

# Path Analysis and Renaming for Predicated Instruction Scheduling

Lori Carter    Beth Simon    Brad Calder    Larry Carter    Jeanne Ferrante

Department of Computer Science and Engineering

University of California, San Diego

San Diego, CA, USA 92093-0114

Phone: 858-534-2466

Fax: 858-534-7029

{lrcarter,esimon,calder,carter,ferrante}@cs.ucsd.edu

## Abstract

*Increases in instruction level parallelism are needed to exploit the potential parallelism available in future wide issue architectures. Predicated execution is an architectural mechanism that increases instruction level parallelism by removing branches and allowing simultaneous execution of multiple paths of control, only committing instructions from the correct path. In order for the compiler to expose and use such parallelism, traditional compiler data-flow and path analysis needs to be extended to predicated code.*

*In this paper, we motivate the need for renaming and for predicates that reflect path information. We present Predicated Static Single Assignment (PSSA) which uses renaming and introduces Full-Path Predicates to remove false dependences and enable aggressive predicated optimization and instruction scheduling. We demonstrate the usefulness of PSSA for Predicated Speculation and Control Height Reduction. These two predicated code optimizations used during instruction scheduling reduce the dependence length of the critical paths through a predicated region. Our results show that using PSSA to enable speculation and control height reduction reduces execution time from 12% to 68%.*

*Key Words: predicated execution, static single assignment, instruction scheduling, renaming*

# 1 Introduction

The Explicitly Parallel Instruction Computing (EPIC) architecture has been put forth as a viable architecture for achieving the *instruction level parallelism* (ILP) needed to keep increasing future processor performance [8, 17]. Intel’s application of EPIC architecture technology can be found in their IA-64 architecture whose first instantiation is the Itanium processor [1]. An EPIC architecture issues wide instructions, similar to a VLIW architecture, where each instruction contains many operations.

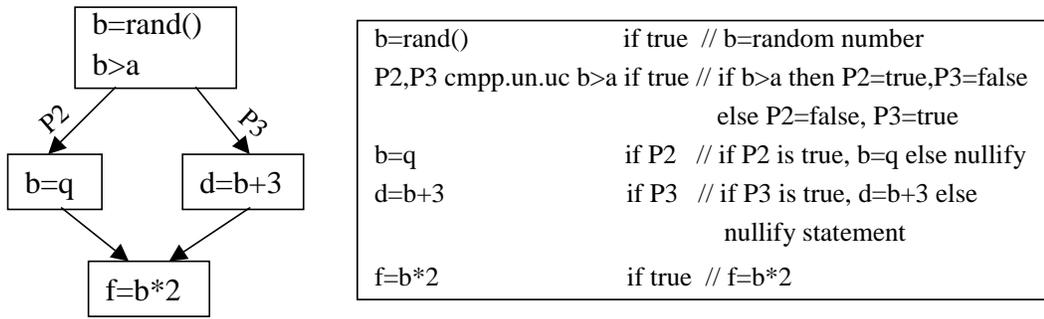
One of the new features of the EPIC architecture is its support for *predicated execution* [24], where each operation is guarded by one of the predicate registers available in the architecture. An operation is committed only if the value of its guarding predicate is true.

One advantage of predicated execution is that it can eliminate hard-to-predict branches by combining both paths of a branch into a single path. Another advantage comes from using predication to combine several smaller basic blocks into one larger hyperblock [22]. This provides a larger pool from which to draw ILP for EPIC architectures.

A significant limitation to ILP is the presence of control-flow and data-flow dependences. Static Single Assignment (SSA) is an important compiler transformation used to remove false data dependences across basic block boundaries in a control flow graph [12]. Removing these false dependences reveals more ILP, allowing better performance of optimizations like instruction scheduling. Without performing SSA, the benefit of many optimizations on traditional code is limited.

Eliminating false dependences is equally important and a more complex task for predicated code, since multiple control paths are merged into a single predicated region. However, the control-flow and data-flow analysis needed to support predicated compilation is different than traditional analysis used in compilers for superscalar architectures. A sequential region of predicated code contains not only data dependences, but also *predicate dependences*. A predicate dependence exists between every operation and the definition(s) of its guarding predicate. Our technique introduces a chain of predicate dependences which represents a unique control path through the original code.

We describe a predicate-sensitive implementation of SSA called *Predicated Static Single Assignment* (PSSA). PSSA introduces *Full-Path Predicates* to extend SSA to handle predicate dependences and the multiple control paths that are merged together in a single predicated region. We demonstrate that PSSA allows effective predicated scheduling by (1) eliminating false dependences along paths via renaming, (2) creating full-path predicates, and (3) providing path-sensitive data-flow analysis. We show the benefit of using PSSA to perform Predicated Speculation and Control Height Reduction during instruction scheduling. Using PSSA allows these two optimizations, when applied together, to schedule all operations at their *earliest schedulable cycle*. In our implementation, the earliest schedulable cycle takes into consideration true data dependences and load/store constraints. In this paper we expand



a) Original Control Flow Graph

b) Predicated Hyperblock

Figure 1: Short code example showing the transformation from non-predicated code to predicated hyperblock.

upon work we presented in [11] by including additional benchmarks and by motivating the need for renaming and for predicates that reflect path information above and beyond what is available from traditional If-converted code.

The paper is organized as follows. Section 2 describes predicated execution. Section 3 motivates the need for predicate-sensitive analysis and full-path predicates. Section 4 presents Predicated Static Single Assignment. Section 5 shows how PSSA can enable aggressive Predicated Speculation and Control Height Reduction. Section 6 reports the increased ILP and reduced execution times achieved by applying our algorithms to predicated code. Section 7 summarizes related work. Section 8 discusses using PSSA within the IA-64 framework, and Section 9 describes our future work. Finally, Section 10 summarizes the contributions of this paper.

## 2 Predicated Execution

Predicated execution is a feature designed to increase ILP and remove hard-to-predict branches. It has also been used to support software pipelining[14, 25]. Machines with hardware to support predicated code include an additional set of registers called predicate registers. The process of predication replaces branches with compare operations that set predicate registers to either true or false based on the comparison in the original branch. Each operation is then associated with one of these predicate registers which will hold the value of the operation’s *guarding predicate*. The operation will be committed only if its guarding predicate is true<sup>1</sup>. This process of replacing branches with compare operations and associating operations with a predicate defined by that compare is called *If-Conversion* [5, 24].

Our work uses the notion of a hyperblock [22]. A *hyperblock* is a predicated region of code consisting of a straight-line sequence of instructions with a single entry point and possibly multiple exit points. Branches with both targets in the hyperblock are eliminated and converted to predicate definitions using If-conversion. All remaining

<sup>1</sup>One exception is the unconditional definition of a predicate. This is discussed later in the section.

branches have targets outside the hyperblock. Consequently, there are no cyclic control-flow or data-flow dependences within the hyperblock. The selection of instructions to be included in the hyperblock is based on program profiling of the original basic blocks which includes information such as execution frequency, basic block size, operation latencies, and other characteristics [22].

A typical code section to include in a hyperblock is one that contains a hard-to-predict (unbiased) branch [21], as shown in Figure 1. After If-conversion, the Control Flow Graph (CFG) in Figure 1(a), which is comprised of four basic blocks, results in the predicated hyperblock shown in Figure 1(b). All operations in the hyperblock are now guarded, either by a predicate register set to the constant value of true, or by a register that can be defined as either true or false by a `cmpp` (compare and put (result) in predicate) operation. Operations guarded by the constant true, such as the operation  $f = b * 2$  in Figure 1, will be executed and committed regardless of the path taken. Operations guarded by a predicate register, such as the operation  $b = c$ , will be put into the pipeline, but only committed if the value of the operation's guarding predicate ( $P2$  for this operation) is determined to be true.

In what follows, we describe three types of operations that can be included in a hyperblock – `cmpp` operations, the predicate OR operation, and normal (non-predicate-defining) operations.

As defined in the Trimaran System [2] (which supports EPIC computing via the Playdoh ISA [19]), guarding predicates are assigned their values via `cmpp` operations [8]. Consider an operation  $B, C \text{ cmpp.un.ac } a > c \text{ if } A$  as an example. The `cmpp` operation can define one or two predicates. This operation will define predicates B and C. The first tag (`.un`) applies to the definition of the first predicate B and the second tag (`.ac`) to C. The first character of a tag defines how the predicate is to be defined. The character `u` means that the predicate will unconditionally get a value, whether the guarding predicate (A in this case) is true or false. If A is false, then B is set to false. Otherwise, A is true and the value of B depends upon the evaluation of  $a > c$ .

The character `a` in the second tag (`.ac`) indicates that the full definition of the related predicate C is contingent on the value of A, the evaluation of  $a > c$ , AND the prior value of C. If A is false, the value of predicate C does not change. If A is true and C has previously been set false then C remains false. Additionally, the second character of a tag defines whether the normal (`n`) result of the condition ( $a > c$ ) or the complement (`c`) of the condition must be true to make the related predicate true. If A is true and C is true and  $!(a > c)$  is true then the new value of C will be true<sup>2</sup>. For a complete definition of `cmpp` statements see the Playdoh architecture specification [19].

In our implementation of PSSA, we introduce a new OR operation currently not defined by Trimaran. The *predicate* OR operation defines block predicates by taking the logical OR of multiple predicates. For example, consider an operation  $G = \text{OR}(A, B, C) \text{ if true}$  (where A, B and C are predicates, each defining a unique path to G). If any one of them has the value of true, G will receive a value of true, otherwise G will be assigned false.

---

<sup>2</sup>Conversely, if A is true and C is true and  $!(a > c)$  is false, then the new value of C will be false.

When scheduling, we make the reasonable assumption that the definition of a predicate is available for use as a source for another operation, or as a guard to a subsequent `cmp` operation in the cycle following its definition. When used as a guard for all other operations, the predicate definition is available for use in the same cycle as it is defined.

We refer to all other operations, which do not define predicates, as *normal* operations. Normal operations include assignments, arithmetic operations, branches, and memory operations.

### 3 Motivation for Predicate-Sensitive Analysis

A major task for the scheduler of a multi-issue machine is to find independent instructions. Unfortunately, predication introduces additional dependences that traditional code doesn't have to consider. In Figure 1(b), there is a dependence between the definition of the guarding predicate  $P2$  and its use in the statement  $b=c$  if  $P2$ . Since predication combines multiple basic blocks, it introduces false dependences between disjoint paths. For example, in Figure 1(b), in the absence of predicate dependence information, we would infer a dependence between the definition of  $b$  in  $b=c$  if  $P2$  and the use of  $b$  in  $d=b+3$  if  $P3$ . However, these two statements are guarded by *disjoint* predicates. Therefore, only one of the predicates ( $P2$  or  $P3$ ) can possibly be true; only one of the statements will actually be committed and no dependence does in fact exist.

Johnson et. al. [18] devised a scheme to determine the disjointness of predicates using the predicate partition graph. This analysis allowed more effective register allocation as live ranges across predicated code could be more accurately determined [15]. Their approach was limited to describing disjointness with restricted path information. Path information that extended across join points was not collected. In Figure 2, the predicate partition graph would determine that the following pairs of predicates are disjoint:  $G$  and  $H$ ,  $B$  and  $C$ ,  $D$  and  $\neg D$ . However, no information regarding the relationship between  $D$  and  $G$  or  $D$  and  $H$  would be available.

This “cross-join” information is needed to provide the scheduler full flexibility in scheduling statements such as  $y=t+r$ . If path information is not available, then  $y=t+r$  is guarded on true and the scheduler correctly assumes this statement is dependent on  $t=r$  and  $( )$ ,  $t=t-s$ ,  $r=5+x$ , and  $r=x+8$ . However, since there are two possible definitions of each operand, there are 4 combinations of operands that could in fact cause the definition of  $y$  – each executable via one (or more) paths of execution through the region. If 4 versions of this statement could be made (one for each combination of the operands), then each could be scheduled at the minimum dependence length for that version. While disjointness information can maintain information regarding paths since the most recent join, we will need to combine path information across joins to remove unnecessarily conservative scheduling dependences. Figure 2(b) shows the cross-join path information that would be needed to guard each assignment of  $y$  so that the

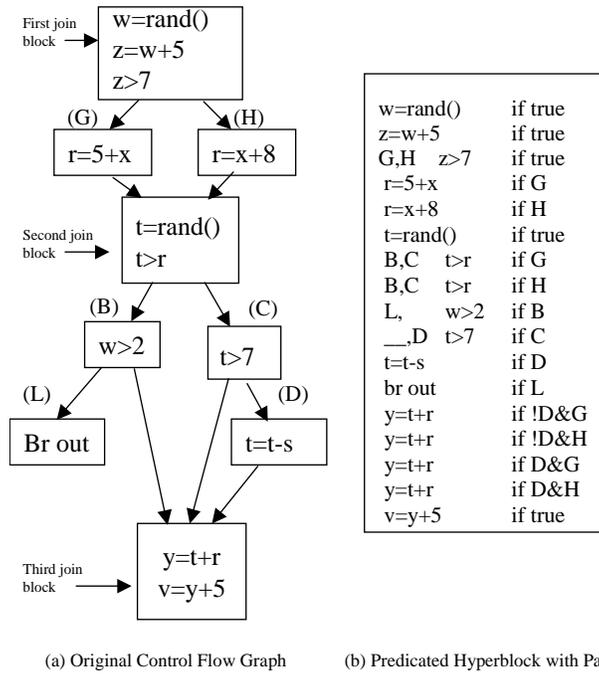


Figure 2: Code is duplicated when more than one definition reaches a use to maintain maximum flexibility for the scheduler. In (b), the statement  $y=t+r$  is duplicated for each pair of definitions that may reach this statement. Each copy is guarded by the predicates that defined the path along which those definitions would occur.

scheduler can know the precise dependences for each copy. This will allow the most flexibility in scheduling each statement.

Although precise dependence information can be determined from guarding predicate relations, we will also show that renaming techniques can be of additional use to achieve greater scheduling flexibility. By renaming variables that have more than one definition in a region, we will maintain path information even after optimizations which change the guarding predicate of a statement have been applied.

## 4 Predicated Static Single Assignment (PSSA)

Techniques such as renaming [4] and Static Single Assignment (SSA) [13, 12] have proved useful in eliminating false dependences in traditional code [31]. Removing false dependences allows more flexibility in scheduling since data independent operations can move past each other during instruction scheduling.

In non-predicated code, SSA assigns each target of an assignment operation a unique variable. At join nodes a  $\phi$ -function may need to be inserted if multiple definitions of a variable reach the join. The  $\phi$ -functions determine which version of the variable to use and assign it to an additional renamed version. This new variable is used to represent the merging of the different variable names. Figure 3 shows the simple example from Figure 1 in SSA form. In the assignment  $b3 \rightarrow \phi(b1, b2)$ , the variable  $b3$  represents the reaching definition of  $b$  which is to be

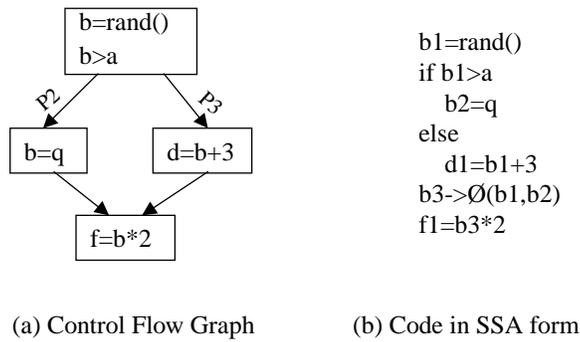


Figure 3: Static Single Assignment

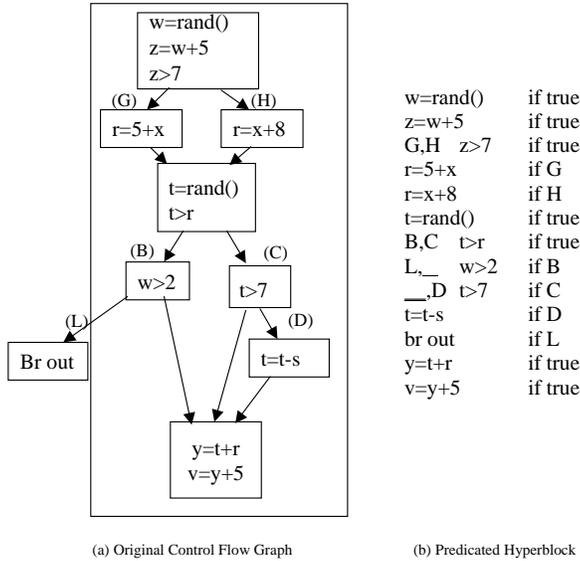


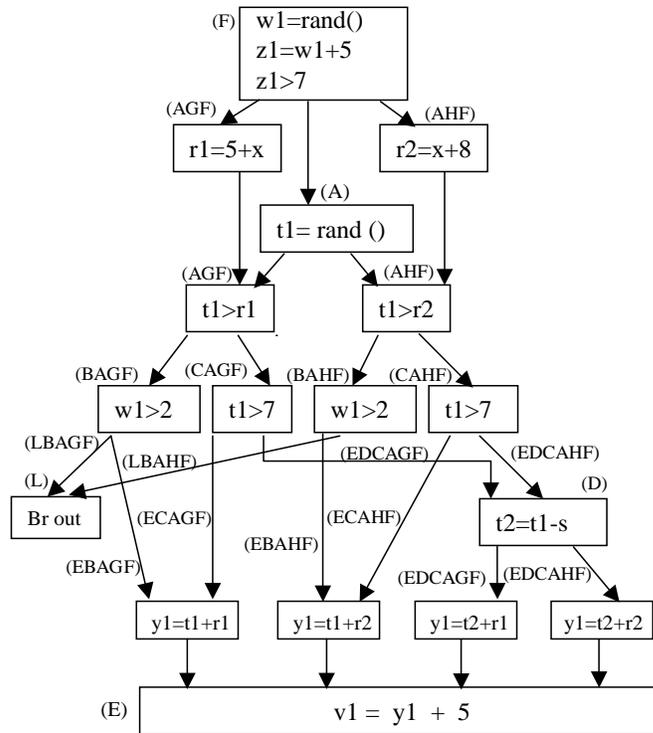
Figure 4: Extended example of transformation from non-predicated CFG to predicated hyperblock

used after the join of definition b1 or b2.

As discussed in section 3, eliminating false dependences is equally important and a more complex task for predicated code, since multiple control paths are merged. To address this problem we developed a predicate-sensitive implementation of SSA called *Predicated Static Single Assignment (PSSA)*.

PSSA seeks to accomplish the same objectives as SSA for a predicated hyperblock. First, it must assign each target of an assignment operation in the hyperblock a unique variable. Second, at points in the hyperblock where multiple paths come together it must summarize under what conditions each of the multiple definitions of a variable reaches that join. The second objective is accomplished through the creation of full-path predicates and path-sensitive analysis.

Consider the sample predicated code shown in Figure 4 using traditional hyperblock predication [22]. In this predicated example, all branches have been replaced (except the one leaving the hyperblock) with predicate-defining operations using If-conversion. The predicates that are defined in this example correspond to the two edges exiting



(a) PSSA dependence graph

F=true	if true	1
w1=rand()	if F	1
z1=w1+5	if F	2
AGF,AHF	cmpp.un.uc z1>7	if F
r1=5+x	if AGF	3
r2=x+8	if AHF	3
A=OR(AGF,AHF)	if true	4
t1=rand()	if A	4
BAGF,CAGF	cmpp.un.uc t1>r1	if AGF
BAHF,CAHF	cmpp.un.uc t1>r2	if AHF
LBAGF,EBAGF	cmpp.un.uc w1>2	if BAGF
LBAHF,EBAHF	cmpp.un.uc w1>2	if BAHF
ECAGF,EDCAGF	cmpp.un.uc t1>7	if CAGF
ECAHF,EDCAHF	cmpp.un.uc t1>7	if CAHF
D=OR(EDCAGF,EDCAHF)	if true	7
t2=t1-s	if D	7
L=OR(LBAHF,LBAGF)	if true	7
br out	if L	7
y1=t1+r1	if EBAGF	6
y1=t1+r2	if EBAHF	6
y1=t1+r1	if ECAGF	6
y1=t1+r2	if ECAHF	6
y1=t2+r1	if EDCAGF	8
y1=t2+r2	if EDCAHF	8
E=OR(EBAGF,EBAHF,ECAGF, ECAHF,EDCAGF,EDCAHF)	if true	7
v1=y1+5	if E	9

(b) PSSA-transformed code

Figure 5: The PSSA dependence graph shows the flow of data and control through the PSSA-transformed code. Blocks labeled with *full-path* predicates (indicated by multiple letters) contain statements that are only executed along that path. Blocks labeled with *block* predicates (single letters) contain statements that will be executed along several paths.

each conditional branch in the CFG in Figure 4. Figure 5 shows this example after PSSA has been applied and displays a graph showing the post-PSSA dependence relationships.

The PSSA transformation has 2 phases: pre- and post-optimization. Hyperblocks are converted to PSSA form before optimization. After optimization, PSSA inserts clean-up code on edges leaving the hyperblock, copying renamed variables back to their original names and then removes any unused predicate definitions.

## 4.1 Converting to PSSA Form

When converting to PSSA form, each operation is processed in turn beginning at the top of the hyperblock and proceeding to the end. *Control PSSA* is applied to predicate-defining operations, and *Normal PSSA* is applied to all other operations.

We first describe Normal PSSA. If the operation is an assignment, the variable defined is renamed. The third operation in Figure 5(b),  $z1 = w1 + 5$ , is an example. All operands are adjusted to reflect previously renamed variables (e.g.  $w$  becomes  $w1$ ). If the operation is part of a join block, multiple versions of the operands may be live. The first operation ( $y = t + r$ ) in the third join block of Figure 2(a) provides an example. Here, the operation will be duplicated for each path leading to the join and the correct operand versions for each path will be used in the duplicate statement as seen in Figure 5 (in the multiple definitions of  $y1$ ). The duplicates are guarded by the full-path predicate (described below) associated with the path along which the operands are defined. Though there are 6 definitions of  $y1$  (only 4 are unique), there is only one definition of  $y1$  on any given path. These definitions are predicated on disjoint predicates; only one of them can possibly be true, and only one of them will be committed.

We next describe Control PSSA. The single `cmpp` operation that defined one or two block predicates (such as the definitions of B and C in Figure 4) is replaced by one or more `cmpp` operations, each associated with a particular path leading to that block. As can be seen in Figure 5(b) there are now two `cmpp` operations: one defining BAGF and CAGF, and one defining BAHF and CAHF. These new predicates are called *full-path predicates* (FPPs). Each FPP definition has the appropriate operand versions for its path and each is guarded by the FPP that defined the path prior to reaching the new block. For example, the `cmpp` defining BAGF and CAGF is predicated on AGF. A FPP specifies the unique path along which an operation is valid for execution, enabling PSSA to provide correct guarding predicates for the duplicate statements previously described.

In the example in Figure 2 we pointed out that the definitions of  $y1$  needed guarding predicates that captured information about paths of execution. The first definition of  $y1$  needed to be guarded by a predicate representing a path of execution through block G but not block D. In addition, the predicate needs to reflect that the execution actually reached the block of the statement in question (E in this case). Register  $y1$  would be incorrectly modified if, for example, the branch out of the hyperblock is taken and block E is never reached. The new FPP EBAGF represents

the precise conditions for correct execution.

In addition to the `cmpp` statements added to define FPPs, `cmpp` statements are included to rename join blocks whose statements were originally predicated on `true`. A and E and their associated FPPs are examples. The operations in Figure 4(b) predicated on `true`, are predicated on F, A and E in the PSSA version of the code shown in Figure 5. This is necessary to maintain exact path information.

Clearly, this has the potential to cause an exponential amount of code duplication. It might seem more reasonable to follow the example of SSA and insert  $\phi$ -functions at join points to resolve multiple definitions. For example, an implementation of  $\phi$ -functions resolving `r` and `t` in the definition of `y1` could be:

```
(1) r=r1 if G
(2) r=r2 if H
(3) t=t1 if true
(4) t=t2 if D
(5) y1=r+t if true
```

While this would have the advantage of decreasing duplication, it does not eliminate the need for predicate-sensitive analysis. Predicate relationship information is still needed to determine the reaching definitions and associated predicates, and to determine the order of the copy operations. For example, both of the statements (3) and (4) defining `t` in the previous sequence could be committed. The literal predicate `true` is always true, and predicate `D` could be true as well. For the use of `t` in (5) to get the correct definition, statement (4) cannot be executed before statement (3). Moreover, other side effects that degrade performance are introduced. Most important is that the insertion of  $\phi$ -functions adds data dependences. For example, a true dependence is introduced between the definition of `t1` and its use in (3). In addition, false dependences are re-introduced. An example is the output dependence between the two definitions of `t`. Thus, SSA and the usual  $\phi$ -function implementation does not give the desired scheduling flexibility.

*Block predicates* are also important to the PSSA transformation. PSSA uses predicate OR statements to redefine the block predicates as the union of the FPPs associated with the paths that reach the block. PSSA does not simply duplicate every path through the hyperblock. Duplication only occurs when necessary to remove false dependences. When there is only one version of all operands reaching a statement, only one version of the statement is required. This is the case with `v1=y1+5` in Figure 5. The variable `y1` is the only version live in node E. This statement is guarded by E, a block predicate created by taking the logical OR of EBAGF, EBAHF, ECAGF, ECAHF, EDCAGF, and EDCAHF. As long as control reaches node E, regardless of the path taken, we will execute and commit the statement `v1=y1+5`.

## 4.2 Post-Optimization Clean-up

After optimization is applied to code in PSSA form, a clean-up phase is run to remove unnecessary code and to assure consistent code outside of the hyperblock.

The PSSA implementation described in this paper generates `cmp` statements for every path and block. These are entered into the PSSA data structure that maintains information about the relationships between the predicates they define, which provides maximum flexibility during optimization. However, some of these FPP definitions may not be used, and the corresponding `cmp` operations will be discarded, reducing the code size significantly.

Finally, to assure correct execution following the hyperblock, PSSA inserts copy operations assigning the original variable names to all renamed definitions that are live out of the hyperblock. These are placed on the appropriate exit of the hyperblock. For example, the exit branch guarded by `L` in Figure 4 would include `t = t1 if L if t` was live out of the hyperblock at this exit.

## 5 Hyperblock Scheduling Optimizations

In this section, we describe how PSSA enables Predicated Speculation (PSpec) and Control Height Reduction (CHR) for aggressive instruction scheduling. PSpec allows operations to be executed before their guarding predicates are determined and CHR allows the guarding predicates to be determined as soon as possible, reducing the number of operations that need to be speculated. Used together with PSSA, we demonstrate that we can schedule the code at its earliest schedulable cycle, assuming a machine with unlimited resources.

### 5.1 Predicated Speculation

This section describes how to perform speculation on PSSA-transformed code. In general, speculation is used to relieve constraints which control dependences place on scheduling. One can speculatively execute operations from the likely-taken path of a highly-predictable branch, by scheduling those operations before their controlling branch [20]. Similarly, Predicated Speculation (PSpec) will schedule a normal operation above the `cmp` operation it is dependent upon, optimizing a hyperblock's execution time.

PSpec handles placement of the speculated predicated operation in a uniform manner. PSpec schedules a normal operation at its earliest schedulable cycle. When speculating an operation, the operation is scheduled earlier than the operation it is control dependent on, and is predicated on true. We assume that any exceptions raised by the speculated operations will be taken care of using architecture features such as poison bits [7].

F=true		if true	1
w1=rand()		if F	1
z1=w1+5		if F	2
AGF, AHF	cmpeq.un.uc z1>7	if F	3
r1=5+x		if true	1
r2=x+8		if true	1
A=OR(AGF, AHF)		if true	4
t1=rand()		if true	1
BAGF,CAGF	cmpeq.un.uc t1>r1	if AGF	4
BAHF,CAHF	cmpeq.un.uc t1>r2	if AHF	4
LBAGF,EB AGF	cmpeq.un.uc w1>2	if BAGF	5
LBAHF,EB AHF	cmpeq.un.uc w1>2	if BAHF	5
ECAGF, EDCAGF	cmpeq.un.uc t1>7	if CAGF	5
ECAHF, EDCAHF	cmpeq.un.uc t1>7	if CAHF	5
D=OR(EDCAGF,EDCAHF)		if true	6
t2=t1-s		if true	2
L=OR(LBAHF,LB AGF)		if true	6
br out		if LBAGF	5
br out		if LBAHF	5
y1=t1+r1		if true	2
y2=t1+r2		if true	2
y3=t1+r1		if true	2
y4=t1+r2		if true	2
y5=t2+r1		if true	3
y6=t2+r2		if true	3
E= OR(EBAGF,E BAHF,ECA GF, ECAHF,EDCAGF,EDC AHF)		if true	6
v1=y1+5		if true	3
v2=y2+5		if true	3
v3=y3+5		if true	3
v4=y4+5		if true	3
v5=y5+5		if true	4
v6=y6+5		if true	4

Figure 6: Extended code example after PSpec optimization has been applied. Statements (other than first statement) predicated on true have been speculated.

---

```

PSpec(normal_op)
{
  if (normal_op.guarding_predicate not defined by
      normal_op.earliest_schedulable_cycle)
  {
    if (multiple defs of normal_op.target exist
        {
          rename(normal_op.target);
          update_uses(normal_op.target);
        }
    normal_op.schedule(earliest_schedulable_cycle);
    normal_op.set_predicate(true);
  }
  else
  {
    normal_op.schedule(earliest_schedulable_cycle);
  }
}

```

---

Figure 7: Basic PSpec Algorithm.

### 5.1.1 Instruction Scheduling with Speculation

To demonstrate the usefulness of PSSA in enabling PSpec, Figure 6 shows the code from Figure 5 after the PSpec optimization has been applied. The assignments to `r1` and `r2` are examples of speculated operations. Notice that based on dependences, they could both be scheduled at cycle one which would have been impossible without renaming.

During predicated speculation, each operation is considered sequentially, beginning with the first instruction in the hyperblock. If it is a normal, non-store operation, PSpec compares its earliest schedulable cycle with the cycle in which its guarding predicate is currently defined. If the operation can be scheduled earlier than its guarding predicate, the operation is predicated on true and scheduled at its earliest schedulable cycle.

Recall that PSSA has not performed full renaming, so further renaming may be required by PSpec. An example is the definition of `y1` in Figure 5. If we speculate any of the definitions of `y1` by predicated them on true without renaming, incorrect code can result. Consequently, we must rename the operations being speculated. The results of applying this to the 6 definitions of `y1` (now `y1`, `y2`, `y3`, `y4`, `y5`, and `y6`) appear in Figure 6. Speculation and renaming may require the duplication of operations using the definition being speculated, since there may now be multiple reaching definitions. When speculating `y1`, the operation `v1=y1+5` had to be duplicated and guarded on the appropriate FPP (though in Figure 6 these statements are shown after they, too, have been speculated). This is possible because PSSA previously created all the necessary FPPs and path information.

If the guarding predicate has been defined by the operation's earliest schedulable cycle, we do not apply PSpec.

It is again scheduled at its earliest schedulable cycle, but guarded by the guarding predicate assigned by PSSA. The instruction  $z1=w1+5$  is an example. The algorithm for PSpec instruction scheduling is shown in Figure 7.

Using PSpec, the hyperblock can now be scheduled in 6 cycles as compared to 9 cycles in Figure 5. Since PSpec is applied whenever the definition of the operation's guarding predicate occurs later than the earliest schedulable cycle of the operation, we could reduce the number of operations that need to be speculated by moving the definition of the guarding predicates earlier. The goal of the next optimization, Control Height Reduction, is to allow predicates to be defined as early as possible.

### 5.1.2 Branches and Speculation

We chose not to PSpec branches. Therefore, a branch statement's earliest schedulable cycle is the one in which its guarding predicate is known. However, if a branch has been predicated on its block predicate by PSSA (because it does not have multiple operand versions reaching it) then it may be unnecessarily delayed in scheduling by waiting for that block predicate to be computed. As shown in Figure 6, we may choose to duplicate this statement, much as we do in normal PSpec, and guard the execution of these duplicates on their respective FPPs, instead of predicating the single instruction on its block predicate.

## 5.2 Control Height Reduction

Control Height Reduction (CHR) eases control constraints between multiple control statements. CHR allows successive control operations on the control path to be scheduled in the same cycle, effectively reducing control dependence height. For example, in the code in Figure 6, the control comparisons for  $z1>7$  and  $t1>r1$  are scheduled in cycles 3 and 4, respectively. However, the second comparison is only waiting for the definition of its guarding predicate AGF.

To schedule it earlier, consider the PSSA dependence graph in Figure 5. The definition of BAGF (defined by the condition  $t1>r1$ ), is control dependent on the definition of AGF (defined by the condition  $z1>7$ ). We could define BAGF directly as the logical AND of the conditions  $z1>7$  and  $t1>r1$  removing the dependence on the definition of AGF. This AND expression could be scheduled in cycle 3.

Control Height Reduction was proposed in [27]. It was successfully used to reduce the height of control recurrences found in loops when applied to superblocks. A *superblock* is a selected trace of basic blocks through the control flow graph containing only one path of control [26]. The path-defining aspects of PSSA allow our algorithm to effectively apply CHR to predicated hyperblocks, since the full-path predicates expose all of the original, separate paths throughout the hyperblock.

F=true		if true	1
w1=rand()		if F	1
z1=w1+5		if F	2
<i>AGF,A HF</i>	<i>cmpp.un.uc z1&gt;7</i>	<i>if F</i>	3
r1=5+x		if true	1
r2=x+8		if true	1
A=OR(AGF/AHF)		if true	4
t1=rand()		if true	1
<i>BAGF, CAGF</i>	<i>cmpp.an.an z1&gt;7</i>	<i>if true</i>	3
<i>BAGF, CAGF</i>	<i>cmpp.an.ac t1&gt;r1</i>	<i>if true</i>	3
<i>BAHF, CAHF</i>	<i>cmpp.ac.ac z1&gt;7</i>	<i>if true</i>	3
<i>BAHF, CAHF</i>	<i>cmpp.an.ac t1&gt;r2</i>	<i>if true</i>	3
LBAGF,EB AGF	cmpp.an.an z1>7	if true	3
LBAGF,EB AGF	cmpp.an.an t1>r1	if true	3
LBAGF,EB AGF	cmpp.an.ac w1>2	if true	3
LBAHF,EB AHF	cmpp.ac.ac z1>7	if true	3
LBAHF,EB AHF	cmpp.an.an t1>r2	if true	3
LBAHF,EB AHF	cmpp.an.ac w1>2	if true	3
ECAGF,EDC AGF	cmpp.an.an z1>7	if true	3
ECAGF,EDC AGF	cmpp.ac.ac t1>r1	if true	3
ECAGF, EDC AGF	cmpp.an.ac t1>7	if true	3
ECAHF,EDC AHF	cmpp.ac.ac z1>7	if true	3
ECAHF,EDC AHF	cmpp.ac.ac t1>r2	if true	3
ECAHF, EDC AHF	cmpp.an.ac t1>7	if true	3
D=OR(EDCAGF,EDCAHF)		if true	4
t2=t1-s		if true	2
L=OR(LBAHF,LBAGF)		if true	4
br out		if LBAGF	3
br out		if LBAHF	3
y1=t1+r1		if true	2
y2=t1+r2		if true	2
y3=t1+r1		if true	2
y4=t1+r2		if true	2
y5=t2+r1		if true	3
y6=t2+r2		if true	3
E=OR(EBAGF,E BAHF,ECA GF, ECAHF,EDCAGF,E DCAHF)		if true	4
v1=y1+5		if EBAGF	3
v1=y2+5		if EBAHF	3
v1=y3+5		if ECAGF	3
v1=y4+5		if ECAHF	3
v1=y5+5		if EDCAGF	4
v1=y6+5		if EDCAHF	4

Figure 8: Extended example after PSpec and CHR optimizations have been applied. Cmpp instructions displayed in italics define predicates that are not used after optimization. Therefore, the statements can be removed from the final code.

Schlansker et. al. [28] recently expanded on their previous research, applying speculation prior to attempting height reduction. Speculation is needed to remove dependences between the branch conditions that need to be combined to accomplish the reduction. However, in that work, speculation was limited to operations that would not overwrite a live register or memory value if speculated, since they did not use renaming. In Figure 5, the `cmpp` operation defining `BAGF` and `CAGF` is shown scheduled at cycle 5 due to dependences on `t1` and `r1`. PSSA allows us to apply PSpec and schedule these definitions in cycle 1, making the `cmpp` available for CHR as shown in Figure 8.

### 5.2.1 Instruction Scheduling with PSpec and CHR

During instruction scheduling, PSpec is performed as described in Section 5.1.1. During the same sequential pass through instructions, for each *control* operation (`cmpp`), CHR is performed if possible.

Recall that the operations in Figure 5 are scheduled in the order given in the PSSA hyperblock. Like PSpec, CHR compares an operation’s earliest schedulable cycle with when it must be scheduled if it waited for its guarding predicate to be defined. If it does not need to wait on the definition of its guarding predicate, it is simply scheduled at its earliest schedulable cycle. Without PSpec, the definition of `BAGF` was waiting on the definition of `t1` and `r1`. With PSpec, it is only waiting on the definition of its guarding predicate. Therefore, it is beneficial to control height reduce.

By ANDing the condition of the current definition with the condition that defined its guarding predicate, we can schedule this definition earlier. If the definition of the guarding predicate involved conditions that were ANDed as well, all of the conditions must be included, so the number of `cmpp` statements needed to define the current operation increases. The `.a` tag on each of these `cmpp` statements indicates that all of them are required for the final definition.

Consider the operations `z1>7`, `t1>r1` and `t1>7` in Figure 5. We control height reduce these operations in Figure 8, since they are all schedulable in cycle 3 based on our scheduling constraints. The definition of `ECAGF` now describes the combination of `z1>7` being true AND `t1>r1` having a value of false AND `t1>7` having a value of true. We implement this logical AND using the `.ac` and `.an` qualifiers. The definition of `ECAGF` requires that the conditions `z1>7` and `t1>7` and the condition `!(t1>r1)` evaluate to true for the FPP to get a value of true. If any one of the requirements are not met, the FPP will be set to false. The compares can be performed in the same cycle [19], allowing multiple links in a control path to be defined simultaneously. The algorithm for CHR is found in Figure 9.

Using PSpec and CHR on PSSA-transformed code results in the 4 cycle schedule shown in Figure 8. Note that the operations shown in italics can be removed in a post-pass because these operations define predicates that are never used. Using predicated speculation and control height reduction together on PSSA-transformed code allows every operation to be scheduled at its earliest schedulable cycle.

---

```

CHR(cmpp_op)
{
  if (cmpp_op.guarding_pred defined
      by cmpp_op.earliest_schedulable_cycle)
  {
    cmpp_op.schedule(cmpp_op.earliest_schedulable_cycle)
  }
  /* Apply Control Height Reduction */
  else
  {
    while (more_stmts_defining(cmpp_op.guarding_pred))
    {
      next_def=next_defining_stmt(cmpp_op.guarding_pred)
      copy=duplicate(next_def)
      copy.schedule(next_def.get_scheduling_time())
      copy.predicate_on(next_def.get_guarding_pred())
      copy.set_define(cmpp_op.get_pred_defined())
      copy.set_tag_to(a)
    }
    cmpp_op.schedule(next_def.get_scheduling_time())
    cmpp_op.predicate_on(next_def.get_guarding_pred())
    cmpp_op.set_tag_to(a)
  }
}

```

---

Figure 9: Basic Control Height Reduction Algorithm.

## 6 Results

We have implemented algorithms to perform PSSA, CHR and PSpec on hyperblocks in the Trimaran System (Version 2.00). We collect profile-based execution weights for operations in the codes and schedule operations with an assumed one-cycle latency in order to calculate execution time. Additionally, we conservatively assume that a load is dependent on all prior stores along a given path, and that a store is dependent on prior stores as well. We also ensure that all instructions along a path leading to a branch out of the hyperblock are executed prior to exiting the hyperblock.

Figure 10 shows normalized execution time when applying our optimizations for several Trimaran benchmarks: `fib`, `mm`, `wc`, `fir`, `wave`, `nbradar` (a Trimaran media benchmark), `qsort`, `alvinn` (from SPEC FP92), `compress` (from SPEC INT95), and `li` (from SPEC INT95). These codes are described in the Trimaran Benchmark Certification [2]. The original execution times are created from the default Trimaran settings, with the exception that the architecture issue rate is set to 16. Execution time is estimated by summing together the frequency of execution of each hyperblock multiplied by the number of cycles it takes to execute the hyperblock assuming a perfect memory system. Infinite results do not restrict the number of operations issued per cycle.

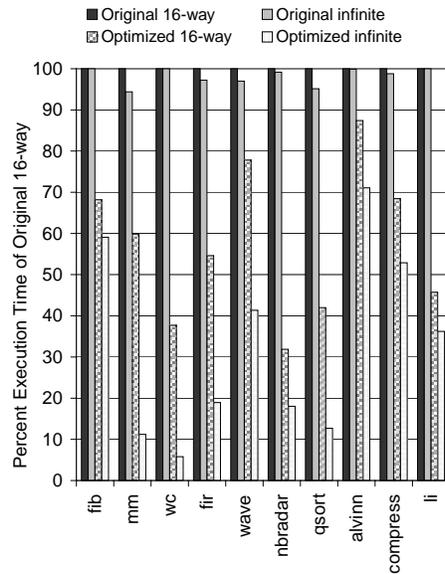


Figure 10: Executed cycles normalized to the number of cycles to execute the original code produced by Trimaran for a 16 issue machine.

16-way results are obtained by dividing each cycle which has been scheduled with more than 16 operations into  $\text{ceiling}(\text{total\_operations\_scheduled\_in\_cycle} / 16)$  cycles. The results are normalized to the original schedule generated by Trimaran for a 16-issue machine and scheduled 16-way. The optimized results show the performance after applying PSSA, PSpec, and CHR. The results show that using PSSA with PSpec and CHR results in a significant reduction in executed cycles.

Figure 11 shows the average number of operations executed per cycle for the configurations examined in Figure 10. In comparing the two graphs for the 16-way results, 3-4 times as many instructions are issued per cycle after applying PSSA, PSpec, and CHR, and this resulted in a reduction in execution time ranging from 12% to 68%. Since PSpec and CHR as applied to PSSA code have the effect of removing the restrictions of control dependence, the optimized infinite results provide a picture of "best case" instruction level parallelism. Inspection of the optimized infinite results of `alvinn`, `compress`, and `li` show that, given current hyperblock formation, peak IPC is somewhat limited.

The renaming required by PSSA and PSpec also significantly increases register pressure. Trimaran's ISA (Playdoh) supports 4 register files: general purpose, floating point, branch, and predicate [2, 19]. Figure 12 shows the average number of live registers for the original code and the optimized code using PSSA, PSpec and CHR. The average live register results are weighted by the frequency of hyperblock execution. For example, `matrix multiply` has on average 17 live general purpose registers in the original code, and 54 live general purpose registers after optimization. Though the increase in utilization of all these register files is notable, the weighted average utilization mostly still remains within the reported IA-64 register file sizes (128 general purpose, 128 floating point, 8 branch,

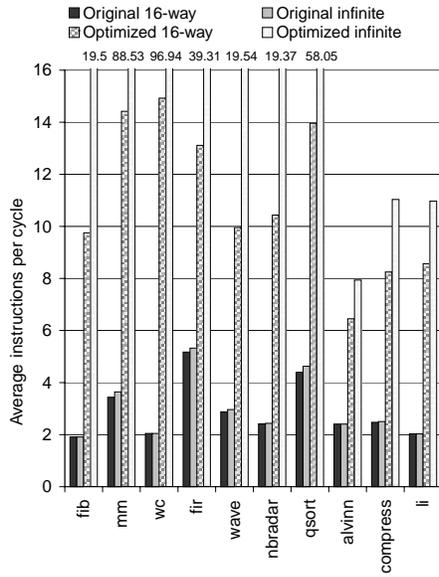


Figure 11: Weighted average number of operations scheduled per cycle for hyperblocks when using PSSA with Predicated Speculation and Control Height Reduction. Note that several of the "Optimized infinite" results are greater than 16 – the issue width simulated in these experiments.

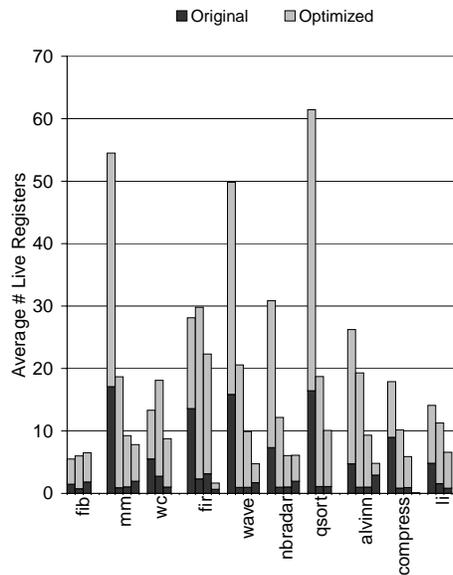


Figure 12: Weighted average register pressure in hyperblocks when using PSSA with Predicated Speculation and Control Height Reduction. Shown from left to right for each benchmark is the general purpose file, predicate file, branch file, and floating point file (zero utilization for some benchmarks).

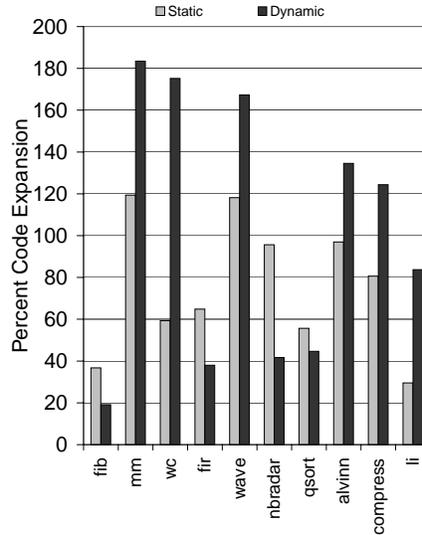


Figure 13: Static and Dynamic Code Expansion normalized to original code size. Dynamic code expansion indicates an increase in the working set size to be supported by the instruction cache.

and 64 predicate) [3].

Additionally, PSSA combined with aggressive PSpec and CHR significantly increases code size – both static and dynamic. Aggressive and resource insensitive application of CHR and PSpec aims to reduce cycles required to schedule at the cost of duplicated code specialized for particular paths (in the case of PSpec) or duplicated code for faster computation of predicates (in the case of CHR). Figure 13 shows both the static and dynamic code expansion of the PSSA, PSpec, and CHR optimized code over the original code. We calculate static code expansion by comparing the number of static operations in the optimized code with the number of static operations in the original code. Dynamic code expansion is measured similarly, with the exception that each static operation is weighted by the number of times that it is executed (as calculated by Trimaran’s profile-based region weights). This “dynamic code expansion” is intended to capture the run-time effect that the introduced duplicated code will have on the memory system. Dynamic code expansion indicates an increase in the working set size to be supported by the instruction cache.

## 7 Related Work

Predicated execution presents challenges and prospects that researchers have addressed in a variety of ways. Mahlke et. al. [21] showed that predicated execution can be used to remove an average of 27% of the executed branches and 56% of the branch mispredictions. Tyson also found similar results and correlated the relationship between

predication and branch prediction [29].

In an effort to relieve some of the difficulties related to applying compiler techniques to predicated code, Mahlke et. al. [22] defined the hyperblock as a single-entry, multiple-exit structure to help support effective predicated compilation. These hyperblocks are formed via selective If-conversion [5, 24] – a technique that replaces branches with predicate define instructions. The success of predicated execution can depend greatly on the region of the code selected to be included in the predicated hyperblock. August et. al. [9] relates the pitfalls and potentials of hyperblock formation heuristics that can be used to guide the inclusion or exclusion of paths in a hyperblock. Warter et. al. [30] explore the use of Reverse If-conversion for exposing scheduling opportunities in architectures lacking support for predicated execution as well as for re-forming hyperblocks to increase efficiency for predicated code [9, 30].

The challenges of doing data-flow and control-flow analysis on hyperblocks have also been addressed. Since hyperblocks include multiple paths of control in one block, traditional compiler techniques are often too conservative or inefficient when applied to them. Methods of predicate-sensitive analysis have been devised to make traditional optimization techniques more effective for predicated code [15, 18]. The work presented in [11] (and expanded upon in this work) extended the localized predicate-sensitive analysis presented in [15, 18] to complete path analysis through the hyperblock. Path-sensitive analysis has previously been found useful for traditional data-flow analysis [6, 10, 16]. We use this specialized path information to accomplish PSSA (a predicate-sensitive form of SSA [13, 12]) which enables Predicated Speculation and Control Height Reduction for hyperblocks that have previously been examined only in the presence of the single path of control found in superblocks [26, 27, 28].

Moon and Ebcioğlu [23] have implemented selective scheduling algorithms, which can schedule operations at their earliest possible cycle for non-predicated code. Our work extends theirs for predicated code, by allowing earliest possible cycle scheduling using predicated renaming with full-path predicates.

## 8 Implementing PSSA in IA-64

Implementing PSSA using the IA-64 ISA [3] would be straightforward with the exception of the predicate OR statement we introduced. We found this OR statement to be very useful in efficiently combining path information in order to eliminate unnecessary code expansion. If this instruction were not explicitly added to IA-64 then it could be implemented by transferring the predicate register file into a general register using the move from predicate instruction in IA-64. The general purpose masking instruction would then be used to mask all but the bits corresponding to the sources of the predicate OR instruction. A result of zero evaluates to false, and anything else evaluates to true.

IA-64, unlike the Playdoh ISA, places limits on compare instructions. For example, conditions that are included

in logical AND compare statements can only compare a variable to zero. Specifically, the statement `LBAGF, EBAGF cmp .and .and t1>r1 if true` in Figure 8 would not be permitted. In implementing CHR, we would have to transform the prior expression into the following 2 statements (expressed in IA-64 notation) [3] :

```
(PR[0]) cmp.unc.gt PR[TEMP_1],PR[TEMP_2] = GR[t1],GR[r1]
(PR[TEMP_2]) cmp.and.gt PR[LBAGF],PR[EBAGF] = GR[0],GR[0]
```

## 9 Future Work

When constructing a hyperblock schedule for a specific processor implementation, resource limits will mandate how many operations can be performed in each cycle. Architectural characteristics such as issue width, resource utilization, number of available predicate registers, and number of available rename registers all need to be considered when creating an architecture-specific schedule. The goal of a hyperblock scheduler is to reduce the execution-height while taking these architectural features into consideration.

In this paper, our goal was to show that PSSA provided an efficient form of renaming and precise path information to allow all operations to be scheduled at their earliest schedulable cycle. We are currently examining different PSSA representations to reduce code duplication and the number of full-path predicates created. Since various control paths through a hyperblock may have different true data dependence heights, it may provide no advantage to speculate operations that are not on the critical path through the hyperblock. PSSA could concentrate on only the critical paths through the hyperblock, reducing code duplication. For non-critical paths, it may be advantageous in PSSA to implement  $\phi$ -functions combining different variable names, instead of maintaining renamed variables for each full-path in the hyperblock. At a point in the hyperblock where all paths join, copy operations could be used to return renamed definitions to original names. Path definitions could then be restarted at this point. This would reduce the amount of duplication required for a given operation to use correctly renamed variables. Our future research concentrates on these issues and creating a more efficient implementation of PSSA.

## 10 Conclusions

This paper extended [11], where Predicated Static Single Assignment was first introduced. It motivated the need for renaming and for predicate analysis that extends across all paths of the hyperblock. It demonstrated how Predicated Static Single Assignment (PSSA), a predicate-sensitive implementation of SSA that implements renaming using full-path predicates, can be used to eliminate false dependences for predicated code. We showed the benefit of using

PSSA to enable Predicated Speculation (PSpec) and Control Height Reduction (CHR) during scheduling. Predicated Speculation allows operations to be executed at their earliest schedulable cycle, even before their guarding predicates are determined. Control Height Reduction allows guarding predicates to be defined as soon as possible, reducing the amount of speculation needed.

By maintaining information about each of the original control paths in a hyperblock, PSSA can provide information that allows precise placement of renamed and speculated code, and allows the correct, renamed values to be propagated to subsequent operations. The renaming used by PSSA allows more aggressive speculation, as overwriting live values is no longer a concern. In addition, PSSA supports Control Height Reduction along every control path using full-path predicates, reducing control dependence depth throughout the hyperblock.

Our experiments show that PSSA is an effective tool for optimizing predicated code. We gave extended experiments that show using PSSA with PSpec and CHR results in a reduction in executed cycles ranging from 12% to 68% for a 16 issue machine.

## Acknowledgments

We would like to thank the Compiler and Architecture Research Group at Hewlett Packard, University of Illinois' IMPACT Group, and New York University's ReaCT-ILP Group for providing Trimaran. We specifically appreciate the time and patience of Rodric Rabbah, Scott Mahlke, Vinod Kathail, and Richard Johnson in answering many questions regarding the Trimaran system. In addition, we would like to thank Scott Mahlke for providing useful comments on this paper. This work was supported in part by NSF CAREER grant No. CCR-9733278, a National Defense Science and Engineering Graduate Fellowship, a research grant from Intel Corporation, and equipment support from Hewlett Packard and Intel Corporation.

## References

- [1] Intel Press Release. Merced processor and IA-64 architecture., 1998. <http://developer.intel.com/design/processor/future/iaa64.htm>, 1998.
- [2] Trimaran, An Infrastructure for Research in Instruction Level Parallelism, 1998. <http://www.trimaran.org>.
- [3] IA-64 Application Developer's Architecture Guide, Revision 1.0, 1999.
- [4] A. V. Aho, R. Sethi, and J. D. Ullman. *Compilers : Principles, Techniques and Tools*. Addison-Wesley, 1986.
- [5] J.R. Allen, K. Kennedy, C. Porterfield, and J. Warren. Conversion of control dependence to data dependence. In *Proceedings of the 10th ACM Symposium on Principles of Programming Languages*, December 1994.

- [6] G. Ammons and J. R. Larus. Improving data-flow analysis with path profiles. *ACM SIGPLAN Notices*, 33(5):72–84, May 1998.
- [7] D. I. August, D. A. Connors, S. A. Mahlke, J. W. Sias, K. M. Crozier, B. Cheng, P. R. Eaton, Q. B. Olaniran, and W. W. Hwu. Integrated predicated and speculative execution in the IMPACT EPIC architecture. In *Proceedings of the 25th Intl. Symp. on Computer Architecture*, July 1998.
- [8] D. I. August, K. M. Crozier, J. W. Sias, P. R. Eaton, Q. B. Olaniran, D. A. Connors, and W. W. Hwu. The IMPACT EPIC 1.0 Architecture and Instruction Set reference manual. Technical Report IMPACT-98-04, IMPACT, University of Illinois, Feb 1998.
- [9] D. I. August, W. Hwu, and S. A. Mahlke. A framework for balancing control flow and predication. In *30th Annual Intl. Symp. on Microarchitecture*, December 1997.
- [10] T. Ball and J. R. Larus. Efficient path profiling. In *Proceedings of the 29th Annual International Symposium on Microarchitecture*, pages 46–57, December 2–4, 1996.
- [11] L. Carter, B. Simon, B. Calder, L. Carter, and J. Ferrante. Predicated static single assignment. In *Proceedings of International Conference on Parallel Architectures and Compilation Techniques*, October 1999.
- [12] R. Cytron, J. Ferrante, B. K. Rosen, M. K. Wegman, and F. K. Zadeck. An efficient method of computing static single assignment form. In *16th Annual ACM Symposium on Principles of Programming Languages*, pages 25–35, 1989.
- [13] R. Cytron, J. Ferrante, B. K. Rosen, M. N. Wegman, and F. K. Zadeck. Efficiently computing static single assignment form and the control dependence graph. *ACM Transactions on Programming Languages and Systems*, 13(4):451–490, October 1991.
- [14] James C. Dehnert and Ross A. Towle. Compiling for the Cydra 5. *The Journal of Supercomputing*, 7(1-2):181–227, May 1993.
- [15] D. M. Gillies, D. R. Ju, R. Johnson, and M. Schlansker. Global predicate analysis and its application to register allocation. In *Proceedings of the 29th Annual Intl. Symp. on Microarchitecture*, pages 114–125, December 1996.
- [16] R. Gupta, D. A. Berson, and J. Z. Fang. Path profile guided partial dead code elimination using predication. In *Proceedings of the 1997 International Conference on Parallel Architectures and Compilation Techniques*, pages 102–113, November 10–14, 1997.
- [17] L. Gwennap. Intel, HP make EPIC disclosure. *Microprocessor Report*, 11(14):1–9, October 1997.
- [18] R. Johnson and M. Schlansker. Analysis techniques for predicated code. In *Proceedings of the 29th Annual Intl. Symp. on Microarchitecture*, pages 100–113, December 1996.
- [19] V. Kathail, M. S. Schlansker, and B. R. Rau. HPL PlayDoh architecture specification: Version 1.0. Technical Report HPL-93-80, HP Labs, Feb 1994.

- [20] P. G. Lowney, S. M. Freudenberger, T. J. Karzes, W. D. Lichtenstein, R. P. Nix, J. S. O'Donnell, and J. C. Ruttenberg. The Multiflow Trace Scheduling compiler. *The Journal of Supercomputing*, 7(1-2):51–142, May 1993.
- [21] S. A. Mahlke, R. E. Hank, R. A. Bringmann, J. C. Gyllenhaal, D. M. Gallagher, and W. W. Hwu. Characterizing the impact of predicated execution on branch prediction. In *Proceedings of the 27th Annual Intl. Symp. on Microarchitecture*, pages 217–227, December 1994.
- [22] S. A. Mahlke, D. C. Lin, W. Y. Chen, R. E. Hank, and R. A. Bringmann. Effective compiler support for predicated execution using the hyperblock. In *Proceedings of the 25th Annual Intl. Symp. on Microarchitecture*, pages 45–54, December 1992.
- [23] S. Moon and K. Ebcioglu. Parallelizing nonnumerical code with selective scheduling and software pipelining. *ACM Transactions on Programming Languages and Systems*, 19(6):853–898, November 1997.
- [24] J. C. H. Park and M. Schlansker. On Predicated Execution. Technical Report HPL-91-58, HP Labs, May 1991.
- [25] B. Ramakrishna Rau, David W. L. Yen, Wei Yen, and Ross A. Towle. The Cydra 5 departmental supercomputer. *Computer*, 22(1):12–35, January 1989.
- [26] M. Schlansker and V. Kathail. Critical path reduction for scalar programs. In *Proceedings of the 28th Annual International Symposium on Microarchitecture*, pages 57–69, November 29–December 1, 1995.
- [27] M. Schlansker, V. Kathail, and S. Anik. Height reduction of control recurrences for ILP processors. In *Proceedings of the 27th Annual Intl. Symp. on Microarchitecture*, pages 40–51, December 1994.
- [28] M. Schlansker, S. Mahlke, and R. Johnson. Control CPR: A branch height reduction optimization for EPIC architectures. In *Proceedings of the 1999 ACM SIGPLAN Conference on Programming Language Design and Implementation*, May 1999.
- [29] G. S. Tyson. The effects of predicated execution on branch prediction. In *Proceedings of the 27th Annual International Symposium on Microarchitecture*, pages 196–206, November 30–December 2, 1994.
- [30] N. J. Warter, S. A. Mahlke, W. W. Hwu, and B. R. Rau. Reverse if-conversion. In *Proceedings of the SIGPLAN '93 Conference on Programming Language Design and Implementation*, pages 290–299, June 1993.
- [31] M. Wolfe. *High Performance Compilers for Parallel Computing*. Addison-Wesley, Redwood City, CA, 1996.