Fault-tolerant Typed Assembly Language

Frances Perry† Lester Mackey† George A. Reis† Jay Ligatti‡ David I. August† David Walker†

†Departments of Computer Science and Electrical Engineering
Princeton University
{frances, lmackey, gareis, august, dpw}@cs.princeton.edu

‡Department of Computer Science and Engineering
University of South Florida
ligatti@cse.usf.edu

Abstract

A *transient hardware fault* occurs when an energetic particle strikes a transistor, causing it to change state. Although transient faults do not permanently damage the hardware, they may corrupt computations by altering stored values and signal transfers. In this paper, we propose a new scheme for provably safe and reliable computing in the presence of transient hardware faults. In our scheme, software computations are replicated to provide redundancy while special instructions compare the independently computed results to detect errors before writing critical data. In stark contrast to any previous efforts in this area, we have analyzed our fault tolerance scheme from a formal, theoretical perspective. To be specific, first, we provide an operational semantics for our assembly language, which includes a precise formal definition of our fault model. Second, we develop an assembly-level type system designed to detect reliability problems in compiled code. Third, we provide a formal specification for program fault tolerance under the given fault model and prove that all well-typed programs are indeed fault tolerant. In addition to the formal analysis, we evaluate our detection scheme and show that it only takes 34% longer to execute than the unreliable version.

Categories and Subject Descriptors D.3.1 [*Programming Languages*]: Formal Definitions and Theory; B.8.1 [*Performance and Reliability*]: Reliability, Testing, and Fault-Tolerance

General Terms Languages, Reliability, Theory, Verification

Keywords transient hardware faults, soft faults, fault tolerance, type systems, typed assembly language

1. Introduction

A *transient fault* or *soft error* is a temporary hardware failure that alters a signal transfer, a register value, or some other processor component. While transient faults are temporary, they corrupt computations and have led to costly failures in high-end systems in recent years. For example, in 2000 there were reports that transient faults caused crashes at a number of Sun's major customer sites, including America Online and eBay [2]. Later, Hewlett Packard admitted multiple problems in the Los Alamos Labs supercomputers due to transient faults [7]. Finally, Cypress Semiconductor has confirmed "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" [28].

Unfortunately, while soft errors can already cause substantial reliability problems, current trends in hardware design suggest that fault rates will increase in the future. More specifically, faster clock rates, increasing transistor density, decreasing voltages and smaller feature sizes all contribute to increasing fault rates [1, 11, 21]. Due to a combination of these factors, fault rates in modern processors have been increasing at a rate of approximately 8% per generation [3].

These trends are well known in the architecture and compiler communities, and, consequently, many solutions to the threat of soft errors have been proposed. At a high level, all of these solutions involve adding redundancy to computations in one way or another, but the specifics vary substantially. For instance, there are proposals involving hardware-only solutions such as error-correcting codes, watchdog co-processors [6] and redundant hardware threads [4, 9, 16, 25] as well as software-only techniques that use both single and multiple cores [12, 13, 17, 18, 20, 24]. Broadly speaking, if the technique can scale, hardware-only solutions are more efficient for a single, fixed reliability policy, but software-only solutions are more flexible (they may be deployed exactly when, where, and to the degree needed) and less costly in terms of hardware. In an attempt to gain some of the best of both worlds, researchers have also recently proposed hybrid software-hardware solutions involving strong fault tolerance mechanisms implemented in hardware but controlled by the software running on the processor [19].

Software-only and hybrid hardware-software techniques also possess at least one further, little-mentioned drawback — *they may not actually work*. To be fair, many of these techniques appear extremely promising. However, as far as we are aware, the published transient fault-tolerance techniques come with no rigorous proofs that they guarantee any particular reliability properties. In general, researchers satisfy themselves with presenting an algorithm for fault-tolerance and leave the audience to judge for themselves whether or not the algorithm is correct. In fact, the literature does not even precisely define what it might mean for an assembly-level program to be fault tolerant. This paper tackles this gaping hole in the existing literature by defining a new hybrid hardware-software technique for tolerating transient faults, and, unlike any previous work, actually proving it has strong fault-tolerance properties.

The specification and proof of fault tolerance comes in several stages. First, before proving any particular properties, it is necessary to define a fault model precisely. Most of the current literature uses the *Single Event Upset (SEU) Model*, which states that only one fault may occur during execution [16, 19, 26]. However, the details of exactly where and when faults may occur are usually given in English. We also assume the SEU model, but we specify exactly where by including faulty transitions as formal rules in the operational semantics of our assembly language.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PLDI'07 June 11–13, 2007, San Diego, California, USA.

Copyright © 2007 ACM 978-1-59593-633-2/07/0006...$5.00.

Second, it is necessary to state precisely what "fault tolerance" actually means. Abstractly, a program is fault-tolerant if no fault can change the observable behavior of a program. More concretely, we assume our system operates in the presence of a memory-mapped output device, and hence a program is *not* fault-tolerant if a fault can cause a deviation in the sequence of values written to memory. We formalize this property more precisely as a mathematical theorem that relates faulty and non-faulty executions of a program.

Third, it is necessary to provide a technique for actually proving that specific programs are fault tolerant relative to the fault model. Our proof technique is presented in the form of a type system. All well-typed programs satisfy variants of the standard progress and preservation lemmas, even in the presence of transient faults, as well as the stronger fault tolerance property mentioned above. In addition to being theoretically important as a proof technique for fault tolerance, the type system can be used to debug compilers that intend to generate reliable code. If the output from these compilers type check, their code will have strong fault tolerance guarantees. In the past, researchers have proposed testing compiler outputs using fault injection techniques that randomly insert errors into programs. However, using a type checker in this case is a much better idea. In principle, a conventional testing technique would need to test all combinations of features *in conjunction with all combinations of faults*, causing an explosion in the number of test cases, and yet still failing to achieve perfect fault coverage in practice. By using the type checker we have designed, one achieves perfect fault coverage relative to the fault model without needing to increase the compiler test suite.

The rest of this paper presents the details of our hybrid hardware-software fault-tolerance technique. Section 2 presents the syntax and operational semantics of the new, idealized assembly language we have designed for fault tolerance. It is a RISC-based architecture with special instructions to facilitate reliable communication with memory and to detect control-flow faults. Section 3 presents the key principles and formal definitions for the fault-tolerant assembly language type system ($\mathrm{TAL_{FT}}$ for short). Though the typing rules are specific to our particular setting, the underlying principles are more general; we believe many of these principles will apply to reasoning about related fault-tolerant systems. Our innovative combination of a TAL-like type-theory with concepts from classical Hoare Logics is a particularly general and important technical contribution. Section 4 describes the key theorems we have proven including Progress, Preservation, "No False Positives," and Fault Tolerance. Section 5 provides empirical evidence that our new hybrid solution to fault tolerance is feasible for many applications by measuring performance results on simulated hardware. Related work is discussed in more detail in Section 6. Due to space considerations, some of the technical details and all of the proofs have been omitted. A companion technical report [15] contains the complete specification of our system and a relatively detailed proof outline.

2. The Faulty Hardware

The faulty 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, since recovery is largely orthogonal to detection, we omit the former, focusing only on the latter in this paper.

The general strategy of every fault-tolerant program is to maintain two redundant and independent threads of computation, a *green* (G) computation and a *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. This methodology has been show in the literature as an effective implementation of fault tolerance [13, 18], and we expand on this style of implementation by formalizing the fault model and coverage.

The execution of assembly programs is specified using a small-step operational semantics that maps *machine states* (Σ) 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 meta variable a ranges over all sorts of registers, and meta variable r ranges only over general-purpose registers (r_1, r_2, ...). 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 *destination register*, d. 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.[1]

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.[2] 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 (Σ) 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 four components are as discussed above, and ir is either an instruction i to be executed, or "$\cdot$" indicating the next instruction should be fetched from code memory. Figure 1 summarizes the syntax of machine states. Here and elsewhere in the paper, we use overbar notation to indicate a sequence of objects.

2.1 The Fault Model

The operational semantics is designed both to model proper execution of machine instructions and to make perfectly explicit, precise, and transparent all of our assumptions about when and where faults may occur. The central operational judgment has the form $\Sigma_1 \longrightarrow^{s}_{k} \Sigma_2$, which expresses a single step transition from state Σ_1 to state Σ_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 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.

[1] In contrast, the tags on instruction opcodes, to be introduced momentarily, *do* have an effect on evaluation.

[2] Address 0 is not considered a valid code address.

$$
\begin{array}{llll}
\textit{colors} & c & ::= & G \mid B \\
\textit{colored values} & v & ::= & c\,n \\
\textit{registers} & r & ::= & r_n \\
\textit{general regs} & a & ::= & r \mid d \mid pc_c \\
\textit{register file} & R & ::= & \cdot \mid R, a \to v \\
\textit{code memory} & C & ::= & \cdot \mid C, n \to i \\
\textit{value memory} & M & ::= & \cdot \mid M, n \to n \\
\textit{store queue} & Q & ::= & \overline{(n,n)} \\
\textit{ALU ops} & op & ::= & add \mid sub \mid mul \\
\textit{instructions} & i & ::= & op\ r_d, r_s, r_t \mid op\ r_d, r_s, v \\
 & & \mid & ld_c\ r_d, r_s \mid st_c\ r_d, r_s \mid mov\ \ r_d, v \\
 & & \mid & bz_c\ \ r_z, r_d \mid jmp_c\ \ r_d \\
\textit{inst register} & ir & ::= & i \mid \cdot \\
\textit{state} & \Sigma & ::= & (R, C, M, Q, ir) \mid fault
\end{array}
$$

Figure 1. Syntax of instructions and machine states.

Different fault-tolerance techniques protect different components of machines. In the literature, the protected areas are usually inside the *Sphere of Replication* (SoR) [16]. 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.

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

$$\frac{\begin{array}{c} Q_1 = \overline{(n_1, n_1')}, (m_1, m'), \overline{(n_2, n_2')} \\ Q_2 = \overline{(n_1, n_1')}, (m_2, m'), \overline{(n_2, n_2')} \end{array}}{(R, C, M, Q_1, ir) \longrightarrow_1 (R, C, M, Q_2, ir)}\ (\textit{Q-zap1})$$

$$\frac{\begin{array}{c} Q_1 = \overline{(n_1, n_1')}, (m, m_1'), \overline{(n_2, n_2')} \\ Q_2 = \overline{(n_1, n_1')}, (m, m_2'), \overline{(n_2, n_2')} \end{array}}{(R, C, M, Q_1, ir) \longrightarrow_1 (R, C, M, Q_2, ir)}\ (\textit{Q-zap2})$$

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. Since the color tag is fictional (has no effect on run-time behavior), this poses no limitation on the fault model. Rules Q_1*-zap* and Q_2*-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-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 equivalent to correct execution of an instruction composed with a fault either immediately before or afterwards. 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.

An important benefit of our formal model is that there is actually some 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.

Instruction Fetch:

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

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

Basic Instructions:

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

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

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

Figure 2. Operational rules for basic instructions.

2.2 Instruction Semantics

The syntax of machine instructions was presented along with the rest of the components of our abstract machine in Figure 1. The semantics is described formally by the inference rules in Figures 2, 3, and 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 $find(Q, n)$ 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 alternatively fetching an instruction from code memory and executing that instruction. 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). Fortunately, however, 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. Fault recovery is an orthogonal issue to fault detection, so we leave it unspecified here. 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*) are completely standard. The first arithmetic operation $op\ \ r_d, r_s, r_t$ performs *op* on the values in r_s and r_t, storing the result in r_d. The second arithmetic operation uses a constant operand v in addition to r_s and r_d. 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 our 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.

$\boxed{\Sigma \longrightarrow^{s}_{k} \Sigma'}$

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

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

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

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

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

Figure 3. Selected operational rules for memory instructions.

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. This is the conundrum of the *Time-Of-Check-Time-Of-Use* (TOCTOU) fault.

To avoid TOCTOU faults, our machine possesses a modified store buffer (the queue Q), which is similar to the store buffer used in previous hardware [16] and hybrid [19] fault tolerant systems. In addition, there are two special storage instructions, each tagged with a color. The green store instruction $st_G\ r_d, r_s$ places the address-value pair $(R_{val}(r_d), R_{val}(r_s))$ on the front of the queue (rule *st_G-queue*). The blue store instruction $st_B\ r_d, r_s$ 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_B-mem*). If the pairs are different, the hardware signals a fault. Failure rules appear in Appendix A.1. Since green stores must always come before blue stores, instruction scheduling is somewhat constrained. As we will show later in Section 5, we have evaluated the performance both 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 instruction 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, and, of course, 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, 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. 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 or might result in loading some arbitrary value. Failure rules that model both possibilities appear in Appendix A.1.

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 [6], but that makes both versions of the control flow explicit as in software-only control flow protection [12, 18].

To achieve an unconditional jump, one executes a *jmp_G* instruction first and a related *jmp_B* instruction at some point in the future. A *$jmp_G\ r_1$* moves the destination address from r_1 into the special destination register d (rule *jmp_G*). 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 *$jmp_B\ r_2$* instruction, the contents of r_2 are compared to the contents of the destination register and if they are equal, control jumps to that location (rule *jmp_B*). If the addresses are different, the hardware detects a fault (see rule *jmp_B-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 5 will show that this scheduling constraint has only a minimal performance impact.

$\boxed{\Sigma \longrightarrow_k^s \Sigma'}$

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

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

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

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

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

$$\frac{\begin{array}{c}R_{val}(d) = 0 \qquad R_{val}(r_z) = 0 \\ R' = R{+}{+}[d \mapsto R(r_d)]\end{array}}{(R, C, M, Q, bz_G\ r_z, r_d) \longrightarrow_0 (R', C, M, Q, \cdot)}\ (bz_G\text{-}taken)$$

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

Figure 4. Selected operational rules for control flow instructions.

The following code illustrates a typical control-flow transfer.

```
1 ldG r1, r2
3 ldB r3, r4
2 jmpG r1
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 d (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 d, 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*). The rules for the associated failure cases appear in Appendix A.1. 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).

Static Expressions

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

Types

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

Contexts

heap typing	Ψ	::=	$\cdot \mid \Psi, n : b$
static context	Θ	::=	$\Delta; \Gamma; \overline{(E_d, E_s)}; E_m$

Figure 5. TAL_{FT} type syntax.

3. 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 a memory-mapped output device can never read a corrupt value and make it visible to a user. We call this property "fault tolerance."

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. Of course, as with other sorts of typed assembly language or proof-carrying code, this fact presents no particular difficulty in practice — it is easy for a compiler to generate sufficient "typing hints" to make type reconstruction trivial. For the reader's reference, the objects used in the type system are presented in Figure 5.

3.1 Static Expressions

Our "type system" is actually a combination of two theories, one being a relatively simple type theory for assembly, inspired by previous work on TAL [8], and the second being a Hoare Logic, designed to enforce the more precise invariants required for strong fault tolerance. The latter component requires we define a language of *static expressions* for reasoning about values and storage.

For the purposes of this paper, the static expressions are drawn from the standard theory of arithmetic and arrays used in many classical Hoare Logics (*c.f.*, Necula's thesis [10]). These static expressions are classified as either integers (kind κ_{int}) or memories (kind κ_{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 Δ is a mapping from variables to kinds, and the judgment $\Delta \vdash E : \kappa$ classifies expression E as having kind κ. The judgment $\Delta \vdash S : \Delta'$ holds when the substitution S maps variables in $Dom(\Delta')$ to values well-formed in Δ 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 $[\![E]\!]$ supplies the denotation of the closed static expression E as either an integer or a memory, depending on its kind. The definitions for $[\![E]\!]$ and $\Delta \vdash E_1 = E_2$ are shown in Appendix A.2, and the remaining judgments are defined in the companion technical report [15].

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

$$\frac{}{\Psi \vdash\ n : \mathit{int}}\ (\textit{int-t}) \qquad \frac{}{\Psi \vdash\ n : \Psi(n)}\ (\textit{base-t})$$

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

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

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

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

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

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

Figure 6. Value Typing.

3.2 Value Typing

Since 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 $\Theta \rightarrow \mathit{void}$ are pointers to code (continuations). One must satisfy the precondition Θ 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 kind of singleton type.

Value Typing Judgment. The value typing judgment has the form $\Psi; \Delta \vdash^{Z}\ v : t$, where Ψ maps heap addresses to basic types, and Δ contains the free expression 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 $(E' = 0 \Rightarrow \langle G, \Theta \rightarrow \mathit{void}, E'_r\rangle)$. When the static expression E' is equal to zero, values of this type also have type $\langle G, \Theta \rightarrow \mathit{void}, E'_r\rangle$. 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 $\langle c, \mathit{int}, E_2\rangle$ 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 Γ_2 be a supertype of the corresponding register in Γ_1. Note that here is no required relationship between the special registers d, pc_G, and pc_B. The rules for these judgments appear in the companion technical report [15].

3.3 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 much the same manner as previous typed assembly languages [8]. 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 Ψ, which only contains invariant heap typing assumptions, Θ contains fine-grained context-sensitive information about the current state of memory and the register file. More specifically, Θ consists of the following subcontexts: (1) Δ, which describes the free expres-

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

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

$$\frac{\Gamma(r_s)=\langle c,\mathit{int},E'_s\rangle \qquad \Gamma(r_t)=\langle c,\mathit{int},E'_t\rangle}{\Psi;(\Delta;\Gamma;\overline{(E_d,E_s)};E_m)\vdash\ op\ \ r_d,r_s,r_t\ \Rightarrow(\Delta;\Gamma{+}{+}[r_d\mapsto\ \langle c,\mathit{int},E'_s\ op\ E'_t\rangle];\overline{(E_d,E_s)};E_m)}\ (\textit{op2r-t})$$

$$\frac{\Gamma(r_s)=\langle c,\mathit{int},E'_s\rangle}{\Psi;(\Delta;\Gamma;\overline{(E_d,E_s)};E_m)\vdash\ op\ \ r_d,r_s,c\,n\ \Rightarrow(\Delta;\Gamma{+}{+}[r_d\mapsto\ \langle c,\mathit{int},E'_s\ op\ n\rangle];\overline{(E_d,E_s)};E_m)}\ (\textit{op1r-t})$$

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

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

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

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

$$\frac{\begin{array}{c}\Gamma(r_d)=\langle B,b\ \mathit{ref},E''_d\rangle \quad \Gamma(r_s)=\langle B,b,E''_s\rangle\\ \Delta\vdash\ E'_d=E''_d \quad \Delta\vdash\ E'_s=E''_s\end{array}}{\Psi;(\Delta;\Gamma;\overline{(E_d,E_s)},(E'_d,E'_s);E_m)\vdash\ st_B\ r_d\ r_s\ \Rightarrow(\Delta;\Gamma{+}{+};\overline{(E_d,E_s)};\mathit{upd}\ E_m\ E'_d\ E'_s)}\ (\textit{st}_B\textit{-t})$$

$$\frac{\begin{array}{c}\Gamma(d)=\langle G,\mathit{int},0\rangle \qquad \Gamma(r_z)=\langle G,\mathit{int},E_z\rangle\\ \Gamma(r_d)=\langle G,\Theta\to\mathit{void},E'_d\rangle \qquad \Theta=(\Delta';\Gamma';\overline{(E'_d,E'_s)};E'_m) \qquad \Gamma'(d)=\langle G,\mathit{int},0\rangle\end{array}}{\Psi;(\Delta;\Gamma;\overline{(E_d,E_s)};E_m)\vdash\ bz_G\ r_z\ r_d\ \Rightarrow(\Delta;\Gamma{+}{+}[d\mapsto\ E_z=0\Rightarrow\langle G,\Theta\to\mathit{void},E'_d\rangle];\overline{(E_d,E_s)};E_m)}\ (\textit{bz}_G\textit{-t})$$

$$\frac{\begin{array}{c}\Gamma(r_d)=\langle G,\Theta\to\mathit{void},E_{rd'}\rangle \qquad \Theta=(\Delta';\Gamma';\overline{(E'_d,E'_s)};E'_m)\\ \Gamma(d)=\langle G,\mathit{int},0\rangle \qquad \Gamma'(d)=\langle G,\mathit{int},0\rangle\end{array}}{\Psi;(\Delta;\Gamma;\overline{(E_d,E_s)};E_m)\vdash\ jmp_G\ r_d\ \Rightarrow(\Delta;\Gamma{+}{+}[d\mapsto\ \langle G,\Theta\to\mathit{void},E_{rd'}\rangle];\overline{(E_d,E_s)};E_m)}\ (\textit{jmp}_G\textit{-t})$$

$$\frac{\begin{array}{c}\Gamma(r_z)=\langle B,\mathit{int},E_z\rangle\\ \Gamma(r_d)=\langle B,(\Delta';\Gamma';\overline{(E'_d,E'_s)};E'_m)\to\mathit{void},E_r\rangle\\ \Gamma(d)=E'_z=0\Rightarrow\langle G,(\Delta';\Gamma';\overline{(E'_d,E'_s)};E'_m)\to\mathit{void},E'_r\rangle\\ \Delta\vdash\ E_z=E'_z\\ \Delta\vdash\ E_r=E'_r\\ \exists S.\Delta\vdash\ S:\Delta'\\ S(\Gamma')(d)=\langle G,\mathit{int},0\rangle\\ S(\Gamma')(pc_G)=\langle G,\mathit{int},E'_r\rangle\\ S(\Gamma')(pc_B)=\langle B,\mathit{int},E_r\rangle\\ \Delta\vdash\Gamma\le\ S(\Gamma')\\ \Delta\vdash\ \overline{(E_d,E_s)}=S(\overline{(E'_d,E'_s)})\\ \Delta\vdash\ E_m=S(E'_m)\end{array}}{\Psi;(\Delta;\Gamma;\overline{(E_d,E_s)};E_m)\vdash\ bz_B\ r_z\ r_d\ \Rightarrow\ (\Delta;\Gamma{+}{+};\overline{(E_d,E_s)};E_m)}\ (\textit{bz}_B\textit{-t})$$

$$\frac{\begin{array}{c}\Gamma(d)=\langle G,(\Delta';\Gamma';\overline{(E'_d,E'_s)};E'_m)\to\mathit{void},E'_r\rangle\\ \Gamma(r_d)=\langle B,(\Delta';\Gamma';\overline{(E'_d,E'_s)};E'_m)\to\mathit{void},E_r\rangle\\ \Delta\vdash\ E_r=E'_r\\ \exists S.\Delta\vdash\ S:\Delta'\\ S(\Gamma')(d)=\langle G,\mathit{int},0\rangle\\ S(\Gamma')(pc_G)=\langle G,\mathit{int},E'_r\rangle\\ S(\Gamma')(pc_B)=\langle B,\mathit{int},E_r\rangle\\ \Delta\vdash\Gamma\le\ S(\Gamma')\\ \Delta\vdash\ \overline{(E_d,E_s)}=S(\overline{(E'_d,E'_s)})\\ \Delta\vdash\ E_m=S(E'_m)\end{array}}{\Psi;(\Delta;\Gamma;\overline{(E_d,E_s)};E_m)\vdash\ jmp_B\ r_d\ \Rightarrow\mathit{void}}\ (\textit{jmp}_B\textit{-t})$$

Figure 7. Instruction Typing.

sion variables appearing in the other context-sensitive objects, (2) Γ, 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 Θ', 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 Γ++ adds one to the static expression associated with each program counter register in Γ. The expression $\overline{upd}\ E_m\ \overline{(E_d, E_s)}$ is $(upd\ (...(upd\ E_m\ E_{d_k}\ E_{s_k})...)\ E_{d_1}\ E_{s_1})$ when $\overline{(E_d, E_s)} = ((E_{d_1}, E_{s_1}), ..., (E_{d_k}, E_{s_k}))$. Figure 7 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 both values 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 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.

Typing Control-Flow Instructions. While the typing rules for control-flow instructions have many premises, they 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.

The simplest rule involves the green unconditional jump. This instruction is just a move from register r_d to the special destination register d. The type of register d is updated to the type of r_d (obeying both principles 1 and 2). The rule contains constraints that d must be equal to 0 in both Γ and Γ' since the hardware resets the destination register to 0 after a jump.

The blue unconditional jump is a true jump. According to principle 1, it checks the standard typing invariants needed to ensure safety in any typed assembly language, including (1) that the jump target has 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 Δ from the code type (see the final seven premises).

The typing of the conditional branches is quite similar to that of unconditional jumps. One difference is that the bz_G instruction is now a conditional move as 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.

3.4 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 Σ are shown in Figure 8 and explained below.

Register File Typing. The judgment $\Psi \vdash^Z R : \Gamma$ states that the register file R has the register file type Γ under heap typing Ψ and a zap tag Z. The contents of each register must have the type given to that register by Γ. Each program counter must have the appropriate color, and the program counters must compute 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 Ψ. 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 Θ' (meaning that control may fall through to the next instruction) then the subsequent instruction must be well typed using Θ' as its precondition.

Memory Typing. The judgment $\Psi \vdash M : E_m$ states that given heap typing Ψ the value memory M is well-formed and can be described by the static expression E_m. The static expression E_m must have kind κ_{mem}, and M must be the denotation of E_m. Each location in the domain of M must have a type $b\ ref$ 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 Ψ 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 $b\ ref$ 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 κ_{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 Σ is well-typed under zap tag Z. This judgment holds when Σ is a five-tuple (R, C, M, Q, ir), and these elements are each well-typed and consistent with each other. Note that Σ is not well-typed when it is the fault state $fault$.

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

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

$\boxed{\Psi \vdash C}$

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

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

$$\frac{\begin{array}{c}\cdot \vdash E_m : \kappa_{mem} \qquad [\![E_m]\!] = M \\ \forall \ell \in Dom(M).\ \Psi \vdash \ell : b\,ref \ \wedge\ \Psi \vdash M(\ell) : b\end{array}}{\Psi \vdash M : E_m} \ (M\text{-}t)$$

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

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

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

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

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

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

Figure 8. Machine State Typing.

4. Formal Results

In order to prove properties of our type system, we extend our single-step transition $\Sigma_1 \longrightarrow^s_k \Sigma_2$ from Section 2 to a sequence of n transitions containing exactly k faults $\Sigma_1 \xrightarrow{n}{}^s_k \Sigma_2$, where n is greater than or equal to zero, and k is still either 0 or 1.

4.1 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 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 1 (Progress)
1. *If* $\vdash \Sigma$ *then* $\Sigma \longrightarrow^s_0 \Sigma'$ *and* $\Sigma' \neq fault$.
2. *If* $\vdash^c \Sigma$ *then* $\Sigma \longrightarrow^s_0 \Sigma$.

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* $\Sigma \longrightarrow^s_0 \Sigma'$ *and* $\Sigma' \neq fault$ *then* $\vdash^Z \Sigma'$.
2. *If* $\vdash \Sigma$ *and* $\Sigma \longrightarrow^s_1 \Sigma'$ *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* $\forall\, n.\ \Sigma \xrightarrow{n}{}^s_0 \Sigma'$ *and* $\vdash \Sigma'$.

4.2 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 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.

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 9.

Using the similarity relations, we can state and prove the fault tolerance theorem for well-typed programs precisely. Assume that machine state Σ is well-typed under the empty zap tag, and non-faulty execution of Σ for n steps results in a state Σ' 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 Σ'_f, then Σ' simulates Σ'_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.

$\boxed{v_1\ sim^Z\ v_2}$

$$\frac{}{C\ n\ sim^Z\ C\ n}\ (\textit{sim-val}) \qquad \frac{}{C\ n\ sim^C\ C\ n'}\ (\textit{sim-val-zap})$$

$\boxed{R\ sim^Z\ R'}$

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

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

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

$$\frac{G\ n_1\ sim^Z\ G\ n_1' \quad G\ n_2\ sim^Z\ G\ n_2' \quad Q\ sim^Z\ Q'}{((n_1, n_2), Q)\ \ sim^Z\ ((n_1', n_2'), Q')}\ (\textit{sim-Q})$$

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

$$\frac{R\ sim^Z\ R' \qquad Q\ sim^Z\ Q'}{(R, C, M, Q, ir)\ sim^Z\ (R', C, M, Q', ir)}\ (\textit{sim-}\Sigma)$$

Figure 9. Similarity of Machine States.

Theorem 4 (Fault Tolerance)
If $\vdash \Sigma$ *and* $\Sigma \xrightarrow{n}_0 {}^{s} \Sigma'$ *then either* $\Sigma \xrightarrow{(n+1)}_1 {}^{s'} \Sigma'_f$ *or* $\exists m \leq (n+1)\ .\ \Sigma \xrightarrow{m}_1 {}^{s'}$ *fault, and*

1. *For all derivations* $\Sigma \xrightarrow{(n+1)}_1 {}^{s'} \Sigma'_f$ *where* $\Sigma'_f \neq$ *fault.* $s' = s$ *and* $\exists\ c.\ \Sigma'\ sim^c\ \Sigma'_f$.
2. *For all derivations* $\Sigma \xrightarrow{m}_1 {}^{s'}$ *fault where* $m \leq (n+1)$. s' *is a prefix of* s.

5. Performance

To better understand how TAL_{FT} can be applied to real world situations, we simulated the TAL_{FT} 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_{FT}. While not an exact representation of the performance of TAL_{FT}, simulating the performance of TAL_{FT} applied to this architecture will give guidance as to the feasibility of this system in a real architecture.

To evaluate the performance impact of our techniques, a version of the VELOCITY compiler [23] 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.

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 10 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, we find that 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 [18] that show the degradation due to redundant code to be less than double.

As alluded to in Section 2.2, Figure 10 compares the performance degradation both 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, our compiler was modified to produce code that had more flexibility in the scheduling of the green and blue versions. We then simulated a more aggressive hardware implementation that could correlate the original and redundant memory operations regardless of the executed order. As expected, this version has better performance (in most cases) than the unconstrained code. 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%. Although the colored ordering restriction of TAL_{FT} may seem costly, removing this restriction provides only a small improvement.

6. Related Work

Fault tolerance based on software replication is a well-populated field with decades of history. TAL_{FT} differs from previous approaches in that it provides a type-theoretic framework for obtaining strong guarantees about the reliability of machine code.

Most closely related to TAL_{FT} is our previous work on λ_{zap}, a highly abstract type-theoretic model for studying the basic principles of fault tolerance in the lambda calculus [26]. There are two important distinctions between TAL_{FT} and λ_{zap}. First, λ_{zap}, working at the level of the lambda calculus, is very far removed from real machine code. For instance, it lacks a program counter, a register file, memory, and load or store instructions. Memory references in particular constitute a key challenge in the current technical work. Second, the properties of the λ_{zap} type system are relatively weak compared with the properties of the current type system. The "end-to-end" fault tolerance property proven for λ_{zap} depends not only on the type system but also the nature of the translation from the ordinary simply-typed lambda calculus. In contrast, the type system of TAL_{FT} is much stronger, capable of ensuring a strong fault tolerance property independently of the process that compiles the code.

Also closely related to TAL_{FT} is the original TAL system, which first applied strong type checking to machine code to guarantee its safety [8]. TAL operates under the assumption of nonfaulty hardware and therefore ignores the major issues of reliability on which this paper has focused.

There have been various implementations of software-only, hardware-only, and hybrid techniques for transient fault mitigation. Hardware techniques have a long history of using very localized bit-level techniques like error correcting codes or parity bits additions. These techniques are efficient for storage structures like memory, but are costly or impossible to apply to other processor elements like pipeline latches or arithmetic units. Higher level techniques are used when protection is necessary for larger segments of the processor. These techniques include the duplication of coarse-grained structures such as functional units, processor cores [5, 22, 27], or hardware contexts [9, 16, 25].

To provide protection when the hardware costs of these approaches are prohibitive, software-only approaches have been proposed as alternatives [12, 13, 17, 18, 20, 24]. While software-only systems are cheaper to deploy and can be configured after deployment, they cannot achieve the same performance or reliability as

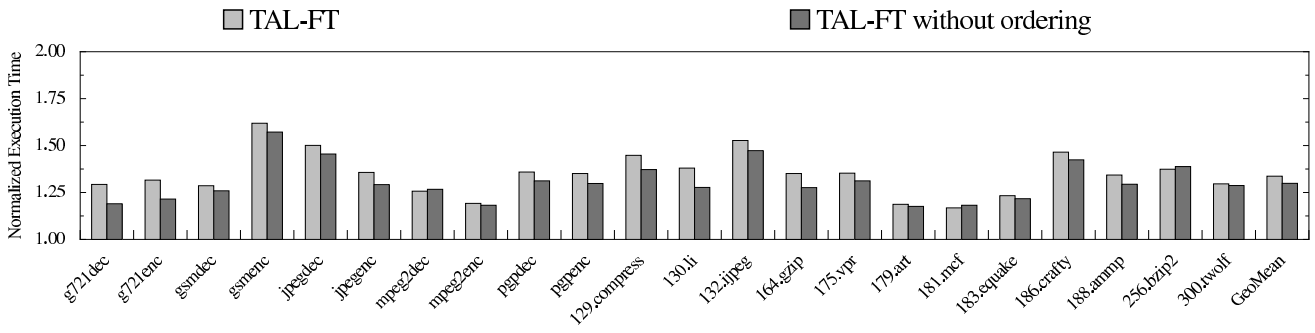

Figure 10. Performance Normalized to Unprotected Version.

hardware-based techniques, since they have to execute additional instructions and are unable to examine microarchitectural state. Despite these limitations, software-only techniques have shown promise, in the sense that they can significantly improve reliability with reasonable performance overhead [12, 13, 18].

TAL_{FT} attempts to exploit the benefits of both sorts of systems by using a hybrid approach to fault tolerance. There have been previous hybrid approaches to transient fault tolerance, some focusing solely on control-flow protection [14] and recently others looking at full processor protection [19]. This work differs from those previous approaches because regardless of the type of implementation, software, hardware, or hybrid, none of those previous approaches have given rigorous formal proofs of the correctness of their systems.

7. Conclusions

In conclusion, transient faults are already a significant cause for concern at major semiconductor manufacturers and threaten to be more so in the coming years and decades. This paper takes one step forward for the science of fault tolerance by presenting a principled and practical hybrid software-hardware scheme for detecting transient faults. More specifically, we identify four general principles for verifying correctness of fault tolerant systems and capture these in an assembly language type system. From a theoretical perspective, the type system acts as a sound proof technique for verifying reliability properties of programs. From a practical perspective, it can be used as a debugging aid within a compiler, strictly dominating any conventional testing technique. Our two main formal results show that a single fault affecting observable behavior in a well-type 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.

Acknowledgments

This research is funded in part by NSF awards CNS-0627650, CNS-0615250, and CCF 0633268. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

References

[1] 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.

[2] 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.

[3] 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.

[4] 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.

[5] R. W. Horst, R. L. Harris, and R. L. Jardine. Multiple instruction issue in the NonStop Cyclone processor. In *Proceedings of the 17th International Symposium on Computer Architecture*, pages 216–226, May 1990.

[6] A. Mahmood and E. J. McCluskey. Concurrent error detection using watchdog processors-a survey. *IEEE Transactions on Computers*, 37(2):160–174, 1988.

[7] 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 Labratory's ASC Q computer. *IEEE Transactions on Device and Materials Reliability*, 5(3):329–335, September 2005.

[8] 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.

[9] 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.

[10] G. C. Necula. *Compiling with Proofs*. PhD thesis, Carnegie Mellon University, 1998.

[11] 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.

[12] 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.

[13] 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.

[14] J. Ohlsson and M. Rimen. Implicit signature checking. In *International Conference on Fault-Tolerant Computing*, June 1995.

[15] 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.

[16] 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.

[17] 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.

[18] 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.

[19] 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.

[20] P. P. Shirvani, N. Saxena, and E. J. McCluskey. Software-implemented EDAC protection against SEUs. In *IEEE Transactions on Reliability*, volume 49, pages 273–284, 2000.

[21] 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.

[22] 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.

[23] 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.

[24] R. Venkatasubramanian, J. P. Hayes, and B. T. Murray. Low-cost on-line fault detection using control flow assertions. In *Proceedings of the 9th IEEE International On-Line Testing Symposium*, pages 137–143, July 2003.

[25] 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.

[26] D. Walker, L. Mackey, J. Ligatti, G. Reis, and D. I. August. Static typing for a faulty lambda calculus. In *ACM International Conference on Functional Programming*, Portland, Oregon, Sept. 2006.

[27] 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.

[28] J. F. Ziegler and H. Puchner. *SER - History, Trends, and Challenges: A Guide for Designing with Memory ICs*. 2004.

A. Appendix

A.1 Failure Rules

Operational rules omitted from Figures 2, 3, and 4.

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

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

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

$$\frac{\begin{array}{c} R_{val}(r_s) \notin Dom(M) \\ R' = R{+\!+}[r_d \mapsto B\ n] \end{array}}{(R,C,M,Q,ld_B\ r_d,r_s) \longrightarrow_0 (R',C,M,Q,\cdot)} \ (ld_B\text{-}rand)$$

$$\frac{}{(R,C,M,\cdot,st_B\ r_d,r_s) \longrightarrow_0 fault} \ (st_B\text{-}queue\text{-}fail)$$

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

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

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

$$\frac{\begin{array}{c} R_{val}(r_z) = 0 \\ R_{val}(r_d) \neq R_{val}(d) \text{ or } R_{val}(d) = 0 \end{array}}{(R,C,M,Q,bz_B\ r_z,r_d) \longrightarrow_0 fault} \ (bz_B\text{-}taken\text{-}fail)$$

A.2 Semantics of Static Expressions

$\boxed{[\![E]\!]}$

$$\begin{array}{lcl} [\![n]\!] & = & n \\ [\![E_1\ op\ E_2]\!] & = & [\![E_1]\!]\ op\ [\![E_2]\!] \\ [\![emp]\!] & = & \cdot \\ [\![sel\ E_m\ E_n]\!] & = & [\![E_m]\!]([\![E_n]\!]) \\ [\![upd\ E_m\ E_1\ E_2]\!] & = & [\![E_m]\!][\ [\![E_1]\!] \mapsto [\![E_2]\!]\ ] \end{array}$$

$\boxed{\Delta \vdash E_1 = E_2}$

$$\frac{\begin{array}{c} \Delta \vdash E_1 : \kappa_{int} \qquad \Delta \vdash E_2 : \kappa_{int} \\ \forall S.\, \cdot \vdash S : \Delta \implies [\![S(E_1)]\!] = [\![S(E_2)]\!] \end{array}}{\Delta \vdash E_1 = E_2} \ (E\text{-}eq)$$

$$\frac{\begin{array}{c} \Delta \vdash E_1 : \kappa_{int} \qquad \Delta \vdash E_2 : \kappa_{int} \\ \forall S.\, \cdot \vdash S : \Delta \implies [\![S(E_1)]\!] \neq [\![S(E_2)]\!] \end{array}}{\Delta \vdash E_1 \neq E_2} \ (E\text{-}neq)$$

$$\frac{\begin{array}{c} \Delta \vdash E_1 : \kappa_{mem} \qquad \Delta \vdash E_2 : \kappa_{mem} \\ Dom([\![S(E_1)]\!]) = Dom([\![S(E_2)]\!]) \\ \forall \ell \in Dom([\![S(E_1)]\!]).\, [\![S(E_1)]\!](\ell) = [\![S(E_2)]\!](\ell) \end{array}}{\Delta \vdash E_1 = E_2} \ (E\text{-}mem\text{-}eq)$$