_0^{s_f} \Sigma'_f$  and  $\Sigma' \text{ sim}^c \Sigma'_f$  and  $s_f \stackrel{\max(\text{Dom}(\Sigma.C))}{\simeq} s$ , or
2. Exists  $m \leq n$ .  $\Sigma_f \xrightarrow{m}_0^{s_f} \text{fault}$  and  $s_f \stackrel{\max(\text{Dom}(\Sigma.C))}{\preceq} s$ .

*Proof* By induction on the structure of  $\Sigma \xrightarrow{n}_0^s \Sigma'$ . (Modifications: Minor changes to support the new output comparisons.) ■

**Lemma 48 (Fault Similarity)**

If  $\Sigma \xrightarrow{1} \Sigma_f$ , then  $\exists c. \Sigma \text{ sim}^c \Sigma'$ .

*Proof* By case analysis on the definition of  $\Sigma \xrightarrow{1} \Sigma'$ . (Modifications: Minor changes to support the changes to the simulation relation.) ■

**Theorem 49 (Fault Tolerance)**

If  $\vdash \Sigma$  and  $\Sigma \xrightarrow{n}_0^s \Sigma'$  then either  $\Sigma \xrightarrow{(n+1)}_1^{s'} \Sigma'_f$

or  $\exists m \leq (n+1) . \Sigma \xrightarrow{m}_1^{s'} \text{fault}$ , and

1. For all derivations  $\Sigma \xrightarrow{(n+1)}_1^{s'} \Sigma'_f$  where  $\Sigma'_f \neq \text{fault}$ .  
 $s' \stackrel{\max(\text{Dom}(C))}{\simeq} s$  and  $\exists c. \Sigma' \text{ sim}^c \Sigma'_f$ .
2. For all derivations  $\Sigma \xrightarrow{m}_1^{s'} \text{fault}$  where  $m \leq (n+1)$ .  
 $s' \stackrel{\max(\text{Dom}(C))}{\preceq} s$ .

*Proof* By case analysis on the definition of  $\Sigma \xrightarrow{n} {}^s_0\Sigma'$ . (*Modifications: Minor changes to support the new output comparisons.*) ■

# Appendix E

## MiniC to ETAL<sub>FT</sub> Translation

This appendix provides the complete rules for the translation from MiniC to ETAL<sub>FT</sub> and proof sketches of the formal results. A condensed version is given in Section 3.4, and the full proofs are available in [49].

### E.1 Overview

In order to simplify the translation, it uses the designated registers  $pc_G$ ,  $pc_B$ ,  $sp_G$ ,  $sp_B$ , and  $gd$  and then as many fresh temporary registers  $t_1, t_2, \dots$  as needed. Many of these temporary registers can be easily removed by coalescing move instructions. Section 3.6 discusses how to support register allocation if the number of temporary registers is greater than the number of actual registers.

The translation uses a simplistic calling convention. In order to support fault tolerance, all function arguments and return values need to be duplicated. Arguments are passed on the stack, with the last argument pushed on first. There will be two assembly-level arguments for each MiniC argument, and the green argument of each pair is always

below the corresponding blue argument on the stack. Below the arguments are the blue and green copies of the return address. When a function returns, it pops the return address and all the arguments and pushes two copies of the return value.

On function entry, each function loads all the arguments into temporary registers. When making a function call, all temporary registers that correspond to local variables are spilled to the stack before the arguments and return address are pushed. After the call returns, the registers for the local variables are restored and then the return values are moved to their destinations.

At a high level, the translation works by passing around a code memory  $C$  and continually accumulating new instructions onto the end. The notation  $C@i$  is used to append instruction  $i$  onto the end of code memory  $C$ . In addition, many judgments track the following additional information:

$n$  *the number of temporary registers required so far.  $\text{freshReg}(n)$  generates two fresh temporary registers and increments  $n$  appropriately.*

$V$  *a mapping from MiniC variables  $x$  to pairs of registers  $(r_g, r_b)$ .  $V_g(x)$  refers to the first component of the pair and  $V_b(x)$  refers to the second.*

$B$  *a mapping from function variables to addresses in code memory.*

Because ETAL<sub>FT</sub> has no notion of termination, the program "completes" by jumping to a designated label  $l_{halt}$  which contains four instructions that create an infinite loop.

## E.2 Translating Typing Information

### E.2.1 Type Translation

The type translation  $\llbracket \tau \rrbracket = b$  in Figure E.1 translates a MiniC type  $\tau$  into a ETAL<sub>FT</sub> base type  $b$ . The translation of integers and references is straightforward. The only thing to note is that the translation of a MiniC reference is a fully initialized ETAL<sub>FT</sub> reference.

The translation of the function types into code types is quite a bit more complicated. First, the two components of the function type are used to generate the type of the stack on entry to the function (defined in a moment) and the base type  $b_\tau$  of the return value.

Next, we choose a number of fresh expression variables. At function entry, the memory is described by  $\alpha_m$ , the program counters by  $\alpha_p$ , and the return addresses by  $\alpha_r$ .  $\alpha'_m$  and  $\alpha_\tau$  will be used to describe the return memory and return values.

The contexts with a subscript  $r$  describe the return state of the function. When the function returns, the two program counters will contain the return address. The arguments and return address will have been removed from the stack and the return values pushed on in their place.

On function entry, the top of the stack contains two copies of the return address, followed by the stack type constructed using the function arguments. The stack pointers point to the top of the stack, the program counters are equal, and the destination register is zero.

### E.2.2 Context Generation

The typing information can be used to generate a number of different ETAL<sub>FT</sub> contexts.



$$\boxed{\llbracket \tau \rrbracket = b}$$

$$\overline{\llbracket int \rrbracket = int} \quad (trans-int)$$

$$\overline{\llbracket b \text{ ref} \rrbracket = b \text{ ref}^1} \quad (trans-ref)$$

$$\begin{aligned} \llbracket A \rrbracket &= \alpha_{\ell-2*n} : \langle G, b_1, \alpha_1 \rangle \quad :: \alpha_{\ell-2*n-1} : \langle B, b_1, \alpha_1 \rangle \\ &\quad :: \dots \\ &\quad :: \alpha_{\ell-2} : \langle G, b_n, \alpha_n \rangle \quad :: \alpha_{\ell-1} : \langle B, b_n, \alpha_n \rangle \\ &\quad :: \alpha_{\ell} : \alpha_{\sigma} \end{aligned}$$

$$\llbracket \tau \rrbracket = b_{\tau}$$

$\alpha_m, \alpha_p, \alpha_r, \alpha_{\tau}, \alpha'_m$  fresh

$$\Delta_r = \alpha'_m : \kappa_{mem}, \alpha_{\tau} : \kappa_{int}$$

$$\begin{aligned} \Gamma_r &= pc_G \rightarrow \langle G, int, \alpha_r \rangle, sp_G \rightarrow \langle G, sptr, (\alpha_{\ell-2}) \rangle, \\ &\quad pc_B \rightarrow \langle B, int, \alpha_r \rangle, sp_B \rightarrow \langle B, sptr, (\alpha_{\ell-2}) \rangle, \\ &\quad gd \rightarrow \langle G, int, 0 \rangle \end{aligned}$$

$$\zeta_r = (\alpha_{\ell-2}) : \langle G, b_{\tau}, \alpha_{\tau} \rangle \quad :: (\alpha_{\ell-1}) : \langle B, b_{\tau}, \alpha_{\tau} \rangle \quad :: \alpha_{\ell} : \alpha_{\sigma}$$

$$\Theta_r = (\Delta_r, \Gamma_r, (), \alpha'_m, \zeta_r)$$

$$\begin{aligned} \Delta &= \alpha_m : \kappa_{mem}, \alpha_p : \kappa_{int}, \alpha_r : \kappa_{int}, \alpha_{\ell} : \kappa_{int}, \alpha_{\sigma} : \kappa_{\sigma}, \\ &\quad \alpha_1 : \kappa_{int}, \dots, \alpha_n : \kappa_{int} \end{aligned}$$

$$\begin{aligned} \Gamma &= pc_G \rightarrow \langle G, int, \alpha_p \rangle, sp_G \rightarrow \langle G, sptr, \alpha_{\ell-2*n-2} \rangle, \\ &\quad pc_B \rightarrow \langle B, int, \alpha_p \rangle, sp_B \rightarrow \langle B, sptr, \alpha_{\ell-2*n-2} \rangle, \\ &\quad gd \rightarrow \langle G, int, 0 \rangle \end{aligned}$$

$$\begin{aligned} \zeta &= (\alpha_{\ell-2*n-2}) : \langle G, \Theta_r \rightarrow void, \alpha_r \rangle \\ &\quad :: (\alpha_{\ell-2*n-1}) : \langle B, \Theta_r \rightarrow void, \alpha_r \rangle \quad :: \llbracket A \rrbracket \end{aligned}$$

---


$$\llbracket A \rightarrow \tau \rrbracket = (\Gamma, \Delta, (), \alpha_m, \zeta) \rightarrow void \quad (trans-(X \rightarrow \tau))$$

Figure E.1: Type Translation.

The type of the stack on function entry is generated from the function arguments  $A$ . The arguments appear in reverse order. There are two copies of each argument, one green and one blue. The base type for each of these two copies is given by the type translation.

$$\boxed{\llbracket A \rrbracket = \zeta}$$

$$\frac{\alpha_\ell, \alpha_\sigma \text{ fresh}}{\llbracket \cdot \rrbracket = \alpha_\ell : \alpha_\sigma} \text{ (gen-}\zeta\text{-base)}$$

$$\frac{\llbracket \tau \rrbracket = b_\tau \quad \llbracket A \rrbracket = E_s : \sigma \quad \alpha_\tau \text{ fresh}}{\llbracket x : \tau, A \rrbracket = (E_s - 2) : \langle G, b_\tau, \alpha_\tau \rangle :: (E_s - 1) : \langle B, b_\tau, \alpha_\tau \rangle :: E_s : \sigma} \text{ (gen-}\zeta\text{)}$$

Using the function arguments  $A$  and the corresponding stack type  $\zeta$ , we can generate the register file type that results after loading each of these arguments into the register specified by  $V$ .

$$\boxed{\llbracket A; \zeta \rrbracket_V = \Gamma}$$

$$\frac{}{\llbracket \cdot; \alpha_\ell : \alpha_\sigma \rrbracket_V = \cdot} \text{ (gen-}\Gamma\text{-args-emp)}$$

$$\frac{\llbracket A; E_s : \sigma \rrbracket_V = \Gamma \quad \Gamma' = \Gamma, V_g(x) \rightarrow t_G, V_b(x) \rightarrow t_B}{\llbracket x : \tau, A; (E_s - 2) : t_G :: (E_s - 1) : t_B :: E_s : \sigma \rrbracket = \Gamma'} \text{ (gen-}\Gamma\text{-args)}$$

Given a variable context  $X$  and a mapping  $V$  from each variable in  $X$  to a pair of registers,  $\llbracket X \rrbracket_V$  is the register file type that contains typing information for each of these registers. Each pair of registers is described by a different expression variable.

$$\boxed{\llbracket X \rrbracket_V = \Gamma}$$

$$\frac{}{\llbracket \cdot \rrbracket_V = \cdot} \text{ (gen-}\Gamma\text{-X-}\cdot\text{)}$$

$$\frac{\llbracket X \rrbracket_V = \Gamma \quad \llbracket \tau \rrbracket = b_\tau \quad \alpha_\tau \text{ fresh}}{\llbracket x : \tau, X \rrbracket = \Gamma, V_g(x) \rightarrow \langle G, b_\tau, \alpha_\tau \rangle, V_b(x) \rightarrow \langle B, b_\tau, \alpha_\tau \rangle} \text{ (gen-}\Gamma\text{-X)}$$

Finally, the judgment  $\llbracket A \rightarrow \tau; L \rrbracket_V = \Theta$  can be used to generate the static context used within the body of a function. It includes the information about function entry generated by  $\llbracket A \rightarrow \tau \rrbracket$ , plus the additional registers corresponding to the arguments and local definitions in  $L$ .

$$\boxed{\llbracket A \rightarrow \tau; L \rrbracket_V = \Theta}$$

$$\frac{\begin{array}{l} \llbracket A \rightarrow \tau \rrbracket = (\Delta, \Gamma, (), E_m, \zeta) \\ \Gamma' = \Gamma, \llbracket L \rrbracket_V, \llbracket A; \zeta \rrbracket_V \end{array}}{\llbracket A \rightarrow \tau; L \rrbracket_V = (\Delta, \Gamma', (), E_m, \zeta)} \text{ (gen-}\Theta\text{)}$$

### E.3 Extending Code Memory

As the translation progresses, it continually adds new code to the code memory  $C$ . We define a series of judgments below that can be used to state that the in-progress code memory has the properties we desire.

As a reminder, the original code typing judgment is shown below. It requires that for all addresses in the code memory, if code type  $\Theta \rightarrow \text{void}$  is assigned to that address, then  $\Theta$  can be used to type check the instruction in the address. In addition, if type checking the instruction gives a modified static context  $\Theta'$ , then the type assigned to the

next location is  $\Theta' \rightarrow \text{void}$ . In other words, the only time two adjacent instructions do not need to agree on the intermediate type, is when the first instruction is a blue jump.

$$\boxed{\Psi \vdash C}$$

$$\frac{\begin{array}{l} 0 \notin \text{Dom}(C) \\ \forall n \in \text{Dom}(C). \Psi(n) = \Theta \rightarrow \text{void} \wedge \Psi; \Theta \vdash C(n) \Rightarrow RT \wedge \\ (RT = \Theta' \text{ implies } \Psi(n+1) = \Theta' \rightarrow \text{void}) \end{array}}{\Psi \vdash C} \quad (C-t)$$

### E.3.1 Invariants Between Functions

The judgment  $B; F \vdash C : P$  describes a code memory that contains the translations of the functions in  $F$ . The code memory should be well-typed using  $C$ -t. In addition, the mapping  $B$  should map each function in  $F$  to an address that has the ETAL<sub>FT</sub> type corresponding to the function type. The only items in the type  $\Psi$  are those for the actual addresses in  $C$ . And finally, the designated label  $l_{\text{halt}}$  that contains the code for an infinite loop, should have the type provided.

$$\boxed{B; F \vdash C : P}$$

$$\frac{\begin{array}{l} \Psi \vdash C \\ \forall f \in F. \Psi(B(f)) = \llbracket F(f) \rrbracket \\ \text{dom}(\Psi) = \text{dom}(C) \\ \Psi(l_{\text{halt}}) = ( (ap : \kappa_{\text{int}}, \alpha_m : \kappa_{\text{mem}}, \alpha_\ell : \kappa_{\text{int}}, \alpha_\sigma : \kappa_\sigma) , \\ (pc_G \rightarrow \langle G, \text{int}, l_{\text{halt}} \rangle, pc_B \rightarrow \langle B, \text{int}, l_{\text{halt}} \rangle \\ sp_G \rightarrow \langle G, \text{sptr}, \alpha_\ell \rangle, sp_B \rightarrow \langle B, \text{sptr}, \alpha_\ell \rangle \\ gd \rightarrow \langle G, \text{int}, 0 \rangle), \\ (, \alpha_m, \alpha_\ell : \alpha_\sigma) \end{array}}{B; F \vdash C : \Psi} \quad \text{wf-}F$$

### E.3.2 Partial Code Memories

While we are in the process of generating code, we are often in the situation where the last instruction added is in the middle of a code block. Since the remaining instructions have not yet been added, the code memory is not well-typed according to  $C$ - $t$ .

Instead, we define the judgment below.  $\Psi \vdash C : \Theta$  states that all the addresses in code memory  $C$  are well-typed in the usual manner, but that the last instruction results in static context  $\Theta$ .  $\Psi$  already contains the typing information for the next instruction that will be added.

$$\boxed{\Psi \vdash C : \Theta}$$

$$\frac{\begin{array}{l} 0 \notin \text{dom}(C) \\ l_m = \max(\text{dom}(C)) \\ \text{dom}(\Psi) = \text{dom}(C) \cup (l_m + 1) \\ \forall l \in \text{dom}(C). \\ \quad \Psi(l) = \Theta \rightarrow \text{void} \wedge \\ \quad \Psi; \Theta \vdash C(l) \Rightarrow RT \wedge \\ \quad RT = \Theta' \implies \Psi(l+1) = \Theta' \rightarrow \text{void} \\ \Psi(l_m + 1) = \Theta \rightarrow \text{void} \end{array}}{\Psi \vdash C : \Theta} \quad C\text{-partial-}\Theta$$

### E.3.3 Invariants within a Function

In addition, there are some extra invariants which should hold within the body of a function. Before defining these, we need to define two sub-judgments.

A static context  $\Theta$  is a subtype of another static context  $\Theta'$  when all the corresponding elements have a subtyping relationship.

$$\boxed{\Theta \leq \Theta'}$$

$$\frac{\begin{array}{l} \exists S. \Delta \vdash S : \Delta' \\ \Delta \vdash \Gamma \leq S(\Gamma') \\ \Delta \vdash \overline{(E_d, E_s)} = S(\overline{(E'_d, E'_s)}) \\ \Delta \vdash E_m = S(E'_m) \\ \Delta \vdash \zeta \leq S(\zeta') \end{array}}{\Delta, \Gamma, \overline{(E_d, E_s)}, E_m, \zeta \leq \Delta', \Gamma', \overline{(E'_d, E'_s)}, E'_m, \zeta'} \Theta\text{-subtp}$$

The judgment  $X; V \vdash \Theta$  wf states that the static context  $\Theta$  is consistent with program variables in  $X$  and the mapping  $V$ . In other words, for every variable in  $X : \tau$ , the two registers in  $V(x)$  have base type  $\llbracket \tau \rrbracket$ , the first is green and the second is blue, and they are described by equivalent expressions.

$$\boxed{X; V \vdash \Theta \text{ wf}}$$

$$\frac{\begin{array}{l} \Theta = (\Delta, \Gamma, seq, E_m, \zeta) \\ \forall x : \tau \in X. \\ V(x) = (r, r') \wedge \\ \Gamma(r) = \langle G, \llbracket \tau \rrbracket, E \rangle \\ \Gamma(r') = \langle B, \llbracket \tau \rrbracket, E' \rangle \\ \Delta \vdash E = E' \end{array}}{X; V \vdash \Theta \text{ wf}} \Theta\text{-V-wf}$$

The judgment  $f; V; B; F; A; L \vdash C : \Psi$  should hold for any code memory  $C$  used in the process of translating function  $f$ . There are five components to this judgment

- The code memory is well formed and ends with some context  $\Theta$ .
- $\Theta$  is consistent with the arguments and local variables.

- $\Theta$  is a subtype of the type obtained by translating the function type and the local declarations. (In other words, all the information from the function entry still holds, though there may be additional information about temporary registers.)
- All functions in the function context have correctly typed start addresses.
- The halt label  $l_{halt}$  has the appropriate type.

$$\boxed{f;V;B;F;A;L \vdash C : \Psi}$$

$$\begin{array}{l}
\Psi \vdash C : \Theta \\
(A \cup L); V \vdash \Theta \text{ wf} \\
\Theta \leq \llbracket F(f); L \rrbracket_V \\
\forall f' \in F. \Psi(B(f)) = \llbracket F(f') \rrbracket \\
\Psi(l_{halt}) = ( \quad ap : \kappa_{int}, \alpha_m : \kappa_{mem}, \alpha_\ell : \kappa_{int}, \alpha_\sigma : \kappa_\sigma \quad ), \\
\quad ( pc_G \rightarrow \langle G, int, l_{halt} \rangle, \quad pc_B \rightarrow \langle B, int, l_{halt} \rangle \\
\quad \quad sp_G \rightarrow \langle G, sptr, \alpha_\ell \rangle \quad sp_B \rightarrow \langle B, sptr, \alpha_\ell \rangle \\
\quad \quad gd \rightarrow \langle G, int, 0 \rangle ), \\
\quad ( ), \alpha_m, \alpha_\ell : \alpha_\sigma ) \\
\hline
f;V;B;F;A;L \vdash C : \Psi \quad \text{wf-fun-body}
\end{array}$$

### E.3.4 Adding Instructions to Code Memory

The Code Addition Lemma states that if we have a code memory that is well-typed ending in  $\Theta$  and a sequence of instructions, the first of which is well-typed given  $\Theta$  and who agree on the interleaving static contexts, then the sequence of instructions can be appended to the code memory, and the result will be the static context generated by the last instruction.

**Lemma 50 (Code Addition)**

If  $\Psi \vdash C : \Theta_1$  and  $\Psi; \Theta_1 \vdash i_1 : \Theta_2$  and  $\Psi; \Theta_2 \vdash i_2 : \Theta_3$  and ... and  $\Psi; \Theta_n \vdash i_n : \Theta'$  then  $\Psi' \vdash C @ i_1 @ \dots @ i_n : \Theta'$  where  $l_m = \max(\text{dom}(\Psi))$  and  $\Psi' = \Psi', l_m + 1 \rightarrow (\Theta_2 \rightarrow \text{void}), l_m + 3 \rightarrow (\Theta_3 \rightarrow \text{void}), \dots, l_m + n \rightarrow (\Theta' \rightarrow \text{void})$

*Proof* By inversion and reconstruction of  $C$ -*partial-t*. ■

The Block Extension Lemma is similar except the added sequence of instructions begins a new block instead of appending on to a block in progress.

**Lemma 51 (Block Extension)**

If  $\Psi \vdash C$  and  $\Psi; \Theta_1 \vdash i_1 : \Theta_2$  and  $\Psi; \Theta_2 \vdash i_2 : \Theta_3$  and ... and  $\Psi; \Theta_n \vdash i_n : \Theta'$  then  $\Psi' \vdash C @ i_1 @ \dots @ i_n : \Theta'$  where  $l_m = \max(\text{dom}(\Psi))$  and  $\Psi' = \Psi', l_m + 1 \rightarrow (\Theta_1 \rightarrow \text{void}), \dots, l_m + n \rightarrow (\Theta_n \rightarrow \text{void}), l_m + n + 1 \rightarrow (\Theta' \rightarrow \text{void})$

*Proof* Similar to Lemma 50; deconstruct  $C$ - $t$  to build first  $C$ -*partial-t*. ■

## E.4 Value Translation

The value translation appends code to the existing code memory that moves the value into the two registers  $r$  and  $r'$ . For a constant value  $n$ ,  $n$  is moved into two fresh registers. For variables, it simply returns the existing register pair, and does not add any new instructions. References are created by allocating a new memory location, initializing that location, and then returning the two registers contain the new address.

$$\boxed{\llbracket X \vdash v : \tau \rrbracket C n V = C' n' r r'}$$



$$\frac{\begin{array}{l} (n', r, r') = \text{freshReg}(n) \\ C' = C @ \text{mov } r \ n @ \text{mov } r' \ n \end{array}}{\llbracket X \vdash n : \text{int} \rrbracket C \ n \ V = C' \ n' \ r \ r'} \quad (\text{trans-v-n})$$

$$\frac{}{\llbracket X \vdash x : X(x) \rrbracket C \ n \ V = C \ n \ V_g(x) \ V_b(x)} \quad (\text{trans-v-x})$$

$$\frac{\begin{array}{l} \llbracket X \vdash v : \tau \rrbracket C \ n \ V = C_v \ n_v \ r_v \ r'_v \\ \llbracket \tau \rrbracket = b_\tau \\ (n', r, r') = \text{freshReg}(n_v) \\ C' = C @ \text{malloc}[b_\tau] \ r \ r' @ \text{st}_G \ r \ r_v @ \text{st}_B \ r' \ r'_v \end{array}}{\llbracket X \vdash \text{ref } v : b \ \text{ref} \rrbracket C \ n \ V = C' \ n' \ r \ r'} \quad (\text{trans-v-ref})$$

The Value Translation Lemma states that after adding the instructions to translate the value, the types of the two destination registers are based on the type of the value. The resulting register file type is a subtype of the starting register file type because additional temporaries may be added in rule *trans-v-ref*.

**Lemma 52 (Value Translation)**

If  $\llbracket X \vdash v : \tau \rrbracket C \ n \ V = C' \ n' \ r \ r'$  and  $\Psi \vdash C : (\Delta, \Gamma, \text{seq}, E_m, \zeta)$  and  $X; V \vdash \Theta$  wf then  $\exists E'_m, \Psi'. \Psi' \vdash C' : (\Delta, \Gamma', \text{seq}, E_m, \zeta)$  and  $\Gamma' \leq \Gamma$  and  $\Gamma'(r) = \langle G, \llbracket \tau \rrbracket, E \rangle$  and  $\Gamma'(r') = \langle B, \llbracket \tau \rrbracket, E' \rangle$  and  $\Delta \vdash E = E'$

*Proof* By case analysis on  $\llbracket X \vdash v : \tau \rrbracket C \ n \ V = C' \ n' \ r \ r'$ . ■

## E.5 Value List Translation

The value list translation  $\llbracket X \vdash vs : X \rrbracket C n V n_p = C' n'$  is used to push the arguments to a call onto the stack. It assumes that the stack space has already been allocated.

$$\boxed{\llbracket X \vdash vs : X \rrbracket C n V n_p = C' n'}$$

$$\frac{}{\llbracket X \vdash \cdot : \cdot \rrbracket C n V n_p = C n} \text{ (trans-vs-.)}$$

$$\frac{\begin{array}{l} \llbracket X \vdash v : \tau \rrbracket C n V = C_v n_v r_v r'_v \\ C'_v = C_v @ \text{sst}(2 * n_p) r_v @ \text{sst}(2 * n_p + 1) r'_v \\ \llbracket X \vdash v, vs : x : \tau, X \rrbracket C'_v n_v V (n_p + 1) = C' n' \end{array}}{\llbracket X \vdash v, vs : x : \tau, X \rrbracket C n V n_p = C' n'} \text{ (trans-vs)}$$

The Value List Translation Lemma states that after the translation, the resulting stack type contains the types of the arguments in the appropriate slots. The translation may generate additional temporary registers while getting the values to be stored, and so the resulting register file type may contain extra information.

### Lemma 53 (Value List Translation)

If  $\llbracket X \vdash vs : ps \rrbracket C n V P_{num} = C' n'$  and  $vs = x_1 : \tau_1, \dots, x_n : \tau_n$  and  $\Psi \vdash C : (\Delta, \Gamma, seq, E_m, \zeta)$  and  $\zeta = E_l : t :: \dots :: E_l + 2P_{num} : t_1 :: \dots :: E_l + 2(P_{num} + n - 1) : t_n :: \zeta''$  then  $\Psi \vdash C' : (\Delta, \Gamma', seq, E'_m, \zeta')$  where  $\zeta' = E_l : t :: \dots :: E_l + 2P_{num} : \langle G, \llbracket \tau_1 \rrbracket, E_1 \rangle :: \dots :: E_l + 2(P_{num} + n - 1) : \langle B, \llbracket \tau_n \rrbracket, E_n \rangle :: \zeta''$  and  $\Gamma' \leq \Gamma$ .

*Proof* By induction on the structure of  $\llbracket X \vdash vs : ps \rrbracket C n V P_{num} = C' n'$ . The recursive case uses Lemma 52 and then calls the Induction Hypothesis. ■

## E.6 Local Declaration Translation

The local declaration translation is a recursive judgment. For each declaration, the value is obtained and then moved into two fresh registers. The value mapping  $V$  is extended to map the variables to these registers.

$$\boxed{\llbracket F; A; L \vdash lds : L' \rrbracket C n V = C' n' V'}$$

$$\frac{}{\llbracket F; A; L \vdash \cdot : L \rrbracket C n V = C n V} \text{ (trans-lds-empty)}$$

$$\begin{array}{l}
 \llbracket (A \cup L) \vdash v : \tau \rrbracket C n V = C_v n_v r_v r'_v \\
 (n_{ld}, r, r') = \text{freshReg}(n_v) \\
 V_{ld} = V[x \mapsto (r, r')] \\
 C_{ld} = C_v @ \text{mov } r \ r_v @ \text{mov } r' \ r'_v \\
 \llbracket F; A; L[x : \tau] \vdash lds : L' \rrbracket C_{ld} n_{ld} V_{ld} = C' n' V' \\
 \hline
 \llbracket F; A; L \vdash \tau x = v; lds : L' \rrbracket C n V = C' n' V' \text{ (trans-lds)}
 \end{array}$$

The Local Declaration Translation Lemma states that the code for the translation of local variables  $lds$  consistently adds the types of those declarations.

### Lemma 54 (Local Decl. Translation)

$$\begin{array}{l}
 \text{If } \llbracket F; A; L \vdash lds : L' \rrbracket C n V = C' n' V' \\
 \quad f; V; B; F; A; L \vdash C : \Psi \\
 \text{then } \exists \Psi'. f; V'; B; F; A; L' \vdash C' : \Psi'
 \end{array}$$

*Proof* By induction on the structure of  $\llbracket F; A; L \vdash lds : L' \rrbracket C n V = C' n' V$ . The recursive case uses Lemma 52 and then calls the Induction Hypothesis. ■

$$\boxed{\llbracket F;A;L \vdash s \text{ wf} \rrbracket_f C n V B = C' n'}$$

$$\frac{\begin{array}{l} \llbracket F;A;L \vdash s_1 \text{ wf} \rrbracket_f C n V B = C_1 n_1 \\ \llbracket F;A;L \vdash s_2 \text{ wf} \rrbracket_f C_1 n_1 V B = C_2 n_2 \end{array}}{\llbracket F;A;L \vdash s_1; s_2 \text{ wf} \rrbracket_f C n V B = C_2 n_2} \text{ (trans-s-seq)}$$

$$\frac{\begin{array}{l} \llbracket (A \cup L) \vdash v : \tau \rrbracket C n V = C_v n_v r_v r'_v \\ C' = C_v @ \text{mov } V_G(x) r_v @ \text{mov } V_B(x) r'_v \end{array}}{\llbracket F;A;L \vdash x = v \text{ wf} \rrbracket_f C n V B = C' n_v} \text{ (trans-s-assign)}$$

$$\frac{\begin{array}{l} \llbracket (A \cup L) \vdash v_1 : \text{int} \rrbracket C n V = C_1 n_1 r_1 r'_1 \\ \llbracket (A \cup L) \vdash v_2 : \text{int} \rrbracket C_1 n_1 V = C_2 n_2 r_2 r'_2 \\ C' = C_2 @ \text{op } V_G(x) r_1 r_2 @ \text{op } V_B(x) r'_1 r'_2 \end{array}}{\llbracket F;A;L \vdash x = v_1 \text{ op } v_2 \text{ wf} \rrbracket_f C n V B = C' n_2} \text{ (trans-s-op)}$$

$$\frac{\begin{array}{l} \llbracket (A \cup L) \vdash v : \tau \text{ ref} \rrbracket C n V = C_v n_v r r' \\ C' = C_v @ \text{ld}_G V_G(x) r @ \text{ld}_B V_B(x) r' \end{array}}{\llbracket F;A;L \vdash x = !v \text{ wf} \rrbracket_f C n V B = C' n_v} \text{ (trans-s-deref)}$$

$$\frac{\begin{array}{l} \llbracket (A \cup L) \vdash v_1 : \tau \text{ ref} \rrbracket C n V = C_1 n_1 r_1 r'_1 \\ \llbracket (A \cup L) \vdash v_2 : \tau \rrbracket C_1 n_1 V = C_2 n_2 r_2 r'_2 \\ C' = C_v @ \text{st}_G r_1 r_2 @ \text{st}_B r'_1 r'_2 \end{array}}{\llbracket F;A;L \vdash v_1 := v_2 \text{ wf} \rrbracket_f C n V B = C' n_2} \text{ (trans-s-update)}$$

Figure E.2: Statement Translation - Basic Statements.

$$\begin{aligned}
& l_{start} = \max(\text{dom}(C)) + 1 \\
& \llbracket (A \cup L) \vdash v : \text{int} \rrbracket C \ n \ V = C_v \ n_v \ r_v \ r'_v \\
& (n_c, re, re', rs, rs') = \text{freshReg}(n_v) \\
& l_{fixend} = \max(\text{dom}(C_v)) + 1 \\
& C_c = C_v \ @ \ \text{mov } re \ l_{halt} \ @ \ \text{mov } re' \ l_{halt} \\
& \quad \quad \quad @ \ \text{bz}_G \ r_v \ re \quad @ \ \text{bz}_B \ r'_v \ re' \\
& \llbracket F; A; L \vdash s \ \text{wf} \rrbracket_f C_c \ n_c \ V = C_s \ n_s \\
& C'_s = C_s \ @ \ \text{mov } rs \ l_{start} \ @ \ \text{mov } rs' \ @ \ l_{start} \\
& \quad \quad \quad @ \ \text{jmp}_G \ rs \quad @ \ \text{jmp}_B \ rs' \\
& l_{end} = \max(\text{dom}(C'_s)) + 1 \\
& C' = C'_s \ \begin{array}{l} [l_{fixend} \quad \mapsto \ \text{mov } re \ l_{end}] \\ [l_{fixend} + 1 \quad \mapsto \ \text{mov } re' \ l_{end}] \end{array} \\
& n' = n_s \\
\hline
\llbracket F; A; L \vdash \text{while } v \ \text{do } s \ \text{wf} \rrbracket C \ n \ V \ B = C' \ n' \quad (\text{trans-s-while})
\end{aligned}$$

Figure E.3: While Statement Translation.

$$\begin{aligned}
& l_c = \max(\text{dom}(C)) + 1 \\
& \llbracket (A \cup L) \vdash v : \text{int} \rrbracket C \ n \ v = C_v \ n_v \ r_v \ r'_v \\
& (n_c, rt, rf, rjt, rjf, rt', rf', rjt', rjf', re, re') = \text{freshReg}(n_v) \\
& l_{fixtrue} = \max(\text{dom}(C_v)) + 1 \\
& C_c = C_v \ @ \ \text{mov } rt \ l_{halt} \ @ \ \text{mov } rt' \ l_{halt} \\
& \quad \quad \quad @ \ \text{bz}_G \ r_v \ rt \quad @ \ \text{bz}_B \ r'_v \ rt' \\
& \llbracket F; A; L \vdash s_2 \ \text{wf} \rrbracket_f C_c \ n_c \ V = C_f \ n_f \\
& l_{fixjoin} = \max(\text{dom}(C_f)) + 1 \\
& C'_f = C_f \ @ \ \text{mov } rjf \ l_{halt} \ @ \ \text{mov } rjf' \ l_{halt} \\
& \quad \quad \quad @ \ \text{jmp}_G \ rjf \quad @ \ \text{jmp}_B \ rjf' \\
& l_{true} = \max(\text{dom}(C'_f)) + 1 \\
& \llbracket F; A; L \vdash s_1 \ \text{wf} \rrbracket_f C'_f \ n_f \ V = C_t \ n_t \\
& l_{join} = \max(\text{dom}(C_t)) + 1 \\
& C'_t = C_t \ @ \ \text{mov } re \ l_{join} \ @ \ \text{mov } re' \ l_{join} \\
& \quad \quad \quad @ \ \text{jmp}_G \ re \quad @ \ \text{jmp}_B \ re' \\
& C' = C'_t \ \begin{array}{l} [l_{fixjoin} \quad \mapsto \ \text{mov } rjf \ l_{join}] \\ [l_{fixjoin} + 1 \quad \mapsto \ \text{mov } rjf' \ l_{join}] \\ [l_{fixtrue} \quad \mapsto \ \text{mov } rt \ l_{true}] \\ [l_{fixtrue} + 1 \quad \mapsto \ \text{mov } rt' \ l_{true}] \end{array} \\
& n' = n_t \\
\hline
\llbracket F; A; L \vdash \text{if } v \ \text{then } s_1 \ \text{else } s_2 \ \text{wf} \rrbracket_f C \ n \ V \ B = C' \ n' \quad (\text{trans-s-if})
\end{aligned}$$

Figure E.4: If Statement Translation.

$$\begin{aligned}
& \mathit{argspace} = (\mathit{size}(vs) * 2) + 2 \\
& \mathit{vspace} = \mathit{size}(\mathit{dom}(V)) * 2 \\
& \mathit{spillspace} = \mathit{vspace} + \mathit{argspace} \\
& C_{\mathit{temps}} = @ \mathit{salloc} \mathit{argspace} \\
& \quad @ \mathit{sst} (\mathit{argspace} + 0) V_g(x_1) \\
& \quad @ \mathit{sst} (\mathit{argspace} + 1) V_g b(x_1) \\
& \quad @ \dots @ \mathit{sst} (\mathit{argspace} + \mathit{vspace} - 2) V_g(x_n) \\
& \quad @ \mathit{sst} (\mathit{argspace} + \mathit{vspace} - 1) V_b(x_n) \\
& \llbracket (A \cup L) \vdash vs : ps \rrbracket C_{\mathit{temps}} n V 1 = C_{vs} n_{vs} \\
& (n_{\mathit{call}}, rf, rf', ra, ra') = \mathit{freshReg}(n_{vs}) \\
& \mathit{retaddr} = \max(\mathit{dom}(C_{vs})) + 9 \\
& C_{\mathit{call}} = C_{vs} @ \mathit{mov} ra \mathit{retaddr} \quad @ \mathit{sst} 0 ra \\
& \quad @ \mathit{mov} ra' \mathit{retaddr} \quad @ \mathit{sst} 1 ra' \\
& \quad @ \mathit{mov} rf B(g) \quad @ \mathit{mov} rf' B(f) \\
& \quad @ \mathit{jmp}_G rf \quad @ \mathit{jmp}_B rf' \\
& \quad @ \mathit{sld}_G 2 V_g(x_1) \\
& \quad @ \mathit{sld}_G 3 V_b(x_1) \\
& \quad @ \dots \\
& \quad @ \mathit{sld}_B (\mathit{vspace}) V_g(x_n) \\
& \quad @ \mathit{sld}_B (\mathit{vspace} + 1) V_b(x_n) \\
& \quad @ \mathit{sld}_G V_g(x) 0 \quad @ \mathit{sld}_B V_b(x) 1 \\
& \quad @ \mathit{sfree} (\mathit{vspace} + 2) \\
\hline
& \llbracket F; A; L \vdash x = g(vs) \ \mathit{wf} \rrbracket_f C n V B = C_{\mathit{call}} n_{\mathit{call}} \quad (\mathit{trans-s-call})
\end{aligned}$$

Figure E.5: Function Call Translation.

## E.7 Statement Translation

The statement translations are shown in Figures E.2, E.3, E.4, and E.5. The rules that do not involve control flow transfers are very straightforward.

The rules for `while` and `if` are quite a bit more involved. The main complication comes from the backpatching required to fill in the correct addresses for jump targets. Because the code is generated sequentially, those instructions that jump forward do not know the correct address to use at the point they are generated. Instead, they temporarily jump to  $l_{halt}$  and are later patched to jump to the correct address. The call instruction generates a lot of code, but this code just implements the calling convention discussed earlier.

The Statement Translation Lemma says that translating a statement  $s$  has no effect on the main typing information. Again, there may be changes to additional temporary registers and the description of memory, but the final type is still consistent with the function arguments and local declarations.

### Lemma 55 (Statement Translation)

If  $\llbracket F;A;L \vdash s \text{ wf} \rrbracket_f C n V B = C' n'$  and  $f;V;B;F;A;L \vdash C : \Psi$   
 then  $\exists \Psi'. f;V;B;F;A;L \vdash C' : \Psi'$ .

*Proof* By induction on the structure of  $\llbracket F;A;L \vdash s \text{ wf} \rrbracket_f C n V B = C' n$  and Lemma 52. ■

## E.8 Function Translation

The function translation depends on two auxiliary judgments to generate the function entry and exit sequences.

### E.8.1 Prologue and Epilogue

The prologue generates code that loads the arguments from the stack into fresh registers and updates the variable mapping appropriately.

$$\boxed{genPrologue(C, n, V, n_p, A) = (C', n', V')}$$

$$\frac{}{genPrologue(C, n, V, n_p, \cdot) = (C', n', V')} \text{ (gen-prologue-emp)}$$

$$\frac{\begin{array}{l} (n_x, r, r') = freshReg(n) \\ C_x = C @ sld_G(2 * n_p + 2) r @ sld_B(2 * n_p + 3) r' \\ genPrologue(C_x, n_x, V[x \mapsto (r, r')], (n_p + 1), A) = (C', n', V') \end{array}}{genPrologue(C, n, V, n_p, (x : \tau, A)) = (C', n', V')} \text{ (gen-prologue)}$$

Before the prologue code,  $C$  is well-typed with regard to the set of functions defined. Afterwards,  $C'$  is well-typed with regard to the body of function  $f$  (though there are no local declarations yet).

#### Lemma 56 (Prologue Generation)

If  $genPrologue(ps) C n V P_{num} = C' n' V'$  and  $F; B \vdash C$

then  $\exists \Psi'. f; V'; B; F[f : ps \rightarrow \tau]; ps; \cdot \vdash C' : \Psi'$ .

*Proof* By induction on the structure of  $genPrologue(ps) C n V P_{num} = C' n' V'$ . ■



The epilogue generates code to free the function arguments and return addresses and push the return values.

$$\boxed{\text{genEpilogue}(X \vdash v : \tau, C, n, V, A) = (C', n')}$$

$$\begin{array}{l} \llbracket X \vdash v : \tau \rrbracket C n V = C_v n_v r_v r'_v \\ (n', r, r') = \text{freshReg}(n_v) \\ C' = C \quad @ \text{sld}_G 0 r \quad @ \text{sld}_B 1 r' \\ \quad @ \text{sfree} 2 * \text{size}(A) + 2 \quad @ \text{salloc} 2 \\ \quad @ \text{sst} 0 r_v \quad @ \text{sst} 1 r'_v \\ \quad @ \text{jmp}_G r \quad @ \text{jmp}_B r' \\ \hline \text{genPrologue}(X \vdash v : \tau, C, n, V, A) = (C', n') \quad (\text{gen-epilogue}) \end{array}$$

Before the epilogue,  $C$  is well-typed with regard to the body of function  $f$ . Afterwards,  $C'$  is well-typed with regard to the set of functions.

**Lemma 57 (Epilogue Generation)**

If  $\text{genEpilogue}((A \cup L) \vdash v : \tau, C, n, A) = (C', n')$  and  $f; V; B; F; A; L \vdash C : \Psi$  then  $\exists \Psi'. B; F \vdash C : \Psi'$ .

*Proof* By deconstruction of (wf-fun-body) and the appropriate instruction typing rules. ■

## E.8.2 Function Translation Judgment

The function translation judgment generates the function prologue, translates the local declaration, translates the function body, and generates the epilogue for each function. The maximum number of registers required by the set of functions is the maximum required by any single function.

$$\boxed{\llbracket F \vdash fds : F' \rrbracket C B = C' B' n}$$

$$\frac{}{\llbracket F \vdash \cdot : F \rrbracket C B = C B} \text{ (trans-fds-empty)}$$

$$\frac{\begin{array}{l} f_{addr} = \max(\text{dom}(C)) + 1 \\ \text{genPrologue}(C, 0, V, 0, A) = C_p n_p V_p \\ \llbracket F[f : A \rightarrow \tau]; A; \cdot \vdash lds : L \rrbracket C_p n_p V_p = C_{lds} n_{lds} V_{lds} \\ \llbracket F[f : A \rightarrow \tau]; A; L \vdash s \text{ wf} \rrbracket_f C_{lds} n_{lds} B[f \mapsto f_{addr}] = C_s n_s \\ \text{genEpilogue}((A \cup L) \vdash v : \tau, C_s, n_s, V_{lds}, A) = (C', n_r) \\ \llbracket F[f : A \rightarrow \tau] \vdash fds : F' \rrbracket C B[f \mapsto f_{addr}] = C' B' n_{fds} \\ n = \max(n_r, n_{fds}) \end{array}}{\llbracket F \vdash \tau f(A) \{lds; s; \text{return } v\} fds : F' \rrbracket C B = C' B' n'} \text{ (trans-fds)}$$

The Function Translation Lemma states that the code memory before and after translating a set of functions is well-typed with regards to those functions which are defined.

**Lemma 58 (Function Translation)**

If  $\llbracket F \vdash fds : F' \rrbracket C B = C' B' n$  and  $B; F \vdash C : \Psi$

then  $\exists \Psi'. B'; F' \vdash C' : \Psi'$

*Proof* By induction on the structure of  $\llbracket F \vdash lds : F' \rrbracket C B = C' B' n$ . The recursive case uses Lemmas 56 (Prologue Generation), 54 (Local Declaration Translation), 55 (Statement Translation), and 57 (Epilogue Generation) and then calls the induction hypothesis. ■

## E.9 Program Translation

The final judgment  $\llbracket \vdash p \text{ wf} \rrbracket = C \ n \ l$  takes a MiniC program  $p$  and generates a code memory  $C$ , a maximum number of registers  $n$ , and a start address  $l$ . It initializes the code memory to contain four instructions that create an infinite loop

$$\boxed{\llbracket \vdash p \text{ wf} \rrbracket = C \ n \ l}$$

$$\begin{array}{l} C_{start} = l_{start} \mapsto \text{mov } t_0 \ l_{halt}, \quad l_{start} + 1 \mapsto \text{mov } t_1 \ l_{halt} \\ \quad \quad \quad l_{start} + 2 \mapsto \text{jmp}_G \ t_0, \quad l_{start} + 3 \mapsto \text{jmp}_B \ t_1 \\ \llbracket \cdot \vdash fds : F \rrbracket C_{start} \ 0 \ \cdot = C_{fds} \ B_{fds} \ n_{fds} \\ l_{main} = \max(\text{dom}(C_{fds})) + 1 \\ \llbracket F; \cdot; \vdash lds : L \rrbracket C_{fds} \ 0 \ \cdot = C_{lds} \ n_{lds} \ V_{lds} \\ \llbracket F; \cdot; L \vdash s \text{ wf} \rrbracket_f C_{lds} \ n_{lds} \ V_{lds} \ B_{fds} = C_s \ n_s \\ \text{genEpilogue}(L \vdash v : \tau, C_s, n_s, V_{lds}, ()) = C' \ n_r \\ n' = \max(n_r, n_{fds}) \\ \hline \llbracket \vdash fds; lds; s; \text{return } v \text{ wf} \rrbracket = C' \ n' \ l_{main} \quad (\text{trans-}p) \end{array}$$

Finally, the Translation Theorem states that the result of translating a well-formed program  $p$  can be used to create a well-typed ETAL<sub>FT</sub> machine state. The code memory is just the code memory returned by the translation. The contents of the heap is empty, and the stack contains two pointers to the designated label containing the termination code. The register file is built by calling the function  $buildR(n)$  to generate a blank register with  $n$  temporary registers. The two program counters are set to the start address generated by the translation, and the stack pointers are set to the top of the stack.

**Theorem 59 (Translation)**

*If*  $\llbracket \vdash fds;lds;\zeta \text{ wf} \rrbracket = C \ n \ l_{start}$

*then*  $\vdash (R, C, M, (), \cdot)$

*where*  $st = \min(\text{dom}(C)) - 3$

$$R = \text{buildR}(n), \ pc_G \mapsto l_{start}, \ pc_B \mapsto l_{start}, \ sp_G \mapsto st, \ sp_B \mapsto st, \ gd \mapsto 0$$

$$M = st + 2 \mapsto 0, \ st + 1 \mapsto l_{halt}, \ st \mapsto l_{halt}$$

*Proof* By Lemmas 58 (Function Translation), 54 (Local Declaration Translation), 55 (Statement Translation), and 57 (Epilogue Generation). ■

# Appendix F

## TAL<sub>CF</sub> Formal Results

This appendix expands on the formal results for TAL<sub>CF</sub> presented in Section 4.4. We provide proof sketches for all lemmas and theorems. The complete proofs appear in the companion technical report [49].

### F.1 Type Safety

#### F.1.1 Typing Lemmas

First, we briefly explain the main lemmas used to prove type safety.

Expression equality is transitive. Conditional expression  $E_z ? E_f : E_t$  is equal to either  $E_t$  or  $E_f$  depending on the value of  $E_z$ .

#### **Lemma 60 (Expression Equality)**

1. *If  $\Delta \vdash E_1 = E_2$  and  $\Delta \vdash E_2 = E_3$  then  $\Delta \vdash E_1 = E_3$*
2. *If  $\Delta \vdash \sigma_1 = \sigma_2$  and  $\Delta \vdash \sigma_2 = \sigma_3$  then  $\Delta \vdash \sigma_1 = \sigma_3$*

3. If  $\Delta \vdash E_z = 0$  then  $\Delta \vdash E_z ? E_f : E_t = E_t$ .

4. If  $\Delta \vdash E_z \neq 0$  then  $\Delta \vdash E_z ? E_f : E_t = E_f$ .

**Proof:** By the definition of  $\Delta \vdash E = E$  and  $\Delta \vdash \sigma = \sigma$  and the definition of  $\llbracket E \rrbracket$ .

Substituting an expression of kind  $\kappa$  for a free variable of kind  $\kappa$  preserves typing. Applying a substitution  $S$  that provides substitutions for a number of free variables also preserves typing.

**Lemma 61 (Substitution)**

1. If  $\Delta, x : \kappa \vdash E' : \kappa'$  and  $\Delta \vdash E : \kappa$  then  $\Delta \vdash E'[E/x] : \kappa'$

2. If  $\Delta, x : \kappa \vdash E_1 = E_2$  and  $\Delta \vdash E : \kappa$  then  $\Delta \vdash E_1[E/x] = E_2[E/x]$

3. If  $(\Delta, x : \kappa); \Psi \vdash^Z v : t$  and  $\Delta \vdash E : \kappa$  then  $\Delta; \Psi \vdash^Z v : t[E/x]$

4. If  $(\Delta, x : \kappa); \Psi; \Gamma; \sigma; E_i; \tau \text{ opt} \vdash b$  and  $\Delta \vdash E : \kappa$   
then  $\Delta; \Psi; \Gamma[E/x]; \sigma[E/x]; E_i[E/x]; \tau \text{ opt}[E/x] \vdash b$

5. If  $(\Delta_1, \Delta_2) \vdash E' : \kappa'$  and  $\Delta_1 \vdash S : \Delta_2$  then  $\Delta_1 \vdash S(E') : \kappa'$

6. If  $(\Delta_1, \Delta_2) \vdash E_1 = E_2$  and  $\Delta_1 \vdash S : \Delta_2$  then  $\Delta_1 \vdash S(E_1) = S(E_2)$

7. If  $(\Delta_1, \Delta_2); \Psi \vdash^Z v : t$  and  $\Delta_1 \vdash S : \Delta_2$  then  $\Delta_1; \Psi \vdash^Z v : S(t)$

8. If  $(\Delta_1, \Delta_2); \Psi; \Gamma; \sigma; E_i; \tau \text{ opt} \vdash b$  and  $\Delta_1 \vdash S : \Delta_2$   
then  $\Delta_1; \Psi; S(\Gamma); S(\sigma); S(E_i); S(\tau \text{ opt}) \vdash b$

**Proof:**

1. By induction on the structure of  $\Delta, x : \kappa \vdash E' : \kappa'$

2. By case analysis on the structure of  $\Delta, x : \kappa \vdash E_1 = E_2$  using Part 1.

3. By case analysis on the structure of  $(\Delta, x : \kappa); \Psi \vdash^Z v : t$  using Parts 1 and 2.
4. By induction on the structure of  $(\Delta, x : \kappa); \Psi; \Gamma; \sigma; E_i; \tau \text{ opt} \vdash b$  using Parts 1-3. The case for rule (*recovernz-t*) divides into two subcases depending on if  $E_z = 0$  and uses rule (*recovernz-eq-t*) or rule (*recovernz-neq-t*) as appropriate.
- 5-8. By induction on the size of  $\Delta_2$ , using Parts 1-4 respectively.

If a value has a type  $t$  and this type is a subtype of  $t'$ , then the value can also be given type  $t'$ .

**Lemma 62 (Subtyping)**

If  $\Delta \vdash t \leq t'$  and  $\Delta; \Psi \vdash^Z v : t$  then  $\Delta; \Psi \vdash^Z v : t'$

**Proof:** By induction on the derivation of  $\Delta; \Psi \vdash^Z v : t$ . Each case uses inversion of the subtyping rules and Lemma 60 (Expression Equality).

If a value  $c \ n$  has type  $\langle c', \tau', E' \rangle$  under zap tag  $Z$ , then our knowledge about  $n$  depends both on the base type  $\tau'$  and also the relationship between  $Z$  and  $c'$ . If  $Z = CF$  and the color in the type is  $G$  or  $B$ , then the judgment may be been derived using rule (*val-zap-CF-t*), so we know nothing about  $n$ . However, we do know that the expression has kind  $\kappa_{int}$ . If  $Z$  is  $c'$ , then the judgment may have been derived by rule (*val-zap-c-t*), so again we only know that the expression  $E'$  has kind  $\kappa_{int}$ . The remaining case applies when  $Z$  is not equal to the color in the type  $c'$  and when either  $Z \neq CF$  or the color in the type is neither  $G$  nor  $B$ . In this case the judgment must have been derived using rule (*val-t*), so we know the color tag on the value is equal to the color tag in the type and the expression

$E'$  is equal to the value  $n$ . In addition, if  $\tau'$  is a code type, we also know that  $n$  is a valid code address.

**Lemma 63 (Canonical Forms)**

If  $\cdot; \Psi \vdash^Z c n : \langle c', \tau', E' \rangle$  and  $\vdash C : \Psi$  then

1. If  $Z = CF$  and ( $c' = B$  or  $c' = G$ ), then  $\cdot \vdash E' : \kappa_{int}$ .
2. If  $Z = c'$  then  $c = c'$  and  $\cdot \vdash E' : \kappa_{int}$ .
3. If  $Z \neq c'$  and ( $Z \neq CF$  or ( $c' \neq B$  and  $c' \neq G$ )) then
  - $c = c'$
  - $\cdot \vdash E' = n$
  - $\tau' = \forall[\Delta](\Gamma, \sigma) \implies n \in \text{Dom}(C)$

**Proof:** By inspection of  $\cdot; \Psi \vdash^Z c n : \langle c', \tau', E' \rangle$ .

If a value has a type  $t$  under zap tag  $Z$ , then that value also has type  $t$  under any zap tag  $Z'$  that is a supertype of  $Z$ .

**Lemma 64 (Color Weakening)**

If  $\Delta; \Psi \vdash^Z v : t$  and  $Z \leq Z'$  then  $\Delta; \Psi \vdash^{Z'} v : t$

**Proof:** By case analysis of  $\Delta; \Psi \vdash^Z v : t$  and the definition of  $Z \leq Z'$ .

## F.1.2 Type Safety

We have proven that the TAL<sub>CF</sub> type system is sound using the standard notion of Progress and Preservation. Progress asserts that machine states well-typed under the empty zap tag



can take a step to another ordinary machine state. States that are well-typed under any zap can also take a step, but this step may reach any state, including `recover(h)` or `hwerror(h)`.

**Theorem 65 (Progress)**

1. If  $\vdash \Sigma$  then  $\Sigma \longrightarrow_0 \Sigma'$ .
2. If  $\vdash^Z \Sigma$  then  $\Sigma \longrightarrow_0 \mathcal{F}$ .

**Proof:** *The proof for each part is by case analysis on the current block *b* of  $\Sigma$  using Lemma 60 (Expression Equality) and Lemma 63 (Canonical Forms).*

Preservation states that execution preserves typing. States well-typed under the empty zap tag continue to be so after taking a non-faulty step. States typed under any zap also remain well-typed after a non-faulty step, but the zap tag may escalate to a supertype. This elevation might occur at control flow transfers. A zap tag of *B* or *G* becomes *CF* whenever the corruption has spread to the operands being used in the transfer. This way the block that results from the transfer can be well-typed under *CF* even when control has transferred to a totally unexpected block. The intention register is always the only orange value that is live across control flow transfers, and we have already seen that it is well-typed even when a control fault has occurred. Finally, a state is well-typed under the empty zap tag and takes a faulty step, then the resulting state is well-typed under some color *c*.

**Theorem 66 (Preservation)**

1. If  $\vdash \Sigma$  and  $\Sigma \longrightarrow_0 \Sigma'$  then  $\vdash \Sigma'$
2. If  $\vdash^Z \Sigma$  and  $\Sigma \longrightarrow_0 \Sigma'$  then  $\exists Z' . \vdash^{Z'} \Sigma'$  and  $Z \leq Z'$ .
3. If  $\vdash \Sigma$  and  $\Sigma \longrightarrow_1 \Sigma'$  then  $\exists c . \vdash^c \Sigma'$

**Proof:** The proof for each part is by case analysis on the corresponding single step judgment using Lemma 60 (Expression Equality), Lemma 63 (Canonical Forms), and Lemma 64 (Color Weakening). Cases for the jump and branch rules also use Lemma 61 (Substitution) and Lemma 62 (Subtyping).

## F.2 Fault Tolerance Results

We first present a handful of definitions and lemmas relating machine states to other states, and then use these to formally state and prove the Fault Tolerance Theorem.

### F.2.1 Machine State Simulation

We say that a faulty value simulates a fault-free value under color  $c$  if the values are equal when they are not colored by  $c$ .

$$\frac{}{c' n \text{ sim}^c c' n} \text{ (sim-val)} \quad \frac{}{c n \text{ sim}^c c n'} \text{ (sim-val-zap)}$$

A faulty machine state  $\Sigma_f$  simulates a fault-free state  $\Sigma$  under color  $c$  if  $\Sigma_f$  is well-typed under  $c$ ,  $\Sigma$  is well-typed under the empty zap tag, and the two states are identical modulo the values in registers colored  $c$ .

$$\frac{\vdash (C, h, R, b) \quad \vdash^c (C, h, R_f, b) \quad \forall r. R_f(r) \text{ sim}^c R(r)}{(C, h, R_f, b) \text{ sim}^c (C, h, R, b)} \text{ (sim-}\Sigma\text{)}$$

### F.2.2 Block Execution

The Block Step Lemma states that given a non-faulty computation and a corresponding faulty version  $\Sigma_f$ , if the non-faulty computation can take a non-faulty step to some other

state *in the same block*, then the faulty computation will either also take a step within the current block or will take a single step to the recover state.

**Lemma 67 (Block Step)**

If  $\Sigma_f \text{ sim}^c (C, h, R, b)$  and  $(C, h, R, b) \longrightarrow_0 (C, h, R', b')$  then either

1.  $\Sigma_f \longrightarrow_0 \Sigma'_f$  and  $\Sigma'_f \text{ sim}^c (C, h, R', b')$ , or
2.  $\Sigma_f \longrightarrow_0 \text{recover}(h)$

**Proof:** By case analysis of  $(C, h, R, b) \longrightarrow_0 (C, h, R', b')$  and Theorem 66 (Preservation).

In order to reason about block execution, we extend the single step relation  $\Sigma \longrightarrow_k \Sigma'$  from Section 4.2 to create the judgment  $\Sigma \rightsquigarrow_k \mathcal{F}$  which states that  $\mathcal{F}$  is the result of executing the *current block* of  $\Sigma$  while incurring  $k$  faulty transitions. Execution proceeds up to the control-flow transfer statement at the end of the current block or the recover state if the block terminates prematurely by transitioning to recovery code. For example, if  $\Sigma = (C, h, R, i_1; \dots; i_n; \text{jmp } r_t)$ , then either  $\mathcal{F} = (C, h, R', \text{recover}(h))$  or  $\mathcal{F} = (C, h, R', \text{jmp } r_t)$ .

$$\begin{array}{c}
 \boxed{\Sigma \rightsquigarrow_k \mathcal{F}} \\
 \frac{(C, h, R, b) \longrightarrow_0 \text{recover}(h)}{(C, h, R, b) \rightsquigarrow_0 \text{recover}(h)} \text{ (blk-eval-recover)} \\
 \frac{}{(C, h, R, \text{jmp } r_t) \rightsquigarrow_0 (C, h, R, \text{jmp } r_t)} \text{ (blk-eval-jmp)} \\
 \frac{}{(C, h, R, \text{brz } r_z r_t) \rightsquigarrow_0 (C, h, R, \text{brz } r_z r_t)} \text{ (blk-eval-brz)} \\
 \frac{(C, h, R, b) \longrightarrow_{k_1} (C, h, R', b') \quad (C, h, R', b') \rightsquigarrow_{k_2} \mathcal{F}}{(C, h, R, b) \rightsquigarrow_{(k_1+k_2)} \mathcal{F}} \text{ (blk-eval-sequence)}
 \end{array}$$

The Block Execution Lemma states that given a faulty computation  $\Sigma_f$  that simulates a non-faulty computation, the result of executing the faulty block will either simulate

the result of executing the non-faulty block, or executing the faulty block will result in  $\text{recover}(h)$ .

**Lemma 68 (Block Execution)**

If  $\Sigma_f \text{sim}^c (C, h, R, b)$  and  $(C, h, R, b) \rightsquigarrow_0 \Sigma'$  then either

1.  $\Sigma_f \rightsquigarrow_0 \Sigma'_f$  and  $\Sigma'_f \text{sim}^c \Sigma'$ , or
2.  $\Sigma_f \rightsquigarrow_0 \text{recover}(h)$

**Proof:** By induction on the structure of  $(C, h, R, b) \rightsquigarrow_0 \Sigma'$  and Lemma 67 (Block Step).

### F.2.3 Fault Recovery

The CF Fault Step Lemma states that once a control flow fault has occurred, execution will either step within the same block or will step to recovery code.

**Lemma 69 (CF Fault Step)**

If  $\vdash^{CF} (C, h, R, b)$  then either

1.  $(C, h, R, b) \longrightarrow_0 (C, h, R', b')$  and  $\vdash^{CF} (C, h, R', b')$
2.  $(C, h, R, b) \longrightarrow_0 \text{recover}(h)$ .

**Proof:** By case analysis on the structure of  $b$  using Theorem 66 (Preservation).

The Fault Recovery Lemma states that once a control flow fault has occurred, control will always reach recovery code before exiting the current block.

**Lemma 70 (CF Fault Block Execution)**

If  $\vdash^{CF} (C, h, R, b)$  then  $(C, h, R, b) \rightsquigarrow_0 \text{recover}(h)$ .

**Proof:** By induction on the length of  $b$  and Lemma 69 (CF Fault Step).

**F.2.4 Block Transitions**

In order to reason about transitions *between* blocks, we define the judgment  $\Sigma \Longrightarrow^\ell \Sigma'$  whenever  $(C, h, R, b) \longrightarrow_0 (C, (h, \ell), R', b')$ . In other words, control transfers from the end of one block to the beginning of another block  $\ell$  in a single step.

$$\boxed{\Sigma \Longrightarrow^\ell \Sigma'}$$

$$\frac{(C, h, R, b) \longrightarrow_0 (C, (h, \ell), R', b')}{(C, h, R, b) \Longrightarrow^\ell (C, (h, \ell), R', b')} \text{ (trans-eval)}$$

The Block Transition Lemma states that whenever a non-faulty computation transitions to a new block, the corresponding faulty computation will either (1) transition to the same block and continue to be indistinguishable from the non-faulty computation, (2) trigger a hardware error, or (3) transition to an incorrect block where the error will be detected before control leaves the incorrect block.

**Lemma 71 (Block Transition)**

If  $\Sigma_f \text{sim}^c (C, h, R, b)$  and  $(C, h, R, b) \Longrightarrow^\ell \Sigma'$  then either

1.  $\Sigma_f \Longrightarrow^\ell \Sigma'_f$  and  $\Sigma'_f \text{sim}^c \Sigma'$
2.  $\Sigma_f \longrightarrow_0 \text{hwerror}(h)$
3.  $\Sigma_f \Longrightarrow^{\ell'} \Sigma'_f$  and  $\Sigma'_f \rightsquigarrow_0 \text{recover}(h, \ell')$

**Proof:** By case analysis of the structure of  $(C, h, R, b) \Longrightarrow^\ell \Sigma'$  and Lemma 70 (CF Fault Block Execution).

## F.2.5 Program Execution

The judgment  $\Sigma \Longrightarrow_k^h \mathcal{F}$  states that machine state  $\Sigma$  executes through a sequence of blocks  $h$  to reach state  $\mathcal{F}$  while incurring  $k$  faulty transitions. In other words, if  $\Sigma = (C, h_1, R, b)$ , then  $\mathcal{F}$  is either the regular state  $(C, (h_1, h), R', \text{jmp } r_t)$ , the regular state  $(C, (h_1, h), R', \text{brz } r_z r_t)$ , the hardware error state  $\text{hwerror}(h_1, h)$ , or the recovery state  $\text{recover}(h_1, h)$ .

$$\boxed{\Sigma \Longrightarrow_k^h \Sigma'}$$

$$\frac{\Sigma \rightsquigarrow_k \mathcal{F}}{\Sigma \Longrightarrow_k^{()} \mathcal{F}} \text{ (prog-exec-blk)}$$

$$\frac{\Sigma \Longrightarrow_k^h \Sigma' \quad \Sigma' \longrightarrow_0 \text{hwerror}(h', h)}{\Sigma \Longrightarrow_k^h \text{hwerror}(h', h)} \text{ (prog-exec-seq-hwerror)}$$

$$\frac{\Sigma \Longrightarrow_{k_1}^h \Sigma' \quad \Sigma' \Longrightarrow^\ell \Sigma'' \quad \Sigma'' \rightsquigarrow_{k_2} \mathcal{F}}{\Sigma \Longrightarrow_{(k_1+k_2)}^{(h, \ell)} \mathcal{F}} \text{ (prog-exec-seq-trans-blk)}$$

The Faulty Execution Lemma states that if a faulty execution  $\Sigma_f$  simulates a non-faulty execution  $\Sigma$  under some color  $c$ , then  $\Sigma_f$  behaves in one of four possible ways with regards to  $\Sigma$ . (1) Executing  $\Sigma_f$  results in the same sequence of blocks  $h$  as executing  $\Sigma$  and the resulting faulty state simulates the corresponding non-faulty state under the same color  $c$ . (2) Executing  $\Sigma_f$  results in an attempt to transfer control to an invalid address outside the domain of code memory and triggers a hardware fault. Prior to the occurrence of the hardware fault, the execution of  $\Sigma_f$  visits the same blocks as the execution of  $\Sigma$ .

(3) While executing  $\Sigma_f$ , a fault is detected and control is transferred to recovery code even though no incorrect blocks have been visited. This situation can be caused by a fault affecting the intention register or the checking code. (4) While executing  $\Sigma_f$ , control veers off course to a block that is not visited in the execution of  $\Sigma$ . In this case, the checking code in the invalid block catches the error and transfers control to the recovery code.

**Lemma 72 (Faulty Execution)**

If  $\Sigma_f \text{ sim}^c \Sigma$  and  $\Sigma \Longrightarrow_0^h \Sigma'$  then either:

1.  $\Sigma_f \Longrightarrow_0^h \Sigma'_f$  and  $\Sigma'_f \text{ sim}^c \Sigma'$
2.  $\Sigma_f \Longrightarrow_0^{h_f} \text{hwerror}(h', h_f)$  and  $h_f$  is a prefix of  $h$
3.  $\Sigma_f \Longrightarrow_0^{h_f} \text{recover}(h', h_f)$  and  $h_f$  is a prefix of  $h$
4.  $\Sigma_f \Longrightarrow_0^{h_f} \text{recover}(h', h_f)$  and  $h_f = (h_1, l')$  and  $h = (h_1, l, h_2)$

**Proof:** By induction on the structure of  $\Sigma \Longrightarrow_0^h \Sigma'$ , Lemma 68 (Block Execution) and Lemma 71 (Block Transition).

## F.2.6 The Fault Tolerance Theorem

A program is fault-tolerant if any execution of the program with a single fault behaves in one of four possible ways with regards to the original, non-faulty computation. (1) The faulty computation visits the same sequence of blocks as the original and the resulting faulty state simulates the corresponding original state under some color  $c$ . (2) The faulty

computation attempts to transfer control to an invalid address outside the domain of code memory and triggers a hardware fault. Prior to the occurrence of the hardware fault, the faulty computation visits the same blocks as the original computation. (3) The faulty computation detects a fault in software and jumps to recovery code even though no incorrect blocks have been visited. This situation can be caused by a fault affecting the intention register or the checking code. (4) The faulty computation veers off course to a block that does not match the corresponding block in the original computation. In this case, the checking code in the invalid block catches the error and transfers control to the recovery code.

**Theorem 73 (Fault Tolerance)**

If  $\vdash \Sigma$  and  $\Sigma \Longrightarrow_0^h \Sigma'$  then at least one of the following cases applies and all derivations  $\Sigma \Longrightarrow_1^{h_f} \mathcal{F}$  where  $\text{length}(h_f) \leq \text{length}(h)$  fit one of these cases:

1.  $\Sigma \Longrightarrow_1^h \Sigma'_f$  and  $\exists c . \Sigma'_f \text{sim}^c \Sigma'$
2.  $\Sigma \Longrightarrow_1^{h_f} \text{hwerror}(h', h_f)$  and  $h_f$  is a prefix of  $h$
3.  $\Sigma \Longrightarrow_1^{h_f} \text{recover}(h', h_f)$  and  $h_f$  is a prefix of  $h$
4.  $\Sigma \Longrightarrow_1^{h_f} \text{recover}(h', h_f)$  and  $h_f = (h_1, l')$  and  $h = (h_1, l, h_2)$

**Proof:** By case analysis on the structure of  $\Sigma \Longrightarrow_0^h \Sigma'$ . In essence, arbitrarily divide the computation  $\Sigma \Longrightarrow_0^h \Sigma'$  into two pieces with some  $\Sigma''$  as the intermediate state. Use one of the fault rules to step  $\Sigma''$  to  $\Sigma''_f$ . If  $c$  is the color of the value that faults, then  $\Sigma''_f \text{sim}^c \Sigma''$ . Then use Lemma 72 (Faulty Execution) with  $\Sigma''_f \text{sim}^c \Sigma''$  and the remainder of the computation from  $\Sigma''$  to  $\Sigma'$  to determine what happens after the fault. Use these



*results and the first half of the computation to show that one of the four cases applies to the entire computation containing a single fault.*

# Appendix G

## TAL<sub>CF</sub> Translation

This appendix gives additional details for the translation from while loops to TAL<sub>CF</sub> presented in Section 4.5. We provide proof sketches for all lemmas and theorems. The complete proofs appear in the companion technical report [49].

### G.1 A Simple While Loop Language

The while loop language statements consist of simple assignment, subtraction, if statements, while loops, and sequences of statements. As all the variables in this language contain integers, the well-formedness judgment  $V \vdash s$  simply enforces that all variables  $v$  in  $s$  exist in the variable context  $V$ .

$$\begin{aligned} s ::= & v := n \mid v_d := v_a - v_b \\ & \mid \text{if } 0 \ v_z \text{ then } s_1 \text{ else } s_2 \mid \text{while } v_z \neq 0 \text{ do } s \\ & \mid s_1; s_2 \end{aligned}$$

## G.2 Checking Code and Exit Code Macros

The translation rules and lemmas make use of the following macros that implement the protocol from Section 4.1. These macros make use of two temporary registers:  $t_g$  and  $t_b$ .

Macro “check  $\ell$ ” generates the checking code at the entry of block  $\ell$  to check that control has correctly transferred to this block. Macro “intendjump  $\ell_t$ ” sets the intention and then executes the jump to target block  $\ell_t$ . Finally, macro “intendzbrz  $r_z \ell_t \ell_f$ ” uses register  $r'_z$  to conditionally set the intention to fall through to block  $\ell_f$  or branch to block  $\ell_t$ , and then uses register  $r_z$  to execute the conditional branch.

$$\text{check } \ell \equiv \begin{array}{l} \text{movi } t_g \ O \ \ell; \text{sub } t_g \ t_g \ r_i; \\ \text{recovernz } t_g \end{array}$$

$$\text{intendjump } \ell_t \equiv \begin{array}{l} \text{movi } t_b \ B \ \ell_t; \text{intend } t_b; \\ \text{movi } t_g \ G \ \ell_t; \text{jmp } t_g \end{array}$$

$$\text{intendzbrz } r_z \ \ell_t \ \ell_f \equiv \begin{array}{l} \text{movi } t_b \ B \ \ell_f; \text{intend } t_b; \\ \text{movi } t_b \ B \ \ell_t; \text{intendz } r'_z \ t_b; \\ \text{movi } t_g \ G \ \ell_t; \text{brz } r_z \ t_g \end{array}$$

## G.3 Translating Variable Context $V$

Since all variables are considered live at all program points, every assembly-level instruction block will have essentially the same signature. If the context  $V$  contains variable  $v_1, \dots, v_n$ , then each block in the translation requires  $2n + 3$  registers: a green copy  $r_k$  and a blue copy  $r'_k$  for each variable  $v_k$ , the intention register  $r_i$ , and two temporary registers  $t_g$  and  $t_b$ . (We have made no effort to optimize this translation, it merely serves to demonstrate the theoretical expressiveness of the target language.)

The function  $\llbracket V \rrbracket_\ell$  generates the code type of the code at label  $\ell$ . The generated  $\Delta$  contains all the expression variables that are need in  $\Gamma$  and  $\sigma$ . Each label will have a slightly different history typing  $\sigma$ , since the sequence ends with the current label. Register file typing  $\Gamma$  gives types to each of the  $2n + 3$  registers. Register  $r_k$  is green and  $r'_k$  is blue. Both registers have basic type *int* and are described by the same expression variable  $x_k$ , which enforces that they are equal on entry to the block. Register  $r_i$  is orange, has basic type *check*, and is described by expression variable  $x_i$ . Again, since we have not optimized the translation, we will assume that during block transitions  $t_g$  always contains a green value and  $t_b$  always contains a green value. They may hold values of other colors during the body of a block.

$$\boxed{\llbracket V \rrbracket_\ell = \forall[\Delta](\Gamma, \sigma)}$$

$$\begin{array}{l} \text{choose fresh variables } x_1, \dots, x_n, x_h, x_{r_i}, x_g, x_b, x_o \\ \Delta = x_1 : \mathbf{K}_{int}, \dots, x_n : \mathbf{K}_{int}, x_h : \mathbf{K}_{hist}, x_{r_i} : \mathbf{K}_{int}, x_g : \mathbf{K}_{int}, x_b : \mathbf{K}_{int} \\ \sigma = x_h \circ \ell \\ \Gamma = \{ r_1 : \langle G, int, x_1 \rangle, r'_1 : \langle B, int, x_1 \rangle, \dots, r_n : \langle G, int, x_n \rangle, r'_n : \langle B, int, x_n \rangle, \\ \quad r_i : \langle O, check, x_{r_i} \rangle, \\ \quad t_g : \langle G, int, x_g \rangle, t_b : \langle B, int, x_b \rangle \} \\ \hline \llbracket v_1, \dots, v_n \rrbracket_\ell = \forall[\Delta](\Gamma, \sigma) \quad (\text{trans-}V) \end{array}$$

The function  $\text{Gen}\Psi(V, L)$  computes the heap typing  $\Psi$  that maps each label in  $L$  to its corresponding type

$$\boxed{\text{Gen}\Psi(V, L)}$$

$$\begin{array}{l} \text{Gen}\Psi(V, \cdot) = \cdot \\ \text{Gen}\Psi(V, (L, \ell)) = \text{Gen}\Psi(V, L), \ell \mapsto \llbracket V \rrbracket_\ell \end{array}$$

## G.4 Partial Translations

A 4-tuple of objects  $(L, C, \vec{i}, \ell)$  is used to track the code generated during the translation.  $C$  is the code memory that contains all blocks generated so far.  $L$  contains labels that may be referred to by blocks in  $C$  but whose corresponding blocks have not yet been generated.  $\ell$  is the label that will be assigned to the block that is currently being generated.  $\vec{i}$  contains the list of instructions for this block that have been generated so far. The instructions for checking the checking code and exit code are not included and will be added when the block is added to  $C$ .

The judgment  $V; \Psi_1 \vdash C : \Psi_2$  is used to type code memory  $C$  as it is being generated. There are two disjoint heaping typings:  $\Psi_1$  contains labels that may be referenced by  $C$  but whose corresponding code blocks may not have been generated yet, and  $\Psi_2$  contains the types for the blocks that have already been generated. Both  $\Psi_1$  and  $\Psi_2$  map each label  $\ell$  to  $\llbracket V \rrbracket_\ell$ . In addition, each label in  $\Psi_2$  has a type that can be used to type check corresponding block.

$$\boxed{V; \Psi_1 \vdash C : \Psi_2}$$

$$\begin{array}{l}
Dom(\Psi_1) \cap Dom(\Psi_2) = \emptyset \\
\forall \ell \in Dom(\Psi_1) . \Psi_1(\ell) = \llbracket V \rrbracket_\ell \\
Dom(C) = Dom(\Psi_2) \\
\forall \ell \in Dom(\Psi_2) . \\
\quad \Psi_2(\ell) = \llbracket V \rrbracket_\ell = \forall [\Delta] ( (\Gamma, r_i \mapsto \langle O, check, x_i \rangle), x_h \circ \ell ) \\
\quad \wedge \Delta; (\Psi_1 \cup \Psi_2); (\Gamma, r_i \mapsto \langle O, check, x_i \rangle); (x_h \circ \ell); x_i; (\Psi_1 \cup \Psi_2)(\ell + 1) \\
\quad \vdash C(\ell) \\
\hline
V; \Psi_1 \vdash C : \Psi_2 \quad (C-wf)
\end{array}$$

Judgment  $V \vdash \vec{i}$  wf states that  $\vec{i}$  is a sequence of pairs of instructions that perform duplicate moves and subtractions. For example, the following is a well-formed list of instructions.

$$\text{movi } r_3 \text{ } G \ 3; \text{ movi } r'_3 \text{ } B \ 3; \text{ sub } r_4 \ r_5 \ r_6; \text{ sub } r'_4 \ r'_5 \ r'_6; \dots$$

Using these definitions, we say a partial translation  $(L, C, \vec{i}, \ell)$  is well-formed when the code memory  $C$  is well-formed using the heap typings calculated from the label  $\ell$  and the labels in  $L$  and the labels already in the domain of  $C$ . In addition, the instruction list  $\vec{i}$  is well-formed.

$$\boxed{V \vdash (L, C, \vec{i}, \ell) \text{ wf}}$$

$$\frac{\begin{array}{l} \Psi_1 = \text{Gen}\Psi(V, (L, \ell)) \\ \Psi_2 = \text{Gen}\Psi(V, \text{Dom}(C)) \\ V; \Psi_1 \vdash C : \Psi_2 \\ V \vdash \vec{i} \text{ wf} \end{array}}{V \vdash (L, C, \vec{i}, \ell) \text{ wf}} \quad (\text{partial-trans- wf})$$

The Block Construction Lemma says that the instruction list  $\vec{i}$  from a well-formed partial translation can be used to construct a block by adding checking code to the beginning and exit code to the end. The exit code can refer to any existing label  $\ell'$  as the jump target. The exit code can be a conditional branch only if the fall-through block  $\ell + 1$  exists. The new code memory formed by adding this new block is also well-formed.

**Lemma 74 (Block Construction)**

If  $V \vdash (L, C, \vec{i}, \ell)$  wf then  $\forall \ell' \in ((L, \ell) \cup \text{Dom}(C))$ .

1.  $\text{Gen}\Psi(V, L) \vdash C[\ell \mapsto \text{check } \ell; \vec{i}; \text{intendjump } \ell'] : \text{Gen}\Psi(V, (\text{Dom}(C), \ell))$

2. If  $l+1 \in ((L, \ell) \cup \text{Dom}(C))$

then  $\text{Gen}\Psi(V, L) \vdash C[\ell \mapsto \text{check } \ell; \vec{i}; \text{intendzbrz } r_z \ell' \ell + 1] : \text{Gen}\Psi(V, (\text{Dom}(C), \ell))$

**Proof:** Using the macro definitions, the definition of  $V \vdash (L, C, \vec{i}, \ell)$  wf, and instruction typing rules from Section 4.3.2.

## G.5 Translating Statements

The main translation judgment  $\llbracket V \vdash s \rrbracket (L, C, \vec{i}, \ell) = (L', C', \vec{i}', \ell')$  extends the existing partial translation  $(L, C, \vec{i}, \ell)$  with the translation of statement  $s$ .

The statement translation rules are shown in Figure G.1. Translating simple assignment and subtraction statements simply adds pairs of assembly instructions to the end of the current instruction sequence. Sequencing two statements uses the partial translation from the first statement to translate the second.

Translating `if0` statements requires the addition of new blocks:  $\ell_f$  contains the fall-through branch,  $\ell_t$  contains the true branch, and  $\ell_m$  is where the two branches merge. The function  $\text{NumBlock}(s)$  calculates the number of blocks generated by the translation of  $s$ . The current block  $b_\ell$  contains checking code, the code  $\vec{i}$  generated for the block so far, and ends with a conditional branch to  $\ell_t$  (and an automatic fall-through to  $\ell_f$ ). The new label  $\ell_t$  is the starting point for the code generated for the true branch  $s_1$ . The ending label of this code  $\ell'_t$  finishes by merging back to the common block at  $\ell_m$ . The translation of the false branch is similar. The final code memory contains all blocks generated by either branch as well as the blocks ending each branch by jumping to the merge block  $\ell_m$ . The label in the resulting partial translation is  $\ell_m$ .

Translating `while` statements also requires the addition of new blocks. The current block at  $\ell$  is terminated with an unconditional jump to a beginning block at  $\ell_b$  that tests

$$\boxed{\llbracket V \vdash s \rrbracket (L, C, \vec{i}, \ell) = (L', C', \vec{i}', \ell')}$$

$$\frac{\vec{i}' = \vec{i}; \text{movi } r_k \ G \ n; \text{movi } r'_k \ B \ n}{\llbracket V \vdash v := n \rrbracket (L, C, \vec{i}, \ell) = (L, C, \vec{i}', \ell)} \quad (t\text{-assign})$$

$$\frac{\vec{i}' = \vec{i}; \text{sub } r_k \ r_a \ r_b; \text{sub } r'_k \ r'_a \ r'_b}{\llbracket V \vdash v_d := v_a - v_b \rrbracket (L, C, \vec{i}, \ell) = (L, C, \vec{i}', \ell)} \quad (t\text{-sub})$$

$$\frac{\begin{array}{l} \llbracket V \vdash s_1 \rrbracket (L, C, \vec{i}, \ell) = (L_1, C_1, \vec{i}_1, \ell_1) \\ \llbracket V \vdash s_2 \rrbracket (L_1, C_1, \vec{i}_1, \ell_1) = (L_2, C_2, \vec{i}_2, \ell_2) \end{array}}{\llbracket V \vdash s_1; s_2 \rrbracket (L, C, \vec{i}, \ell) = (L_2, C_2, \vec{i}_2, \ell_2)} \quad (t\text{-seq})$$

$$\ell_f = \ell + 1$$

$$\ell_t = \ell_f + \text{NumBlocks}(s_1)$$

$$\ell_m = \ell_t + \text{NumBlocks}(s_2)$$

$$b_\ell = \text{check } \ell; \vec{i}; \text{intendzbrz } r_z \ \ell_f \ \ell_t$$

$$\llbracket V \vdash s_1 \rrbracket ((L, \ell_f), C[\ell \mapsto b_\ell], \cdot, \ell_t) = (L'_t, C'_t, \vec{i}'_t, \ell'_t)$$

$$b'_t = \text{check } \ell'_t; \vec{i}'_t; \text{intendjump } \ell_m$$

$$\llbracket V \vdash s_2 \rrbracket ((L, \ell_t), C[\ell \mapsto b_\ell], \cdot, \ell_f) = (L'_f, C'_f, \vec{i}'_f, \ell'_f)$$

$$b'_f = \text{check } \ell'_f; \vec{i}'_f; \text{intendjump } \ell_m$$

$$C' = (C'_t \cup C'_f)[\ell'_t \mapsto b'_t][\ell'_f \mapsto b'_f]$$

$$\frac{}{\llbracket V \vdash \text{if0 } v_z \text{ then } s_1 \text{ else } s_2 \rrbracket (L, C, \vec{i}, \ell) = (L, C', \cdot, \ell_m)} \quad (t\text{-if})$$

$$\ell_b = \ell + 1$$

$$\ell_s = \ell_b + 1$$

$$\ell_e = \ell_s + \text{NumBlocks}(s)$$

$$C'' = C[\ell \mapsto \text{check } \ell; \vec{i}; \text{intendjump } \ell_b] \\ [\ell_b \mapsto \text{check } \ell_b; \text{intendzbrz } r_z \ \ell_e \ \ell_s]$$

$$\llbracket V \vdash s \rrbracket ((L, \ell_e), C'', \cdot, \ell_s) = (L'_s, C'_s, \vec{i}'_s, \ell'_s)$$

$$C' = C'_s[\ell'_s \mapsto \text{check } \ell'_s; \vec{i}'_s; \text{intendjump } \ell_b]$$

$$\frac{}{\llbracket V \vdash \text{while } v_z \neq 0 \text{ do } s \rrbracket (L, C, \vec{i}, \ell) = (L'_s, C', \cdot, \ell_e)} \quad (t\text{-while})$$

Figure G.1: Translation of While Programs.



the condition and branches to an ending label  $\ell_e$  if the condition fails. Otherwise it falls through to the block at  $\ell_s$  which contains the translation of  $s$  and terminates with a jump back to the beginning block. The label in the resulting partial translation is  $\ell_e$ .

The Statement Translation Lemma says that given a well-formed partial translation  $(L, C, \vec{i}, \ell)$ , translating a statement  $s$  results in another well-formed partial translation. In addition, the new set of undefined labels  $L'$  is equal to that in the original partial translation.

**Lemma 75 (Statement Translation)**

*If  $\llbracket V \vdash s \rrbracket (L, C, \vec{i}, \ell) = (L', C', \vec{i}', \ell')$  and  $V \vdash (L, C, \vec{i}, \ell)$  wf then  $V \vdash (L', C', \vec{i}', \ell')$  wf and  $L = L'$*

**Proof:** *Using the definition of  $V \vdash (L, C, \vec{i}, \ell)$  wf and Lemma 74 (Block Construction).*

## G.6 The Translation Theorem

To translate a statement  $s$  as a stand-alone program, it is translated as in the previous section with 1 as the starting label. Because there is no halt instruction in TAL<sub>CF</sub>, code is added to the last block in the translation to create an infinite loop at label  $\ell_{halt}$ . The function  $\text{InitRegFile}(V)$  creates an initial register file that maps each register used to translate  $V$  to 0.

The assembly language program corresponding to  $s$  is the TAL<sub>CF</sub> state consisting of the generated code memory, a history with only the first location, an initial register file, and code to jump to the first label in code memory. If the original statement is well-formed, then the translation is well-typed.

**Theorem 76 (Translation)**

If  $\llbracket V \vdash s \rrbracket(\cdot, \cdot, \cdot, 1) = (\cdot, C, \vec{i}, \ell)$  then

$\vdash (C', 0, \text{InitRegFile}(V), \text{intendjump } 1)$

where  $C' = [\ell \mapsto \text{check } \ell; \vec{i}; \text{intendjump } \ell_{halt}] [\ell_{halt} \mapsto \text{check } \ell_{halt}; \text{intendjump } \ell_{halt}]$

**Proof:** Using Lemma 75 (Statement Translation), Lemma 74 (Block Construction), the block typing rules from Section 4.3.3, and the machine state typing rules from Section 4.3.4.

# Bibliography

- [1] M. Abadi, M. Budiu, Úlfar Erlingsson, and J. Ligatti. Control-flow integrity: Principles, implementations, and applications. In *ACM Conference on Computer and Communications Security*, Nov. 2005.
- [2] M. Abadi, M. Budiu, Úlfar Erlingsson, and J. Ligatti. A theory of secure control flow. In *International Conference on Formal Engineering Methods*, Nov. 2005.
- [3] A. W. Appel. Foundational proof-carrying code. In *Sixteenth Annual IEEE Symposium on Logic in Computer Science*, pages 247–258. IEEE, 2001.
- [4] H. Bar-El, H. Choukri, D. Naccache, M. Tunstall, and C. Whelan. The sorcerer’s apprentice guide to fault attacks. In *Proceedings of the IEEE*, volume 94, February 2006.
- [5] R. C. Baumann. Soft errors in advanced semiconductor devices-part I: the three radiation sources. *IEEE Transactions on Device and Materials Reliability*, 1(1):17–22, March 2001.
- [6] R. C. Baumann. Soft errors in commercial semiconductor technology: Overview and scaling trends. In *IEEE 2002 Reliability Physics Tutorial Notes, Reliability Fundamentals*, pages 121\_01.1 – 121\_01.14, April 2002.

- [7] E. Biham and A. Shamir. Differential fault analysis of secret key cryptosystems. pages 513–525, 1997.
- [8] J. Blömer, M. Otto, and J.-P. Seifert. A new crt-rsa algorithm secure against bellcore attacks. In *CCS '03: Proceedings of the 10th ACM conference on Computer and communications security*, pages 311–320, New York, NY, USA, 2003. ACM.
- [9] D. Boneh, R. A. DeMillo, and R. J. Lipton. On the importance of checking cryptographic protocols for faults. *Lecture Notes in Computer Science*, 1233:37–51, 1997.
- [10] E. Borin, C. Wang, Y. Wu, and G. Araujo. Software-based transparent and comprehensive control-flow error detection. In *CGO '06: Proceedings of the International Symposium on Code Generation and Optimization*, pages 333–345, Washington, DC, USA, 2006. IEEE Computer Society.
- [11] S. Borkar. Designing reliable systems from unreliable components: the challenges of transistor variability and degradation. In *IEEE Micro*, volume 25, pages 10–16, December 2005.
- [12] C.-L. Chen and M. Y. B. Hsiao. Error-correcting codes for semiconductor memory applications: A state-of-the-art review. *IBM Journal of Research and Development*, 28(2):124–134, 1984.
- [13] J. Chen. A typed intermediate language for compiling multiple inheritance. In *POPL '07: Proceedings of the 34th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 25–30, New York, NY, USA, 2007. ACM.
- [14] J. Chen, C. Hawblitzel, F. Perry, M. Emmi, J. Condit, D. Coetzee, and P. Pratikaki. Type-preserving compilation for large-scale optimizing object-oriented compilers.

- In *ACM Conference on Programming Language Design and Implementation*, June 2008.
- [15] J. Chen and D. Tarditi. A simple typed intermediate language for object-oriented languages. In *POPL '05: Proceedings of the 32nd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 38–49, New York, NY, USA, 2005. ACM.
- [16] J. Chen, D. Wu, A. W. Appel, and H. Fang. A provably sound TAL for back-end optimization. In *PLDI '03: Proceedings of the ACM SIGPLAN 2003 conference on Programming language design and implementation*, pages 208–219, New York, NY, USA, 2003. ACM.
- [17] K. 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.
- [18] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pages 137–150, 2004.
- [19] P. Dusart, G. Letourneux, and O. Vivolo. Differential fault analysis on A.E.S, 2003.
- [20] M. Elsman. Fault-tolerant voting in a simply-typed lambda calculus. Technical Report ITU-TR-2007-99, IT University of Copenhagen, June 2007.
- [21] X. Feng and Z. Shao. Modular verification of concurrent assembly code with dynamic thread creation and termination. In *ICFP '05: Proceedings of the tenth ACM SIGPLAN international conference on Functional programming*, pages 254–267, New York, NY, USA, 2005. ACM.

- [22] X. Feng, Z. Shao, Y. Dong, and Y. Guo. Certifying low-level programs with hardware interrupts and preemptive threads. In *PLDI '08: Proceedings of the 2008 ACM SIGPLAN conference on Programming language design and implementation*, pages 170–182, New York, NY, USA, 2008. ACM.
- [23] M. Gomaa, C. Scarbrough, T. N. Vijaykumar, and I. Pomeranz. Transient-fault recovery for chip multiprocessors. In *Proceedings of the 30th annual international symposium on Computer architecture*, pages 98–109. ACM Press, 2003.
- [24] S. Govindavajhala and A. Appel. Using memory errors to attack a virtual machine. In *Proceedings of the 2003 Symposium on Security and Privacy*, pages 153–165, May 2003.
- [25] R. Hamming. Error-detecting and error-correcting codes. In *Bell System Technical Journal*, volume 29(2), pages 147–160, 1950.
- [26] S. Hareland, J. Maiz, M. Alavi, K. Mistry, S. Walsta, and C. Dai. Impact of CMOS process scaling and SOI on the soft error rates of logic processes. In *Proceedings of the Symposium on VLSI Technology*, pages 73–74, 2001.
- [27] C. Hawblitzel, H. Huang, L. Wittie, and J. Chen. A garbage-collecting typed assembly language. In *TLDI '07: Proceedings of the 2007 ACM SIGPLAN international workshop on Types in languages design and implementation*, pages 41–52, New York, NY, USA, 2007. ACM.
- [28] P. Hazucha, T. Karnik, J. Maiz, S. Walstra, B. Bloechel, J. Tschanz, G. Dermer, S. Hareland, P. Armstrong, and S. Borkar. Neutron soft error rate measurements in a 90-nm cmos process and scaling trends in sram from 0.25- $\mu\text{m}$  to 90-nm

- generation. In *Proceedings of the 2003 Electron Devices Meeting*, pages 21.5.1–21.5.4, Dec 2003.
- [29] M. Y. B. Hsiao, W. C. Carter, J. W. Thomas, and W. R. Stringfellow. Reliability, availability, and serviceability of IBM computer systems: A quarter century of progress. *IBM Journal of Research and Development*, 25(5):453–465, 1981.
- [30] L. Jia, F. Spalding, D. Walker, and N. Glew. Certifying compilation for a language with stack allocation. In *IEEE Symposium on Logic in Computer Science*, pages 407–416, June 2005.
- [31] T. Karnik, B. Bloechel, K. Soumyanath, V. De, and S. Borkar. Scaling trends of cosmic ray induced soft errors in static latches beyond 0.18. In *Proceedings of the Symposium on VLSI Technology*, pages 61–62, 2001.
- [32] L. Lamport and S. Merz. Specifying and verifying fault-tolerant systems. In *ProCoS: Proceedings of the Third International Symposium Organized Jointly with the Working Group Provably Correct Systems on Formal Techniques in Real-Time and Fault-Tolerant Systems*, pages 41–76, London, UK, 1994. Springer-Verlag.
- [33] C. Lin, A. McCreight, Z. Shao, Y. Chen, and Y. Guo. Foundational typed assembly language with certified garbage collection. In *TASE '07: Proceedings of the First Joint IEEE/IFIP Symposium on Theoretical Aspects of Software Engineering*, pages 326–338, Washington, DC, USA, 2007. IEEE Computer Society.
- [34] J.-L. Lin, M. H. Dunham, and M. A. Nascimento. A survey of distributed database checkpointing. *Distributed and Parallel Databases*, 5(3):289–319, 1997.
- [35] A. Mahmood and E. J. McCluskey. Concurrent error detection using watchdog processors—a survey. *IEEE Transactions on Computers*, 37(2):160–174, 1988.

- [36] A. Meixner, M. E. Bauer, and D. Sorin. Argus: Low-cost, comprehensive error detection in simple cores. In *MICRO '07: Proceedings of the 40th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 210–222, Washington, DC, USA, 2007. IEEE Computer Society.
- [37] S. E. Michalak, K. W. Harris, N. W. Hengartner, B. E. Takala, and S. A. Wender. Predicting the number of fatal soft errors in Los Alamos National Laboratory's ASC Q computer. *IEEE Transactions on Device and Materials Reliability*, 5(3):329–335, September 2005.
- [38] G. Morrisett, K. Crary, N. Glew, and D. Walker. Stack-based Typed Assembly Language. *Journal of Functional Programming*, 12(1):43–88, Jan. 2002.
- [39] G. Morrisett, D. Walker, K. Crary, and N. Glew. From System F to Typed Assembly Language. In *Twenty-Fifth ACM Symposium on Principles of Programming Languages*, pages 85–97, San Diego, Jan. 1998.
- [40] G. Morrisett, D. Walker, K. Crary, and N. Glew. From System F to Typed Assembly Language. *ACM Transactions on Programming Languages and Systems*, 3(21):528–569, May 1999.
- [41] S. S. Mukherjee, M. Kontz, and S. K. Reinhardt. Detailed design and evaluation of redundant multithreading alternatives. In *Proceedings of the 29th Annual International Symposium on Computer Architecture*, pages 99–110. IEEE Computer Society, 2002.
- [42] G. Necula. Proof-carrying code. In *Twenty-Fourth ACM Symposium on Principles of Programming Languages*, pages 106–119, Paris, 1997.



- [43] G. Necula and P. Lee. Safe kernel extensions without run-time checking. In *Proceedings of Operating System Design and Implementation*, pages 229–243, Seattle, Oct. 1996.
- [44] G. C. Necula. Proof-carrying code. In *Proceedings of the 24th ACM Symposium on Principles of Programming Languages*, pages 106–119, Paris, Jan. 1997.
- [45] G. C. Necula. *Compiling with Proofs*. PhD thesis, Carnegie Mellon University, 1998.
- [46] T. J. O’Gorman, J. M. Ross, A. H. Taber, J. F. Ziegler, H. P. Muhlfeld, I. C. J. Montrose, H. W. Curtis, and J. L. Walsh. Field testing for cosmic ray soft errors in semiconductor memories. In *IBM Journal of Research and Development*, pages 41–49, January 1996.
- [47] N. Oh, P. P. Shirvani, and E. J. McCluskey. Control-flow checking by software signatures. In *IEEE Transactions on Reliability*, volume 51, pages 111–122, March 2002.
- [48] N. Oh, P. P. Shirvani, and E. J. McCluskey. Error detection by duplicated instructions in super-scalar processors. In *IEEE Transactions on Reliability*, volume 51, pages 63–75, March 2002.
- [49] F. Perry. Reasoning about software in the presence of transient faults – complete proofs. Technical Report TR-831-08, Princeton University, 2008.
- [50] F. Perry, C. Hawblitzel, and J. Chen. Simple and flexible stack types. In *International Workshop on Aliasing, Confinement, and Ownership*, July 2007.

- [51] F. Perry, L. Mackey, G. A. Reis, J. Ligatti, D. I. August, and D. Walker. Fault-tolerant typed assembly language. In *International Symposium on Programming Language Design and Implementation (PLDI)*, June 2007.
- [52] F. Perry, L. Mackey, G. A. Reis, J. Ligatti, D. I. August, and D. Walker. Fault-tolerant typed assembly language. Technical Report TR-776-07, Princeton University, 2007.
- [53] F. Perry and D. Walker. Reasoning about control flow in the presence of transient faults. Technical Report TR-799-07, Princeton University, 2007.
- [54] F. Perry and D. Walker. Reasoning about control flow in the presence of transient faults. In *International Static Analysis Symposium*, July 2008.
- [55] R. Phelan. Addressing soft errors in ARM core-based SoC. ARM White Paper, December 2003.
- [56] G. Piret and J.-J. Quisquater. A differential fault attack technique against spn structures, with application to the AES and KHAZAD. In *CHES*, pages 77–88, 2003.
- [57] S. K. Reinhardt and S. S. Mukherjee. Transient fault detection via simultaneous multithreading. In *Proceedings of the 27th Annual International Symposium on Computer Architecture*, pages 25–36. ACM Press, 2000.
- [58] G. A. Reis. *Software Modulated Fault Tolerance*. PhD thesis, Department of Electrical Engineering, Princeton University, Princeton, NJ, 2008.
- [59] G. A. Reis, J. Chang, and D. I. August. Automatic instruction-level software-only recovery methods. In *IEEE Micro Top Picks*, volume 27, January 2007.

- [60] G. A. Reis, J. Chang, N. Vachharajani, R. Rangan, and D. I. August. SWIFT: Software implemented fault tolerance. In *Proceedings of the 3rd International Symposium on Code Generation and Optimization*, March 2005.
- [61] G. A. Reis, J. Chang, N. Vachharajani, R. Rangan, D. I. August, and S. S. Mukherjee. Design and evaluation of hybrid fault-detection systems. In *Proceedings of the 32th Annual International Symposium on Computer Architecture*, pages 148–159, June 2005.
- [62] Z. Shao. An overview of the FLINT/ML compiler. In *Workshop on Types in Compilation*, Amsterdam, June 1997. ACM.
- [63] P. Shivakumar, M. Kistler, S. W. Keckler, D. Burger, and L. Alvisi. Modeling the effect of technology trends on the soft error rate of combinational logic. In *Proceedings of the 2002 International Conference on Dependable Systems and Networks*, pages 389–399, June 2002.
- [64] D. P. Siewiorek. Fault tolerance in commercial computers. *Computer*, 23(7):26–37, 1990.
- [65] T. J. Slegel, R. M. Averill III, M. A. Check, B. C. Giamei, B. W. Krumm, C. A. Krygowski, W. H. Li, J. S. Liptay, J. D. MacDougall, T. J. McPherson, J. A. Navarro, E. M. Schwarz, K. Shum, and C. F. Webb. IBM’s S/390 G5 Microprocessor design. In *IEEE Micro*, volume 19, pages 12–23, March 1999.
- [66] D. Tarditi, G. Morrisett, P. Cheng, C. Stone, R. Harper, and P. Lee. TIL: A type-directed optimizing compiler for ML. In *ACM Conference on Programming Language Design and Implementation*, pages 181–192, Philadelphia, May 1996.

- [67] M. Treaster. A survey of fault-tolerance and fault-recovery techniques in parallel systems, 2005.
- [68] M. Tremblay and Y. Tamir. Support for fault tolerance in VLSI processors. In *Proceedings of the IEEE International Symposium on Circuits and Systems*, volume 1, pages 388–392, May 1989.
- [69] S. Triantafyllis, M. J. Bridges, E. Raman, G. Ottoni, and D. I. August. A framework for unrestricted whole-program optimization. In *ACM SIGPLAN 2006 Conference on Programming Language Design and Implementation*, pages 61–71, June 2006.
- [70] D. Tullsen, S. Eggers, and H. Levy. Simultaneous multithreading: Maximizing on-chip parallelism. In *Proceedings of the 22nd International Symposium on Computer Architecture*, June 1995.
- [71] T. N. Vijaykumar, I. Pomeranz, and K. Cheng. Transient-fault recovery using simultaneous multithreading. In *Proceedings of the 29th Annual International Symposium on Computer Architecture*, pages 87–98. IEEE Computer Society, 2002.
- [72] B. D. Vito and R. Butler. Provable transient recovery for frame-based, fault-tolerant computing systems, 1992.
- [73] D. Walker, L. Mackey, J. Ligatti, G. A. Reis, and D. I. August. Static typing for a faulty lambda calculus. In *ACM International Conference on Functional Programming*, Portland, Oregon, Sept. 2006.
- [74] A. Wood. Data integrity concepts, features, and technology. White Paper, Tandem Division, Compaq Computer Corporation, 1999.

- [75] Y. Yeh. Triple-triple redundant 777 primary flight computer. In *Proceedings of the 1996 IEEE Aerospace Applications Conference*, volume 1, pages 293–307, February 1996.
- [76] J. F. Ziegler and H. Puchner. *SER - History, Trends, and Challenges: A Guide for Designing with Memory ICs*. 2004.