

Using Register Lifetime Predictions to Protect Register Files Against Soft Errors

Pablo Montesinos, *Member, IEEE*, Wei Liu, *Member, IEEE*, and Josep Torrellas, *Fellow, IEEE*

Abstract—Device scaling and large integration increase the vulnerability of microprocessors to transient errors. One of the structures where errors can be most harmful is the register file — a storage structure that is read very frequently.

To increase the resistance of register files to soft errors, this paper presents the *ParShield* architecture. *ParShield* is based on two observations: (i) the data in a register is only useful for a small fraction of the register’s lifetime, and (ii) not all registers are equally vulnerable. *ParShield* selectively protects registers by generating, storing, and checking the ECCs of only the most vulnerable registers while they contain useful data. In addition, it stores a parity bit for all the registers, re-using the ECC circuitry for parity generation and checking. *ParShield* has no SDC AVF and a small average DUE AVF of 0.040 and 0.010 for the integer and floating-point register files, respectively. *ParShield* consumes on average only 81% and 78% of the power of a design with full ECC for the SPECint and SPECfp applications, respectively. Finally, *ParShield* has no performance impact and little area requirements.

Index Terms—B.8.0 Performance and Reliability, C.1.0 Processor Architectures

I. INTRODUCTION

WITH increased chip integration levels, reduced supply voltages, and higher frequencies, soft errors are becoming a serious threat for high-performance processors. Such errors can be due to a variety of events, most notably the impact of high-energy particles [1], [2], [3]. Since soft errors can result in program visible errors [4], there have been proposals for several architectural designs that protect different structures of the processor, such as caches, memories, and datapaths [5], [6], [7], [8].

One of the critical structures to protect in a processor is the register file. It is a sizable structure that stores architectural state. It often stores data for long periods of time and is read frequently, which increases the probability of spreading a faulty datum to other parts of the machine. For these reasons, some commercial processors protect their register files with either parity [9], [10] or error correcting codes (ECC) [11]. Protecting the register file with only parity enables error detection but not correction. In this case, when the error is detected, recovery is only possible by invoking a high-level operation at the operating system (OS) or application level. Since the software might not always be able to recover from the error, the application may need to terminate. Full ECC support, on the other hand, enables on-the-fly detection and correction of errors. However, it does so at a cost in power and possibly performance.

This work was supported in part by the National Science Foundation under grants CHE-0121357 and CCR-0325603; DARPA under grant NBCH30390004; DOE under grant B347886; and gifts from IBM and Intel. A shorter version of this paper appeared in the 37th International Conference on Dependable Systems and Networks (DSN), in June 2007.

A cost-effective protection mechanism for soft errors in register files should have no performance impact, keep the remaining Architectural Vulnerability Factor (AVF) [12] to a small value, consume modest power, and use little area. To design such mechanism, we make two key observations on the use of registers in general-purpose processors. The first one is that the data stored in a physical register is not always useful. A soft error in a physical register while the data is not useful will not have any impact on the processor’s architectural state. Consequently, we only need to protect a register when it contains useful data. The second observation is that not all the registers are equally vulnerable to soft errors. A small set of long-lived registers account for a large fraction of the time that registers need to be protected. The contribution of most of the other registers to the vulnerable time is very small.

Based on these two key observations, this paper proposes *ParShield*, a novel architecture that provides cost-effective protection for register files against soft errors. *ParShield* relies on the Shield concept, which selectively protects a subset of the registers by generating, storing, and checking the ECCs of only the most vulnerable registers while they contain useful data. Such support reduces the AVF of the integer register file by an average of 73% to 0.040, and the AVF of the floating-point register file by an average of 85% to 0.010. *ParShield* also includes a parity bit for all the registers and re-uses the ECC circuitry for parity generation and checking as well. As a result, *ParShield* has no Silent Data Corruption (SDC) AVF (all single-bit errors are detected), has a Detected Unrecoverable Error (DUE) AVF as low as Shield’s AVF, and consumes on average only 81% and 78% of the power of a design with full ECC for the SPECint and SPECfp applications, respectively. Moreover, *ParShield* has no performance impact and little area requirements.

The paper is organized as follows. Section II describes the motivation of this work; Sections III and IV describe the design and the implementation of *ParShield*; Section V describes possible compiler support; Sections VI and VII evaluate *ParShield*; and Section VIII describes related work.

II. MOTIVATION: ASSIGNING RELIABILITY RESOURCES

A. Register Lifetime

Modern out-of-order processors use register renaming with a large number of physical registers to support many in-flight instructions [13], [14]. After the processor decodes an instruction with a destination register, it allocates a free physical register, creating a new *register version*. Later, the instruction is executed and its result is written to the corresponding physical register. Subsequent instructions that use that value are renamed to read from that physical register. The register version is kept until the instruction that redefines the corresponding architectural register retires — this is necessary to handle precise exceptions. Note that,




Fig. 1. Physical register holding different register versions (top), and lifetime of its $n + 1$ th version (bottom).




Fig. 2. ACE periods of two register versions.

while a register version can be read multiple times, it is written to at most *once*.

As Fig. 1 shows, a physical register successively holds different versions during program execution. Each of these versions is assigned to a potentially different architectural register. In between versions, the physical register remains in *NotAllocated* status.

Fig. 1 also illustrates a register version lifetime. It lasts from register allocation to deallocation, and can be divided into three different periods: from the time the register is allocated until the register is written to; from the time the register is written to until it is read for the last time; and from the last read until the register is deallocated. We call these periods *PreWrite*, *Useful*, and *PostLastRead*, respectively. Note that the register stores valuable data only during *Useful* periods.

B. Register ACE Analysis

Errors are usually classified as undetected or detected. The former are known as Silent Data Corruption (SDC), while the latter are usually referred to as Detected Unrecoverable Errors (DUE) [12]. Errors for which detection and recovery succeeds are not treated as errors.

A structure's *Architectural Vulnerability Factor* (AVF) is the probability that a fault in that structure will result in an error [12]. The SDC AVF and DUE AVF are the probabilities that a fault causes an SDC or a DUE error, respectively. In general, if a structure is protected by an error detection mechanism, its SDC AVF is zero. If the structure has error detection and correction capabilities, its DUE AVF is zero. In this work, we assume that the AVF for a register file is the average AVF of all its bits.

Mukherjee *et al.* [12] proposed the concept of *Architecturally Correct Execution* (ACE) to compute a structure's AVF. ACE analysis divides a bit's lifetime into ACE and un-ACE periods. A bit is in ACE state when a change in its value will produce an error. The AVF for a single bit is the fraction of time that it is in ACE state. To calculate the total time a bit is in ACE state, we start by assuming that its whole lifetime is in ACE state, and then we remove the fraction that can be proven un-ACE. The fraction left is an upper bound on the ACE time.

As an example, Fig. 2(a) and Fig. 2(b) show two register versions and their ACE and un-ACE periods. In both cases, a free physical register R is allocated at time t_a and deallocated at time t_d . During its *PreWrite* period, it remains in un-ACE state. At

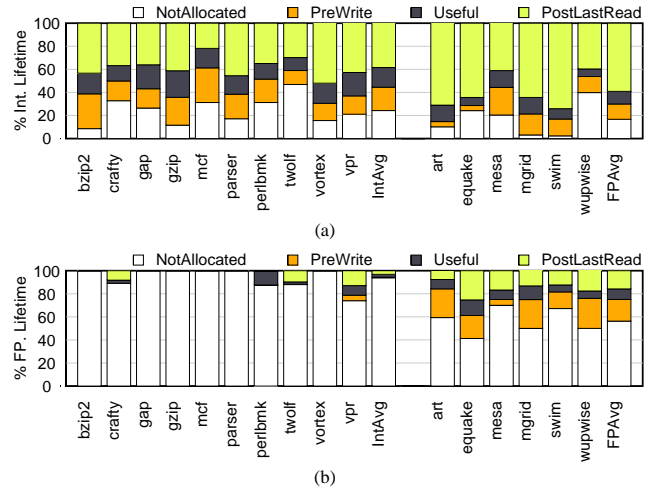


Fig. 3. Integer (a) and floating-point (b) register lifetime breakdown.

time t_w , R is written to and, if it will be consumed at least once, it switches to ACE state. Fig. 2(a) depicts a register version that is never read, so it remains in un-ACE state for its whole lifetime. Fig. 2(b) shows a register version that is consumed n times, so it enters ACE state at t_w and remains in it until it is read for the last time at t_{r_n} . A register version is in un-ACE state during its *PostLastRead* period.

There is one case where a register is read and it should still remain in un-ACE state. This is when the reader instructions are eventually squashed and, therefore, are never committed. For simplicity, however, in this work we do not consider it; if a register will be read, it is ACE.

C. Two Key Observations

Our analysis of SPECint and SPECfp 2000 applications for an out-of-order superscalar processor with 128 integer and 64 floating-point physical registers (Section VI) enables us to make two key observations.

The Combined Useful Time of All the Registers is Small.

We observe that the time a register version is in *Useful* state is only a *small fraction* of the register's lifetime. Fig. 3(a) shows the average fraction of time that an integer physical register spends in *PreWrite*, *Useful*, *PostLastRead* and *NotAllocated* status for both SPECint and SPECfp applications. As shown in the figure, on average only 18% and 12% of the register lifetime is *Useful* for SPECint and SPECfp applications, respectively. Fig. 3(b) shows that the *Useful* period is even smaller in the case of floating-point registers, namely 9% for SPECfp. Thus, there is no need to provide protection during the entire lifetime of a register. Note that since floating-point registers are not used much in SPECint applications, the data on SPECint applications in Fig. 3(b) is not representative. We only show it for completion.

Fig. 4(a) and Fig. 4(b) show the average number of physical registers that are in *Useful* state at a given time in the register files. For the integer one, the average is 19 registers out of 128 for SPECint; it is 17 out of 128 for SPECfp. In the case of the the floating-point register file, the average is only 5 out of 64 for SPECfp.

Overall, we conclude that it is possible to reduce the vulnerability of the register file by only protecting a subset of carefully chosen registers at a time.




Fig. 4. Average number of integer (a) and floating-point (b) physical registers in useful state.




Fig. 5. Short- and long-lived integer register version characterization.

A Few Long-Lived Register Versions Provide Much of the Total Useful Time. The second observation is that not all the register versions are equally vulnerable to soft errors. A small set of long-lived versions account for a large fraction of the time that registers need to be protected. For this section only, we say that a register version is *short-lived* if by the time it is written, an instruction that reads or writes the same architectural register has been renamed. We call the other versions *long-lived*.

To see this effect, consider Fig. 5 and Fig. 6. For each SPEC application, Fig. 5(a) and Fig. 6(a) show the percentage of long- and short-lived register versions allocated in the integer and floating-point register files, respectively. On average, less than 10% of the integer register versions are long-lived in SPECint and SPECfp (Fig. 5(a)). For floating-point register versions in SPECfp (Fig. 6(a)), the fraction of long-lived versions is less than 5% on average.

Fig. 5(b) and Fig. 6(b) show the percentage of the useful lifetime that long- and short-lived versions contribute to, in the integer and the floating-point register files, respectively. For integer register versions (Fig. 5(b)), on average 40% and 46% of the contribution in SPECint and SPECfp, respectively, comes




Fig. 6. Short- and long-lived floating-point register version characterization.

from the few long-lived register versions. For floating-point register versions (Fig. 6(b)), on average 40% of the contribution in SPECfp comes from long-lived versions. Therefore, it is cost-effective to give higher protection priority to these long-lived register versions.

III. PARSHIELD: PROTECTING THE REGISTER FILE

To provide cost-effective protection for register files, we propose *ParShield*. *ParShield* is composed of (i) the *Shield* structures and (ii) the parity support. In this section, we describe the architecture, focusing mostly on *Shield*.

A. *Shield* Concept

ParShield relies on the *Shield* concept, which involves using ECCs to selectively protect only the subset of most vulnerable registers while they contain useful data. *Shield* supports three operations on one such register: (i) when the register is written, *Shield* generates and saves the ECC of the written data, (ii) when the register is read, *Shield* checks whether the register contents are still valid, and (iii) *Shield* keeps the ECC of the data until the register is read for the last time. *Shield* assumes a single-bit fault model.

Fig. 7 shows the *Shield* architecture. It adds three hardware components to a traditional register file for an out-of-order processor: a table that stores the ECCs of some registers, a set of ECC generators and a set of ECC checkers. The ECC table is organized as a CAM. It protects the most vulnerable register versions in the register file. Each entry protects one register version and consists of: (i) a *tag* with the physical register number, (ii) a *parity* bit for the tag, (iii) the *ECC* bits of the data in the register, and (iv) a set of *Status* bits that are used during the replacement of the ECC table entries.

When a physical register is about to be written, a request for protection is sent to *Shield*. If *Shield* decides to protect the register, it tries to allocate an entry for that version. The entries in the ECC table are not pre-allocated during the register renaming stage because there is no need to protect a register version during *PreWrite* time. Once an entry has been successfully allocated in the ECC table, an ECC generator calculates the ECC of the register data in parallel with the register write operation.

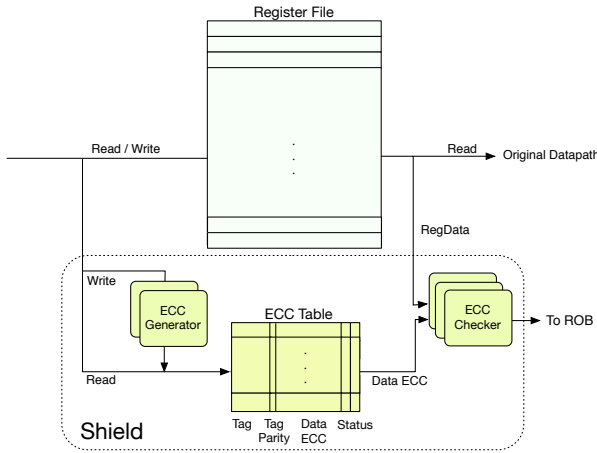


Fig. 7. Shield architecture.

When a physical register is read, the register file sends its data to both the datapath and Shield (Fig. 7). Shield checks whether there is an entry in the ECC table whose tag matches the physical register number. If so, Shield checks the tag’s parity, and sends the corresponding ECC to an ECC checker to verify the data’s integrity. If an error is detected in the tag — thanks to the parity bit — the corresponding entry in the ECC table is invalidated and the processor proceeds. If the ECC checker detects an error in the register data, the processor stalls and takes the following actions to recover from the error: (i) it fixes the register data, (ii) it flushes the reorder buffer (ROB) from the oldest instruction that reads the register version, and (iii) it flushes the whole ECC table and resumes. Finally, if there is no error, the ECC checker signals no error and the processor proceeds.

Each entry of the ROB is augmented with two Finish bits, one for each of the two potential source operand registers. These bits are set if, when the corresponding operand was read, it either completed the ECC check or was not protected by Shield. Single-operand instructions have one bit always set.

When an instruction reaches the head of the ROB and is ready to retire, the Finish bits are checked. If at least one of the bits is not set, the instruction cannot retire; it has to wait for the ECC checker to finish and set the bit(s), or for the ROB to get full. In the latter case, the instruction is retired without taking the Finish bits into consideration in order to minimize performance degradation. Our experiments show that the ROB provides enough slack for the ECC checker to verify the integrity of the data without affecting the IPC.

An entry in the ECC table is deallocated and assigned to another register version when the Shield replacement algorithm decides to evict it or when a new version of the same physical register is written and sent to Shield for protection. When an entry is evicted from the ECC table, its associated register version will no longer be protected.

B. Entry Allocation and Replacement

When Shield receives a request for protection for a physical register version, it tries to allocate an entry in the ECC table. If there is an entry in the table protecting a previous version of the same physical register, Shield re-assigns the entry to the new version. Otherwise, Shield attempts to pick a free table entry. Since the table is much smaller than the register file, there




Fig. 8. Predicting short-lived registers.

may be no free entry, and a decision has to be made to either replace an existing entry in the ECC table or abort the allocation. Entry replacement has to be done carefully. Replacing a recently-allocated entry that protects a long-lived register to accommodate a new one that will protect a short-lived register increases the vulnerability of the system. Therefore, Shield needs to predict the lifespan of register versions.

1) *Predicting Short- and Long-lived Registers*: When Shield considers evicting an entry from the ECC table, it does not know whether the register version that it protects is still in its useful time or not. Shield’s goal is to evict the entry that contributes the least to the overall register file’s AVF. Since a long-lived register contributes more to the register file’s AVF than a short-lived one, Shield tries to evict short-lived registers. To this end, Shield extends the short-lived register predictor proposed by Ponomarev *et al.* [15]. In the following, we first describe their approach and then how we augment it.

Fig. 8 illustrates how Ponomarev *et al.*’s short-lived register predictor works. It maintains a bit vector, called *Renamed*, that has one bit per physical register. In Fig. 8, under *original code*, a LOAD instruction loads into architectural register R1. After R1 is used in the ADD instruction, the MUL instruction overwrites R1. Therefore, the MUL is a *renamer* for the LOAD. In the *renamed code*, R1 has been renamed to P12. When the MUL is renamed, it sets the bit *Renamed*[12]. If by the time the LOAD loads the data into P12, the *Renamed*[12] bit is set, P12 is considered to be short-lived.

Although this predictor is simple and often effective, it is limited. Specifically, suppose that the ADD is the only consumer of R1. In this case, we would want to consider P12 to be a short-lived register. However, if the LOAD has loaded the data before the MUL is renamed, P12 will not be predicted as a short-lived register.

To extend the capability of Ponomarev *et al.*’s algorithm, we reformulate the *Renamed* vector. We call it the *Events* vector, and it has two bits for each physical register, namely *Events.Renamed* and *Events.Used*. The rules for the *Events.Renamed* bit are identical to the Ponomarev *et al.*’s scheme. The new *Events.Used* bit is set when renaming an instruction that consumes the physical register. Based on the *Used* and *Renamed* bits, when we are about to write to a physical register, we predict the register’s lifespan as one of the following four types: 1) *long-lived*, if both *Used* and *Renamed* bits are reset, 2) *dead*, if only the *Renamed* bit is set, 3) *short-lived*, if only the *Used* bit is set, and 4) *ultrashort-lived*, if both *Used* and *Renamed* bits are set.

Fig. 9 shows an example of how our proposed predictor works. In the *original code*, the MUL and DIV instructions act as the renamers for the first and second LOAD instructions, respectively. Therefore, the *Events.Renamed* bits of P12 and P7 are both set. Our algorithm also sets the *Events.Used* bit of P12 because R1 is used by the ADD. However, since R2 is never used, the *Events.Used* bit of P7 remains reset. The four possible combinations of the *Used* and *Renamed* bits are shown




Fig. 9. Predicting the lifespan of physical register versions.

TABLE I
EFFECTIVE REPLACEMENT POLICY IN ECC TABLE.

Prediction	Entries that can be replaced
Long	Free, Ultrashort, Short, Long
Short	Free, Ultrashort, Short
Ultrashort	Free, Ultrashort

in Fig. 9(b). *Dead* register versions are not protected by Shield.

2) *Entry Replacement*: For each protection request, Shield receives the register’s *Events* bits along with the register number and data. If there is neither an entry protecting the same physical register nor a free entry in the ECC table, Shield has to replace an existing entry or abort the allocation. Note that ECC table entries can only be allocated when a register version is written. Once a register version loses its ECC table entry, it cannot get a new one, and remains unprotected for the rest of its lifetime.

We propose a replacement policy that we call *Effective*. It uses the expected lifespan of a register version to select the entry to replace. It works as follows: when a victim entry is needed, Shield tries to select an entry that protects a register version with a shorter or same expected lifespan than the one to be protected. If such an entry is unavailable, Shield aborts the allocation. Table I shows the types of entries that can be replaced for a register version according to its prediction. For example, if a register version is predicted as short, it tries to replace an entry marked as free, ultrashort, and then short — in this priority order.

We also dynamically adjust the entry type to reflect the fact that the expected lifespan gets shorter after reads. When an ultrashort or short entry is read, the type changes to free or ultrashort, respectively. The type of long entries is never changed since these entries tend to have long lifespans and may be read many times during their lifetime.

The information about the type of register version that an ECC table entry contains is kept in two Status bits (Fig. 7). The four possible states of the Status bits are *long*, *short*, *ultrashort*, and *free*.

C. Entry Deallocation

An ECC table entry protects a given register version until the replacement algorithm reassigns the entry to another version. After a register version is read for the last time, it is effectively stale, and it is useless to protect it anymore. Ideally, Shield would like to know the time of the last read to a register version so that it can deallocate the entry then. However, Shield has no way of knowing whether a read is the last one. Therefore, it is possible to have stale entries in the ECC table. These stale entries hurt the efficiency of Shield because they protect nothing and occupy resources. The situation is worse if these stale entries are

marked as long, because they have less chance of being replaced compared to short ones.

To remove stale entries from the ECC table — especially the ones marked as long — we send explicit signals (called *eviction* signals) to the ECC table to indicate which entries just became stale. When the ROB sends a signal to release a physical register, this same signal is also forwarded to the ECC table as an eviction signal. If the ECC table has the corresponding entry, it marks it as free.

D. Error Recovery

Since we use single error correction with double error detection (SEC-DED) codes in this study, Shield allows the processor to recover from transient single-bit errors in the register file and detect double-bit errors. Although the processor may also recover from some double-bit errors, in this paper we only focus on single-bit errors.

When an ECC checker detects that the register data read by instruction I has a single-bit error, the processor stalls and enters recovery mode. First, the checker fixes the error and writes the corrected data back to the physical register (say P). Second, Shield examines the ROB looking for the oldest instruction that reads P and flushes that instruction as well as the others that follow it. Note that only flushing I and the instructions that follow it is not enough to recover from the error. Imagine that an instruction J older than I reads P after I did, but before the error is detected by the ECC checker. The data that J reads has already been corrupted. Consequently, the processor has to flush from the oldest instruction that reads P . The ECC table is then flushed so no entry in the ECC table protects one of the registers that were removed from the ROB. Finally, the processor can resume.

If an error occurs in a mispredicted path, Shield will still recover from it for simplicity.

E. AVF of a Register File with Shield

Fig. 10 shows different physical register versions, their associated ECC table entry, and the time during which they are vulnerable to errors (ACE cycles). In Fig. 10(a), Shield cannot allocate an entry for the register version and, therefore, the register is in ACE state during its whole useful lifetime. In Fig. 10(b), an entry was protecting the register version but is evicted before the version is read for the first time. As a result, the register is in ACE state during its whole useful lifetime. In Fig. 10(c), the entry is evicted after the register version is read at least once but before its last read. Consequently, Shield only protects the version until the read just before the eviction. After that read, the register is in ACE state. Finally, in Fig. 10(d), the ECC entry remains allocated for the whole useful lifetime of the register version. The version is completely protected and is never in ACE state.




Fig. 10. Computing the AVF of different physical register versions.




Fig. 11. ACE and un-ACE periods for fully and partially bypassed register versions.

However, the longer an entry protects a dynamically dead register, the less efficient Shield is. By using the eviction signal described in Section III-C, we are able to mitigate this effect. Using the four cases in this figure, we can compute the ACE cycles of each register version. Since the AVF of a physical register is the fraction of ACE cycles, we can then easily compute the AVF of each physical register and of the whole register file.

To calculate the overall AVF of the system, we also have to take into account the possibility of a bit flip in the ECC table or the Finish bits in the ROB. The tag in the ECC table is protected by the parity bit, and therefore a bit flip in this field can be detected. Shield then deallocates the damaged entry from the ECC table. A bit flip in the ECC field can be easily detected and corrected during the integrity check. A bit flip in the Status bits will not affect the correctness of the system — only the efficiency of Shield. Thus, the AVF of the ECC table is 0.

Finally, the AVF of the Finish bits is also 0, assuming a single-bit error model. If any of the Finish bits flips to 0, the corresponding instruction will stay longer at the head of the ROB, but will eventually retire when the ROB gets full. If any of the Finish bits flips to 1, the instruction might retire before it is actually checked. However, since we are assuming that only one error can occur at a time, no other error can occur and the register data has to be correct.

F. ParShield: Shield Plus Full Register Parity

Finally, we extend Shield with storage for a parity bit for all the physical registers in the processor, and re-use the ECC circuitry for parity generation and checking as well. The result is the complete *ParShield* architecture. With the parity bit, *ParShield* reduces the SDC AVF to zero (all errors are detected) — although the DUE AVF is equal to the AVF of plain Shield (the exposure to non-correctable errors remains the same as in plain Shield). Moreover, this is accomplished at a very small cost in hardware and power.

Specifically, consider when a register write sends a protection request to the ECC table. While *ParShield* is checking if it should generate the data's ECC and enter them in the table — depending on the type of register version and the current contents of the ECC table — *ParShield* uses one ECC generator to compute the data's parity and store it in a Parity bit vector. Such operation takes a small fraction of the time taken by the generation of the full ECC.

In the same way, consider when a register read sends a request to the ECC table. While *ParShield* is checking if the ECC table

TABLE II
PERCENTAGE OF BYPASSED REGISTER READS IN SPECINT AND SPECfp

Suite	Bypassed Register Reads (%)
SPECint	26.3
SPECfp	32.2

contains the corresponding entry, *ParShield* reads the Parity bit vector and uses one ECC checker to check the parity. Again, this operation takes little time. Moreover, computing and checking the parity bits consumes much less power than computing and checking ECCs.

IV. IMPLEMENTATION ISSUES

A. Bypass Network

Processors use the bypass network to send results from one functional unit to another so that dependent instructions can execute back to back. Table II shows that, on average, 26% (SPECint) and 32% (SPECfp) of all register reads take their data from the bypass network, thus not accessing the register file. We therefore need to include register bypassing in our model of AVF.

We distinguish two kinds of bypasses for a register version: (i) all its consumers read the value from the bypass network, and (ii) some of the consumers read it from the register file while others read it from the bypass network. We refer to the former as *full bypass* and to the latter as *partial bypass*.

Calculating the AVF for register versions that are fully bypassed is straightforward. Since the data stored in the register file is never used, their AVFs are zero. On the other hand, partially bypassed versions need to be protected from the time the data is written until their last non-bypassed read.

Fig. 11(a) shows an example where an ADD instruction generates a version of register P1, which is then read by subsequent MUL and SUB instructions. We assume that the SUB is P1's last use. Fig. 11(b) shows the ACE and un-ACE periods if the result is fully bypassed. Neither the MUL nor the SUB accesses the register file. Therefore, this P1 version remains un-ACE during its entire lifetime. In Fig. 11(c), the SUB instruction reads P1 from the register file and, therefore, P1 remains ACE until the SUB executes.

B. Accessing the ECC Table

The ECC table needs fewer ports than the register file. The reason is that the table is not as performance-critical as the register

TABLE III

AVERAGE NUMBER OF READS AND USEFUL LIFETIME (MEASURED IN CYCLES) OF REGISTER VERSIONS IN CERTAIN ARCHITECTURAL REGISTERS.

App.	Architectural Register							
	\$at or \$1		\$v0 or \$2		\$gp or \$28 (in millions)		\$ra or \$31	
	# of Reads	Useful Lifetime	# of Reads	Useful Lifetime	# of Reads	Useful Lifetime	# of Reads	Useful Lifetime
bzip2	1.0	10.2	0.9	4.9	14.6	561.3	1.0	39.3
crafty	1.0	10.1	1.1	7.3	1.1	575.9	1.0	25.2
gap	1.0	6.9	1.1	11.2	19.8	1686.1	1.0	7.8
gzip	1.0	5.0	1.2	9.0	10.1	668.4	1.0	73.0
mcf	1.0	6.5	1.0	17.8	0	0	1.0	8940.1
parser	1.0	17.7	1.0	12.9	12.2	1033.8	1.0	18.6
perlbnk	0.0	0	1.0	7.3	0	0	1.0	8.4
twolf	1.0	49.3	1.0	20.5	16.9	1487.5	1.0	239.7
vortex	1.0	5.7	1.4	6.7	0	0	1.0	11.9
vpr	1.0	14.4	1.0	5.5	1.5	82.7	1.0	72.5
IntAvg	0.9	12.6	1.1	10.3	7.6	609.6	1.0	943.65

file and, therefore, does not need to be multiported for the worst case — twice as many read ports as the issue width and as many write ports as the issue width. Consequently, we reduce the number of ports and perform the ECC generation and checking off the critical path. If necessary, we support some small queueing of requests, which does not affect performance because instructions often wait in the ROB for a long period before committing.

In reality, there is rarely any queueing. One reason is that many instructions have fewer than two register operands. Moreover, many reads obtain their data from the bypass network and, therefore, do not access the ECC table. In addition, many register writes do not create an entry in the ECC table. Specifically, *dead* versions (Section III-B.1), *ultrashort* versions that find the table full with non-*ultrashort* versions, and *short* versions that find the table full with *long* versions, skip ECC generation and table update.

In some cases, ParShield adds additional updates to the ECC table tags. These are caused by the eviction signals (Section III-C). However, these signals are infrequent. They are only sent when the physical register to be freed was predicted as *long* — Fig. 20 shows that, on average, less than 10% of the register versions are predicted as *long*. The rationale is that *short* and *ultrashort* register versions are aged automatically and evicted from the ECC table at a much faster pace than *long* ones — typically before the register is freed.

C. Using More Architectural Knowledge to Improve Efficiency

It is possible to extend the algorithm of Section III-B.1 that predicts the lifespan of register versions by leveraging the usage patterns of certain architectural registers. The goal is to improve the prediction accuracy, given that some architectural registers have a specific purpose and, therefore, special usage patterns.

Table III shows the average number of reads and useful lifetimes of register versions in certain architectural registers of the MIPS ISA [16] for the SPECint application suite.

Columns 2-5 describe the \$at and the \$v0 architectural registers. The former is used by the assembler, while the latter is utilized as a return value register. It can be seen that both of them have very short useful lifetimes and are typically read only

MUL r1<-r5,r4	MUL r1{L}<-r5,r4	MUL (r1,e1)<-r5,r4
...
ADD r2<-r1,r3	ADD r2<-r1,r3	ADD r2<-(r1,e1),r3
...
ST r1->[r6]	ST r1{last}->[r6]	ST (r1,e1)->[r6]
		FREE (e1)
(a)	(b)	(c)

Fig. 12. Original code (a), code for compiler-assisted ECC-table (b), and code for compiler-managed ECC-table (c).

once. Even though it could be beneficial to deallocate the ECC table entries that protect them after one read, we observe that they are likely to be predicted as *ultrashort* and evicted from the table very quickly.

Columns 8-9 show that register \$31, or *return address pointer* register, is also read once per allocation but has a longer lifetime. However, because it is reallocated shortly after it is read, explicitly deallocating it from the ECC has little effect. As a result, we did not explicitly deallocate the entries that protect the physical mappings of registers \$at, \$v0 and \$ra.

The *global pointer* register (\$gp) is used to reference global variables. Unlike the other registers in Table III, it is read many times and has a long useful lifetime. It is important, then, that the ECC table protects it for as long as possible. Therefore, in the ECC table, we pin the entry that protects the physical mapping of the global pointer until it receives an eviction signal.

V. USING THE COMPILER TO IMPROVE PARSHIELD'S EFFICIENCY

The ParShield implementation presented in this paper is a hardware-only approach. The hardware is responsible for allocating and deallocating entries in the ECC table and for deciding which register versions to protect based on the prediction of version lifetime. Such hardware-only approach is fully transparent to software and, therefore, neither software modification nor recompilation is required. However, when changes in software or re-compilation are possible, a more cost-effective hardware-software co-design approach may be possible.

ParShield's lifetime prediction algorithm depends on the size of the instruction window. A larger window would allow the processor to obtain more information about the register versions' lifetime. Unfortunately, scaling the microarchitectural structures in an out-of-order processor has been proven difficult [17]. On the other hand, the compiler observes the entire program and, by using dataflow and interprocedural analysis, can determine the lifetime of most register versions. When the compiler identifies a register version as long-lived, it can either (i) mark it and let the hardware decide for how long or whether to allocate an entry for it in the ECC table, or (ii) manage the allocation and deallocation of the ECC table entry for that register version. We refer to the first approach as *compiler-assisted* and to the second one as *compiler-managed*.

As an example, consider the code in Fig. 12(a), and assume that the compiler determines that r1 holds a long-lived register version. In the compiler-assisted approach (Fig. 12(b)), the compiler annotates the MUL instruction indicating that the register version that MUL creates is likely to be long. However, it is up to the hardware to decide for how long or whether to protect

```
LD r3<-[r4] % cache miss
MOV r2<-r5
ADD r1<-r2,r3
```

Fig. 13. Example of long-lived register version ($r2$) due to a long-latency instruction (LD).

<pre>ADD r5<-r1,r2 MUL r6<-r7,r5 1000s insts ... SUB r1<-r6,r5</pre>	<pre>ADD r5<-r1,r2 MUL r6<-r7,r5 ... ST r5->mem 1000s insts ... LD r5<-mem SUB r1<-r6,r5</pre>
(a)	(b)

Fig. 14. Original code (a) and spilling a long, vulnerable register version to memory (b).

$r1$'s version in the ECC table. In the compiler-managed approach (Fig. 12(c)), the ECC table is fully controlled by software. The compiler assigns ECC table entry $e1$ to $r1$, and informs the hardware that it must access $e1$ to verify $r1$'s integrity for the ADD and ST instructions.

However, compiler-based register lifetime prediction also faces some issues due to the lack of dynamic execution information. For example, consider the code in Fig. 13. The compiler predicts $r2$'s version as short-lived. Now suppose that the LD instruction causes a cache miss, therefore taking hundreds of cycles until data returns. As a result, $r2$ stays in the register file for a long time. Since it was predicted as short-lived, an ECC table entry may not have been allocated, therefore hurting the register file's AVF. This happens because the compiler lacks dynamic information. Previous work [18] has shown that it is possible to statically identify the instructions that depend on a delinquent load. In the example, that would mean that $r2$ would be considered long-lived.

Another limitation of the hardware-only approach is that it is often impossible to determine whether an instruction is the last one reading from a register version. This is not true in the case of the compiler, which, in many cases, can detect the last use of a register version. When this happens, the compiler marks the read instruction as the last read (Fig. 12(b)) or introduces a new instruction that causes the ECC table to deallocate the entry (Fig. 12(c)).

Finally, other compiler techniques might help reduce the register file AVF and can be used with both *compiler managed* and the *compiler assisted* approaches. For example, register spilling could be modified so that long, vulnerable register versions are sent to memory—which is likely to be ECC-protected—instead of keeping them in the register file. Consider the code in Fig. 14(a). Suppose that only the MUL and SUB instructions read from $r5$, and that they are very far apart in the instruction stream. As a result, $r5$'s stays in the register file in useful state for a long time. In many cases, the compiler could detect such situations and send $r5$ to memory right after the MUL instruction, and bring it back right before the SUB. Fig. 14(b) illustrates this case.

TABLE IV
REGISTER CONFIGURATIONS EVALUATED

Configuration	Description
Baseline	Register files with no parity or ECC protection
Shield	Baseline + Shield (Section III-A)
ParShield	Shield + parity for all registers (Section III-F)
FullECC	Baseline + ECC for all registers

VI. EVALUATION METHODOLOGY

We use SESC, a cycle-accurate execution-driven simulator [19] to model the processor and memory system architecture of Table V. The architecture is a MIPS-like 3-issue out-of-order processor with two levels of caches, a 128-entry integer register file with 6 read and 3 write ports, and a 64-entry floating-point register file with 4 read and 2 write ports.

We evaluate the performance and the power of this architecture with the register file configurations of Table IV. *Baseline* is the architecture with no protection for the register files. *Shield* is Baseline plus the Shield architecture of Section III-A. As shown in Table V, the ECC table for the integer register file has 32 entries and 3 read and 3 write ports; the ECC table for the floating-point register file has 16 entries and 2 read and 2 write ports. The number of ECC generators and checkers is the same as the number of write and read ports in the ECC table, respectively. *ParShield* is Shield plus the parity bit for all registers (Section III-F). Finally, *FullECC* is Baseline plus ECC for all the 128 integer registers and 64 floating-point registers. In all cases, 8-bit ECC codes are used to protect the 64-bit registers.

We evaluate the architectures with SPECint and SPECfp 2000 applications running the *Ref* data set. All of the applications are included except those that are not supported by our current framework. The applications are compiled using *gcc-3.4* with *-O3* optimization enabled. After skipping the initialization (typically 1-6 billion instructions), each application executes around 1 billion instructions.

Since applications do not run to completion, we are unable to determine whether or not a register is in ACE state when the simulation finishes. For example, if a simulation ends right after t_w in Fig. 2(a), we would not know if the period after the write is ACE or un-ACE. To handle these edge effects, we use the *cooldown* technique that was proposed by Biswas *et al.* [20]. During the cooldown interval, we track the registers that were live at the moment that the simulation stopped. This helps us determine if a register was in ACE or un-ACE state.

VII. EVALUATION

In this section, we first examine the AVF results and the power and area consumption, then perform a sensitivity analysis, and finally examine register lifespan prediction.

A. AVF Results

We compare the AVF of Baseline to that of Shield with different replacement policies in the ECC table: *Random*, *LRU*, *Effective* (proposed in Section III-B.2), and *OptEffective*. The latter augments *Effective* with the pinning optimization described in Section IV-C. Recall that ParShield has an SDC AVF equal to zero (all errors are detected) and a DUE AVF equal to Shield's AVF. Finally, the AVF of FullECC is zero.

TABLE V
PROCESSOR AND MEMORY SYSTEM MODELED. CYCLE COUNTS ARE IN PROCESSOR CYCLES

Processor		Register File		Cache & Memory		ParShield	
Frequency	4 GHz	Integer:		L1 Cache:		Integer:	
Fetch/Issue/Retire	6/3/3	Entries	128	Size, assoc, line	16KB, 4, 64B	ECC table entries, width	32, 18 bits
ROB size	126	Width	64 bits	Latency	2 cycles	R/W ports	3/3
I-window	68	R/W ports:	6/3	L2 Cache:		ECC latency	4 cycles
LD/ST queue	48/42	FP:		Size, assoc, line	1MB, 8, 64B	FP:	
Mem/Int/FP unit	2/3/2	Entries	64	Latency	12 cycles	ECC table entries, width	16, 17 bits
Branch predictor:		Width	64 bits	Memory:		R/W ports	2/2
Mispred. Penalty	14 cycles	R/W ports:	4/2	Latency	500 cycles	ECC latency	4 cycles
BTB	2K, 2-way						

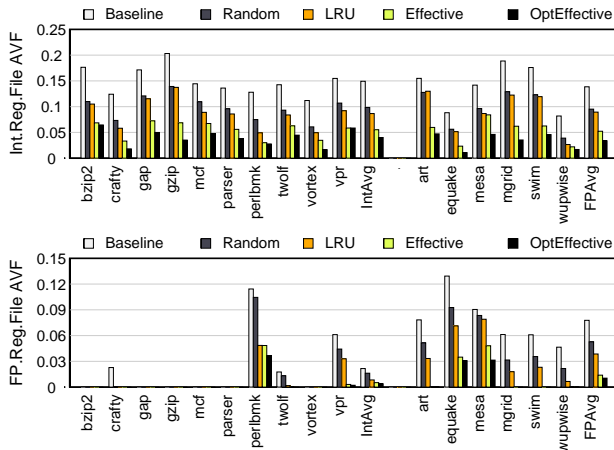


Fig. 15. Integer (top) and floating-point (bottom) register file AVFs.

Fig. 15 and shows the AVFs of the integer (AVF_{int}) and floating-point (AVF_{fp}) register files for the described configurations. The AVFs are shown for all simulated SPECint and SPECfp applications. Since there are almost no floating-point operations in the SPECint applications, we do not discuss the AVF_{fp} for SPECint, and only show it in Fig. 15 for completeness.

Fig. 15 shows that, for all applications and on both register files, *Effective* and, especially *OptEffective*, have an AVF much lower than *Baseline*. For example, *Effective* reduces the AVF_{int} for SPECint by 63% on average and the AVF_{fp} for SPECfp by 42% on average relative to *Baseline*. *OptEffective* reduces the AVF_{int} for SPECint by up to 84% (on average 73%) and the AVF_{fp} for SPECfp by up to 100% (on average 85%) relative to *Baseline*. The resulting average AVF_{int} for SPECint is 0.040 and the average AVF_{fp} for SPECfp is 0.010. As expected, *Random* and *LRU* perform worse than chosen policies.

In general, *Shield* works slightly better for the floating-point register file because it has a smaller fraction of registers in useful state than the integer one. In addition, it is easier to predict the lifespan of floating-point registers. As shown in Fig. 15, *Shield* reduces the AVF_{fp} to nearly zero for *art*, *mgrid*, *swim* and *wupwise*.

B. Power and Area Consumption

Register files consume a significant fraction of the power in modern processors. For example, one estimate suggests that the




Fig. 16. Integer register file power consumption.

integer register file consumes around 14% of the dynamic power in the processor [21]. We use CACTI 4.2 [22] to estimate the dynamic and static power of storage structures such as the register file, the ECC table, and the ECC and parity bit-fields. We use HSpice [23] models to estimate the dynamic and static power of the ECC logic.

In Fig. 16, we show the total power (dynamic plus static) consumed in the integer register file for the different register configurations. For each application, the bars are normalized to *Baseline*. We do not include data for the floating-point register file because, as explained before, many of our applications do not use it much and, therefore, the average differences between configurations are small.

The figure shows that *FullECC* consumes on average 100% more power than *Baseline* for both SPECint and SPECfp applications. This is due to the combination of the ECC generators and checkers, and the additional storage for the ECC bits. With *Shield*, the average power is only 78% and 74% of *FullECC* for SPECint and SPECfp, respectively. This results mainly from the fewer ECC generators and checkers, and the fewer ECC operations performed — although the tags and ports in the ECC table are a significant source of power consumption.

Fig. 16 also shows that *ParShield* consumes only slightly more power than *Shield*. The difference is small because the parity bits consume little power to generate, store and check. Overall, with *ParShield*, the average power is 81% and 78% of *FullECC* for SPECint and SPECfp, respectively. Both *ParShield* and *Shield* are more power-efficient than *FullECC*.

Finally, we estimate the area of the register file and the additional ECC and parity structures using CACTI 4.2. The area of the ECC logic is not added because it is negligible. Adding up the contributions of both the integer and the floating-point register files, we find that *FullECC* uses 4.9% more area than *Baseline*. Moreover, *Shield* and *ParShield* use 15.7% and 17.6% more area




Fig. 17. Impact of the number of ECC table ports on the IPC.




Fig. 18. Impact of the number of ECC table ports on the integer register file AVF.




Fig. 19. Impact of the ECC table size on the integer register file AVF.

than *Fullecc*, respectively. This area increase is tolerable, given the power savings provided.

C. Sensitivity Analysis

To gain insight into the operation of Shield, we examine the impact of the number of read and write ports in the ECC table. Fig. 17 and Fig. 18 show the IPC (Instructions Per Cycle) and the integer register file AVF, respectively, when we use x read ports and y write ports (xR, yW) in the ECC table. For each application, the bars are normalized to the (6R, 4W) configuration.

Fig. 17 shows that, as we go from (6R, 4W) — the same number of ports and type as in the integer register file — to (3R, 3W) — our design choice — the IPCs remain constant. This is because Shield performs the ECC generation and checks off the critical path, queues requests when necessary, and leverages the slack given by the ROB. However, when we further reduce the number of read ports, performance suffers. With (2R, 3W), the average IPC decreases by 6.9% for SPECint and by 4.2% for SPECfp. With (1R, 3W), there is a 50% performance penalty because now the ECC table becomes a major bottleneck.

Fig. 18 shows the impact of the number of ECC table ports on the AVF of the integer register file. With our design choice (3R, 3W), the AVF changes negligibly over (6R, 4W). With (2R, 3W) and (1R, 3W), however, the AVF increases noticeably.

Finally, Fig. 19 shows the impact of the ECC table size on the AVF of the integer register file. Our design choice has 32 entries. Recall from Fig. 4 that, on average, there are only 19 registers




Fig. 20. Breakdown of the types of register versions predicted.




Fig. 21. Average cumulative distribution of useful lifespan for Ultrashort, Short, and Long registers in SPECint.




Fig. 22. Comparing the AVF for different register lifespan prediction algorithms.

with useful state in SPECint and 17 in SPECfp. From Fig. 19, we see that a 16-entry ECC table has a high AVF. The 24-entry table is better, although it still has a 68% and 80% higher AVF than our chosen 32-entry configuration for SPECint and SPECfp, respectively. Using a larger 64-entry ECC table reduces the AVF further but at a very high hardware cost.

Overall, from this section and the previous one, we see that the 32-entry (3R, 3W) ECC table offers a good tradeoff between performance, AVF, and power.

D. Register Lifespan Prediction

Finally, we examine the accuracy of the register lifespan predictor used in Shield (*OptEffective*). Fig. 20 shows the fraction of integer register versions that *OptEffective* predicts of each type, for both SPECint and SPECfp. We can see that, on average for SPECint, it predicts over 60% as Ultrashort and only 10% as Long. For SPECfp, the average fraction of Ultrashort predictions is over 80%. Fig. 21 shows the average cumulative distribution of the useful lifespan for the register versions predicted as Ultrashort, Short, and Long in SPECint. From the figure, we see that *OptEffective* correctly separates the three types of register versions. Indeed, most register versions predicted as Long have over 30 cycles of useful lifespan, while 95% of the register versions predicted as Ultrashort have less than 10 cycles of useful lifespan.

However, *OptEffective* does not possess oracle knowledge. Fig. 22 compares Shield's AVF using *OptEffective* and an oracle

algorithm (*Oracle*) for register lifespan prediction. For completeness, it also shows the AVF with Ponomarev *et al.*'s predictor (Section III-B). The figure shows that, on average, *OptEffective* is very close to *Oracle*. The difference in AVF largely results from the fact that, for some parts of the applications, there are more registers in useful state than entries in the ECC table.

VIII. RELATED WORK

Fully protected register files. Traditional fault-tolerant designs protect the entire register file with parity or ECC. The IBM S/390 G5 [11] uses duplicated, lockstepped pipelines to ensure that only correct data updates the ECC-protected structure that stores the architectural state of the processor. The ERC32 [9] is a SPARC processor with parity-protected registers and buses that also performs program control flow, which imposes extra overhead. Like the ERC32, the Intel Montecito [10] also utilizes parity to protect the whole register file. Both the ERC32 and the Intel Montecito require software intervention to recover when a fault is detected. Our ParShield design uses parity to detect all single-bit errors, and the Shield concept to recover from most single-bit errors — by selectively protecting the register versions that contribute the most to the overall vulnerability of the register file. Processor performance is not affected. The result is a design with a very cost-effective tradeoff between DUE AVF and power.

Partially protected register files. Memik *et al.* [24] proposed the duplication of actively-used physical registers in unused register locations. While their approach can enhance reliability with minimal performance degradation, it can only detect errors, but not recover from them. Yan *et al.* [25] proposed using the compiler to assign the most vulnerable variables to a set of ECC-protected registers. ParShield does not need to re-compile the programs because it offers a hardware-only solution.

Register lifetime analysis. Lozano and Gao [26] used the compiler to identify short-lived variables and prevented their values from going to the register file, thus reducing register pressure. Sangireddy and Somani [27] reduced the access time to the register file by exploiting useless periods in the register lifetime. Ponomarev *et al.* [15] used a small dedicated register file to cache short-lived operands to reduce the energy consumption in the ROB and the architectural register file. ParShield differs from all these proposals in that it proposes a hardware scheme to distinguish between ultrashort-, short- and long-lived operands, and exploits the difference to enhance the register file reliability.

IX. FUTURE WORK

The ParShield architecture described in this paper has focused on protecting the register file. However, other structures in the processor can also benefit from similar selective protection mechanisms. For example, the instruction window is a large structure that often holds unprotected instructions for long periods of time. Instead of protecting the entire instruction window, a ParShield-based protection mechanism would prioritize protecting long-lived instructions. Our future work includes studying the instruction window, load/store queue, and other processor structures that are amenable to the same protection scheme.

X. CONCLUSIONS

Register files are vulnerable to soft errors because they are large and contain architectural state. Registers often store data

for long periods of time and are read frequently, which increases the probability of spreading a faulty datum. A cost-effective protection mechanism for soft errors in register files should have no performance impact, keep the remaining AVF to a small value, consume modest power, and use little area.

In this paper, we have proposed one such mechanism, namely the *ParShield* design. ParShield relies on the Shield concept, which selectively protects a subset of the registers by generating, storing, and checking the ECCs of only the most vulnerable registers while they have useful data. Shield reduces the AVF of the integer register file by an average of 73% to 0.040, and the AVF of the floating-point register file by an average of 85% to 0.010. ParShield extends Shield with a parity bit for all the registers and the re-use of the ECC circuitry for parity generation and checking as well. As a result, ParShield has no SDC AVF (all single-bit errors are detected), has a DUE AVF as low as Shield's AVF, and consumes on average only 81% and 78% of the power of a design with full ECC for the SPECint and SPECfp applications, respectively. Moreover, ParShield has no performance impact and little area requirements.

ACKNOWLEDGMENT

The authors thank the anonymous reviewers and the members of the I-ACOMA group at the University of Illinois at Urbana-Champaign for their invaluable comments. Special thanks go to Brian Greskamp and Radu Teodorescu for their feedback on the paper.

REFERENCES

- [1] E. Czeck and D. Siewiorek, "Effects of Transient Gate-level Faults on Program Behavior," in *International Symposium on Fault-Tolerant Computing*, June 1990.
- [2] W. McKee, H. McAdams, E. Smith, J. McPherson, J. Janzen, J. Ondrusek, A. Hyslop, D. Russell, R. Coy, D. Bergman, N. Nguyen, T. Aton, L. Block, and V. Huynh, "Cosmic Ray Neutron Induced Upsets as a Major Contributor to the Soft Error Rate of Current and Future Generation DRAMs," *1996 IEEE Annual International Reliability Physics*, 1996.
- [3] J. F. Ziegler, H. W. Curtis, H. P. Muhlfeld, C. J. Montrose, and B. Chin, "IBM Experiments in Soft Fails in Computer Electronics (1978-1994)," *IBM J. Res. Dev.*, 1996.
- [4] N. J. Wang, J. Quek, T. M. Rafacz, and S. J. Patel, "Characterizing the effects of transient faults on a high-performance processor pipeline," in *DSN '04: Proceedings of the 2004 International Conference on Dependable Systems and Networks (DSN'04)*. Washington, DC, USA: IEEE Computer Society, 2004, p. 61.
- [5] S. Kim and A. K. Somani, "Area Efficient Architectures for Information Integrity in Cache Memories," in *International Symposium on Computer Architecture*, 1999.
- [6] J. Ray, J. C. Hoe, and B. Falsafi, "Dual Use of Superscalar Datapath for Transient-fault Detection and Recovery," in *International Symposium on Microarchitecture*, 2001.
- [7] P. Shivakumar, M. Kistler, S. Keckler, D. Burger, and L. Alvisi, "Modeling the Effect of Technology Trends on the Soft Error Rate of Combinational Logic," in *International Conference on Dependable Systems and Networks*, 2002.
- [8] C. Weaver, J. Emer, S. Mukherjee, and S. Reinhardt, "Techniques to Reduce the Soft Error Rate of a High-Performance Microprocessor," in *International Symposium on Computer Architecture*, 2004.
- [9] J. Gaisler, "Evaluation of a 32-bit Microprocessor with Built-In Concurrent Error-Detection," in *International Symposium on Fault-Tolerant Computing*, 1997.
- [10] C. McNairy and R. Bhatia, "Montecito: A Dual-Core, Dual-Thread Itanium Processor," *IEEE Micro*, 2005.
- [11] T. Slegel, I. Averill, R.M., M. Check, B. Giamei, B. Krumm, C. Krygowski, W. Li, J. Liptay, J. MacDougall, T. McPherson, J. Navarro, E. Schwarz, K. Shum, and C. Webb, "IBM's S/390 G5 Microprocessor Design," *IEEE Micro*, vol. 19, 1999.

- [12] S. Mukherjee, C. Weaver, J. Emer, S. Reinhardt, and T. Austin, "A Systematic Methodology to Compute the Architectural Vulnerability Factors for a High-Performance Microprocessor," in *International Symposium on Microarchitecture*, 2003.
- [13] G. Hinton, D. Sager, M. Upton, D. Boggs *et al.*, "The Microarchitecture of the Pentium 4 Processor," *Intel Technology Journal*, 2001.
- [14] E. J. McLellan and D. A. Webb, "The Alpha 21264 Microprocessor Architecture," in *ICCD '98: Proceedings of the International Conference on Computer Design*, 1998.
- [15] D. Ponomarev, G. Kucuk, O. Ergin, and K. Ghose, "Isolating Short-Lived Operands for Energy Reduction," *IEEE Transactions on Computers*, 2004.
- [16] K. C. Yeager, "The MIPS R10000 Superscalar Microprocessor," *IEEE MICRO*, 1996.
- [17] K. Olukotun, B. A. Nayfeh, L. Hammond, K. Wilson, and K. Chang, "The Case for a Single-Chip Multiprocessor," in *International conference on Architectural support for programming languages and operating systems*, March 1996.
- [18] V. Panait, A. Sasturkar, and W. Wong, "Static Identification of Delinquent Loads," in *2nd Annual International Symposium on Code Generation and Optimization with Special Emphasis on Feedback-Directed and Runtime Optimization.*, 2004.
- [19] J. Renau, B. Fraguera, J. Tuck, W. Liu, M. Prvulovic, L. Ceze, S. Sarangi, P. Sack, K. Strauss, and P. Montesinos, "SESC Simulator," 2005. [Online]. Available: <http://sesc.sourceforge.net>
- [20] A. Biswas, P. Racunas, R. Cheveresan, J. Emer, S. S. Mukherjee, and R. Rangan, "Computing Architectural Vulnerability Factors for Address-Based Structures," in *International Symposium on Computer Architecture*, June 2005.
- [21] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-Aware Microarchitecture," in *International Symposium on Computer Architecture*, 2003.
- [22] N. P. J. David Tarjan, Shyamkumar Thoziyoor, "CACTI 4.0," *Tech Report HPL-2006-86*, 2006. [Online]. Available: <http://www.hpl.hp.com/techreports/2006/HPL-2006-86.htm>
- [23] Hspice User's Manual, "Applications and examples," 1996. [Online]. Available: citeseer.ist.psu.edu/328246.html
- [24] G. Memik, M. T. Kandemir, and O. Ozturk, "Increasing Register File Immunity to Transient Errors," in *DATE '05: Conference on Design, Automation and Test in Europe*, 2005.
- [25] J. Yan and W. Zhang, "Compiler-guided register reliability improvement against soft errors," in *EMSOFT '05: Proceedings of the 5th ACM International Conference on Embedded Software*, 2005.
- [26] L. A. Lozano and G. R. Gao, "Exploiting Short-Lived Variables in Superscalar Processors," in *International Symposium on Microarchitecture*, 1995.
- [27] R. Sangireddy and A. K. Somani, "Exploiting Quiescent States in Register Lifetime," in *International Conference on Computer Design (ICCD'04)*, 2004.



Wei Liu Wei Liu is a research scientist in Microprocessor Technology Labs at Intel Corporation. His research interests include computer architecture, compiler support for multi-core, processor reliability and programming languages. Liu has a PhD in computer science from Tsinghua University, P.R.China. He is a member of the IEEE and the ACM.



Josep Torrellas Biography text here.



Pablo Montesinos is a PhD candidate in the Computer Science Department of the University of Illinois at Urbana-Champaign. His research interests include support for deterministic replay of parallel execution, programability, operating systems and processor reliability. Montesinos received his B.S. degree in computer engineering from the Universidad de León in 2001, Spain. In 2003 he was awarded the "La Caixa Fellowship" and joined the I-ACOMA group at the University of Illinois at Urbana-Champaign, from where he received his M.S in 2005. He is a student member of the IEEE and the ACM.