

Compiler-Assisted Multiple Instruction Rollback Recovery Using a Read Buffer

Neal J. Alewine, Shyh-Kwei Chen, *Member, IEEE*,
W. Kent Fuchs, *Fellow, IEEE*, and Wen-mei W. Hwu, *Member, IEEE*

Abstract—Multiple instruction rollback (MIR) is a technique that has been implemented in mainframe computers to provide rapid recovery from transient processor failures. Hardware-based MIR designs eliminate rollback data hazards by providing data redundancy implemented in hardware. Compiler-based MIR designs have also been developed which remove rollback data hazards directly with data-flow transformations. This paper describes compiler-assisted techniques to achieve multiple instruction rollback recovery. We observe that some data hazards resulting from instruction rollback can be resolved efficiently by providing an operand read buffer while others are resolved more efficiently with compiler transformations. The compiler-assisted scheme presented consists of hardware that is less complex than shadow files, history files, history buffers, or delayed write buffers, while experimental evaluation indicates performance improvement over compiler-based schemes.

Index Terms—Fault-tolerance, error recovery, instruction retry, compilers.

I. INTRODUCTION

INSTRUCTION retry is a technique for rapid recovery from transient faults in a computer system. Multiple instruction rollback recovery is particularly appropriate when error detection latencies or when error reporting latencies are greater than a single instruction cycle. When transient processor errors occur, multiple instruction rollback (also referred to as multiple instruction retry or simply instruction retry) can be an effective alternative to system-level checkpointing and rollback recovery [1], [2], [3], [4], [5], [6]. Multiple instruction retry within a sliding window of a few instructions [2], [3], [4], [5], or re-execution of a few cycles [7], can be implemented in parallel with error detection methods for recovery from transient processor errors.

A. Hardware-Based Instruction Rollback

Hardware implemented instruction retry schemes belong to one of two groups:

- 1) full checkpointing and
- 2) incremental checkpointing.

Full checkpointing maintains "snapshots" of the required system state space at regular, or predetermined intervals. Upon

error detection, the system can be rolled back to the appropriate checkpointed system state. Incremental checkpointing maintains changes to the system state in a "sliding window". Upon error detection the system state is restored by undoing, or "backing-out" the system state changes to an instruction previous to the one in which the error occurred.

The issues associated with instruction retry are similar to the issues encountered with exception handling in an out-of-order instruction execution architecture [8]. If an instruction is to write to a register and N is the maximum error detection latency (or exception latency), two copies of the data must be maintained for N cycles. Hardware schemes such as reorder buffers, history buffers, future files [9], and micro-rollback [2] differ in where the updated and old values reside, circuit complexity, processor cycle times, and rollback efficiency.

Table I gives a description of various hardware-based methods to restore the general purpose register file contents during single or multiple instruction rollback. In the VAX 8600 and VAX 9000, errors are detected prior to the completion of a faulty instruction. For most VAX instructions, updates to the system state occur at the end of the instruction. If the error is detected prior to updating the system state, the instruction can be rolled back and reexecuted. If the system state has changed prior to detection of the error, a flag is set to indicate that instruction rollback cannot be accomplished. Redundant data storage is not required for the VAX 8600 and VAX 9000.

The IBM 4341, IBM 3081, IBM patent 4,912,707, IBM patent 4,044,337, and history file all require shadow file structures to maintain redundant data. This data is used to restore the system state during rollback recovery. Shadow file structures can add significant circuit overhead, although the level sensitive scan design [14] of the IBM 4341 and IBM 3081 provides this feature without additional cost over that incurred to obtain testability.¹ The VAX 8600 and VAX 9000 schemes avoid shadow files, however, they require an error detection latency of only one instruction.

The micro-rollback scheme also avoids shadow files by using a delayed write buffer to prevent old data from being overwritten until the error detection latency has expired; ensuring that the new data is fault-free. In a delayed write scheme, the most recent write values are contained in the delayed write buffer, and complex bypass and prioritization circuitry is required to forward this data on subsequent reads. The performance impact introduced by the bypass circuitry is a function of the register file size and the maximum rollback distance [2].

1. The 126 scan rings of the IBM 3081 contains 35,000 bits of data.

Manuscript received June 28, 1993; revised Aug. 22, 1994.
N.J. Alewine is with IBM Corporation, Boca Raton, Fla.
S.-K. Chen is with IBM T.J. Watson Research Center, Yorktown Heights, NY 10598.

W.K. Fuchs, and W.-m.W. Hwu are with the Center for Reliable and High-Performance Computing, Coordinated Science Laboratory, 1308 W. Main St., University of Illinois, Urbana, IL 61801.

IEEECS Log Number C95090.

TABLE I
HARDWARE-BASED SINGLE AND MULTIPLE INSTRUCTION
ROLLBACK SCHEMES

Rollback Scheme	Checkpoint Type	Rollback Distance	Location of Data	
			Primary	Redundant
IBM 4341 [1]	full	single instr.	register file	shadow file
IBM 3081 [2]	full	10-20 instr.	register file	shadow file
VAX 8600 [3]	full	single instr.	register file	not required
IBM patent 4,912,707 [4]	full	variable	register file	shadow file
IBM patent 4,044,337 [5]	incremental	single instr.	register file	shadow files
micro-rollback [6]	incremental	variable	write buffer	register file
history buffer [7]	incremental	variable	register file	history buffer
history file [7]	incremental	variable	register file	shadow file
VAX 9000 [8]	full	single instr.	register file	not required
IBM E/S 9000 [9]	incremental	variable	virtual file	physical file

The history buffer scheme maintains redundant data in a separate push-down array and therefore does not require bypass circuitry [9]. The history buffer does however require an extra register file port which complicates the file design and can impact performance by increasing file access times.

In an effort to increase the register file size while maintaining downward code compatibility relative to the 16 architectural registers, the IBM E/S 9000 introduced a virtual register management (VRM) system [15]. The VRM circuitry dynamically maps the 16 architectural registers into 32 physical registers. When the data in a physical register becomes obsolete, the physical register is released for reassignment as a new virtual register. Although the VRM system was primarily intended to reduce register pressure and therefore improve system performance, it has been extended to provide data redundancy to assist in rollback recovery. In the VRM extension, remapping of a physical register to a new virtual register is postponed until the error detection latency has been exceeded for the data contained in the physical register.

B. Compiler-Based Instruction Rollback

Compiler-based approaches to multiple instruction rollback recovery have also been developed [3], [4]. Compiler-based MIR uses data-flow manipulations to remove data hazards that result from multiple instruction rollback. Rollback data hazards (or just hazards) are identified by *antidependencies*² of length $\leq N$, where N represents the maximum rollback distance. Antidependencies are removed at three levels:

- 1) pseudo-code level, or the code level prior to variables being assigned to physical registers,
- 2) machine-code level, or the code level in which variables are assigned to physical registers, and
- 3) post-pass level, that represents assembler-level code emitted by the compiler.

Compiler-based multiple instruction rollback reduces the requirement for data redundancy logic present in hardware-based instruction rollback approaches.

C. Compiler-Assisted Instruction Rollback

Compiler-based multiple instruction rollback resolves all data hazards using compiler transformations. This paper describes a compiler-assisted instruction rollback scheme that

uses dedicated data redundancy hardware to resolve one type of rollback data hazard while relying on compiler assistance to resolve the remaining hazards. Experimental results indicate that by exploiting the unique characteristics of differing hazard types, the new compiler-assisted MIR design can achieve superior performance to either a hardware-only or compiler-only instruction rollback scheme.

II. ERROR MODEL AND HAZARD CLASSIFICATION

A. Rollback Data Hazard Model

The following four assumptions are used in the general error model:

- 1) the maximum error detection latency is N instructions,
- 2) memory and I/O have delayed write buffers and can rollback N cycles,
- 3) the states of the program counter and program status word (PSW) are preserved by an external recording device or by shadow registers [2], and
- 4) the processor state can be restored by loading the correct contents of the register file, program counter, and PSW.

Given the above assumptions, any error which does not manifest itself as an illegal path in the control-flow graph (CFG) of the program is allowed provided that the following two conditions are satisfied:

- 1) register file contents do not spontaneously change, and
- 2) data can not be written to an incorrect register location.

There are four targeted transient error types:

- 1) processor errors such as those caused by an ALU failure,
- 2) incorrect values read from I/O, memory, the register file, or external functional units such as the floating point unit,
- 3) correct/incorrect values read from incorrect locations within the I/O, memory, or register file, and
- 4) incorrect branch decisions resulting from error types 1, 2, or 3.

B. Hazard Classification

The executable code can be represented as a CFG $G(V, E)$, where V is the set of nodes denoting instructions and E is the set of edges denoting control-flow. If there is a direct control-flow from instruction i , denoted I_i , to I_j , where $I_i \in V$ and $I_j \in V$, then there is an edge $(I_i, I_j) \in E$. Let $d_{min}(I_i, I_j)$ denote the smallest number of instructions along any path from I_i to I_j .

The hazard set H_{regs} of the error model is defined as the set of pseudo registers (or machine registers) whose values are inconsistent during different executions of an instruction sequence due to retry. Two properties³ are used to classify rollback data hazards given the error model of Section II.A.

PROPERTY 1. Hazard register x is an element of H_{regs} iff there exists a sequence of instructions I_1, I_2, \dots, I_N which form a legal walk⁴ in G such that x is live at I_1 , and x is defined during the walk.

3. A formal treatment of these properties, along with proofs, can be found in [17].

4. A walk is a sequence of edge traversals in a graph where the edges visited can be repeated [18].

2. For a complete presentation of data-flow properties and manipulation methods, see Aho et al. [16].

PROPERTY 2. Hazards can be classified as one of two types:

- 1) those that appear as antidependencies of length $\leq N$ in $G(V, E)$, referred to as *on-path hazards*, and 2) those that appear at branch boundaries, referred to as *branch hazards*. These two hazard types may overlap.

An on-path or branch data hazard occurs when I_i defines variable x , and after rollback, I_j uses the corrupted x value prior to its redefinition. Fig. 1 shows an error occurring prior to a branch instruction with error detection occurring after the definition of variable x in I_i . Rollback then occurs with the program sequence beginning above the branch.

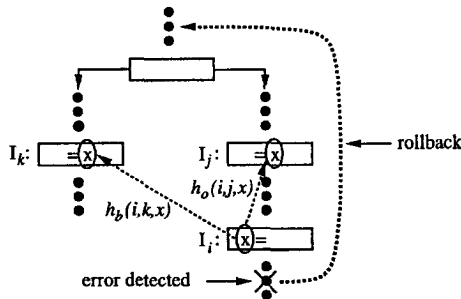


Fig. 1. On-path and branch hazards.

If the post-rollback path is the same as the original path, a use of variable x will occur in I_j prior to the re-definition of x in I_i . This represents an on-path hazard and is denoted as $h_o(i, j, x)$ in Fig. 1. If the post-rollback path is different from the original path (e.g., the error caused an incorrect branch decision during the original program sequence), a use of variable x will occur in I_k prior to a re-definition of x . This second case represents a branch hazard and is denoted as $h_b(i, k, x)$ in Fig. 1.

III. COMPILER-ASSISTED INSTRUCTION ROLLBACK

As shown in Section II, rollback data hazards are of two types: 1) on-path hazards, and 2) branch hazards. Previous work has shown that compiler-driven data-flow manipulations can be used to resolve both on-path [3] and branch [4] hazards. Compiler-assisted multiple instruction rollback described in this section uses hardware to resolve on-path hazards and relies on compiler assistance to resolve the remaining branch hazards.

A. On-Path Hazard Resolution Using a Read Buffer

Fig. 2 shows a hardware scheme to resolve on-path hazards. Each time a register is used, its value appears on the read port and is saved in the read buffer. If a register r_k is defined in I_i and it is an on-path hazard, then r_k must have been read within the last N cycles. In this case, the read buffer will contain the old value and it is permissible to write the new value into the register file. In the event of a rollback of N instructions, the contents of the read buffer are flushed in reverse order and stored back to the register file. For an on-path hazard, the path taken after the rollback will be the same as the path taken prior to rollback and each read of r_k will produce the same value as

before. It is assumed that the read buffer is an integral part of the register file and any error in the system does not corrupt the transfer to the read buffer or its contents.




Fig. 2. Read buffer.

In contrast to a write history buffer which forces a read of r_k prior to writing r_k , the read buffer monitors the register file ports and stores only the values read as part of the normal program flow and, therefore, should not significantly impact the register file performance or processor cycle time. The read buffer is twice the width of a register with a depth of N . This is twice the size of a delayed write buffer, but eliminates the requirement for complex bypassing and prioritization logic.

A.1. Covering On-Path Hazards

In addition to resolving all on-path hazards, the read buffer will resolve some branch hazards. Fig. 3 shows an on-path hazard and a branch hazard both with definitions of x in I_i and uses of x , after rollback, in instructions I_j and I_k , respectively. Note that if path l is initially taken, the read buffer will contain the old value of x and rollback would be successful. However if path m is taken, the read buffer will not contain the old value of x and rollback would be unsuccessful. If only paths such as l exist, the presence of the on-path hazard assures successful rollback or "covers" the branch hazard. In this case, resolution of the branch hazard using compiler techniques is not necessary.




Fig. 3. Covering on-path hazard.

A.2. Post-Pass Transformation

Given the efficiency of the read buffer in resolving on-path hazards, a post-pass transformation on assembler-level code becomes possible as an alternative to nop insertion transformations [3]. The post-pass transformation creates on-path hazards when necessary to assure that all branch hazards are resolved by the read buffer. Given one such branch hazard which defines physical register r_k at instruction I_i , the transformation inserts an $\text{MOV } r_k, r_k$ instruction immediately before I_i . This guarantees that all paths leading to I_i are like path l in Fig. 3.

B. Branch Hazard Resolution

Branch hazards are resolved at three levels: 1) pseudo-level, 2) machine-level, and 3) post-pass level [4]. Pseudo-level hazards are removed by variable renaming, for example, renaming variable x to y in instruction I_i of Fig. 1. Machine-level branch hazards occur when register assignments result in branch hazards that were not present at the pseudo-level. Machine-level hazards are resolved by adding hazard constraints to live range constraints prior to register assignment. Branch hazards which remain after pseudo-level and machine-level transformations are resolved at the post-pass level with read insertions as described in Section III.A.2.

The primary pseudo-level renaming transformation for the removal of branch hazards, involves node splitting [4]. This section presents a one-pass node splitting algorithm which results in marginally reduced code growths and dramatically reduced compile-times relative to previous node splitting algorithms.

Figs. 4a and 4b show a typical data dependence (requiring node splitting) and the node splitting technique, respectively. In Fig. 4a, renaming x in I_i to y will ultimately require the renaming of the *use* register x in I_j to y since multiple definitions of x reach I_k . To break this dependence, the following node splitting criterion is used: If multiple definitions of x reach I_k and x is in the *live_in* set of I_k , I_k will be split into two identical nodes. This “unzipping” is shown in Fig. 4b. Loop protection assures that no loop header is split [3].

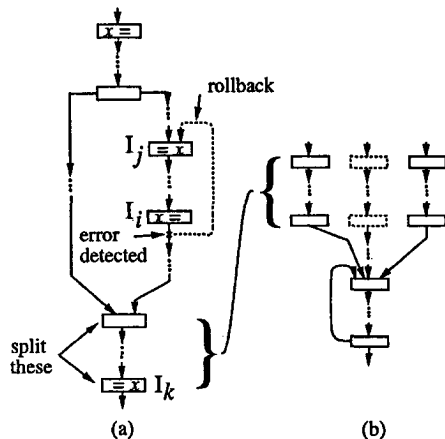


Fig. 4. Node splitting.

B.1. Iterative Node Splitting

Node splitting breaks equivalence relationships which would prevent pseudo register renaming [3], [16]. When two definitions of a hazard variable reach a node in which the hazard variable is live, the node is split. Node splitting to resolve one hazard variable often resolves other unrelated hazard variables. This implies that the hazard set should be recalculated after splitting is performed for each hazard variable. Previous node splitting algorithms use this iterative algorithm to avoid unnecessary node splitting [3].

Fig. 5 demonstrates the effect of the iterative node splitting algorithm on an example subgraph. Node splitting relative to hazard variable x ensures that the definition of x in node n_1 and the definition of x in node n_2 do not both reach the same use of x in node n_5 . Node splitting relative to y ensures that the definition of y in node n_3 and the definition of y in node n_4 do not both reach the same use of y in node n_6 .




Fig. 5. Iterative node splitting relative to hazard variables x and y .

Fig. 5 also shows an optimal subgraph which resolves both hazards with less splitting than produced by the iterative algorithm, indicating that excessive node splitting is possible with the iterative algorithm.

B.2. Node Splitting Using Graph Coloring

To reduce splitting, a node splitting algorithm is developed using the concept of conflicting parents [17]. Ensuring that node n does not have conflicting parents enables resolution of the hazard using variable renaming. The node splitting strategy for a particular node is to group the parents of that node such that elements within a group do not conflict. Each group becomes parent nodes for a duplicate of the original node. For example, if node n has six parent nodes and these nodes can be organized into three nonconflicting groups, then only three total copies of n are required.

Fig. 6 illustrates the use of conflicting parents and graph coloring in node splitting for the QSORT application described in Table III of Section IV.A. Node splitting is performed on pseudo-level code, which for this example is represented by *Lcode* from the IMPACT C compiler [19]. Fig. 6 shows node 48 from the QSORT application. Node 48 has six parent nodes prior to splitting. These nodes can be arranged in a parent conflict graph, where each arc of the graph represents two nodes that conflict. Establishing groups can be achieved by finding the minimum coloring of the parent conflict graph, i.e., coloring the nodes such that no two nodes connected by an arc have the same color. For the example shown in Fig. 6, three colors are sufficient to color the parent conflict graph, resulting in the splitting of node 48 into nodes 48, 48' and 48''. The graph coloring heuristic used for our one-pass node splitting algorithm is a modified version of an algorithm used for register allocation [16].

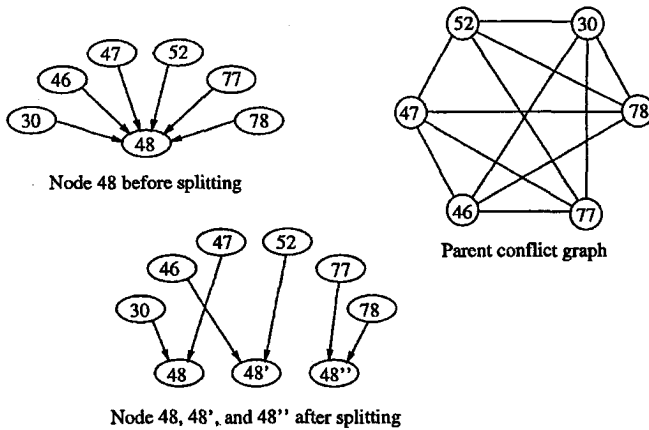


Fig. 6. Node splitting using graph coloring for QSORT.

B.3. One-Pass Node Splitting Algorithm

Both *live_in(n)* and *reaching_out(n)* [16] analyses are required to identify conflicting parent nodes. A one-pass node splitting algorithm becomes possible by precalculating *live_in* and the hazard node set, and then, beginning with the root node, splitting in a topological traversal of the CFG. A topological traversal ensures that when processing node *n*, all ancestors of *n* have been processed and no descendants of *n* have been processed. This latter case ensures that the presplit calculation of *live_in(n)* can be used for parent conflict identification when processing a given node. Unlike *live_in(n)*, *reaching_out(n)* is affected by the splitting of ancestor nodes. Since *reaching_out(n)* is based solely on node *n* and its ancestors, *reaching_out(n)* can be calculated as node splitting proceeds. If a hazard node is split, each duplicate of the node must be added to the hazard node set. Since the root node does not have conflicting parents, a topological traversal of the CFG using the graph coloring node splitting technique ensures that no node in the resulting graph has conflicting parents.

Table II illustrates the improvement of the one-pass node splitting algorithm over the iterative algorithm for the COMPRESS application described in Table III of Section IV.A. The COMPRESS application was compiled on a SPARCserver 490

using the IMPACT C compiler [19] with a rollback distance of 10. Node count values represent pseudo instructions (*Lcode*) created by the IMPACT C compiler before and after splitting. Seven of the 14 COMPRESS functions which required splitting are listed. Algorithm run times represent the overall compile times given each of the two node splitting algorithms.

Table II shows a marginal overall code growth reduction for the one-pass algorithm. Although one function demonstrated a significant code growth reduction (6.7% compared to 75.6%), the function is small and has minimal effect on the overall code size. The improvement in compile-time of the one-pass algorithm is more dramatic, resulting in a speedup factor of 30.2.

TABLE II
NODE SPLITTING COMPARISONS FOR COMPRESS

Orig. Node Cnt.	Iterative Alg.	% Increase	One-pass Alg.	% Increase
547	601	9.9	566	3.5
461	499	8.2	496	7.6
144	147	2.1	147	2.1
181	209	15.5	207	14.4
75	80	6.7	80	6.7
21	28	33.3	27	28.6
45	79	75.6	48	6.7

Iterative Algorithm run time = 614.0 seconds

One-pass Algorithm run time = 20.3 seconds

Speedup = 30.2

C. Performance Enhancement Through Profiling

C.1. Post-Pass Transformation Versus Loop Protection

Some hazards remain after compilation and must be removed using a post-pass transformation. Previous post-pass transformations used nop insertions to increase all anti-dependency distances to $> N$ [3]. Since nop insertion can be costly to performance, previous compiler transformations removed all resolvable hazards, leaving only unresolvable hazards to be removed by the post-pass transformation.

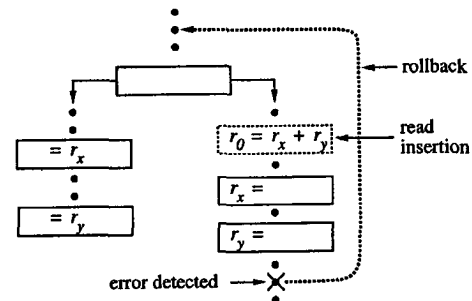


Fig. 7. Post-pass hazard removal using read insertion.

In Section III.A.2, an alternative post-pass transformation was introduced in which nop insertion was replaced by read insertions as the primary hazard removal technique. As illustrated in Fig. 7, up to two branch hazards can be removed by a single read instruction. Figs. 12 and 14 of Section IV.B show performance overhead comparisons between compiler-driven

data-flow manipulations and the post-pass transformation for the PUZZLE and TBL applications described in Table III of Section IV.A. *Comp/PP* indicates that hazards are resolved by the compiler where possible, while the remaining hazards are resolved at the post-pass level. *PP* (post-pass) indicates that compiler transformations have been disabled and that all hazards are removed at the post-pass phase.

For the PUZZLE application, compiler transformations produce better performance than the post-pass transformation alone. For the TBL application, using the post-pass transformation to remove all hazards produces slightly better performance than the combination of compiler and post-pass transformations. Hazard elimination via read insertion introduces a guaranteed but small performance impact due to the longer instruction path length. As demonstrated by the PUZZLE application, pseudo register renaming can eliminate hazards without impacting performance when loop protection is infrequent. The save/restore operations of loop protection can result in more performance impact than read insertion when loop protection is frequent, as demonstrated by results for the TBL application.

Fig. 8 illustrates the potential effect on performance given the following two types of hazard removal:

- 1) hazard removal using register renaming that results in loop protection, and
- 2) hazard removal using read insertion.

If the protected loop of Fig. 8 is executed 20 times and the hazard instruction is executed two times, loop protection would require the execution of 40 additional instructions, where read insertion would require the execution of only two additional instructions. If the loop and hazard instruction execution frequencies were reversed, then read insertion would produce more performance impact than loop protection. As shown in Fig. 8, profiling data can be used to aid in loop protection decisions.

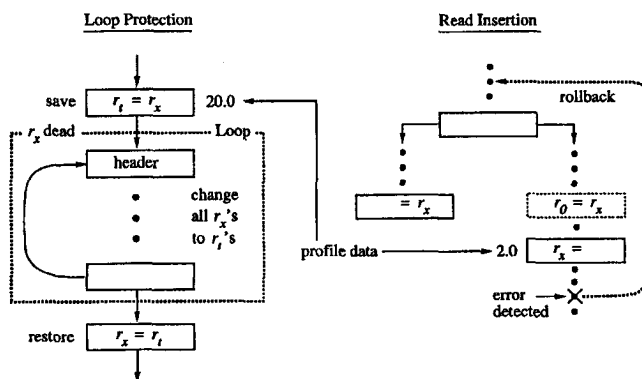


Fig. 8. Loop protection versus read insertion.

C.2. Profiling Effectiveness

Profile data was included in the pseudo-level transformations of Section III.B. The profile data is comprised of both dynamic profile sampling and static prediction. The static prediction is used as a supplement for areas of the application code that are not executed during profile sampling. For static

profiling, a loop is assumed to iterate ten times. Inner loops, therefore, iterate multiples of 10 times depending on the depth of loop nesting. All loop header nodes and hazard nodes are assigned weights based on the profile data.

Protection of loop l due to hazard node n_h is required based on the following condition: if $n_h_weight > 3 * (hdr_node(l)_weight)$, then protect loop l . The constant 3 adjusts the weights to account for both direct and indirect loop protection costs. Direct loop protection costs result from the save/restore instruction pair shown in Fig. 8. Indirect loop protection costs result from:

- 1) an increased number of hazards which in turn required more node splitting and more loop protection, and
- 2) increased register usage due to the save/restore instructions which can result in additional register spills.

Fig. 9 shows the run-time overhead for the TBL application with rollback distances from 1 to 10. *Prof/PP* indicates that profiling data was used in loop protection decisions.




Fig. 9. TBL profile data used for loop protection decisions.

The results show that the use of profile data can improve application performance by postponing some hazard resolutions until the post-pass phase. Using profile data to aid in loop protection decisions did not produce performance equal to that for the post-pass transformation, for the TBL application. As an extension to this work, profile data can be used to aid in register allocation. As discussed in Section III.B, hazards that are present after pseudo register renaming are resolved by adding hazard constraints to live range constraints prior to register allocation. These additional constraints can cause increased register spillage and impact performance. Similar techniques to those developed for loop protection can be used to enhance register allocation decisions.

IV. PERFORMANCE EVALUATION

A. Implementation and Application Programs

The hazard removal algorithms have been implemented in the MIPS processor code generator of the IMPACT C compiler [19]. Transformations resolving pseudo register hazards (loop protection, node splitting, and loop expansion) are invoked just before register allocation. Transformations resolving machine register hazards are invoked after the live range constraints have been generated and before physical register allocation. The nop insertion algorithm, or post-pass algorithm, is called before the assembly code output routine.

Table III lists the eleven application programs used in the

evaluations. The applications were cross-compiled on a SPARCserver 490 and then the compiled program was run on a DECstation 3100. *Static Size* is the number of assembly instructions emitted by the code generator, not including the library routines and other fixed overhead.

TABLE III
APPLICATION PROGRAMS

Program	Static Size	Description
QUEEN	148	eight-queen program
WC	181	UNIX utility
QSORT	252	quick sort algorithm
CMP	262	UNIX utility
GREP	907	UNIX utility
PUZZLE	932	simple game
COMPRESS	1826	UNIX utility
LEX	6856	lexical analyzer
YACC	8099	parser-generator
TBL	8197	table formatting preprocessor
CCCP	8775	preprocessor for gnu C compiler

The results are summarized in Figs. 10 through 14. Each figure contains two plots: The first plot shows the percent of run-time overhead (*Time OH*) of the referenced hazard resolution scheme, and the second plot shows the percent of code growth overhead (*Size OH*) relative to the base values in Table III.

Four hazard resolution techniques were evaluated. *Compiler 1* resolves on-path hazards only, using the compiler-driven data-flow manipulations. *Compiler 2* extends the compiler transformations to resolve both on-path and branch hazards.

PP (post-pass) disables the compiler transformations and relies solely on the post-pass transformation presented in Section III.A.2. *Comp/PP* uses compiler transformations to resolve branch hazards with the techniques described in Section III.B, assumes a read buffer to resolve on-path hazards, and uses the post-pass transformation to remove remaining branch hazards. *Comp/PP* represents the compiler-assisted multiple instruction rollback scheme.

Due to the excessive compile times of the previous *Compiler 1* and *Compiler 2* algorithms for large applications, the evaluations of these schemes were restricted to applications QUEEN, WC, COMPRESS, CMP, PUZZLE, and QSORT. Both *Comp/PP* and *PP* were evaluated for all 11 applications.

B. Performance Analysis

Compiler transformations used for the removal of data hazards can impact performance in several ways. Loop protection inserts save/restore operations at the head and tail of the loop, increasing the path length and, therefore, the run time. Additional arcs in the dependency graph can cause more spill code to be generated, increasing memory references and cache misses. Nop insertion can be costly since up to N nops could be inserted for each unresolved hazard. The insertion of MOV r_k, r_k instructions to create covering on-path hazards in the post-pass transformation also increases path lengths, although typically less than with nop insertions. Finally, the increase in code size, mainly due to loop expansion, may cause more run-time cache misses. The performance numbers shown in




Fig. 10. Run-time and code size overhead for QUEEN.

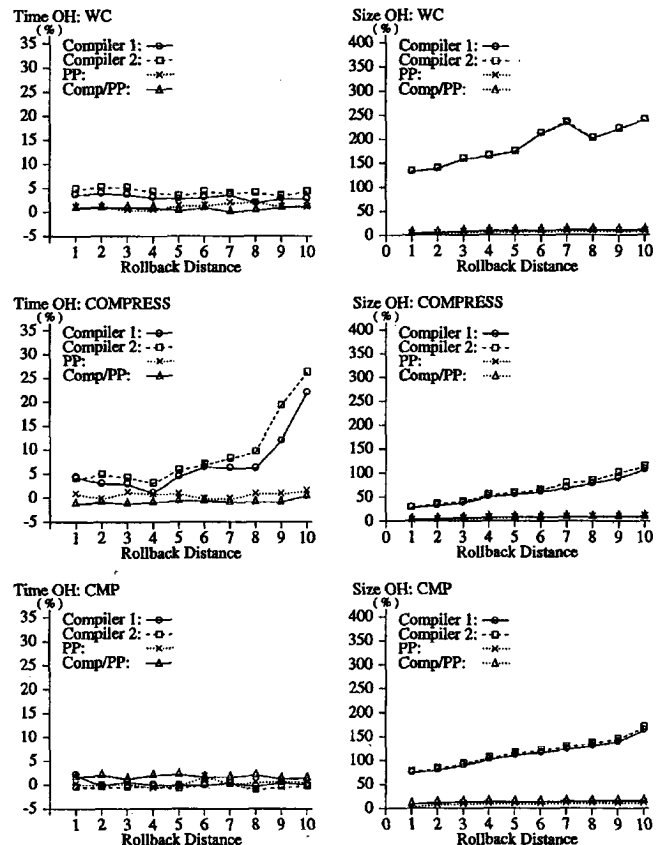


Fig. 11. Run-time and code size overhead for WC, COMPRESS, and CMP.

Figs. 10 through 14 are for execution of the eleven application programs on a DECstation 3100 after they have been compiled with the transforms described.

C. Results: Compiler 2

As can be seen in Figs. 10 through 12, extending the compiler hazard resolution scheme to include branch hazards introduces little incremental performance impact or code growth overhead. Given a rollback distance of 10, resolving both on-path and branch hazards using compiler transformations resulted in a maximum performance impact of 35.4% and an average performance impact of 15.4%. This compares with maximum and average impacts of 32.6% and 12.6%, respectively, for compiler-driven on-path hazard resolution only. The maximum code size overhead measured for the extended compiler-based technique was 372% with an average overhead

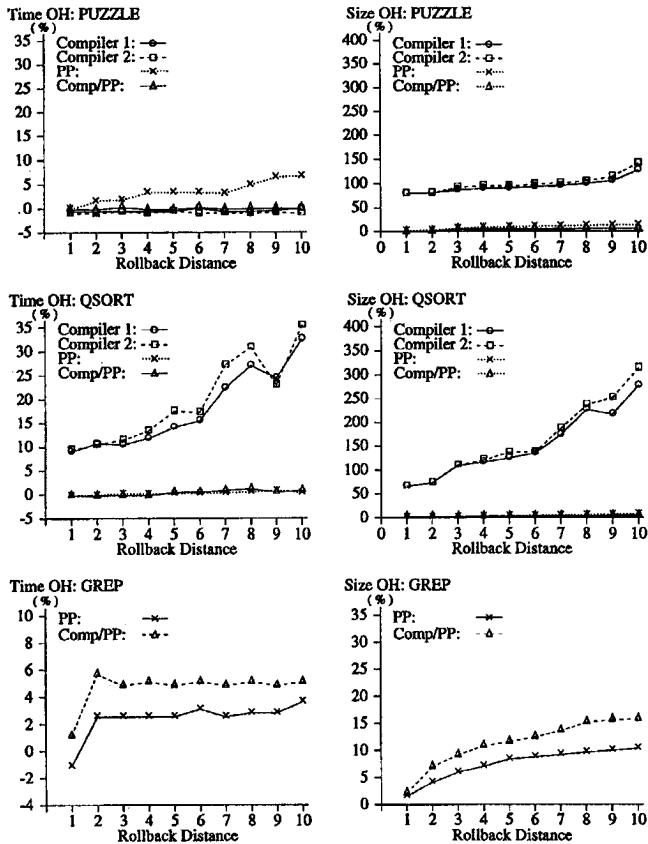


Fig. 12. Run-time and code size overhead for PUZZLE, QSORT, and GREP.

of 225%, for a rollback distance of 10. This compares with a maximum and average overhead of 328% and 207%, respectively, for the unextended compiler-based scheme.

These results indicate a small incremental run-time performance overhead and a small code size overhead given compiler-based branch hazard removal compared to compiler-based on-path hazard removal alone. Three factors account for these small incremental impacts. First, on-path hazards dominate in frequency of occurrence. Second, resolving an on-path hazard at instruction I_i through renaming can sometimes resolve a branch hazard at instruction I_i . Third, resolving on-path hazards with nop insertion may resolve a corresponding branch hazard by increasing the distance between the hazard node and its nearest predecessor branch node.

D. Results: PP

Figs. 10 through 14 show the run-time and code size overheads for each application studied using the read buffer to resolve on-path hazards and the post-pass transformation described in Section III to cover all branch hazards. The results are worst case in that many of the branch hazards could have been resolved with no performance impact using the compiler techniques; instead, they are resolved by the insertion of MOV instructions which cause a guaranteed, although small, performance impact. Given a rollback distance of 10, the post-pass transformation produced a maximum performance impact of 7.7% with an average performance impact of 2.4%, significantly below the levels produced by the compiler-based

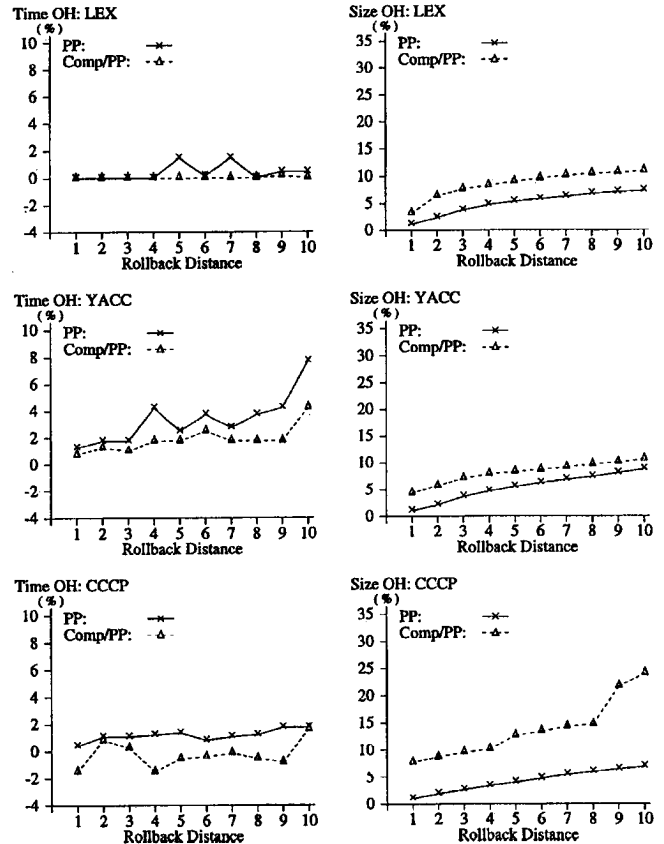


Fig. 13. Run-time and code size overhead for LEX, YACC, and CCCP.




Fig. 14. Run-time and code size overhead for TBL.

scheme. Code growth overhead measurements were correspondingly lower with a maximum overhead of 13.0% and an average overhead of 8.6%.

E. Results: Comp/PP

The compiler-assisted scheme achieved consistently low performance overheads across all applications and slightly better performance than with the post-pass transformation only. Given a rollback distance of 10, the compiler-assisted scheme produced a maximum performance impact of 6.6% with an average performance impact of 2.0%, and a maximum code growth overhead of 51.2% with an average overhead of 15.5%. The run time results of PUZZLE, YACC, and CCCP indicate that compiler techniques are still useful in reducing run-time performance penalties. These compiler techniques, however, have the disadvantage of requiring recompilation and

additional code growth. The primary advantage of the compiler-assisted and post-pass schemes are their utilization of the read buffer to resolve the more frequent on-path hazards.

V. READ BUFFER SIZE REDUCTION

A practical lower bound and average size requirement for the read buffer are established in this section by modifying the design to save only the data required for rollback. The study measures the effect on the performance of ten application programs using six read buffer configurations with varying read buffer sizes. Two alternative configurations are shown to be the most efficient.

Rollback is accomplished with a read buffer by first flushing the read buffer back to the general purpose register file in the reverse order of which the values were saved. Provided that the depth of the dual first-in-first-out (FIFO) read buffers is N , redundant copies of the appropriate register values are available to restore the register file given a rollback of $\leq N$.

The read buffer size requirement of $2N$ is the worst case. The buffer maintains the last N register reads from the register file, assuring data redundancy for all values required. The read buffer may also save data that is not required during rollback. Register reads that must be saved can be determined at compile time. If this information is added to the instruction encoding (e.g., as an extra bit field for source 1 and for source 2), then the read buffer can be designed to save only those values required. As long as the required values are maintained for N cycles, a read buffer size of less than $2N$ is possible.

Fig. 15 illustrates a case in which all register reads do not have to be placed in the read buffer. The register values (denoted $value(r_i)$) that require saving are marked with an “*.” Since only the required values are saved, the read buffer total size can now potentially be less than N . In this case, however, the instruction count must also be saved so that the value can be maintained for at least N cycles. In the event that the read buffer overflows, the oldest value in the buffer must be pushed to memory and a record kept so that during rollback the value can be retrieved from memory. Given a dual FIFO depth of M , memory would serve the function of the remaining $N - M$ of the two FIFOs.

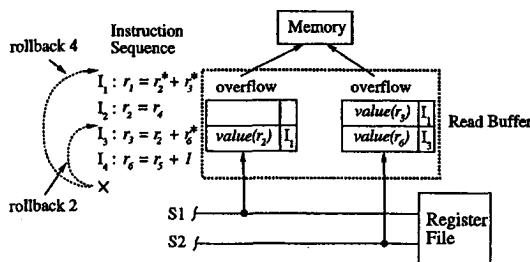


Fig. 15. Read buffer of size $< 2N$.

A. Read Buffer Designs and Evaluation Methodology

Six read buffer configurations were studied. Configuration A1, shown in Fig. 16, has a separate FIFO for each source

bus. Configuration A2 allows access to either FIFO from either source bus. Configuration B1 contains a single FIFO and assumes that both source operands can be written into the single FIFO within the same cycle. This latter *split-cycle-save* assumption is consistent with a register file design that writes during the first half of the cycle and reads during the second half of the cycle [20]. Configuration B2 assumes no split-cycle-save capability. Configuration C contains a single level dual queue to absorb a simultaneous operand save and configuration D extends this design to allow access to either queue from either source bus.

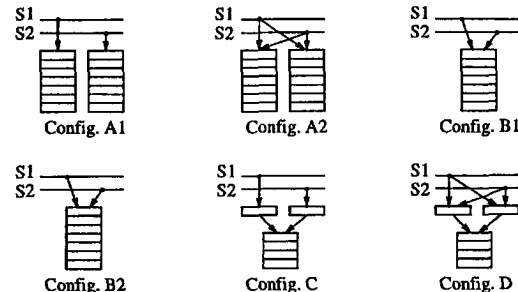


Fig. 16. Read buffer configurations.

The read buffer was simulated at the instruction level. The s-code emitted by the IMPACT C compiler [19] was instrumented with procedure calls to a simulation program containing models for the six read buffer configurations. Branch hazards were removed by the compiler for a rollback distance of 10. Parameters such as which operands require saving in the read buffer were determined at the post-pass level and instrumentation code segments were adjusted to pass this information to the simulation program. Table III lists the ten⁵ application programs used in the evaluations. The applications were cross-compiled on a SPARCserver 490 and run on a DECstation 3100 with read buffer sizes ranging from 0 to 20 (note that 20 represents the maximum read buffer size of $2N$).

B. Evaluation Results

B.1. Detailed Analysis: QUEEN

Fig. 17 shows changes in performance overhead (Cycles OH) for various read buffer sizes and configurations running the QUEEN application. Looking at Fig. 17, configuration A1, it can be seen that significant performance impact is incurred even with a modest reduction in read buffer size. Configuration A1 was consistently the least efficient of the six alternatives across the ten applications studied.⁶ This is due to the fact that the dual FIFO's are dedicated to a single source bus. In many cases saving S1 will cause an overflow because the S1 FIFO is full, even though there is room in the S2 FIFO. Configuration A1 does allow for simultaneous saves of S1 and S2, given sufficient room in each, but this feature does not compensate for the latter inefficiency. Configuration A2 demon-

5. The TBL application was not included in the read buffer size evaluation.

6. An efficient configuration is one with a low performance overhead given a small read buffer size.

strates the improvement gained by allowing either source bus access to either FIFO. Configuration B1 was the most efficient of the six configurations for the QUEEN application. In this configuration a total read buffer size of 13 would produce zero performance impact with a 35% reduction in read buffer size.




Fig. 17. Cycle overhead for QUEEN.

It should be noted that configuration B1 assumes that simultaneous saves of S1 and S2 can be handled within the same cycle. If this latter assumption is invalid, Fig. 17, configuration B2, shows that no less than 9.4% performance impact is achieved regardless of the read buffer size. The “leveling off” of B2 is due to the bottleneck at the single FIFO entry point and not the depth of the FIFO. The flat part of the curve shows the percent of instructions requiring simultaneous saves of S1 and S2 in the QUEEN application.

Fig. 17, configuration C, shows how a single level dual queue placed between the source bus and the single FIFO can alleviate some of the bottleneck effects. The dual queue can absorb a single simultaneous save of S1 and S2, distributing the saves over multiple cycles. A nonzero minimum performance overhead is still present due to cases in which the dual queue has not emptied before the next simultaneous save occurs.

Fig. 17, configuration D, shows the results of an improved queue structure that permits saves from either bus into either queue. This configuration avoids stalls in some cases (e.g., S2 must be saved while the queue dedicated to S2 in configuration C is full and the other queue is empty). Configuration D also has a nonzero minimum performance overhead but gives better performance than configuration C.

The simulation results for QUEEN show that configuration A1 is the least efficient and that given the ability to do split-cycle-saves, configuration B1 is the most efficient. Without the split-cycle-save capability, configuration D is the best of the single FIFO designs resulting in a minimum performance overhead of 4.5%, and configuration A2 is the best of the dual FIFO designs resulting in a 1.7% performance overhead with a read buffer size of 14. For configurations B1, B2, C, and D, a total read buffer size of 13 is sufficient to maximize performance.⁷

B.2. Evaluation of All Application Programs

Results for the other nine application programs are similar to those for QUEEN [17]. The differences between the application results are the points at which the curve “levels off”

(i.e., the buffer size) and, in the case of configurations B2 through D, at what level the performance overhead stabilizes. Table IV summarizes measurements obtained for the ten applications given the two most efficient configurations, A2 and B1. Configuration comparisons are made at read buffer size values that produce low values of performance overhead. Configuration A2 does not level off like configuration D and does not rapidly approach zero like configuration B1. For a better comparison of configurations A2 and B1, Table IV gives the read buffer size value where the performance overhead value drops below 3%. The read buffer size value is referred to as *RB_size* and the performance overhead value is referred to as *OH_level*.

TABLE IV
READ BUFFER SIZE EVALUATION SUMMARY

Program	<i>RB_size</i>		<i>OH_level</i> (%)	
	A2	B1	A2	B1
QUEEN	14	12	1.66	1.36
WC	10	8	0.00	2.54
QSORT	16	15	2.28	0.94
CMP	12	11	0.00	0.00
GREP	10	10	0.18	0.18
PUZZLE	10	9	2.87	0.32
COMPRESS	12	12	2.87	1.12
LEX	12	12	2.73	1.55
YACC	16	15	1.07	0.00
CCCP	12	12	2.34	1.74

It can be seen from Table IV that the read buffer size requirement is roughly the same, per application, regardless of the split-cycle-save assumption (i.e., comparing configurations A2 and B1). The size requirement is application dependent—from 8 for WC to 15 for QSORT and YACC. The measurements show that a considerable reduction in read buffer size is achievable. Given the split-cycle-save assumption and configuration B1, a minimum of 25%, a maximum of 60%, and an average of 42% reduction was achieved. For configuration A2 and no split-cycle-save assumption, a minimum of 20%, a maximum of 50%, and an average reduction of 38% was achieved. The measurements indicate that care should be taken relative to the ultimate selection of read buffer size. Given the steepness of the B1 curve around the *RB_size* value, small decreases in size can produce large performance overheads.

In summary, a dual FIFO with source bus access to each buffer (configuration A2) and the single FIFO with the split-cycle-save capability (configuration B1) consistently outperformed the other four configurations. There were moderate variances between the buffer sizes required for minimum performance impact between the ten applications studied and the performance stabilization value assuming no split-cycle-save capability. Up to a 55% read buffer size reduction was achieved with an average reduction of 39.5% given the most efficient read buffer configuration for the applications. Although significant read buffer size reductions are possible without adversely affecting performance, it should be noted that such an approach requires an additional data-path to the memory unit and more complex recovery logic.

7. Two must be added to each read buffer size value in C and D to account for the queues.

VI. CONCLUDING REMARKS

This paper presented a compiler-assisted multiple instruction rollback scheme that combines compiler-driven data-flow manipulations with dedicated data redundancy hardware to remove data hazards resulting from multiple instruction rollback. Experimental evaluation of the compiler-assisted scheme with a maximum rollback distance of ten showed performance impacts of no more than 6.6% and an average of 2.0%, over the eleven application programs studied. The performance evaluation indicates performance penalties that are lower than for previous compiler-only approaches and hardware complexity that is less than previous hardware-only approaches. Six read buffer configurations were studied to determine the minimum size requirement for general applications. It was found that a significant read buffer size reduction is achievable, but the additional control logic to handle read buffer overflows may limit the overall hardware savings. Current research includes application of compiler-assisted multiple instruction rollback recovery to super-scalar, VLIW, and parallel processing architectures. Extensions of compiler-assisted multiple instruction recovery to speculative execution repair are also under development.


ACKNOWLEDGMENTS

The authors wish to thank C.-C. Jim Li for his help with the compiler aspects of this paper, and Scott Mahlke and William Chen for their invaluable assistance with the IMPACT compiler. We also express our thanks to Janak Patel for his contributions to this research.

This research was supported in part by the National Aeronautics and Space Administration (NASA) under grant NASA NAG 1-613, in cooperation with the Illinois Computer Laboratory for Aerospace Systems and Software (ICLASS), and in part by the U.S. Department of the Navy and managed by the Office of the Chief of Naval Research under Contract N00014-91-J-1283.

REFERENCES


- [1] M.S. Pittler, D.M. Powers, and D.L. Schnabel, "System development and technology aspects of the IBM 3081 processor complex," *IBM J. Research and Development*, vol. 26, pp. 2-11, Jan. 1982.
- [2] Y. Tamir and M. Tremblay, "High-performance fault-tolerant VLSI systems using micro rollback," *IEEE Trans. Computers*, vol. 39, pp. 548-554, Apr. 1990.
- [3] C.-C.J. Li, S.-K. Chen, W.K. Fuchs, and W.-M.W. Hwu, "Compiler-based multiple instruction retry," *IEEE Trans. Computers*, vol. 44, pp. 35-46, Jan. 1995.
- [4] N.J. Alewine, S.-K. Chen, C.-C.J. Li, W.K. Fuchs, and W.-m.W. Hwu, "Branch recovery with compiler-assisted multiple instruction retry," *Proc. 22nd Int'l Symp. Fault-Tolerant Computing*, pp. 66-73, July 1992.
- [5] L. Spainhower, J. Isenberg, R. Chillarege, and J. Berding, "Design for fault-tolerance in system ES/9000 model 9000," *Proc. 22nd Int'l Symp. Fault-Tolerant Computing*, pp. 38-47, July 1992.
- [6] P.M. Kogge, K.T. Truong, D.A. Richard, and R.L. Schoenike, "Checkpoint retry mechanism," U.S. Patent No. 4912707, Mar. 1990. Assignee: International Business Machines Corporation, Armonk, N.Y.
- [7] Y. Tamir, M. Liang, T. Lai, and M. Tremblay, "The UCLA mirror processor: A building block for self-checking self-repairing computing nodes," *Proc. 21st Int'l Symp. Fault-Tolerant Computing*, pp. 178-185, June 1991.
- [8] N.J. Alewine, W.K. Fuchs, and W.-m.W. Hwu, "Application of compiler-assisted rollback recovery to speculative execution repair," *Hardware and Software Architectures for Fault Tolerance*. New York: Springer-Verlag, 1994.
- [9] J.E. Smith and A.R. Pleszkun, "Implementing precise interrupts in pipelined processors," *IEEE Trans. Computers*, vol. 37, pp. 562-573, May 1988.
- [10] M.L. Ciacelli, "Fault handling on the IBM 4341 processor," *Proc. 11th Int'l Symp. Fault-Tolerant Computing*, pp. 9-12, June 1981.
- [11] W.F. Brucker and R.E. Josephson, "Designing reliability into the VAX 8600 system," *Digital Tech. J. Digital Equipment Corp.*, vol. 1, no. 1, pp. 71-77, Aug. 1985.
- [12] G.L. Hicks, D. Howe Jr., and A. Zurla Jr., "Instruction retry mechanism for a data processing system," U.S. Patent No. 4044337, Aug. 1977. Assignee: International Business Machines Corp., Armonk, N.Y.
- [13] D.B. Fite, T. Fossum, and D. Manley, "Design strategy for the VAX 9000 systems," *Digital Tech. J. Digital Equipment Corp.*, vol. 2, no. 4, pp. 13-24, Fall 1990.
- [14] E.B. Eichelberger and T.W. Williams, "A logic design structure for LSI testability," *Proc. 14th Design Automation Conf.*, pp. 462-468, 1977.
- [15] J.S. Liptay, "The ES/9000 high end processor design," *IBM J. Research and Development*, vol. 36, no. 3, May 1992.
- [16] A.V. Aho, R. Sethi, and J.D. Ullman, *Compilers: Principles, Techniques, and Tools*. Reading, Mass.: Addison-Wesley, 1986.
- [17] N.J. Alewine, "Compiler-assisted multiple instruction rollback recovery using a read buffer," PhD thesis, Tech. Rep. CRHC-93-06, Univ. of Illinois at Urbana-Champaign, 1993.
- [18] J.A. Bondy and U. Murty, *Graph Theory with Applications*. London: Macmillan Press Ltd., 1979.
- [19] P. Chang, W. Chen, N. Warter, and W.-m.W. Hwu, "IMPACT: An architecture framework for multiple-instruction-issue processors," *Proc. 18th Ann. Symp. Computer Architecture*, pp. 266-275, May 1991.
- [20] J.L. Hennessy and D.A. Patterson, *Computer Architecture: A Quantitative Approach*. San Mateo, Calif.: Morgan Kaufmann, 1990.




Neal J. Alewine received his BS and MS degrees from Florida Atlantic University in 1980 and 1988, respectively. He received his PhD in electrical engineering at the University of Illinois in 1993.

Dr. Alewine is a senior engineer with the International Business Machines Corp. in Boca Raton, Florida. He is currently the technical staff to a business area manager responsible for a family of coprocessor products used in personal computers and work stations. Since joining IBM in 1980 he has held several positions including designer, first and

second level management, project leader, and program management. He has received IBM outstanding technical achievement and invention achievement awards and was selected as a University Scholar at Florida Atlantic University in 1979. Dr. Alewine's interests include advanced computer architecture, systems design, optimizing compiler design, and fault-tolerant computing.




Shyh-Kwei Chen (S'86-M'94) received the BS degree from the National Taiwan University, Taipei, Taiwan, in 1983, the MS degree from the University of Minnesota, Minneapolis, in 1987, and the PhD from the University of Illinois Urbana-Champaign, all in computer science. He is now with IBM T.J. Watson Research Center, Yorktown Heights, New York. His research interests include parallel processing, fault-tolerant computing, compilers, and parallel debuggers.



W. Kent Fuchs (S'83-M'85-SM'90) received the BSE degree in electrical engineering from Duke University in 1977. In 1984 he received the MDiv degree from Trinity Evangelical Divinity School in Deerfield, Illinois, and in 1985 he received the PhD in electrical engineering from the University of Illinois.

Dr. Fuchs is currently a professor in the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory, University of Illinois. He is the recipient of many awards including the Senior Xerox Faculty Award for Excellence in Research 1993, College of Engineering, University of Illinois, selection as a University Scholar, University of Illinois 1991, appointment as Fellow in the Center for Advanced Studies, University of Illinois 1990, the Xerox Faculty Award for Excellence in Research 1987, University of Illinois, the Digital Equipment Corporation Incentives for Excellence Faculty Award 1986-1988, and the Best Paper Award, IEEE/ACM Design Automation Conference (DAC) 1986, simulation and test category.

Dr. Fuchs was the guest editor of the May 1992 special issue of *IEEE Transactions on Computers* on fault-tolerant computing, and guest coeditor of the April 1992 special issue of *Computer* on wafer-scale integration architectures and algorithms. He has been a member of the editorial board for *IEEE Transactions on Computers* and *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*.



Wen-mei W. Hwu (S'81-M'87) received his PhD degree in computer science from the University of California, Berkeley, in 1987.

Dr. Hwu is an associate professor in the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign. His research interest is in the area of architecture, implementation, and compilation for high-performance computer systems. He is the director of the IMPACT project, which has delivered new compiler and computer architecture technologies to the computer industry since 1987. The IMPACT project has been sponsored by NSF, ONR, and NASA as well as major corporations such as Hewlett-Packard, Intel, SUN, NCR, AMD, and Matsushita. In recognition of his contributions to the areas of compiler optimization and computer architecture, the Intel Corporation named him the Intel Associate Professor at the College of Engineering, University of Illinois, in 1992. He received the National Eta Kappa Nu Outstanding Young Electrical Engineer Award for 1993 and the 1994 Senior Xerox Award for Faculty Research.