

## Branch Target Buffer Design\*

*Chris H. Perleberg*

Department of Electrical Engineering and Computer Science  
University of California, Berkeley  
Berkeley, CA 94720

The Branch Target Buffer (BTB) can reduce the performance penalty of branches in pipelined processors by predicting the path of the branch and caching information used by the branch. This paper discusses two major issues in the design of BTBs with the theme of achieving maximum performance with a limited number of bits allocated to the BTB design. First is the issue of BTB management - when to enter and discard branches from the BTB. Higher performance can be obtained by entering branches into the BTB only when they experience a branch taken execution. A new method for discarding branches from the BTB is examined. This method discards the branch with the smallest expected value for improving performance, outperforming the LRU strategy by a small margin, at the cost of additional complexity.

The second major issue discussed is the question of what information to store in the BTB. A BTB entry can consist of one or more of the following: branch tag (i.e. the branch address), prediction information, the branch target address, and instructions at the branch target. A variety of BTB designs, with one or more of these fields, are evaluated and compared. This study is then extended to multilevel BTBs, in which different levels have different amounts of information per entry. For the specific implementation assumptions used, multilevel BTBs improved performance over single level BTBs only slightly, at the cost of additional complexity. Multi-level BTBs may provide significant performance improvements for other implementations, however.

Design target miss ratios for BTBs are also developed, so that the performance of BTBs for real workloads may be estimated.

December 23, 1989

---

\* The material presented here is based on research supported in part by the National Science Foundation under grant MIP-8713274, by NASA under consortium agreement NCA2-128, by the State of California under the MICRO program, and by the Digital Equipment Corporation, International Business Machines Corporation, Philips Laboratories/Signetics and Apple Computer Corporation.

# Branch Target Buffer Design\*

*Chris H. Perleberg*

Department of Electrical Engineering and Computer Science  
University of California, Berkeley  
Berkeley, CA 94720

## Table of Contents

1.0 Introduction .....	2
2.0 Solutions to the Branch Problem .....	2
3.0 Methodology .....	5
3.1 Trace Driven Simulation .....	5
3.2 Description of Traces .....	5
4.0 BTB Design Optimization .....	8
5.0 Branch Prediction .....	9
5.1 Best Case Prediction Rates .....	10
5.2 Effect of Context Switching on Prediction Rates .....	12
6.0 BTB Management .....	14
6.1 Entering Branches into the BTB .....	14
6.2 Discarding Branches from the BTB .....	15
7.0 Relationship between BTB and Instruction Cache Miss Rates .....	19
7.1 BTB Design Target Miss Ratios .....	21
8.0 Information Stored in the BTB .....	22
8.1 Tradeoffs Between Size and Number of Prediction Bits .....	26
8.2 Multi-Level BTBs .....	26
8.2.1 Results of Simulating Multi-Level BTBs .....	28
9.0 Conclusions .....	30
10.0 Appendix .....	32
11.0 Bibliography .....	57

---

\* The material presented here is based on research supported in part by the National Science Foundation under grant MIP-8713274, by NASA under consortium agreement NCA2-128, by the State of California under the MICRO program, and by the Digital Equipment Corporation, International Business Machines Corporation, Philips Laboratories/Signetics and Apple Computer Corporation.

## 1. Introduction

Most modern computers use pipelining to significantly increase performance. Pipelining divides the execution of each instruction into several pieces, normally called pipeline stages. A typical pipeline might have five pipeline stages, including instruction fetch, instruction decode, operand fetch, execution, and result writeback. It is important that the operation time of each pipeline stage is almost identical, since the rate at which instructions flow through the pipeline is limited by the slowest pipeline stage. In the optimal case, five instructions could be operated on simultaneously in this pipeline, each in a separate pipeline stage. This optimal case, though seldom achieved in practice, would produce a five fold performance gain over the same processor designed without a pipeline. More details about pipelining can be found in [Ramamoorthy76] and [Kogge81].

Branch instructions can reduce the performance of pipelines by interrupting the steady flow of instructions into the pipeline. To execute a branch, the processor must decide whether the branch is taken, calculate the branch's target address, and then, if the branch is taken, fetch instructions from the target address. Branches hurt performance because the pipelined processor needs to know the path of the branch (in order to fetch more instructions) before the path of the branch has been determined. Thus, when a branch occurs, the processor has two choices, either wait for the branch to finish executing before fetching more instructions, or continue fetching instructions, possibly from the wrong path. Either choice hurts performance, although the second choice hurts less, especially if the path of the branch is predicted correctly most of the time.

The Branch Target Buffer (BTB) can be used to reduce the performance penalty of branches by predicting the path of the branch and caching information about the branch [Lee84]. This paper discusses two major issues in the design of BTBs, with the theme of achieving maximum performance with a limited number of bits allocated to the BTB design. First, the question of BTB management, when to enter branches into the BTB, and how to discard branches from the BTB. Second, the question of what information to keep in the BTB. Four types of information have been kept in BTBs, including branch tags, the branch target address, prediction information, and instruction bytes from the branch target address. A BTB contains a subset of these four types of information, and a multi-level BTB contains several distinct subsets, each in a separate level of the BTB. Higher performance BTBs generally contain a larger subset of the possible information types than lower performance BTBs. As a sidenote, a comparison of BTB and instruction cache miss rates is performed, with the purpose of showing that BTB miss rates can be derived from the instruction cache miss rates that exist in literature [SmithAJ87].

This paper is organized as follows. In the following section, a number of solutions to the branch problem are surveyed. The methodology used to generate the results in this paper (trace driven simulation) is then presented, followed by a discussion of optimizations in the design of BTBs. Branch prediction is covered, and the issue of BTB management. A discussion on miss rates of BTBs and instruction caches is followed by a consideration of the issue of what information to keep in a BTB. Conclusions are then provided.

## 2. Solutions to the Branch Problem

This survey of solutions to the branch problem is intended to give a brief explanation of each solution, and provide the names of a few machines that have used each solution. Further discussion is available in [Lee84], [Lilja88], and [McFarling86].

**Static Branch Prediction:** As discussed in the introduction, if the path of a branch can be correctly predicted, the branch penalty can be reduced. Predicting all branches not taken or predicting all branches taken (Lee and Smith found 68 percent of branches to be taken [Lee84]; we found 70 percent of branches to be taken) is the simplest type of static prediction. Predicting

branches based on the direction of the branch (forward or backward) and/or the branch opcode is also possible [SmithJ81]. The National Semiconductor NS32532 uses the branch opcode and the branch direction to predict 71 percent of branches correctly [Maytal89]. A static prediction bit in each branch instruction, used to indicate which path the branch will probably take, is also effective. The bit can be set using heuristics and execution profiling [McFarling86]. The Ridge Computers Ridge-32 [Folger83], the Intel 80960 [Hinton89], and the AT&T CRISP [Ditzel87] all use a static prediction bit. An average prediction rate of 85 percent has been reported for the static prediction bit method [McFarling86].

**Dynamic Branch Prediction:** Dynamic branch prediction predicts branch paths based on the runtime execution history (the sequence of taken and not taken executions, represented with one bit per execution) of each branch. One proposed dynamic method uses a two bit counter state machine which takes the branch execution history as input and outputs the predicted branch path for the next branch execution [SmithJ81]. Another method examines the last  $n$  executions (taken or not taken) of a branch, and using statistics gathered beforehand to calculate the chance of a taken branch, predicts the branch path [Lee84] [Widdoes77]. Widdoes proposes keeping bits with each instruction in the instruction cache for storing branch execution history [Widdoes77]. Machines that use dynamic branch prediction include the S-1 [Hailpern79], the MU5 [Rau77], several TRON microprocessors [Sakamura87], the Edgecore Edge 2000 (Motorola 680X0 compatible) [Manual87] [Walter89], and the NexGen processor (Intel 80386 compatible) [Stiles89a].

**Delayed Branch:** Executing  $n$  instructions after a branch whether the branch is taken or not often allows useful instructions to be executed while the branch path is resolved and branch target address calculated. The  $n$  instructions (usually  $n=1$ ) are called delay instructions. For more flexibility, some machines allow the execution of the delay instructions to be skipped, depending on the value of a "squash" bit in the branch and the path of the branch, taken or not taken [McFarling86]. The delayed branch technique is used in most RISC processors, and requires intelligent compilers for optimal use [DeRosa87] [Gross82] [McFarling86] [Patterson81].

**Loop Buffers:** Loop buffers are small, high speed buffers used for instruction prefetching and fast execution of instruction loops. If a loop is completely contained within a loop buffer, the loop can be executed from the high speed buffer, instead of accessing slower main memory for the instructions. Multiple loop buffers allow multiple non-contiguous loops to execute from the buffers. Machines using loop buffers include the CDC-6600 with eight 60-bit words [Thornton64], the CDC-7600 with twelve 60-bit words [Elrod70], the IBM 360/91 and the IBM 360/195 with eight 64-bit words [Anderson67] [Murphey70], the Cray-1 with four 128-byte buffers [Russell78], and the Cray-XMP with four 256 byte buffers [Hwang84].

**Multiple Instruction Streams:** To avoid the penalty of a wrong branch prediction, both paths of a branch can be fetched and executed, and the results of the incorrect path discarded once the branch is resolved. This method requires duplication of hardware, early calculation of the branch target address, high memory and register file bandwidth, and a method for handling more than one branch in the pipeline. This technique has been used in mainframes, including the IBM 370/168 and the IBM 3033 [IBM73] [IBM78].

**Prefetch Branch Target:** Many computers predict all branches not taken and pay a penalty if the branch is taken. This penalty may be reduced if instructions from the branch target address are fetched while the branch (and the instructions in the not taken branch path) is executing. The IBM 360/91 prefetches two double-words (each 64 bits long) from the target address [Anderson67].

**Prepare to Branch:** The "Prepare-to-Branch" and "Load-Look-Ahead" instructions of the Texas Instruments ASC computer direct the machine to prefetch from the target address of an upcoming branch, rather than from the sequential addresses after the branch [Gaulding75]. The CHoPP supercomputer has a similar instruction [Kartashev90]. Use of these instructions with

branches that are usually taken can increase performance.

**Shared Pipeline Multiprocessor:** An  $n$  stage pipelined processor executing  $n$  independent instruction streams simultaneously, each with a single instruction occupying one of the pipeline stages, eliminates the performance penalty of branches. Within an individual instruction stream, each instruction completes execution before the next instruction begins execution, so branch instructions do not have to wait during execution for prior instructions to finish executing. Each independent instruction stream receives a  $1/n$  fraction of total processor power. The shared pipeline multiprocessor concept was implemented in the Denelcor HEP computer [Jordan83], and is examined in [Shar74].

**Branch Folding:** The AT&T CRISP processor uses a technique called branch folding to reduce branch execution time to zero [Ditzel87]. CRISP uses a decoded instruction cache; each non-branch instruction is transformed into a line of microcode in a cache of decoded instructions. Each entry includes a next address field pointing to the next line of microcode to be executed. Unconditional branches are simply "folded" into the next address field of the previous line of microcode. No branch execution time is necessary once the unconditional branch is in the decoded cache. Conditional branches make use of a second next address field in the microcode and static branch prediction. A bit in the branch instruction predicts whether the branch will be taken or not, and the predicted next address field is selected to find the next line of microcode to execute. If the predicted path is found to be wrong once the condition flag is resolved, the incorrect path results are discarded, and the correct path is executed. Like the unconditional branch, the conditional branch has been folded into the previous line of microcode, resulting in a zero cycle execution time for a correctly predicted branch in the decoded cache. Note that branch folding does not require a decoded instruction cache. Branch folding can be implemented with a branch target buffer that caches instructions from both the taken and not taken branch paths of each branch.

**Branch Target Buffer:** The Branch Target Buffer (BTB) can be used to reduce the performance penalty of branches by predicting the path of the branch and caching information about the branch [Lee84]. Up to four types of information can be cached in a BTB: a tag identifying the branch the information corresponds to (usually the branch address), prediction information for predicting the path of the branch, the branch target address, and instruction bytes from the branch target address. A BTB with all four types of information works as follows. As each instruction is fetched from memory, the instruction address is used to index into the BTB. If a valid BTB entry is found for that address, the instruction is a branch. The branch path is predicted using the branch's prediction information. If the branch is predicted taken, the pipeline is supplied with the cached instructions in the BTB, and the processor begins fetching instructions from the target address, offset by the number of bytes of instructions cached in the BTB. If the branch is predicted not taken, the processor continues fetching sequentially after the branch. After the processor finishes executing the branch, it checks to see if the BTB correctly predicted the branch. If it has, all is well, and the processor can continue sequentially. If the branch was predicted incorrectly (or the branch was taken and the target address changed), the processor must flush the pipeline and begin fetching from the correct branch path. In either case, the branch prediction information and branch target address (if changed) must be updated after the branch. Machines that use a variation of the BTB include: the Advanced Micro Devices Am29000 [AMD88] and the General Electric RPM40 [Lewis88] with branch target caches (no prediction information, just a cache of target instructions); the Mitsubishi M32 with a BTB containing dynamic prediction information, branch tag, and instructions from the target address (no mention of target address being stored) [Yoshida87]; the Edgecore Edge 2000 with a direct map 1024 entry BTB containing one dynamic prediction bit, a branch tag, and the branch target address for each entry [Walter89]; and the NexGen processor with a fully associative LRU replacement BTB containing prediction information, branch tag, branch target address, and instructions from the target address

[Stiles89a] [Stiles89b].

### 3. Methodology

#### 3.1. Trace Driven Simulation

To date, there exists no generally accepted model of program behavior with respect to the performance of BTBs. As a result, we use trace driven simulation to generate the data presented. Trace driven simulation uses a recording (known as a trace) of the execution of a program as input to a software design simulator. Each trace contains a record of the data and instruction addresses issued by the program and the actual instruction bytes fetched by the program. Using a number of these traces as input to a software simulator can give accurate performance numbers, within the accuracy of the simulator itself, and allows a flexibility in varying design parameters that is very difficult to achieve in a hardware prototype. There are, however, reasons to be cautious when dealing with the results of trace driven simulation [SmithAJ85]. Traces tend to be short recordings (less than 1 CPU second) and may not be representative of the expected workload. They often do not include the effects of context switching, operating system code execution, and input/output. In addition, traces are often dependent on the implementation of the computer architecture the trace is recorded on (e.g. number of bytes fetched per instruction fetch), and may not match the assumptions of the simulated design. We have used a wide variety of traces of both small and large programs from a number of different machines and believe that our results are representative of what could be expected in practice.

#### 3.2. Description of Traces

A total of 32 program traces are used in this paper. Four processors are represented in the traces, including the MIPS Co. R2000 (MIPS), the Sun Microsystems SPARC (SPARC), the DEC VAX 11/780 (VAX), and the Motorola 68010 (68k) [DEC77] [Fujitsu87] [Kane89] [Motorola89]. The traces are divided into six workloads, so the individual characteristics of the workloads can be observed in the results [Lee84]. The workloads include COMP (MIPS compiler related traces), FP (MIPS floating point traces), TEXT (MIPS text processing traces), SPARC (SPARC traces), 68k (68010 traces), and VAX (VAX traces). The large number of traces for the MIPS reflects the fast, flexible, and bug free tracer "pixie" that MIPS Co. provides for its computers. The small number of traces for the SPARC is due to difficulty with the SPARC tracer. The 68k traces are described in [Grochowski86]. The VAX traces are described in [Henry83]. In general, the traces, in order of significance (and confidence in the results they generate), are the MIPS, the SPARC, the 68k, and the VAX.

In some of our discussions, we refer to "static" and "dynamic" instructions. The instructions in the execution trace are dynamic instructions. The instructions in the executed program that appear in the trace are static instructions. Thus, with a two instruction loop that is executed 10 times, there are two static instructions, and 20 dynamic instructions.

The BTB simulators used in this paper use a reduced trace as input, which is composed of all control transfer (branch) instructions (which average 19 percent of the "static" instructions), both conditional and unconditional [DEC77] [Fujitsu87] [Kane89] [Motorola89]. As a result, all of the data presented is for BTBs that handle both conditional and unconditional branches.

The individual workloads and traces are described below. For a summary of statistics on the traces and workloads, see Table #1. A complete table of data is presented in the appendix (Table #A-1).

- (1) **COMP:** MIPS Compiler Traces. This workload is comprised of traces of MIPS compiler related programs. The traces are of significant real world programs, as the names, and numbers in Table #1 show. There are seven traces in this workload:

Table #1: Average Workload and Trace Statistics								
	COMP	TEXT	FP	SPARC	VAX	68k	Work Ave	Trace Ave
CPU	MIPS	MIPS	MIPS	SPARC	VAX	68k		
Dynamic Instr	2161788	2428066	2439411	1976181	1045646	681475	1788761	1644401
Static Instr	15737	11696	14157	4860	2832	3978	8877	8619
Object Instr	54516	48287	60658	43235	na	na	na	na
Object Bytes	218064	193148	242632	172940	na	na	na	na
Stat Branch Instr	17.47	19.93	12.29	21.11	26.05	23.69	20.09	20.71
Stat Branch Bytes	17.47	19.93	12.29	21.11	13.06	20.76	17.44	17.23
Stat Cond Instr	9.74	12.97	8.03	11.77	16.04	11.60	11.69	11.90
Stat Uncond Instr	7.73	6.96	4.26	9.34	10.01	12.10	8.40	8.81
Branch Taken	72.84	67.80	69.91	59.29	70.79	69.50	68.36	69.39
Cond Taken	64.48	63.44	59.49	44.18	61.14	56.62	58.23	59.73
Dyn Blk Instr	5.65	5.00	10.38	4.60	3.38	4.27	5.55	5.26
Dyn Blk Bytes	22.61	19.98	41.53	19.72	14.26	11.97	21.68	20.22
Stat Blk Instr	5.86	5.03	8.92	4.75	3.96	4.25	5.46	5.27
Stat Blk Bytes	23.46	20.10	35.68	19.01	17.60	14.25	21.68	20.85
Obj Blk Instr	5.38	4.70	7.08	4.46	na	na	na	na
Obj Blk Bytes	21.52	18.79	26.56	17.85	na	na	na	na
Definitions								
Dynamic Instr	Average number of dynamic instructions							
Static Instr	Average number of static instructions							
Object Instr	Average number of instructions in the entire program (object text segment)							
Object Bytes	Average number of bytes in the entire program (object text segment)							
Stat Branch Instr	Percentage static branch instructions of all static instructions							
Stat Branch Bytes	Percentage bytes of static branch instructions of all static instruction bytes							
Stat Cond Instr	Percentage static conditional branch instructions of all static instructions							
Stat Uncond Instr	Percentage static unconditional branch instructions of all static instructions							
Branch Taken	Percentage of branches that are taken							
Cond Taken	Percentage of conditional branches that are taken							
Dyn Blk Instr	Dynamic basic block size in instructions, including branch instruction							
Dyn Blk Bytes	Dynamic basic block size in bytes, including branch instruction							
Stat Blk Instr	Static basic block size in instructions, including branch instruction							
Stat Blk Bytes	Static basic block size in bytes, including branch instruction							
Obj Blk Instr	Object (text segment) basic block size in instructions, including branch instruction							
Obj Blk Bytes	Object (text segment) basic block size in bytes, including branch instruction							
Work Ave	Average of all the workload average values							
Trace Ave	Average of all of the individual trace values							

Table #1: This table contains statistics on the 32 traces we used. The values are all averages, including the workload averages (COMP, TEXT, FP, SPARC, VAX, and 68k), the overall trace average for the 32 traces (Trace Ave), and the average of the workload averages (Workload Ave). More information can be found in Table #A-1.

- asm MIPS Co. assembler, as01.21
  - ccom MIPS Co. C Compiler front end, ccom1.21
  - fcom MIPS Co. F77 Fortran Compiler front end, fcom1.21
  - ld MIPS Co. link editor, ld1.21
  - ugen MIPS Co. microcode generator, ugen1.21
  - uopt MIPS Co. microcode optimizer, uopt1.21
  - upas MIPS Co. Pascal Compiler front end, upas1.21
- (2) **TEXT:** MIPS Text Traces. This workload is comprised of traces of MIPS text processing programs. There are four traces in this workload:
- emacs GNU Emacs compiled on the MIPS.
  - grep GNU grep version 1.1 compiled on the MIPS.
  - nroff MIPS Co. version of nroff.
  - vi MIPS Co. version of vi.
- (3) **FP:** MIPS Floating Point Traces. This workload is comprised of floating point intensive programs compiled on the MIPS using the MIPS Co. fortran77 compiler. There are four traces in this workload:
- doduc F77 high energy physics program (doduc.f) [Doduc89].
  - integ F77 integration program (integral.f)
  - mold F77 molecular dynamics program (moldyn.f)
  - spice F77 circuit simulator program (spice2g6) [UCB87].
- (4) **SPARC:** SPARC Program Traces. This workload is comprised of three traces, not quite as significant as the MIPS traces.
- ccom2 Old (unknown) version of Sun 4 C Compiler front end.
  - esp2 Trace of espresso, logic minimization program.
  - gbis2 Trace of GNU Bison, 'yacc' like program.
- (5) **VAX:** VAX Program Traces. This workload is comprised of seven traces [Henry83].
- awk Unix awk
  - ls Unix ls
  - otmdl Parser/Constructor, written in Pascal, uses set operations
  - sedx Unix stream editor sedx
  - spic Spice circuit simulator
  - troff Unix troff
  - ymerge Parse table compacter, written in C
- (6) **68k:** 68010 Program Traces. This workload is comprised of seven traces [Grochowski86].
- as Unix assembler as
  - egrep Unix egrep utility
  - fort1 F77 matrix manipulation program
  - fort1a Trace of 'fort1' compiled with optimizer
  - ls Unix ls program with '-R' flag set



stat Trace analyzer 'stat'  
 stata Trace of 'stat' compiled with optimizer

#### 4. BTB Design Optimization

The theme of this paper is to maximize BTB performance with a limited number of bits allocated to the BTB implementation. Describing this problem mathematically helps to clarify the issues. To begin, it is necessary to understand the performance impact of different amounts and types of information stored for each branch in the BTB.

Potential BTB performance increases as the information content of each entry of the BTB increases, and as the number of entries increases. Increased information content requires more bits of storage, as does an increased number of entries. The simplest BTB with the lowest level of performance contains only prediction information. Predicting the branch direction can reduce the branch penalty. A medium complexity BTB which gives equal or higher performance contains the branch tag, prediction information, and the branch target address. With the branch target address available before it is normally calculated, fetching can occur earlier at the branch target. A high complexity BTB with still higher performance contains the branch tag, prediction information, target address, and instruction bytes from the target address. This design can decrease the delay in fetching instructions from the target. Increasing the number of types of information (e.g. prediction information, target address, target instruction bytes) stored for each branch improves performance. In addition, increasing (per entry) the amount (bits of storage) of each type of information also improves potential performance.

An optimization problem exists if a maximum performance BTB is to be designed with a limited number of bits. As both greater numbers of entries and greater amounts of information per entry increase performance, the problem is selecting the number of entries and the amount of information per entry that will maximize performance. Which design provides better performance, a large number of entries with a small amount of information per entry, or a small number of entries with a large amount of information per entry?

The following simple mathematical BTB model is intended to help clarify the problem. This model gives insight into how to maximize the performance of a BTB design that has been added onto a processor that normally predicts all branches not taken. The equation presents the probable savings per branch of an optimal BTB designed with  $N$  storage bits. Some simplifications have been made in setting up this equation, as noted below.

$$\text{Max} \left\{ \sum_i h(i) \left[ t(i)ptt(i)V(i) - (1-t(i))ptnt(i)W(i) \right] \right\} \quad \text{such that} \quad \sum_i \text{num\_bits}(i) \leq N$$

where:

- $N$  = number of storage bits in the BTB
- $h(i)$  = probability that branch  $i$  is referenced (executed)
- $t(i)$  = probability that branch  $i$  will be taken
- $ptt(i)$  = probability of predicting branch  $i$  taken, given that branch  $i$  is taken
- $ptnt(i)$  = probability of predicting branch  $i$  taken, given that branch  $i$  is not taken
- $V(i)$  = cycle savings for correct taken prediction of branch  $i$
- $W(i)$  = cycle cost for incorrect taken prediction of branch  $i$
- $\text{num\_bits}(i)$  = number of bits in BTB allocated to branch  $i$

Note that  $ptt(i)$  and  $prnt(i)$  are zero for branches not in the BTB, so the sum is only effected by branches in the BTB.

The first product in the equation, when summed up, is the expected cycle savings due to the BTB, in the case of a correct taken prediction. The second product in the equation, summed up, is the expected cycle cost of an incorrect taken prediction. The difference of these two sums is the net cycle savings per branch due to the BTB. The case of a correct not taken prediction has no performance penalty with or without the BTB (unless branch folding is used, in which case there may be a savings), so it does not appear in the equation. The case of incorrect not taken prediction is not included, since the cost is the same with or without a BTB (except when the BTB contains target instructions).

The problem of selecting the best BTB design is described by the equation above. For a fixed number of bits, increasing the performance of entries in the BTB is equivalent to increasing  $V(i)$  and decreasing  $W(i)$ , and decreases the number of entries. Increasing the number of entries in the BTB adds more branches into the sum, and decreases the information per entry. Finding the best performance for a limited number of total bits in the BTB requires a careful balance between performance per entry and number of entries.

The simple BTB model supports three concepts used later in this paper.

- (i) *If a branch does not have potential for improving performance, do not enter it into the BTB.* The BTB model indicates that the BTB should be filled with the branches that have the most potential for improving performance. Placing into the BTB branches with no performance value displaces branches that may have performance value.
- (ii) *When it is necessary to discard a branch from the BTB, discard the branch with the least potential performance value.* The BTB should contain the branches with the most potential for improving performance. A replacement strategy based on this concept might discard the branch that is least likely to be referenced AND least likely to be taken. A not taken branch that is in the BTB has the same performance as one that is not in the BTB, so discarding a branch that is not likely to be taken has little penalty.
- (iii) *A multi-level BTB, each level possibly containing different amounts/types of information per entry, may be able to maximize performance by achieving a better balance of number of entries and quantity of information per entry.* Note that  $V(i)$  and  $W(i)$  in the BTB model are functions of  $i$ . The BTB performance may be maximized if  $V(i)$  and  $W(i)$  are not constant. If branches with high performance potential are allocated more bits per entry, and branches that have less performance potential are allocated fewer bits per entry (but more entries), higher overall performance may be achieved. The constraint of a limited number of storage bits for the BTB design exists, so not all branches can be given the large size (higher performance) BTB entries.

## 5. Branch Prediction

In this paper, dynamic branch prediction is used, in which predictions are based on probability of taken branch statistics (as a function of  $n$  bits of branch execution history, known as prediction bits) compiled from all 32 traces. Since every branch has a startup period during which  $n$  bits of execution history are not available (only 1 to  $n-1$  bits are available in this period), statistics were also compiled for partial bit strings of execution history, from 1 to  $n-1$ . The statistics give the measured probability of a taken branch for the 32 traces as a function of  $n$  bits of execution history. As a result, predictions based on these statistics are the best possible predictions (yielding the highest correct prediction rate) for the 32 traces of any fixed, (i.e. non-learning or adaptive) method that only uses  $n$  bits of execution history as input. Table #A-2 (in appendix) gives probability of taken and probability of occurrence statistics for 6 branch history bits for

each of the workloads and for the 32 trace average. The statistics in Table #A-2 are appropriate for BTBs that enter branches on the first execution. Similar statistics are contained in Table #A-3 for BTBs that enter branches on taken branch executions only.

Achieving the highest prediction rate is not the same as achieving the highest performance. Four cases are possible when a prediction is made: predict not taken, branch not taken (zero cycle cost); predict not taken, branch taken (cycle cost  $b$ ); predict taken, branch not taken (cycle cost  $c$ ); and predict taken, branch taken (cycle cost  $d$ ). Maximizing the prediction rate is accomplished by predicting branch taken whenever the probability  $p$  of a taken branch is greater than 50 percent. Maximizing performance (minimizing cycle cost) is accomplished by predicting branch taken when the cost of predicting taken is less than the cost of predicting not taken (i.e. when  $(1-p)*c + p*d < p*b$ ). As this equation requires implementation dependent numbers ( $b,c,d$ ), all prediction rates presented in this paper are based on predicting taken if the probability of taken is greater than 50 percent. Later, in our discussion of multi-level BTBs, for which a single prediction rate is not a useful measure, performance is maximized using a set of cycle costs for the four cases described above.

### 5.1. Best Case Prediction Rates

As a reference point, it is useful to know the realistic upper limit for prediction rates. To determine this upper limit, two infinite size BTB simulators were designed (BTB#1 and BTB#2). Design BTB#1 takes in branches on their first execution and keeps them in the BTB until the end of the trace. Design BTB#2 takes in branches on their first taken execution and keeps them in the BTB until the end of the trace. Both BTBs predict all instructions that are not in the BTB as not taken. Note that the BTB prediction rate is independent of the amount of information per entry in the BTB, although the BTB must contain prediction bits and a branch tag (i.e. this is not a hash table BTB).

Table #2: Branch Prediction Rates (including target change effects)								
Prediction Bits	BTB Design	COMP	TEXT	FP	SPARC	VAX	68k	Average
0	BTB#1	67.21	66.13	61.78	56.80	68.85	64.96	64.29
0	BTB#2	81.27	80.86	81.04	73.04	81.54	76.68	79.07
1	BTB#1	86.30	94.83	85.03	87.51	89.71	85.55	88.15
1	BTB#2	86.30	94.83	85.03	87.51	89.71	85.55	88.15
2	BTB#1	87.81	95.45	86.72	88.69	91.56	86.27	89.42
2	BTB#2	87.81	95.45	86.72	88.69	91.56	86.27	89.42
4	BTB#1	88.66	96.02	87.85	90.61	93.36	88.21	90.79
4	BTB#2	88.66	96.02	87.84	90.61	93.35	88.20	90.78
8	BTB#1	89.19	96.18	89.28	91.78	93.70	89.00	91.52
8	BTB#2	89.19	96.18	89.27	91.77	93.67	88.97	91.51
16	BTB#1	90.14	96.40	89.90	92.93	95.25	91.48	92.68
16	BTB#2	90.13	96.39	89.89	92.91	95.22	91.45	92.67

Table #2: Prediction Rates for BTB#1 and BTB#2, both of infinite size. BTB#1 enters all branches as they are executed. BTB#2 enters branches only as they are taken. These prediction rates are upper limits on the prediction rates possible from finite BTBs.

Table #3: Prediction Rate Reduction due to Target Change						
COMP	TEXT	FP	SPARC	VAX	68k	Average
5.20	1.41	7.78	2.35	1.61	4.18	3.75

Table #3: This table indicates the constant reduction of the prediction rate due to the effect of target address changes.

Table #4: Lee and Smith's Percentage Target Changes [Lee84]						
IBM/CPL	IBM/BUS	IBM/SCI	IBM/SUP	PDP11	CDC6400	Average
4.2	2.1	4.4	1.4	12.0	2.9	4.5

Table #4: Lee and Smith's percentage probability of target change for their workloads [Lee84].

BTB#1 and BTB#2 prediction rates are presented in Table #2 (the full set of results is presented in the appendix, Table #A-4). The prediction rates in each column under a workload name are simple workload averages, and the **Average** column is the average of the workload averages. Table #2 includes the effect of target address changes in the prediction rate by considering a target change occurring during a taken branch to be a misprediction. Note that target changes do not affect a BTB that contains only prediction bits (no target address or instructions are stored). It happens that virtually all target changes occur for branches that are always taken (e.g. subroutine returns) and thus the effect of target changes is to decrease prediction accuracy by a constant amount. Those constants can be found in Table #3, and are identical for BTB#1 and BTB#2. Equivalent constants for the six workloads of Lee and Smith are shown in Table #4 [Lee84]. Figure #1 shows the average prediction rate for BTB#2 (including and not including the effect of target changes) as a function of the number of prediction bits.

The BTB#1 and BTB#2 prediction rates are very similar, except for the case of zero prediction bits. For zero prediction bits, the BTB#2 prediction rate is higher than BTB#1's prediction rate. This is due to the fact that for zero prediction bits, BTB#2 only predicts as taken those branches that are in the BTB. Since BTB#2 only enters branches into the BTB on a taken branch execution, branches that are never taken or are initially not taken are correctly predicted, thus boosting the prediction rate. BTB#1 enters branches into the BTB on the first execution, predicting the branches as taken thereafter, for the zero prediction bit case. Thus, the prediction rate of BTB#1 is approximately the average percentage of taken branches. Note that even with zero prediction bits, a BTB can be useful by caching the branch target and instructions at the branch target; in such a design, all branches in the BTB are predicted as taken.

The prediction rates are very close to those of Lee and Smith [Lee84]. Table #5 contains three columns of prediction rates, one for the average of prediction rates from Lee and Smith, one for the BTB#1 prediction rate (calculated in a similar fashion to Lee and Smith), and one for the BTB#2 prediction rate. These prediction rates do not include the effect of target address changes, so our values are from Table #2 with the **Average** constant from Table #3 added in. Lee and Smith use the percentage of taken branches as the prediction rate for zero prediction bits, basically the same as BTB#1, which is not comparable to the BTB#2 zero prediction bit case for reasons explained above. Lee and Smith generated prediction rates for up to 5 prediction bits. For up to four prediction bits, the prediction rate increases sharply. For more than 4 prediction bits, Figure #1 indicates that the prediction rate grows slowly.

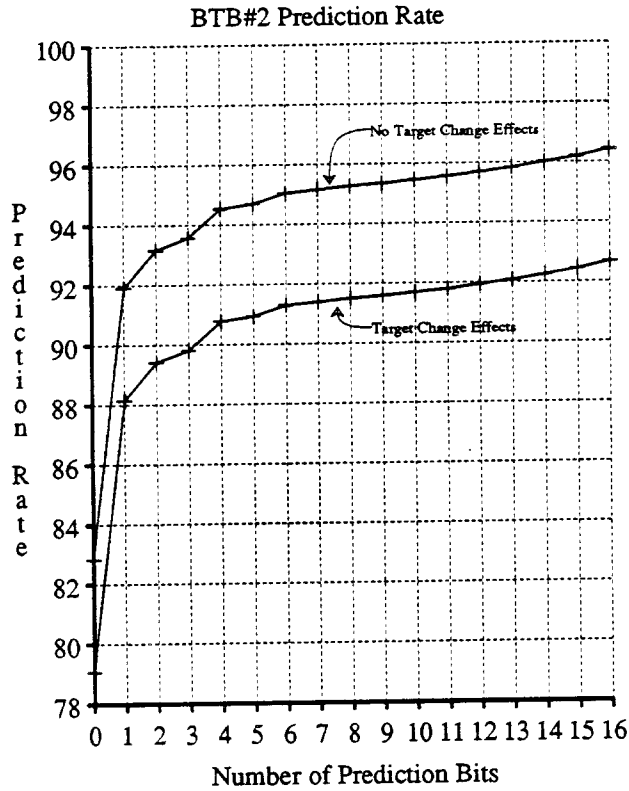


Figure #1: BTB#2 (best case) prediction rates as a function of number of prediction bits.

Prediction Bits	Lee & Smith	BTB#1	BTB#2
0	67.4	68.0	82.8
1	88.7	91.9	91.9
2	92.0	93.2	93.2
3	92.7	93.6	93.6
4	93.3	94.5	94.5
5	93.5	94.7	94.7

Table #5: Comparison of BTB#1 and BTB#2 prediction rates to those of Lee and Smith [Lee84].

### 5.2. Effect of Context Switching on Prediction Rates

Multitasking computer systems frequently experience context switches. During a context switch, the processor stops executing one program and begins executing another. The BTB normally contains invalid data after a context switch, and is flushed. As a result, context switching has a detrimental effect on the prediction rate of the BTB. In order to observe this effect, the infinite size BTB#2 simulator was altered to flush the BTB after every  $N$  instructions. Prediction rates were generated for  $N$  equal to 1000, 3160, 10000, 31600, 100000, 316000, and 1000000.

The results are shown in Figure #2, with curves for 0, 1, 2, 4, 8, and 16 prediction bits (the full set of results are presented in the appendix in Table #A-5). Each curve represents the average prediction rate of the six workload averages. The prediction rates include the effect of branch target changes (i.e. target change on a taken branch is considered a misprediction).




Figure #2: Effect of context switching on the prediction rate of BTB#2. Each of the curves is for BTB#2 with a constant number of prediction bits.

Examining Figure #2, it is apparent that context switching significantly reduces the prediction rate, and correspondingly, the performance benefit of a BTB. With  $N=10000$ , two prediction bits attain a 84 percent prediction rate while with no context switching they attain a 89 percent prediction rate. For context switches that occur less than every 10000 instructions, the effectiveness of adding additional prediction bits to increase the prediction rate is limited. For  $N$  above 10000 instructions, adding prediction bits has a stronger effect on the prediction rate. Some of the traces contain fewer instructions than  $N$  (for some  $N$ ), so that for  $N$  greater than the number of instructions in the trace, the context switching does not effect the prediction rate for the trace.

The curve for zero prediction bits in Figure #2 is interesting. Apparently, with zero prediction bits, context switching can increase the prediction rate. This anomaly is due to the characteristics of the zero prediction bit BTB, discussed in the previous section. The zero prediction bit BTB correctly predicts taken branches that are in the BTB, and not taken branches that are not in the BTB. Context switches clear the BTB of branches that are experiencing not taken executions, allowing them to be predicted correctly. These correct predictions counter the mispredictions it takes to get taken branches back into the BTB, producing a maximum prediction rate for a finite context switch rate.

## 6. BTB Management

Management of BTBs is concerned with the issue of entering and discarding branches from the BTB. Good management schemes significantly improve the performance of a BTB. In this section, the first two optimization concepts presented in the **BTB Design Optimization** section are used to present alternative BTB management schemes.

### 6.1. Entering Branches into the BTB

The first concept presented in the **BTB Design Optimization** section is: *If a branch does not have the potential for improving performance, do not enter it into the BTB.* Two methods for entering branches into the BTB are evaluated here. The first method is to enter branches into the BTB when branches are first executed (design BTB#3). This obvious method has an apparent drawback. If a branch is entered into the BTB on a not taken execution, the BTB will not improve performance while the branch continues not taken executions, and will mispredict the first taken execution (with no performance improvement) that the branch experiences. In addition, if the BTB contains target instruction bytes, entering the branch on a not taken execution requires extra memory bandwidth to fetch the target instruction bytes. The second method for entering branches into the BTB is to enter branches on their first taken execution (design BTB#4). This method avoids entering into the BTB entries of limited usefulness, which may themselves displace more useful entries. Thus, the BTB does not have to deal with branches that are never taken and which have no performance penalty that the BTB can reduce. If the BTB contains target instructions, entering branches on taken execution allows the target instructions to be cached in the BTB at no cost, since they are normally fetched on a taken branch execution.

The two methods of entering branches into the BTB were simulated. Design BTB#3 enters branches into the BTB on the first branch execution, is set associative, and uses LRU (least recently used) replacement. Design BTB#4 enters branches into the BTB on the first branch taken execution, is set associative, and uses LRU replacement. Both BTBs use branch tags to uniquely identify the branches. The BTBs were simulated for 0 to 16 prediction bits, 16 to 2048 entries, and set sizes of 1 to 16. The prediction rates presented are averages of the workload averages, and include the effect of target address changes (i.e. a target change on a taken branch execution is considered an incorrect prediction).

Table #6 presents the prediction rates for the two BTB designs, for set size of 4, and number of prediction bits 0, 1, 2, 4, 8, and 16 (the full set of results is presented in the appendix in Table #A-6). As expected, BTB#4 significantly outperforms BTB#3 in every case, but especially for low numbers of entries in the BTB. Entering branches into the BTB on the first taken execution avoids placing useless information in the BTB that displaces useful information. For the 32 traces, 17 percent of all static branches are never taken, and these never taken branches constitute 12 percent of all dynamic branch executions. These branches never enter BTB#4, but they enter BTB#3, hurting performance. Not only does BTB#4 have higher performance, it also has a simpler design. The BTB is notified to enter a branch whenever a branch is taken that does not exist in the BTB. If the BTB contains target instructions, they can be stored into the BTB as the instruction fetch stage of the processor fetches from the target address of the taken branch.

Table #6 indicates that the number of entries in a BTB has a significant impact on the prediction rate of the BTB. With two prediction bits, prediction rates range from 64 percent with 16 entries to 89 percent with 2048 entries. The growth in prediction rates (as a function of BTB size) levels off after 512 entries, so for the traces used, a 512 entry BTB would have about the same performance as a 2048 entry BTB, at a much lower cost.

Table #6: Comparison of Methods of Entering Branches into a BTB										
			Number of BTB Entries							
Prediction Bits	BTB Design	Set Size	16	32	64	128	256	512	1024	2048
0	BTB#3	4	47.60	50.63	55.58	58.97	61.46	63.18	63.92	64.18
0	BTB#4	4	62.72	69.47	74.68	77.10	78.32	78.96	79.00	79.05
1	BTB#3	4	59.39	67.68	75.31	80.67	84.20	86.58	87.63	88.01
1	BTB#4	4	63.61	71.27	77.92	82.11	85.28	87.10	87.83	88.08
2	BTB#3	4	59.74	68.20	76.22	81.72	85.36	87.81	88.88	89.27
2	BTB#4	4	64.01	72.03	78.88	83.21	86.47	88.33	89.09	89.33
4	BTB#3	4	59.97	68.59	76.89	82.65	86.51	89.11	90.23	90.64
4	BTB#4	4	64.25	72.50	79.57	84.19	87.65	89.64	90.43	90.70
8	BTB#3	4	60.13	68.81	77.31	83.15	87.13	89.79	90.95	91.37
8	BTB#4	4	64.43	72.85	80.01	84.70	88.30	90.34	91.16	91.42
16	BTB#3	4	60.29	69.02	77.61	83.71	87.94	90.82	92.07	92.52
16	BTB#4	4	64.60	73.09	80.41	85.35	89.19	91.40	92.29	92.58

**Table #6:** Comparison of two methods of entering branches into a BTB. Higher prediction rate indicates higher performance. BTB#3 enters branches into the BTB on the first branch execution, while BTB#4 enters branches on the first branch taken execution.

## 6.2. Discarding Branches from the BTB

The second concept presented in the **BTB Design Optimization** section is: *When it is necessary to discard a branch from the BTB, discard the branch with the least performance value.* It is not possible to determine the branch with the least performance value in a real system, but it is possible to find a branch that probably has the least performance value. The Least Recently Used (LRU) algorithm is known by experience to work well for replacing (discarding) entries in caches, although the random algorithm works better with small caches, avoiding the worst case behavior of the LRU algorithm for loops that are larger than the small cache [SmithJ83]. The LRU algorithm replaces the entry that has been accessed least recently. Design BTB#4, described above, uses LRU replacement.

A BTB has an added level of complexity over a cache, since each entry can contain prediction bits. The prediction bits provide more information that can be used in finding the branch with the least performance value. If a branch's prediction bits strongly predict not taken, it may be reasonable to discard the branch, as the BTB does not contribute to the performance of a not taken branch. But if this branch has a high probability of being referenced, it will probably be executed soon, and may predict taken in the future, and thus probably has some performance value. What is needed is an algorithm that replaces the branch that is least likely to be used and least likely to be taken. The Minimum product of Probability of reference and Probability of taken (MPP) algorithm is a good approximation. This algorithm replaces the branch with the minimum product of probability of reference and probability of branch taken. The probability of reference can be obtained empirically as a function of LRU bits that the algorithm maintains. The LRU bits indicate the order in which the branches have been referenced. The probability of branch taken is obtained from the branch prediction statistics. Table #7 gives the probability of reference values for three BTB sizes and a set size of four (LRU bits=0 is the most recently referenced entry in the set, LRU bits=3 is the least recently referenced entry in the set). Table #8



gives the probability of a taken branch as a function of three prediction bits of branch history. Table #9 presents an ordered list of the possible MPP products with a 128 entry 4 way set associative three prediction bit BTB (combination of Tables #7 and #8, disregards partial history strings). Table #A-3 (in appendix) gives probability of taken and probability of occurrence statistics for 6 branch history bits for each of the workloads and for the 32 trace average (these statistics are for BTBs that enter branches on taken executions only).

<b>Table #7: Probability of Reference (Set Size of 4)</b>				
	<b>LRU bits: Order of Reference</b>			
<b>BTB Entries</b>	<b>0</b>	<b>1</b>	<b>2</b>	<b>3</b>
<b>16</b>	.5111	.2199	.1464	.1225
<b>32</b>	.5504	.2123	.1405	.0969
<b>64</b>	.6147	.2030	.1073	.0750
<b>128</b>	.7016	.1729	.0812	.0443
<b>256</b>	.7771	.1401	.0550	.0277
<b>512</b>	.8450	.1060	.0337	.0153
<b>1024</b>	.8967	.0781	.0188	.0064
<b>2048</b>	.9417	.0474	.0084	.0025

**Table #7:** For several sizes of a 4 way set associative BTB, this table gives the probability of reference of the entries in each set as a function of the LRU bits. The most recently accessed entry of the set has LRU bits=0, while the least recently accessed entry of the set has LRU bits=3.

<b>Table #8: Probability of Taken Branch</b>		
<b>Branch History</b> N = Not taken T = Taken	<b>Probability of Occurrence</b>	<b>Probability of Taken Branch</b>
NNN	.1305	.0780
NNT	.0151	.3405
NTN	.0207	.5186
NTT	.0249	.6794
TNN	.0154	.3258
TNT	.0307	.6469
TTN	.0250	.7914
TTT	.7377	.9766

**Table #8:** As a function of a branch's history, this table gives the probability that the next execution of the branch will be taken, and the probability with which this branch history will occur.

The MPP algorithm can be implemented relatively cheaply. A set size of two or four entries is common in caches, so only one or two LRU bits are required. Context switching has a severe impact on the performance of large numbers of prediction bits so one to five prediction

Table #9: Order of Possible MPP Products for 128 Entry 4 Way Set Associative Three Prediction Bit BTB		
LRU Bits	Prediction Bits	MPP Products
0	TTT	.6852
0	TTN	.5552
0	NTT	.4767
0	TNT	.4539
0	NTN	.3638
0	NNT	.2389
0	TNN	.2286
1	TTT	.1689
1	TTN	.1368
1	NTT	.1175
1	TNT	.1118
1	NTN	.0897
2	TTT	.0793
2	TTN	.0643
1	NNT	.0589
1	TNN	.0563
2	NTT	.0552
0	NNN	.0547
2	TNT	.0547
3	TTT	.0433
2	NTN	.0421
3	TTN	.0351
3	NTT	.0301
3	TNT	.0287
2	NNT	.0276
2	TNN	.0230
3	NTN	.0230
3	NNT	.0151
3	TNN	.0144
1	NNN	.0135
2	NNN	.0063
3	NNN	.0035

**Table #9:** Table lists in order of MPP product value the 32 possible products for a 128 entry 4 way set associative BTB with three prediction bits. The products from branch history strings of less than three are not presented. Note that the MPP algorithm would replace the entry (one of four) that has the minimum product value.

bits is reasonable. Probability of reference is a function of the LRU bits, and probability of taken is a function of the prediction bits. Therefore, at most seven bits of input are required to calculate the product for MPP for each entry. A total of four input bits will probably be typical. Four input bits indicate 16 possible values. What is important is not the actual value of each product of probability of reference and probability of taken but the relative ordering of the products. Each branch's product can therefore be encoded into four bits, and can be generated by a simple four bit table with four bits of input. The four bit results of the branches in a set can be quickly compared to find the branch with the minimum product.

To compare the LRU and MPP algorithms, two BTB simulators were constructed. Design BTB#4 is set associative, uses LRU replacement, and enters branches on taken executions. Design BTB#5 is set associative, uses MPP replacement, and enters branches on taken executions. Both BTBs contain branch tags and prediction bits. The BTBs were simulated for 0 to 16 prediction bits, 16 to 2048 entries, and set sizes of 1 to 16. The prediction rates presented are averages of workload averages, and the effect of target changes is included in the prediction rate (i.e. a target change on a taken execution is a misprediction). The full table of prediction rate data is in the appendix in Table #A-6. Table #10 contains prediction rate data for BTB#4 and BTB#5 with set sizes of four and number of predictions bits 0, 1, 2, 4, 8, and 16.

Table #10: Comparison of BTB Replacement Strategies										
			Number of BTB Entries							
Prediction Bits	BTB Design	Set Size	16	32	64	128	256	512	1024	2048
0	BTB#4	4	62.72	69.47	74.68	77.10	78.32	78.96	79.00	79.05
0	BTB#5	4	62.72	69.47	74.68	77.10	78.32	78.96	79.00	79.05
1	BTB#4	4	63.61	71.27	77.92	82.11	85.28	87.10	87.83	88.08
1	BTB#5	4	63.92	71.50	78.14	82.28	85.35	87.13	87.84	88.08
2	BTB#4	4	64.01	72.03	78.88	83.21	86.47	88.33	89.09	89.33
2	BTB#5	4	64.17	72.11	79.02	83.26	86.52	88.35	89.09	89.34
4	BTB#4	4	64.25	72.50	79.57	84.19	87.65	89.64	90.43	90.70
4	BTB#5	4	64.50	72.56	79.68	84.25	87.71	89.65	90.44	90.70
8	BTB#4	4	64.43	72.85	80.01	84.70	88.30	90.34	91.16	91.42
8	BTB#5	4	64.69	72.91	80.11	84.77	88.34	90.35	91.16	91.43
16	BTB#4	4	64.60	73.09	80.41	85.35	89.19	91.40	92.29	92.58
16	BTB#5	4	64.85	73.14	80.46	85.33	89.19	91.39	92.29	92.57

Table #10: This table contains a comparison of the MPP and LRU replacement strategies. BTB#4 uses LRU replacement and BTB#5 uses MPP replacement. The MPP strategy has a slightly higher prediction rate in most cases.

Examining Table #10, the MPP replacement algorithm (BTB #5) outperforms the LRU replacement algorithm (BTB #4) except for a few BTB sizes with 16 prediction bits. However, the prediction rate difference between the two designs is trivial, and probably does not justify the extra complexity of the MPP design.

## 7. Relationship between BTB and Instruction Cache Miss Rates

Smith hypothesized that a BTB and a instruction cache should have similar miss rates if the following conditions are met [SmithAJ89]:

- (i) *The cache line is one instruction in length.*
- (ii) *The cache size is the BTB size (number of entries) multiplied by the average basic block size.*

His reasoning is that a BTB is simply an instruction cache that only holds branches. Assume the basic block size is  $n$  instructions (lines). One branch exists in every basic block (i.e. on average, one of every  $n$  instructions is a branch). Therefore, the instruction cache has  $n$  misses for every miss in the BTB and  $n$  hits for every hit in the BTB. If the two conditions above are met, the miss ratio should be the same.


The hypothesis was tested with a BTB and a cache simulator. The BTB simulator (design BTB#3, the same design as was used previously) is set associative with LRU replacement and enters branches on their first execution (both taken and not taken). The BTB sizes are powers of two. The cache simulator is set associative with LRU replacement and a four byte line size. Cache sizes are not powers of two since the cache size is equal to the BTB size multiplied by the basic block size. The six workloads were fed into the simulators, each workload using its own average basic block size to size the cache. The BTB and cache miss rates can thus be compared for equality within each workload. Table #A-7 (in appendix) contains the resulting miss rates. Each miss rate in the table is the geometric mean of the miss rates of a workload's traces.

Figure #3 displays the cache and BTB#3 miss rates for the COMP, TEXT, and FP workloads with set size of four. Figure #4 displays the cache and BTB#3 miss rates for the SPARC, VAX, and 68k workloads with set size of four. It is immediately apparent that the cache and BTB miss rates are nearly equivalent. The four byte line size worked very well for the MIPS (COMP, TEXT, FP) and SPARC (SPARC) traces, since MIPS and SPARC only have four byte instructions. In Table #1, by dividing **Dyn Blk Bytes** with **Dyn Blk Instr**, we can calculate that the VAX traces have an average instruction length of 4.2 bytes and the 68k traces have an average instruction length of 2.8 bytes. The curves in Figure #4 for the VAX are not compensated for this slightly larger instruction length. The 68k curves in Figure #4, however, have been compensated by multiplying the cache miss rate by a factor of  $4.0/2.8$ . The compensation is necessary since a single miss in the cache reads in  $4.0/2.8$  instructions instead of the assumed 1 instruction, thus lowering the cache miss rate by a factor of  $2.8/4.0$ . Once compensated, the BTB and cache miss rates for the 68k are virtually equivalent.

With set sizes of 1 and 2, the cache and BTB miss rates differ more than for set sizes of 4 or 8. Generally, with set size of one or two, the BTB miss rate is higher than the cache miss rate. The reason for this is simple. Programs tend to have code that is slightly ordered with respect to powers of two, and a cache with a power of two number of sets will have more collisions than one that has a non-power of two number of sets. The effect of collisions is greatest when the set size is small [Hill89].

The BTB#3 design (branches entered on first execution) has been shown to not perform as well as the BTB#4 design (branches entered on first taken execution). It is interesting, therefore, to compare their miss rates. Table #A-8 (in the appendix) contains the miss rate data for BTB#3 and BTB#4. Each value under a workload heading is a geometric mean of the miss rates for the traces in the workload. The average column contains values that are geometric means of the workload geometric mean miss rates.

Figure #5 contains the BTB#3 and BTB#4 miss rates for the COMP, TEXT, and FP workloads (set size of four), and Figure #6 contains the BTB#3 and BTB#4 miss rates for the SPARC, VAX, and 68k workloads (set size of four). The miss rate curves indicate that BTB#4 has




**Figure #3:** BTB#3 and corresponding instruction cache miss rates for the workloads COMP, TEXT, and FP. This figure shows that the BTB and cache miss rates can be computed from one another.

significantly lower miss rates than BTB#3.

The relationship between the BTB#3 and BTB#4 miss rates depends on the percentage of taken branches and the BTB size. BTB#3 enters branches into the BTB on the first execution of the branches, and BTB#4 enters branches on the first taken execution of the branches. Therefore, as the percentage of taken branches increases, the BTB#4 miss rate approaches the BTB#3 miss rate. The effect of BTB size on the relationship between BTB#3 and BTB#4 miss rates can be observed by examining the two cases, infinite BTB size and small BTB size. In an infinite size BTB, BTB#4 will not miss on those branches that are never taken (17 percent of static branch instructions in our traces), so the number of misses for BTB#4 is the number of misses for BTB#3 (number of static branches) minus the number of static branches that are never taken. For our traces, the resulting miss rate for BTB#4 is 83 percent of the BTB#3 miss rate. As BTB size decreases, entering not taken branches generates additional misses by displacing useful branches. The result is that BTB#3's miss rate grows relative to BTB#4's miss rate, as BTB size decreases.

Table #11 contains the BTB#3 and BTB#4 miss rates and their ratio as a function of BTB size for a 4 way set associative BTB. Notice that for BTB#3 miss rates that are greater than 1.0%, the BTB#4/BTB#3 ratio ranges between .542 and .598 (a relatively small range) with an average value of .57. With less than a 1.0% miss rate for BTB#3, the ratio rises sharply. A simple model of the BTB#4/BTB#3 miss rates ratio is therefore .57 for BTB#3 miss rates greater than 1.0%, and undefined for BTB#3 miss rates less than 1.0% (alternatively, this region could be assigned a worst case miss rate ratio of 1.0). This model works reasonably well for set sizes other




**Figure #4:** BTB#3 and corresponding instruction cache miss rates for the workloads SPARC, VAX and 68k. This figure shows that the BTB and cache miss rates can be computed from one another.

than four as well. This model does not take into account the effect of the taken branch percentage, so workloads with taken branch percentages other than 70 percent may need a different model.

### 7.1. BTB Design Target Miss Ratios

Given the relationship that we have obtained between BTB and instruction cache miss ratios, we can derive *design target miss ratios* for BTBs. Design target miss ratios are those to be expected for an "average" workload on a real machine in normal operation. They were defined and created in [SmithAJ85] and [SmithAJ87]; the methodology is described in those papers. To convert instruction cache design target miss ratios to design target miss ratios for BTBs, the average basic block size and the average instruction size in the workloads used are needed. We have combined our data with that from some other published sources. Peuto and Shustek extracted extensive information from seven IBM 370 traces that indicate average basic block size of 24.22 bytes and an average instruction size of 3.68 bytes [Peuto77]. Lee and Smith used four IBM 370 workloads with a total of 15 traces that have an average basic block size of 18.98 bytes, assuming an average instruction length of 3.68 bytes [Lee84]. Our 7 VAX traces indicate an average basic block size of 4.2 bytes and basic block size of 14.26 bytes. The average instructions lengths bracket 4 bytes per instruction, so for simplicity, 4 bytes per instruction is assumed in the conversion calculations. The basic block sizes bracket our 32 trace average of 20.22 bytes so this basic block size is assumed.



**Figure #5:** BTB#3 and BTB#4 miss rates for the workloads COMP, TEXT, and FP as a function of the number of BTB entries. BTB#4 has significantly lower miss rates than BTB#3.

Using data from [SmithAJ87] for a 4-way set associative instruction cache with 4-byte line size, LRU replacement, and cache sizes from 32 bytes to 32768 bytes, BTB miss ratio data was generated. Note that miss ratios for caches with line sizes greater than 4 bytes can be converted to an equivalent 4-byte line size miss ratio using the ratios of miss ratios data in [SmithAJ87]. Figure #7 contains the resulting miss ratios for BTB designs BTB#3 and BTB#4 (four way set associative, LRU replacement). Table #12 contains BTB design target miss ratio data for 2 to 1024 BTB entries.

## 8. Information Stored in the BTB

The information stored in a BTB strongly influences the performance of the design. Four types of information have been stored in BTBs: branch tags (uniquely identifying the branch), prediction information, branch target address, and target instruction bytes. In addition, in a BTB design described below, the instruction bytes immediately following the branch instruction are stored in the BTB, allowing branch folding to occur (eliminating branch execution time). In order to get an idea of the performance potential of the above types of information, we will consider their effect using the five stage pipeline mentioned in the introduction. Seven possible single-level BTB designs will be discussed using this example pipeline. The five pipeline stages are: instruction fetch, instruction decode, operand fetch, execution, and result writeback. All the pipeline stages normally take one cycle to execute.

For a processor with branch instructions that have a simple uniform target address encoding, the target address can be available after the instruction decode stage, if all instructions





Figure #6: BTB#3 and BTB#4 miss rates for the workloads SPARC, VAX, and 68k. BTB#4 has significantly lower miss rates than BTB#3.

Table #11: Miss Rates for BTB#3 and BTB#4 and Ratio of Miss Rates BTB#4/BTB#3				
Entries	Set Size	BTB#3 Miss Rate	BTB#4 Miss Rate	BTB#4/BTB#3
16	4	46.82	25.90	.553
32	4	29.06	15.74	.542
64	4	14.87	8.34	.561
128	4	7.55	4.42	.585
256	4	3.58	2.14	.598
512	4	1.65	0.97	.588
1024	4	0.81	0.52	.642
2048	4	0.53	0.41	.774

Table #11: Table gives the miss rates of BTB#3 and BTB#4 and the ratio of their miss rates as a function of BTB size.

(branch and non-branch, since the instruction type is not determined until after decode) undergo target address calculations during the instruction decode stage. In the following designs, it is assumed that the target address is available after the decode stage. The conditions for a conditional branch are stable at the end of the operand fetch stage (i.e. at the end of the execution stage





**Figure #7:** BTB design target miss ratios, converted from cache design target miss ratios in [SmithAJ87].

of the previous instruction), at which time instructions can be fetched from the target address if the branch is taken. With one cycle random access memory, this pipeline's branch penalty for a taken branch is two cycles.

With the pipeline and assumptions given above, seven BTB designs are discussed below.

**i) Hash Table of Prediction Information:** This simple BTB is accessed with the branch instruction address (modulo the table size) during the instruction decode (or fetch) stage to obtain prediction information. If the decoded instruction is a branch, the predicted branch path is fetched. In the example pipeline, this design saves a cycle for a correctly predicted taken branch (removes the delay between target address available and conditions stable). The prediction information must be updated after every branch execution. Note that the hash table (in which no branch tag exists) allows collisions to occur (multiple branches accessing the same entry) which reduces the prediction rate. The Mitsubishi GMICRO/100 uses a one prediction bit 256 entry BTB of this type [Sakamura87].

**ii) Branch Tag and Prediction Information:** A branch tag identifies a branch instruction with an entry in the BTB before the decoding stage of the processor determines the instruction is a branch. Using a branch tag can remove a cycle of delay (the decode stage time) in the example pipeline if early branch identification is useful. However, in our pipeline, since the target address is not available until after the instruction decode, this BTB design will not reduce the branch penalty more than BTB (i), but it will eliminate collisions (multiple branches accessing the same entry in the hash table) at the cost of supporting the branch tag.

Table #12: Design Target Miss Ratios		
Number BTB Entries	Miss Ratio BTB#3	Miss Ratio BTB#4
2	70.8	40.4
4	65.4	37.3
8	60.7	34.6
16	58.2	33.2
32	54.2	30.9
64	46.6	26.6
128	35.0	20.0
256	23.8	13.6
512	16.4	9.3
1024	12.9	7.4

Table #12: BTB design target miss ratios, converted from cache design target miss ratios in [SmithAJ87].

iii) **Hash Table of Prediction Information and Target Address:** This BTB can remove the delay from the time of identifying the branch instruction to the time of conditions stable for a correctly predicted taken branch. In our pipeline, the branch is identified after the decoding stage when the target address is available, so this BTB has no performance benefit over BTB (i). However, if a delay exists between the time the instruction is identified as a branch and the time of target address available, this BTB will remove that additional delay, unlike BTB (i). For example, in our pipeline, if branches have a very simple encoding, they may be identified before going through the decoding stage, saving an additional cycle over BTB (i) for a correctly predicted taken branch. As with BTB (i), the hash table allows collisions to occur.

iv) **Branch Tag, Prediction Information, and Target Address:** In this BTB, the branch tag improves performance for our pipeline, unlike BTB (ii). This BTB removes the delay from the time of branch identification to the time of conditions stable. Since the branch tag identifies the branch one cycle earlier than the decoding stage, this BTB saves two cycles from the taken branch delay. The Edgecore E2000 uses a 1024 entry direct map BTB of this type, with one prediction bit, branch tag, and target address [Walter89].

v) **Target Instructions:** This BTB is often called a branch target cache. It is basically an instruction cache that caches instructions from the target addresses of branches. It has been used in the Am29000 and the GE RPM40 [AMD88] [Lewis88]. Hill compares it with instruction caches and instruction buffers [Hill87].

vi) **Branch Tag, Prediction Information, Target Address, and Target Instructions:** This BTB removes the delay from identifying a branch to receiving target instructions for a taken branch. In our pipeline, with one cycle random access memory, this BTB does not improve performance over BTB (iv). However, if the memory is static column memory, with a 2 cycle initial column access, and one cycle for same column accesses thereafter, the availability of the target instructions will reduce the taken branch delay by one additional cycle (total of three cycles saved). The target instructions can also reduce the delay in deep pipelines. This type of BTB is used in the NexGen processor [Stiles89a].

vii) **Branch Tag, Prediction Information, Target Address, Target Instructions, and Instructions after the Branch:** This BTB can implement branch folding, eliminating the execution time of branches [Ditzel87]. This BTB identifies a branch during instruction fetch, and instead of decoding the branch, decodes an instruction (taken from the BTB) from the predicted branch path. Extra hardware is required to decode and execute the branch in parallel, to verify the predicted branch path and target address. Simple branch encodings can reduce the decoding hardware required. This BTB allows branches to execute in zero time.

### 8.1. Tradeoffs Between Size and Number of Prediction Bits

Not only does the type of information contained in the BTB have an impact on cost and performance, the quantity of each type of information has an impact as well. In particular, Table #6 shows that significant tradeoffs can be made between size of the BTB (BTB#4) and number of prediction bits. For example, if a BTB with a prediction rate of 88 percent is required, both a one prediction bit 2048 entry BTB and a two prediction bit 512 entry BTB will suffice. If each BTB entry has enough storage bits such that the number of prediction bits is insignificant, then the two prediction bit BTB has the same performance as a one prediction bit BTB which uses four times as many storage bits. If a BTB with a prediction rate of 89 percent is needed, a two prediction bit 1024 entry BTB can be used, or, for about half of the storage bits (assuming the number of prediction bits is small compared to the total number of bits per entry), a four prediction bit BTB with 512 entries will do. The potential cost savings encourage the designer to be careful in selecting a BTB configuration.

### 8.2. Multi-Level BTBs

The third concept presented in the **BTB Design Optimization** section is: *A multi-level BTB, each level possibly containing different amounts/types of information per entry, may be able to maximize performance by achieving a better balance of number of entries and quantity of information per entry.* Within the constraint of a finite number of storage bits allocated to a BTB design, a multi-level BTB may maximize performance.

To test this concept of multi-level BTBs, four BTB designs were simulated. The multi-level BTBs have up to three levels. The first level, the highest performance level, contains a branch tag, prediction bits, target address, and target instruction bytes for each entry. The second level, the medium performance level, contains a branch tag (one design eliminates this), prediction bits, and a target address for each entry. The third level, the low performance level, is a hash table of prediction bits. Higher performance levels require more storage bits than the lower performance levels.

Four Branch Cases	costx			costy		
	level1	level2	level3	level1	level2	level3
Pred. not taken, branch not taken	0	0	0	0	0	0
Pred. not taken, branch taken	4	4	4	4	4	4
Pred. taken, branch not taken	4	4	4	4	4	4
Pred. taken, branch taken	0	2	3	0	1	2

Table #13: The cycle costs for the two cost sets, costx and costy, as a function of BTB level and the four possible branch cases.

Note that a simple prediction rate does not adequately represent the performance of a multi-level BTB. To obtain a measure of performance, implementation dependent cycle costs are required. Each of the three levels has different cycle costs for the four possible cases when a branch executes: predict not taken, branch not taken; predict not taken, branch taken; predict taken, branch not taken; predict taken, branch taken. Two sets of cycle costs, *costx* and *costy*, are used, and appear in Table #13. The predict not taken, branch not taken case has a zero cycle cost for all three levels and both cost sets. For simplicity, the two incorrect prediction cases (predict not taken, branch taken; predict taken, branch not taken) have a 4 cycle cost for both cost sets. The Amdahl 470 V/6 has a 4 cycle penalty when a branch is taken [Amdahl76], and the CLIPPER has a 3 to 5 cycle penalty [Hollingsworth89]. Note that in general, the two incorrect prediction cases need not have the same costs. Normally the sequential instruction stream is available, so an incorrect taken prediction costs less than a incorrect not taken prediction, in the second and third levels.

For the predict taken, branch taken case, the cycle costs are more interesting. In *costx*, a static column memory is assumed, with a three cycle initial column access, and one cycle access in that column thereafter. The target address is available one cycle after fetching the branch, and the conditions (for conditional branches) are stable one cycle after that. Level one entries have a zero cycle penalty by caching 8 bytes of instructions to make up for the extra two cycle delay in the initial three cycle target fetch. This assumes one instruction (4 bytes) executed every cycle. Level two entries have a two cycle penalty, due to the extra two cycle delay in the initial three cycle target fetch. Level three entries have a three cycle penalty, due to the one cycle to calculate the target address, and the two cycle delay in the initial target fetch.

For *costy*, the predict taken, branch taken case is similar. A static column memory is again assumed, with a two cycle initial column access, and one cycle per access in that column thereafter. The target address is available one cycle after branch instruction fetch, and the conditions are stable two cycles after that. The level one entries have a zero cycle penalty by caching 4 bytes of target instructions to eliminate the extra cycle of the initial two cycle target access. The level two entries have a one cycle penalty, the extra cycle on the initial target access. The level three entries have a two cycle penalty, the cycle to calculate the target address plus the extra cycle of the initial target access.

Using the cost sets *costx* and *costy*, four multi-level BTB designs were simulated. All the designs use global LRU replacement where needed so small regular size changes can be made in the various levels. (In a real implementation, levels one and two would be set associative). Branches in lower levels are promoted to the highest level on a branch taken execution so that the target instructions can be fetched without increasing memory demands. Entries at higher levels are demoted one level at a time as they are replaced at the higher levels. The number of prediction bits used in all levels is fixed at two. Table #14 contains the approximate number of bits required by each entry of the levels in the designs. Justification for the bit allocations is provided in the following paragraphs.

BTB#6 contains level one and level two. For *costx*, level one requires approximately 32 bits of branch tag, prediction bits, 32 bits of target address, and 8 bytes of target instructions, for a rough total of 128 bits per entry. For *costy*, level one requires branch tag, prediction bits, target address, and only 4 bytes of target instructions, for a rough total of 96 bits per entry. In BTB#6, level two is implemented as a hash table (no branch tag) containing prediction bits and target address for an approximate total of 32 bits for both *costx* and *costy*. This level two depends on a simple branch instruction encoding that allows the branch to be identified at the end of instruction fetch. Note that this hash table will experience collisions (multiple branches accessing the same entry) that will reduce performance.

	BTB#6		BTB#7		BTB#8		BTB#9	
	costx	costy	costx	costy	costx	costy	costx	costy
level1	128	96	128	96	128	96	128	96
level2	32	32	64	64	64	64		
level3					2	2	2	2

Table #14: Number of storage bits required per entry of the four multi-level BTBs, as a function of cost set and BTB level.

BTB#7 is identical to BTB#6 except that level two contains a branch tag that increases the number of bits per entry to roughly 64.

BTB#8 is a full three level BTB. Levels one and two are identical to BTB#7. Level three is a 1024 entry hash table of two prediction bit entries. Note that collisions do occur in level three.

BTB#9 contains level one (same as BTB#6) and level three, each entry of which requires two bits.

### 8.2.1. Results of Simulating Multi-Level BTBs

Each of the four multi-level BTB designs was simulated with two cost sets (*costx* and *costy*), three values for total number of bits in the BTB (4096, 8192, 16384), and nine different bit allocations between the levels of the BTBs for each of the six possible cost sets and total number of bits combinations. The resulting cycle savings per branch data is given in full in six tables in the appendix (Table #A-9 to Table #A-14). The values in the workload columns are workload averages, and the average column contains the average of the workload averages. Six reduced tables presenting the average cycle savings per branch are presented here (Table #15 to Table #20). Tables #15 and #16 are for 4096 total bits and cost sets *costx* and *costy* respectively. Tables #17 and #18 are for 8192 total bits and cost sets *costx* and *costy* respectively. Tables #19 and #20 are for 16384 total bits and cost sets *costx* and *costy* respectively.

It is important to understand that the results presented are dependent on the cost sets, the bit allocations per entry, and the BTB designs. A multi-level BTB may or may not be practical, strongly depending on the specific implementation.

The cycle savings per branch results for BTB#6 and BTB#7 indicate that it is worth allocating all available bits to the high performance level one, since any other bit allocation between levels one and two increases the cycle savings per branch by at most 1.5 percent, and in most cases makes performance worse. This is true for all three total bit sizes, and over both cost sets, *costx* and *costy*.

BTB#8, the full three level BTB, achieves optimum results with a 1024 entry level three and all the remaining bits allocated to level one. Bits allocated to level two did not improve performance. BTB#8 outperforms BTB#6 and BTB#7 in four of the six cost set/number of bits combinations. In terms of cycle savings per branch, BTB#8 outperforms BTB#6 and BTB#7 by a maximum of 5 percent.

BTB#9 performs the best of all, in one configuration (*costy*, 4096 bits, 32 level one entries, 512 level three entries) outperforming designs with all bits allocated to level one by about 10 percent. For this best case, if we assume one cycle per non-branch instruction, basic block size of

Table# 15: Average Cycle Savings Per Branch, 4096 bit BTB, costx															
BTB#6 Entries			BTB#7 Entries			BTB#8 Entries			BTB#9 Entries			Cycle Savings			
lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	BTB#6	BTB#7	BTB#8	BTB# 9
0	128	0	0	64	0	0	32	1024	0	0	2048	0.903	0.902	0.903	0.522
4	112	0	4	56	0	2	28	1024	4	0	1792	1.298	1.309	1.195	1.115
8	96	0	8	48	0	4	24	1024	8	0	1536	1.411	1.424	1.300	1.294
12	80	0	12	40	0	6	20	1024	12	0	1280	1.515	1.523	1.369	1.435
16	64	0	16	32	0	8	16	1024	16	0	1024	1.561	1.574	1.412	1.556
20	48	0	20	24	0	10	12	1024	20	0	768	1.589	1.605	1.484	1.610
24	32	0	24	16	0	12	8	1024	24	0	512	1.601	1.628	1.509	1.668
28	16	0	28	8	0	14	4	1024	28	0	256	1.599	1.638	1.530	1.679
32	0	0	32	0	0	16	0	1024	32	0	0	1.665	1.665	1.556	1.665

Table# 16: Average Cycle Savings Per Branch, 4096 bit BTB, costly															
BTB#6 Entries			BTB#7 Entries			BTB#8 Entries			BTB#9 Entries			Cycle Savings			
lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	BTB#6	BTB#7	BTB#8	BTB#9
0	128	0	0	64	0	0	32	1024	0	0	2048	1.465	1.420	1.506	1.136
5	112	0	5	56	0	2	28	1024	5	0	1792	1.674	1.645	1.660	1.560
10	96	0	10	48	0	5	24	1024	10	0	1536	1.747	1.719	1.733	1.711
16	80	0	16	40	0	8	20	1024	16	0	1280	1.785	1.759	1.771	1.777
21	64	0	21	32	0	10	16	1024	21	0	1024	1.793	1.774	1.805	1.867
26	48	0	26	24	0	13	12	1024	26	0	768	1.794	1.779	1.829	1.884
32	32	0	32	16	0	16	8	1024	32	0	512	1.802	1.792	1.847	1.945
37	16	0	37	8	0	18	4	1024	37	0	256	1.796	1.781	1.855	1.937
42	0	0	42	0	0	21	0	1024	42	0	0	1.777	1.777	1.867	1.777

**Table #15:** Average cycle savings per branch for the four BTB designs (BTB#6-BTB#9) with a total of 4096 storage bits and cost set costx.

**Table #16:** Average cycle savings per branch for the four BTB designs (BTB#6-BTB#9) with a total of 4096 storage bits and cost set costly.

5.26 instructions, and 70 percent branches taken, BTB#9 speeds up processor performance by 32 percent, and BTB#6 or BTB#7, with all bits allocated to level one (42 level one entries), speed up processor performance by 28 percent. So in this case, the multi-level BTB produces a net processor speedup of 4 percent over a single level BTB. In practice, the net speedup will be less than 4 percent since the average non-branch instruction will not execute in one cycle.

The limited speedup of a multi-level BTB compared to a single level BTB suggests that the additional complexity of the multi-level BTB is not cost effective. If the complexity of the multi-level BTB requires significant chip area, converting this area to storage bits for a simple single level BTB may produce equal or better performance. But multi-level BTBs are strongly dependent on the implementation. BTB#9 produced optimal performance for costly which gives lower BTB levels more performance value per entry than costx. RISC processors, with delayed branches, and very small (usually one cycle) branch penalties, will probably not find that multi-

Table# 17: Average Cycle Savings Per Branch, 8192 bit BTB, costx															
BTB#4 Entries			BTB#5 Entries			BTB#6 Entries			BTB#7 Entries			Cycle Savings			
lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	BTB#4	BTB#5	BTB#6	BTB#7
0	256	0	0	128	0	0	96	1024	0	0	4096	0.969	0.968	0.992	0.529
8	224	0	8	112	0	6	84	1024	8	0	3584	1.498	1.504	1.483	1.304
16	192	0	16	96	0	12	72	1024	16	0	3072	1.680	1.680	1.643	1.567
24	160	0	24	80	0	18	60	1024	24	0	2560	1.758	1.765	1.737	1.685
32	128	0	32	64	0	24	48	1024	32	0	2048	1.826	1.837	1.797	1.811
40	96	0	40	48	0	30	36	1024	40	0	1536	1.859	1.878	1.845	1.869
48	64	0	48	32	0	36	24	1024	48	0	1024	1.891	1.911	1.890	1.932
56	32	0	56	16	0	42	12	1024	56	0	512	1.910	1.940	1.913	1.971
64	0	0	64	0	0	48	0	1024	64	0	0	1.950	1.950	1.932	1.950

Table# 18: Average Cycle Savings Per Branch, 8192 bit BTB, costly															
BTB#4 Entries			BTB#5 Entries			BTB#6 Entries			BTB#7 Entries			Cycle Savings			
lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	BTB#4	BTB#5	BTB#6	BTB#7
0	256	0	0	128	0	0	96	1024	0	0	4096	1.545	1.520	1.595	1.145
10	224	0	10	112	0	8	84	1024	10	0	3584	1.858	1.835	1.881	1.734
21	192	0	21	96	0	16	72	1024	21	0	3072	1.932	1.909	1.967	1.887
32	160	0	32	80	0	24	60	1024	32	0	2560	1.969	1.956	2.011	1.984
42	128	0	42	64	0	32	48	1024	42	0	2048	1.982	1.978	2.047	2.035
53	96	0	53	48	0	40	36	1024	53	0	1536	2.001	1.998	2.068	2.086
64	64	0	64	32	0	48	24	1024	64	0	1024	2.007	2.009	2.087	2.111
74	32	0	74	16	0	56	12	1024	74	0	512	1.995	2.009	2.103	2.109
85	0	0	85	0	0	64	0	1024	85	0	0	2.008	2.008	2.111	2.008

Table #17: Average cycle savings per branch for the four BTB designs (BTB#6-BTB#9) with a total of 8192 storage bits and cost set costx.

Table #18: Average cycle savings per branch for the four BTB designs (BTB#6-BTB#9) with a total of 8192 storage bits and cost set costly.

level BTBs, or any BTB for that matter, significantly improve performance. CISC processors, however, with deep pipelines and larger branch penalties, may find multi-level BTBs useful, and almost certainly will find that a single level BTB significantly improves performance.

## 9. Conclusions

A Branch Target Buffer can reduce the performance penalty of branches in pipelined processors by predicting the path of the branch, and caching information that the branch needs to execute quickly. Using a set of program traces and software BTB simulators, BTB performance data was generated. Prediction rates for infinite size BTBs, using a dynamic prediction method, increased significantly up to four prediction bits, after which the prediction rate increased slowly up to our maximum measured size of 16 prediction bits. Context switching was found to have a significant negative effect on the prediction rates. With a context switch occurring every 10000 instructions, two prediction bits have a prediction rate of 84 percent, as compared to 89 percent

Table# 19: Average Cycle Savings Per Branch, 16384 bit BTB, costx															
BTB#6 Entries			BTB#7 Entries			BTB#8 Entries			BTB#9 Entries			Cycle Savings			
lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	BTB#6	BTB#7	BTB#8	BTB#9
0	512	0	0	256	0	0	224	1024	0	0	8192	1.017	1.044	1.040	0.531
16	448	0	16	224	0	14	196	1024	16	0	7168	1.736	1.761	1.723	1.576
32	384	0	32	192	0	28	168	1024	32	0	6144	1.907	1.922	1.879	1.819
48	320	0	48	160	0	42	140	1024	48	0	5120	1.993	2.004	1.981	1.949
64	256	0	64	128	0	56	112	1024	64	0	4096	2.048	2.055	2.045	2.034
80	192	0	80	96	0	70	84	1024	80	0	3072	2.064	2.066	2.073	2.065
96	128	0	96	64	0	84	56	1024	96	0	2048	2.067	2.076	2.087	2.090
112	64	0	112	32	0	98	28	1024	112	0	1024	2.065	2.083	2.096	2.104
128	0	0	128	0	0	112	0	1024	128	0	0	2.085	2.085	2.104	2.085

Table# 20: Average Cycle Savings Per Branch, 16384 bit BTB, costly															
BTB#6 Entries			BTB#7 Entries			BTB#8 Entries			BTB#9 Entries			Cycle Savings			
lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	lev1	lev2	lev3	BTB#6	BTB#7	BTB#8	BTB#9
0	512	0	0	256	0	0	224	1024	0	0	8192	1.604	1.633	1.644	1.148
21	448	0	21	224	0	18	196	1024	21	0	7168	2.002	2.033	2.031	1.892
42	384	0	42	192	0	37	168	1024	42	0	6144	2.084	2.102	2.113	2.045
64	320	0	64	160	0	56	140	1024	64	0	5120	2.124	2.138	2.159	2.130
85	256	0	85	128	0	74	112	1024	85	0	4096	2.131	2.145	2.175	2.157
106	192	0	106	96	0	93	84	1024	106	0	3072	2.133	2.147	2.182	2.175
128	128	0	128	64	0	112	56	1024	128	0	2048	2.127	2.147	2.188	2.190
149	64	0	149	32	0	130	28	1024	149	0	1024	2.122	2.141	2.192	2.197
170	0	0	170	0	0	149	0	1024	170	0	0	2.138	2.138	2.197	2.138

**Table #19:** Average cycle savings per branch for the four BTB designs (BTB#6-BTB#9) with a total of 16384 storage bits and cost set costx.

**Table #20:** Average cycle savings per branch for the four BTB designs (BTB#6-BTB#9) with a total of 16384 storage bits and cost set costly.

without context switching.

BTB management, when to enter and discard branches from the BTB, was explored. We found that entering branches into the BTB only on a branch taken execution yields better results than entering branches on both taken and not taken executions. Entering branches on taken executions avoids entering into the BTB branches that can't improve performance at the expense of discarding branches that can improve performance. A new method of discarding BTB entries was discussed and compared to the simple LRU replacement strategy. This new method discards branch information that is both not likely to be referenced (i.e. low in the LRU stack), and not likely to be useful (i.e. predicts not taken). This method performed slightly better than simple LRU.

We found a relationship between miss ratios for BTBs and instruction caches. This allowed us to convert miss rates for caches to BTB miss rates and vice versa. Using this capability, design target miss ratios for instruction caches were taken from the literature and converted to



design target miss ratios for BTBs.

Finally, the nature of BTB entries was discussed. The type and amount of information stored in a BTB has a significant impact on the performance of the BTB. Several BTB designs were presented and evaluated, including entries with one or more of the following information types: branch tag, prediction bits, branch target address, and branch target instructions. Varying the amount of individual types of information has a significant impact, as shown by examining the tradeoffs possible between number of prediction bits and number of entries in the BTB. Multi-level BTBs were considered, i.e. BTBs that vary the information content in each level in the BTB. The multi-level BTBs simulated, with specific cycle cost and implementation assumptions, did outperform single level BTBs, but by a small margin. However, depending on the implementation, multi-level BTBs may significantly improve performance.

## **10. Appendix**

Table #A-1: Trace and Workload Information																		
Workload	CPU	Trace Name	Dynamic Instr	Static Instr	Object Instr	Object Bytes	Static Branch instr	Static Branch bytes	Static Cond instr	Static Uncond instr	Branch Taken	Cond Taken	Dyn Blk instr	Dyn Blk bytes	Stat Blk instr	Stat Blk bytes	Obj Blk instr	Obj Blk bytes
COMP	MIPS	asm	2618790	14757	31794	127176	17.31	17.31	10.40	6.92	66.46	55.62	6.27	25.08	5.78	23.12	5.85	23.40
COMP	MIPS	ccom	2347578	16719	49014	196056	21.47	21.47	14.88	6.59	68.24	59.77	5.26	21.04	4.66	18.64	4.73	18.92
COMP	MIPS	foom	2399675	18001	58702	234808	20.94	20.94	13.48	7.46	74.02	68.36	5.07	20.28	4.78	19.12	4.66	18.64
COMP	MIPS	ld	2139419	11201	35950	143800	16.93	16.93	8.73	8.20	75.82	68.57	5.53	22.12	5.91	23.64	5.78	23.12
COMP	MIPS	ugen	2198663	14541	54798	219192	17.50	17.50	8.87	8.63	77.13	67.25	5.76	23.04	5.71	22.84	5.56	22.24
COMP	MIPS	uopt	2005789	8019	82754	331016	13.75	13.75	3.94	9.81	77.44	69.16	5.94	23.76	7.27	29.08	5.25	21.00
COMP	MIPS	upas	2122605	26923	68602	274408	14.40	14.40	7.91	6.50	70.80	62.65	5.74	22.96	6.94	27.76	5.83	23.32
COMP	MIPS	average	2161788	15737	54516	218064	17.47	17.47	9.74	7.73	72.84	64.48	5.65	22.61	5.86	23.46	5.38	21.52
TEXT	MIPS	emacs	2248793	22152	114909	459636	19.95	19.95	13.11	6.84	68.43	63.15	4.44	17.76	5.01	20.04	4.84	19.36
TEXT	MIPS	grep	2659545	4666	14882	59528	18.69	18.69	13.27	5.42	55.77	55.26	5.63	22.52	5.35	21.40	4.85	19.40
TEXT	MIPS	inroff	2524763	7237	17352	69408	19.84	19.84	11.94	7.90	69.72	64.12	5.26	21.04	5.04	20.16	4.79	19.16
TEXT	MIPS	vi	2279164	12728	42006	168024	21.27	21.17	13.55	7.72	77.28	71.24	4.65	18.60	4.70	18.80	4.31	17.24
TEXT	MIPS	average	2428066	11696	48287	193148	19.93	19.93	12.97	6.97	67.80	63.44	5.00	19.98	5.03	20.10	4.70	18.79
FP	MIPS	doduc	2548784	20191	53414	213656	8.35	8.35	5.87	2.47	65.90	59.37	9.08	36.32	11.98	47.92	8.03	32.12
FP	MIPS	integ	2548031	10187	35946	143784	10.05	10.05	6.19	3.86	68.31	51.48	11.55	46.20	9.95	39.80	5.58	22.32
FP	MIPS	imold	2436649	4774	29782	119128	18.98	18.98	12.19	6.79	71.08	61.41	13.60	54.40	5.27	21.08	4.89	19.56
FP	MIPS	spice	2224179	21475	123490	493960	11.79	11.79	7.87	3.92	74.33	65.68	7.30	29.20	8.48	33.92	9.81	39.24
FP	MIPS	average	2439411	14157	60658	242632	12.29	12.29	8.03	4.26	69.91	59.49	10.38	41.53	8.92	35.68	7.08	26.56
SPARC	SPARC	ccom2	2014692	4984	56898	227592	20.89	20.89	10.75	10.13	66.66	49.36	4.76	19.04	4.79	19.16	3.90	15.60
SPARC	SPARC	esp2	1992550	5722	47261	189044	19.82	19.82	11.50	8.32	65.07	54.34	5.03	20.12	5.05	20.20	4.65	18.60
SPARC	SPARC	gbis2	1921300	3875	25546	102184	22.63	22.63	13.06	9.57	46.14	28.85	4.00	20.00	4.42	17.68	4.84	19.36
SPARC	SPARC	average	1976181	4860	43235	172940	21.11	21.11	11.77	9.34	59.29	44.18	4.60	19.72	4.75	19.01	4.46	17.85
VAX	VAX	awk	111806	1084	na	na	25.37	12.57	13.10	12.27	65.15	51.21	3.71	16.46	3.94	16.55	na	na
VAX	VAX	ls	91579	1686	na	na	27.34	15.49	17.44	9.91	66.25	59.29	3.11	9.81	3.66	13.85	na	na
VAX	VAX	otmdl	3044481	1680	na	na	19.88	10.39	12.68	7.20	67.37	59.06	3.41	11.25	5.03	22.19	na	na
VAX	VAX	sedx	123562	883	na	na	32.84	17.42	21.97	10.87	87.19	85.42	2.37	7.55	3.05	12.53	na	na
VAX	VAX	spic	2463599	10064	na	na	22.59	8.92	11.91	10.67	72.23	50.62	4.47	29.44	4.43	26.36	na	na
VAX	VAX	troff	221684	2065	na	na	31.82	14.90	21.89	9.93	67.70	61.10	3.05	14.19	3.14	14.21	na	na
VAX	VAX	ymerge	1262814	2360	na	na	22.54	11.76	13.30	9.24	69.61	61.27	3.51	11.14	4.44	17.49	na	na
VAX	VAX	average	1045646	2832	na	na	26.05	13.06	16.04	10.01	70.79	61.14	3.38	14.26	3.96	17.60	na	na
68K	68K	as	597346	5186	na	na	24.78	20.61	13.11	11.67	62.52	44.56	4.42	12.50	4.04	13.51	na	na
68K	68K	egrep	660000	1410	na	na	20.71	17.03	11.49	9.22	81.98	77.04	4.75	14.99	4.83	16.09	na	na
68K	68K	fortl	916064	5390	na	na	22.06	21.79	9.91	12.15	70.68	58.77	4.78	11.95	4.53	15.22	na	na
68K	68K	fortla	900866	5271	na	na	21.87	20.93	10.13	11.74	70.52	58.58	4.80	12.00	4.57	15.62	na	na
68K	68K	ls	296050	2003	na	na	26.51	21.13	13.78	12.73	66.92	53.32	3.40	9.82	3.77	12.06	na	na
68K	68K	stat	700000	4377	na	na	25.36	22.08	11.08	14.28	67.26	51.99	3.88	11.41	3.94	13.47	na	na
68K	68K	stata	700000	4212	na	na	24.57	21.76	11.68	12.89	66.62	52.09	3.83	11.15	4.07	13.76	na	na
68K	68K	average	681475	3978	na	na	23.69	20.76	11.60	12.10	69.50	56.62	4.27	11.97	4.25	14.25	na	na
Workload Ave			1788761	8877	na	na	20.09	17.44	11.69	8.40	68.36	58.23	5.55	21.68	5.46	21.68	na	na
Trace Ave			1644401	8619	na	na	20.71	17.23	11.90	8.81	69.39	59.73	5.26	20.22	5.27	20.85	na	na
<b>Definitions</b>																		
Dynamic Instr	Number of dynamic instructions in the trace																	
Static Instr	Number of static instructions in the trace																	
Object Instr	Number of instructions in the entire program (object text segment)																	
Object Bytes	Number of bytes in the entire program (object text segment)																	
Static Branch instr	Percentage static branch instructions of all static instructions																	
Static Branch bytes	Percentage bytes of static branch instructions of all static instruction bytes																	
Static Cond instr	Percentage static conditional branch instructions of all static instructions																	
Static Uncond instr	Percentage static unconditional branch instructions of all static instructions																	
Branch Taken	Percentage of branches that are taken																	
Cond Taken	Percentage of conditional branches that are taken																	
Dyn Blk instr	Dynamic basic block size in instructions, including branch instruction																	
Dyn Blk bytes	Dynamic basic block size in bytes, including branch instruction																	
Stat Blk instr	Static basic block size in instructions, including branch instruction																	
Stat Blk bytes	Static basic block size in bytes, including branch instruction																	
Obj Blk instr	Object (text segment) basic block size in instructions, including branch instruction																	
Obj Blk bytes	Object (text segment) basic block size in bytes, including branch instruction																	
Workload Ave	Average of all the workload average values																	
Trace Ave	Average of all of the individual trace values																	

Table #A-1: This table contains the statistics gathered from all 32 traces. Values are presented for individual traces, workload averages, average of workload averages, and the trace average of all 32 traces. Definitions for headings in the columns and rows are given at the bottom of the table.

Table #A-2: Branch History Sequences: Probability of Taken, Probability of Occurrence (For BTBs with 6 pred. bits that enter branches on first execution of branch)														
Bra. History N=Not taken T=Taken	Probability of Next Execution Taken							Probability of Occurrence						
	comp	text	fp	sparc	vax	68k	ave	comp	text	fp	sparc	vax	68k	ave
NNNNNN	.019	.009	.016	.017	.015	.041	.020	.201	.290	.245	.312	.224	.199	.233
NNNNNT	.279	.216	.209	.264	.221	.285	.260	.005	.003	.004	.006	.004	.010	.005
NNNNTN	.180	.070	.092	.115	.170	.121	.133	.004	.003	.004	.006	.003	.008	.005
NNNNTT	.545	.697	.748	.601	.618	.586	.601	.002	.001	.001	.003	.001	.003	.002
NNNTNN	.174	.072	.074	.168	.123	.122	.130	.004	.003	.004	.007	.004	.009	.005
NNNTNT	.257	.246	.169	.271	.266	.271	.260	.001	.000	.001	.001	.001	.002	.001
NNNTTN	.276	.285	.340	.190	.275	.170	.231	.001	.000	.001	.002	.001	.002	.001
NNNTTT	.630	.845	.830	.688	.691	.785	.736	.001	.001	.001	.002	.001	.003	.002
NNTNNN	.233	.114	.099	.258	.293	.193	.212	.004	.003	.004	.007	.004	.009	.005
NNTNNT	.239	.109	.093	.188	.096	.234	.182	.001	.001	.001	.002	.002	.002	.002
NNTNTN	.451	.448	.578	.499	.660	.437	.505	.001	.000	.001	.002	.001	.002	.001
NNTNTT	.405	.585	.615	.492	.460	.322	.412	.001	.000	.000	.001	.000	.001	.001
NNTTNN	.343	.370	.290	.550	.511	.328	.409	.001	.000	.001	.003	.001	.002	.001
NNTTNT	.518	.603	.405	.581	.463	.572	.525	.001	.000	.000	.001	.001	.001	.001
NNTTTN	.427	.377	.298	.283	.436	.293	.361	.001	.000	.000	.001	.001	.001	.001
NNTTTT	.726	.878	.848	.674	.775	.713	.743	.002	.001	.001	.003	.002	.004	.002
NTNNNN	.195	.107	.110	.216	.179	.188	.177	.004	.003	.004	.006	.003	.008	.005
NTNNNT	.203	.179	.230	.169	.204	.236	.211	.001	.000	.001	.002	.002	.002	.002
NTNNTN	.325	.297	.194	.313	.261	.322	.294	.002	.001	.001	.003	.002	.002	.002
NTNNTT	.499	.561	.479	.465	.566	.458	.489	.001	.000	.000	.001	.000	.001	.001
NTNTNN	.366	.342	.373	.510	.492	.405	.423	.001	.000	.001	.002	.001	.002	.001
NTNTNT	.312	.125	.103	.100	.086	.171	.131	.003	.001	.004	.010	.012	.006	.006
NTNTTN	.801	.605	.799	.664	.719	.530	.690	.002	.000	.000	.002	.001	.001	.001
NTNTTT	.524	.710	.615	.589	.777	.702	.659	.001	.000	.000	.002	.001	.002	.001
NTTNNN	.302	.163	.548	.380	.431	.166	.294	.001	.000	.001	.002	.001	.002	.001
NTTNTN	.523	.601	.711	.794	.474	.419	.563	.001	.000	.000	.002	.001	.001	.001
NTTNTT	.723	.649	.651	.795	.804	.546	.707	.001	.000	.001	.001	.001	.001	.001
NTTTNN	.393	.446	.017	.315	.477	.362	.230	.003	.001	.013	.004	.002	.004	.004
NTTTNT	.428	.343	.648	.378	.403	.307	.389	.001	.000	.000	.001	.001	.001	.001
NTTTTN	.856	.871	.807	.784	.828	.654	.803	.004	.001	.001	.003	.002	.002	.002
NTTTTT	.801	.877	.816	.641	.845	.630	.751	.002	.001	.001	.002	.002	.002	.002
TNNNNN	.915	.879	.938	.835	.928	.872	.899	.012	.008	.006	.008	.011	.012	.010
TNNNNT	.192	.121	.091	.177	.142	.166	.158	.005	.003	.004	.006	.004	.010	.005
TNNNTN	.398	.294	.264	.413	.358	.215	.319	.001	.000	.001	.002	.001	.002	.001
TNNNTT	.249	.135	.327	.231	.214	.293	.252	.001	.000	.001	.002	.002	.002	.001
TNNTNN	.575	.653	.602	.497	.465	.756	.610	.001	.000	.001	.001	.001	.001	.001
TNNTNT	.486	.630	.763	.483	.611	.547	.564	.001	.000	.001	.002	.002	.002	.002
TNNTTN	.461	.414	.271	.356	.271	.381	.371	.001	.000	.000	.001	.001	.001	.001
TNNTTT	.480	.301	.450	.264	.464	.415	.396	.001	.000	.000	.002	.001	.001	.001
TNTNNN	.600	.871	.731	.739	.792	.829	.758	.001	.001	.001	.001	.001	.002	.001
TNTNTN	.321	.194	.242	.391	.424	.433	.377	.001	.000	.001	.001	.001	.002	.001

Table #A-2 (cont.): Branch History Sequences: Probability of Taken, Probability of Occurrence (For BTBs with 6 pred. bits that enter branches on first execution of branch)														
Bra. History N=Not taken T=Taken	Probability of Next Execution Taken							Probability of Occurrence						
	comp	text	fp	sparc	vax	68k	ave	comp	text	fp	sparc	vax	68k	ave
TNTNNT	.398	.369	.378	.337	.261	.333	.340	.001	.000	.000	.001	.001	.001	.001
TNTNTN	.806	.911	.891	.918	.934	.855	.898	.003	.001	.004	.010	.011	.006	.006
TNTNTT	.429	.790	.465	.525	.528	.684	.551	.002	.000	.001	.002	.002	.002	.002
TNTTNN	.633	.577	.646	.677	.462	.513	.569	.001	.000	.000	.001	.001	.001	.001
TNTTNT	.693	.881	.974	.767	.625	.828	.833	.004	.001	.013	.005	.002	.004	.004
TNTTTN	.915	.815	.874	.865	.866	.824	.875	.004	.001	.001	.003	.002	.002	.002
TNTTTT	.880	.919	.866	.819	.822	.868	.861	.013	.007	.005	.008	.011	.011	.010
TTNNTN	.260	.136	.137	.358	.310	.141	.219	.002	.001	.001	.003	.001	.003	.002
TTNNTT	.578	.799	.816	.650	.680	.753	.689	.001	.000	.001	.001	.001	.001	.001
TTNTNN	.561	.638	.425	.559	.505	.396	.490	.001	.000	.000	.001	.001	.001	.001
TTNTNT	.614	.765	.648	.358	.678	.728	.619	.001	.001	.001	.003	.001	.002	.001
TTNTTN	.569	.279	.588	.628	.297	.303	.429	.001	.000	.000	.001	.000	.001	.001
TTNTNT	.591	.623	.335	.612	.413	.560	.533	.002	.000	.001	.002	.002	.002	.002
TTNTTN	.834	.906	.980	.878	.842	.892	.914	.003	.001	.013	.004	.002	.004	.004
TTNTTT	.782	.896	.875	.749	.889	.880	.842	.016	.008	.006	.009	.012	.011	.011
TTTNNN	.343	.208	.262	.204	.348	.271	.284	.001	.001	.001	.002	.001	.003	.002
TTTNTT	.686	.855	.773	.716	.878	.796	.787	.001	.001	.000	.001	.001	.002	.001
TTTNTN	.675	.610	.784	.772	.786	.743	.733	.001	.001	.000	.002	.001	.002	.001
TTTNTT	.930	.953	.937	.862	.934	.879	.917	.016	.008	.006	.009	.012	.011	.011
TTTTNN	.418	.516	.168	.364	.564	.387	.418	.001	.001	.001	.002	.002	.003	.002
TTTTNT	.943	.947	.948	.884	.941	.871	.921	.013	.008	.006	.008	.011	.012	.010
TTTTTN	.911	.904	.839	.802	.867	.816	.862	.012	.008	.006	.008	.011	.012	.010
TTTTTT	.981	.989	.991	.986	.984	.981	.984	.616	.624	.622	.485	.609	.575	.595
NNNNN	.022	.010	.017	.020	.017	.046	.023	.205	.293	.249	.318	.228	.209	.238
NNNNT	.305	.225	.215	.303	.249	.273	.271	.006	.003	.005	.009	.005	.012	.007
NNNTN	.198	.078	.119	.147	.183	.156	.160	.006	.003	.004	.008	.005	.011	.006
NNNTT	.555	.688	.697	.570	.555	.637	.604	.003	.001	.002	.004	.002	.005	.003
NNTNN	.252	.148	.242	.243	.266	.210	.231	.006	.003	.005	.009	.006	.011	.007
NNTNT	.353	.333	.212	.315	.269	.316	.310	.002	.001	.001	.003	.002	.003	.002
NNTTN	.362	.293	.387	.233	.351	.260	.302	.002	.001	.001	.004	.001	.003	.002
NNTTT	.618	.856	.799	.709	.743	.800	.745	.002	.001	.002	.004	.002	.005	.003
NTNNN	.251	.124	.118	.275	.316	.234	.240	.006	.003	.004	.008	.005	.011	.006
NTNNT	.300	.180	.170	.242	.150	.264	.234	.002	.001	.002	.004	.002	.003	.002
NTNTN	.688	.822	.848	.855	.907	.747	.828	.004	.002	.005	.011	.013	.007	.007
NTNTT	.423	.732	.500	.517	.513	.576	.515	.003	.001	.001	.003	.002	.003	.002
NTTNN	.455	.431	.434	.583	.493	.393	.464	.002	.001	.001	.004	.002	.003	.002
NTTNT	.664	.840	.960	.737	.593	.791	.798	.004	.001	.013	.005	.003	.005	.005
NTTTN	.820	.748	.717	.718	.749	.651	.749	.005	.001	.001	.004	.002	.003	.003
NTTTT	.864	.914	.863	.785	.817	.830	.841	.014	.008	.007	.010	.013	.015	.012
TNNNN	.213	.113	.115	.262	.209	.173	.188	.006	.003	.005	.009	.005	.012	.007
TNNNT	.340	.365	.528	.328	.348	.384	.367	.002	.001	.001	.003	.002	.003	.002

Table #A-2 (cont.): Branch History Sequences: Probability of Taken, Probability of Occurrence (For BTBs with six pred. bits that enter branches on first execution of branch)														
Bra. History N=Not taken T=Taken	Probability of Next Execution Taken							Probability of Occurrence						
	comp	text	fp	sparc	vax	68k	ave	comp	text	fp	sparc	vax	68k	ave
TNNTN	.397	.377	.224	.373	.316	.343	.344	.002	.001	.002	.004	.002	.003	.002
TNNTT	.570	.731	.593	.384	.656	.634	.581	.002	.001	.001	.004	.002	.003	.002
TNTNN	.447	.315	.425	.543	.445	.370	.425	.002	.001	.001	.002	.002	.003	.002
TNTNT	.426	.239	.138	.199	.127	.277	.216	.005	.002	.005	.012	.013	.008	.008
TNTTN	.822	.857	.974	.819	.792	.799	.866	.004	.001	.013	.006	.003	.005	.005
TNTTT	.764	.886	.857	.725	.879	.856	.825	.017	.008	.006	.011	.013	.013	.012
TTNNN	.325	.194	.352	.278	.380	.229	.287	.002	.001	.002	.004	.002	.004	.003
TTNNT	.608	.778	.741	.765	.723	.634	.683	.002	.001	.001	.004	.002	.003	.002
TTNTN	.700	.622	.711	.783	.796	.681	.721	.003	.001	.001	.003	.002	.003	.002
TTNTT	.846	.904	.305	.697	.881	.750	.738	.018	.009	.019	.013	.013	.015	.015
TTTNN	.421	.471	.281	.368	.527	.368	.409	.002	.001	.001	.004	.002	.004	.003
TTTNT	.922	.939	.929	.857	.927	.842	.900	.017	.008	.006	.011	.013	.013	.012
TTTTN	.895	.901	.835	.767	.863	.784	.844	.014	.008	.007	.010	.013	.015	.012
TTTTT	.980	.988	.991	.983	.983	.979	.983	.629	.632	.628	.493	.620	.587	.605
NNNN	.028	.012	.019	.027	.021	.053	.027	.211	.296	.253	.327	.233	.221	.245
NNNT	.315	.246	.272	.310	.282	.299	.296	.008	.004	.006	.012	.007	.015	.009
NNTN	.259	.136	.148	.214	.227	.203	.212	.008	.004	.006	.012	.007	.014	.009
NNTT	.562	.706	.661	.478	.603	.636	.595	.004	.002	.002	.007	.004	.007	.004
NTNN	.303	.174	.273	.303	.305	.243	.273	.008	.004	.006	.011	.007	.014	.009
NTNT	.403	.260	.150	.219	.143	.287	.235	.007	.002	.006	.015	.015	.011	.009
NTTN	.682	.671	.940	.581	.631	.615	.713	.006	.002	.014	.009	.004	.008	.007
NTTT	.746	.881	.845	.722	.859	.841	.811	.019	.009	.008	.015	.015	.017	.015
TNNN	.273	.140	.179	.276	.331	.232	.253	.008	.004	.006	.012	.007	.015	.009
TNNT	.436	.498	.352	.508	.419	.427	.437	.004	.001	.002	.007	.004	.006	.004
TNTN	.693	.766	.821	.839	.891	.729	.803	.007	.002	.006	.014	.015	.011	.009
TNTT	.791	.892	.313	.662	.830	.721	.709	.021	.009	.020	.017	.015	.018	.017
TTNN	.436	.458	.338	.481	.513	.378	.432	.004	.002	.002	.008	.004	.007	.004
TTNT	.869	.928	.950	.817	.870	.829	.871	.021	.009	.020	.016	.015	.018	.017
TTTN	.876	.883	.816	.753	.847	.763	.826	.019	.009	.008	.014	.015	.017	.015
TTTT	.977	.987	.989	.979	.979	.976	.980	.643	.640	.635	.504	.633	.602	.617
NNN	.037	.013	.023	.035	.030	.064	.035	.219	.300	.259	.339	.239	.235	.254
NNT	.357	.312	.295	.384	.333	.336	.342	.012	.005	.008	.019	.011	.021	.013
NTN	.460	.386	.482	.559	.666	.430	.518	.015	.006	.012	.026	.022	.025	.018
NTT	.752	.863	.353	.606	.786	.697	.685	.025	.011	.022	.024	.019	.025	.022
TNN	.350	.259	.290	.373	.376	.288	.327	.012	.005	.008	.019	.011	.021	.013
TNT	.755	.793	.772	.535	.513	.627	.645	.028	.012	.025	.031	.030	.028	.026
TTN	.827	.853	.896	.685	.800	.718	.790	.025	.011	.022	.024	.019	.025	.022
TTT	.971	.985	.987	.972	.977	.972	.976	.663	.650	.643	.518	.648	.620	.632
NN	.053	.018	.031	.053	.045	.082	.050	.232	.305	.267	.358	.251	.256	.267
NT	.632	.641	.655	.478	.465	.503	.544	.040	.017	.034	.050	.041	.049	.040
TN	.692	.685	.752	.619	.727	.575	.665	.040	.017	.034	.050	.041	.049	.040

<b>Table #A-2 (cont.): Branch History Sequences: Probability of Taken, Probability of Occurrence (For BTBs with six pred. that enter branches on first execution of branch)</b>														
<b>Bra. History</b>	<b>Probability of Next Execution Taken</b>							<b>Probability of Occurrence</b>						
<b>N=Not taken T=Taken</b>	<b>comp</b>	<b>text</b>	<b>fp</b>	<b>sparc</b>	<b>vax</b>	<b>68k</b>	<b>ave</b>	<b>comp</b>	<b>text</b>	<b>fp</b>	<b>sparc</b>	<b>vax</b>	<b>68k</b>	<b>ave</b>
<b>TT</b>	.963	.983	.966	.956	.971	.961	.967	.688	.661	.665	.543	.667	.645	.654
<b>N</b>	.147	.053	.112	.123	.141	.162	.129	.272	.322	.301	.407	.292	.305	.306
<b>T</b>	.945	.975	.951	.916	.942	.929	.943	.728	.678	.699	.593	.708	.695	.694

**Table #A-2:** This table contains the probability of occurrence and probability of next branch execution taken statistics, as a function of up to the past six branch executions. These statistics can be used by BTBs that enter branches on both taken and not taken executions. Statistics are gathered beginning with the first execution of each static branch.

Table #A-3: Branch History Sequences: Probability of Taken, Probability of Occurrence (For BTBs with 6 pred. bits that enter branches on first taken execution of branch)														
Bra. History N=Not taken T=Taken	Probability of Next Execution Taken							Probability of Occurrence						
	comp	text	fp	sparc	vax	68k	ave	comp	text	fp	sparc	vax	68k	ave
NNNNNN	.059	.018	.075	.033	.031	.096	.048	.071	.169	.062	.185	.112	.094	.107
NNNNNT	.279	.206	.208	.258	.214	.283	.256	.005	.003	.005	.007	.004	.011	.006
NNNNTN	.180	.067	.090	.114	.174	.120	.132	.005	.003	.005	.007	.004	.010	.006
NNNNTT	.538	.691	.754	.588	.616	.589	.599	.002	.001	.001	.003	.001	.003	.002
NNNTNN	.173	.069	.072	.168	.125	.120	.129	.005	.003	.005	.008	.005	.010	.006
NNNTNT	.257	.248	.168	.270	.268	.266	.258	.001	.000	.001	.001	.001	.002	.001
NNNTTN	.276	.273	.345	.185	.285	.173	.234	.001	.000	.001	.002	.001	.002	.001
NNNTTT	.625	.842	.835	.680	.684	.789	.735	.002	.001	.001	.002	.001	.003	.002
NNTNNN	.236	.114	.095	.257	.296	.194	.213	.005	.003	.005	.008	.005	.010	.006
NNTNNT	.240	.108	.094	.190	.099	.234	.183	.002	.001	.001	.003	.002	.003	.002
NNTNTN	.449	.448	.580	.497	.664	.434	.504	.002	.000	.001	.002	.001	.002	.001
NNTNTT	.408	.583	.614	.497	.450	.322	.413	.001	.000	.000	.001	.001	.001	.001
NNTTNN	.343	.375	.284	.543	.532	.327	.410	.001	.000	.001	.003	.001	.002	.001
NNTTNT	.519	.603	.390	.576	.456	.569	.520	.001	.000	.000	.001	.001	.001	.001
NNTTTN	.427	.383	.290	.280	.447	.286	.362	.001	.000	.000	.001	.001	.001	.001
NNTTTT	.724	.875	.859	.672	.768	.711	.742	.002	.001	.002	.003	.002	.004	.002
NTNNNN	.196	.107	.110	.215	.185	.186	.178	.005	.003	.005	.007	.004	.009	.006
NTNNNT	.204	.174	.238	.167	.204	.230	.208	.002	.000	.001	.003	.002	.003	.002
NTNNTN	.325	.301	.188	.313	.263	.321	.294	.002	.001	.002	.003	.002	.003	.002
NTNNTT	.500	.560	.468	.471	.575	.459	.491	.001	.000	.000	.001	.000	.001	.001
NTNTNN	.366	.353	.369	.510	.487	.401	.421	.001	.000	.001	.002	.001	.002	.001
NTNTNT	.313	.124	.103	.101	.087	.168	.131	.003	.002	.005	.011	.013	.006	.007
NTNTTN	.800	.604	.794	.662	.719	.536	.692	.002	.000	.001	.002	.001	.001	.001
NTNTTT	.524	.706	.610	.581	.774	.699	.655	.001	.001	.001	.002	.001	.002	.001
NTTNNN	.305	.164	.570	.376	.461	.166	.302	.001	.000	.001	.002	.001	.002	.001
NTTNTN	.524	.607	.720	.793	.469	.418	.564	.001	.000	.001	.003	.001	.001	.001
NTTNTT	.721	.646	.647	.795	.806	.550	.707	.002	.000	.001	.002	.001	.001	.001
NTTTNN	.396	.439	.016	.309	.483	.360	.221	.003	.001	.017	.005	.002	.004	.005
NTTTNT	.431	.350	.677	.379	.414	.307	.396	.001	.000	.000	.001	.001	.001	.001
NTTTTN	.855	.872	.807	.781	.828	.653	.802	.005	.001	.001	.003	.002	.002	.002
NTTTTT	.803	.882	.818	.649	.846	.629	.753	.002	.001	.001	.003	.003	.003	.002
TNNNNN	.914	.880	.939	.834	.926	.872	.898	.014	.009	.007	.009	.012	.014	.012
TNNNNNT	.192	.118	.089	.177	.144	.167	.158	.005	.004	.005	.008	.004	.011	.006
TNNNTN	.400	.290	.263	.410	.357	.213	.318	.001	.000	.001	.003	.001	.002	.001
TNNNTT	.250	.132	.335	.232	.218	.290	.253	.002	.000	.001	.003	.002	.002	.002
TNNTNN	.574	.651	.609	.493	.464	.755	.607	.001	.000	.001	.001	.001	.001	.001
TNNTNT	.484	.629	.775	.484	.604	.552	.566	.002	.001	.002	.003	.002	.003	.002
TNNTTN	.462	.415	.274	.353	.271	.377	.370	.001	.000	.000	.002	.001	.001	.001
TNNTTT	.481	.298	.454	.263	.468	.420	.398	.001	.000	.000	.003	.001	.001	.001
TNTNNN	.602	.871	.726	.742	.787	.826	.754	.001	.001	.001	.002	.001	.002	.001
TNTNTN	.322	.197	.239	.387	.427	.430	.375	.001	.000	.001	.001	.001	.002	.001

Table #A-3 (cont.): Branch History Sequences: Probability of Taken, Probability of Occurrence (For BTBs with 6 pred. bits that enter branches on first taken execution of branch)														
Bra. History N=Not taken T=Taken	Probability of Next Execution Taken							Probability of Occurrence						
	comp	text	fp	sparc	vax	68k	ave	comp	text	fp	sparc	vax	68k	ave
TNTNNT	.401	.365	.392	.336	.265	.327	.342	.001	.000	.001	.001	.001	.001	.001
TNTNTN	.805	.911	.891	.918	.933	.855	.897	.003	.002	.005	.011	.013	.006	.007
TNTNTI	.434	.787	.459	.528	.528	.684	.551	.002	.001	.001	.003	.002	.002	.002
TNTTNN	.630	.581	.656	.673	.457	.512	.568	.001	.000	.000	.001	.001	.001	.001
TNTTNT	.694	.882	.974	.771	.625	.831	.839	.004	.001	.017	.005	.002	.005	.005
TNTTTN	.914	.816	.872	.861	.866	.826	.875	.005	.001	.001	.003	.002	.002	.002
TNTTTI	.877	.921	.871	.818	.822	.869	.861	.015	.009	.007	.009	.013	.013	.012
TTNNTN	.259	.134	.132	.359	.306	.140	.219	.002	.001	.001	.003	.001	.004	.002
TTNNTI	.576	.798	.823	.647	.678	.750	.687	.001	.000	.001	.001	.001	.001	.001
TTNTNN	.558	.638	.432	.561	.507	.398	.493	.001	.000	.000	.001	.001	.001	.001
TTNTNT	.617	.759	.655	.359	.679	.728	.620	.001	.001	.001	.003	.002	.002	.002
TTNTTN	.567	.283	.600	.628	.300	.307	.434	.001	.000	.000	.001	.001	.001	.001
TTNTTI	.590	.619	.329	.612	.416	.568	.535	.002	.001	.001	.003	.002	.002	.002
TTNTTN	.833	.907	.980	.880	.842	.896	.918	.003	.001	.017	.005	.002	.004	.005
TTNTTI	.780	.899	.878	.754	.887	.879	.842	.018	.009	.007	.011	.013	.013	.013
TTTNNN	.343	.208	.262	.202	.350	.265	.282	.002	.001	.001	.003	.001	.003	.002
TTTNTN	.687	.851	.775	.710	.881	.794	.786	.001	.001	.001	.002	.001	.002	.001
TTTNTI	.674	.618	.783	.771	.786	.746	.734	.002	.001	.001	.002	.001	.002	.001
TTTNTN	.929	.954	.935	.866	.932	.878	.916	.018	.009	.007	.011	.013	.013	.013
TTTNTI	.421	.509	.159	.366	.567	.382	.414	.002	.001	.001	.003	.002	.004	.002
TTTNTN	.941	.949	.950	.885	.940	.873	.921	.015	.009	.007	.009	.013	.013	.012
TTTNTI	.908	.907	.836	.805	.865	.817	.861	.014	.009	.007	.009	.012	.014	.012
TTTTNN	.982	.989	.991	.986	.984	.981	.985	.718	.731	.774	.579	.698	.650	.694
TTNNNN	.213	.111	.114	.260	.210	.173	.188	.008	.005	.007	.014	.007	.017	.010
TTNNNT	.340	.363	.545	.324	.352	.375	.365	.003	.001	.002	.005	.003	.005	.003
TTNNTN	.397	.381	.220	.373	.318	.341	.344	.003	.001	.002	.006	.003	.005	.004
TTNNTI	.572	.726	.596	.386	.658	.633	.582	.003	.001	.001	.006	.003	.004	.003
TNTNNT	.446	.319	.428	.542	.442	.370	.425	.003	.001	.001	.004	.002	.004	.003
TNTNTN	.425	.238	.137	.200	.129	.275	.217	.007	.003	.007	.020	.018	.011	.011
TNTNTI	.821	.857	.974	.820	.791	.805	.871	.006	.001	.020	.009	.004	.007	.007
TNTTTN	.762	.889	.860	.728	.877	.855	.824	.023	.012	.009	.018	.018	.019	.017
TTNNTN	.324	.192	.356	.277	.382	.225	.288	.003	.001	.002	.006	.003	.006	.004
TTNNTI	.609	.773	.745	.762	.723	.633	.683	.003	.001	.001	.006	.003	.004	.003
TTNTNN	.698	.628	.705	.782	.797	.683	.722	.004	.001	.001	.005	.003	.004	.003
TTNTNT	.845	.904	.297	.698	.878	.748	.729	.025	.013	.027	.022	.019	.022	.021
TTTNNN	.423	.468	.280	.370	.530	.364	.408	.003	.002	.002	.006	.003	.006	.004
TTTNTN	.920	.941	.932	.857	.925	.843	.900	.023	.012	.009	.018	.018	.019	.017
TTTNTI	.893	.904	.833	.771	.861	.786	.844	.019	.012	.009	.017	.018	.021	.017
TTTTNN	.980	.988	.991	.984	.983	.979	.983	.864	.933	.899	.837	.875	.845	.874
TTNNNT	.274	.139	.179	.274	.332	.230	.252	.011	.006	.008	.020	.010	.022	.013
TTNNTI	.438	.494	.357	.504	.422	.422	.436	.006	.002	.003	.012	.006	.009	.006



Table #A-3 (cont.): Branch History Sequences: Probability of Taken, Probability of Occurrence (For BTBs with six pred. that enter branches on first taken execution of branch)														
Bra. History N=Not taken T=Taken	Probability of Next Execution Taken							Probability of Occurrence						
	comp	text	fp	sparc	vax	68k	ave	comp	text	fp	sparc	vax	68k	ave
TNTN	.692	.768	.819	.838	.890	.729	.802	.009	.003	.008	.024	.021	.015	.013
TNTT	.790	.893	.305	.664	.827	.720	.702	.029	.014	.028	.027	.022	.026	.025
TTNN	.435	.455	.338	.478	.514	.375	.431	.006	.002	.003	.012	.005	.010	.006
TTNT	.868	.930	.952	.818	.868	.830	.872	.029	.014	.029	.027	.022	.026	.025
TTTN	.873	.885	.813	.755	.845	.765	.825	.026	.014	.011	.024	.021	.025	.021
TTTT	.978	.987	.989	.980	.979	.976	.981	.884	.945	.909	.854	.894	.867	.891
TNN	.349	.256	.294	.370	.375	.285	.326	.017	.008	.012	.032	.016	.030	.019
TNT	.753	.796	.779	.536	.513	.627	.647	.038	.017	.036	.051	.042	.041	.038
TTN	.824	.855	.896	.686	.798	.722	.791	.035	.016	.032	.039	.027	.036	.031
TTT	.971	.985	.988	.973	.977	.972	.977	.910	.959	.920	.878	.915	.892	.912
TN	.689	.686	.755	.617	.726	.575	.666	.055	.025	.048	.082	.058	.072	.057
TT	.963	.983	.966	.957	.971	.961	.967	.945	.975	.952	.918	.942	.928	.943
T	.928	.975	.951	.918	.942	.928	.943	1.000	1.000	1.000	1.000	1.000	1.000	1.000

**Table #A-3:** This table contains the probability of occurrence and probability of next branch execution taken statistics, as a function of up to the past six branch executions. These statistics can be used by BTBs that enter branches on taken branch executions only. Statistics are gathered beginning with the first taken execution of each static branch. Partial history strings (less than six) all start with a taken execution, as is expected.

Table #A-4: Branch Prediction Rates (including target change effects)								
Prediction Bits	BTB Design	COMP	TEXT	FP	SPARC	VAX	68k	Average
0	BTB#1	67.21	66.13	61.78	56.80	68.85	64.96	64.29
0	BTB#2	81.27	80.86	81.04	73.04	81.54	76.68	79.07
1	BTB#1	86.30	94.83	85.03	87.51	89.71	85.55	88.15
1	BTB#2	86.30	94.83	85.03	87.51	89.71	85.55	88.15
2	BTB#1	87.81	95.45	86.72	88.69	91.56	86.27	89.42
2	BTB#2	87.81	95.45	86.72	88.69	91.56	86.27	89.42
3	BTB#1	88.16	95.65	87.05	89.13	91.93	86.95	89.81
3	BTB#2	88.16	95.65	87.05	89.13	91.93	86.95	89.81
4	BTB#1	88.66	96.02	87.85	90.61	93.36	88.21	90.79
4	BTB#2	88.66	96.02	87.84	90.61	93.35	88.20	90.78
5	BTB#1	88.76	96.08	87.90	91.00	93.49	88.41	90.94
5	BTB#2	88.75	96.08	87.89	91.00	93.48	88.39	90.93
6	BTB#1	88.90	96.12	89.22	91.31	93.54	88.68	91.30
6	BTB#2	88.90	96.12	89.21	91.31	93.52	88.66	91.29
7	BTB#1	88.98	96.15	89.22	91.63	93.65	88.81	91.41
7	BTB#2	88.96	96.14	89.23	91.65	93.63	88.79	91.40
8	BTB#1	89.19	96.18	89.28	91.78	93.70	89.00	91.52
8	BTB#2	89.19	96.18	89.27	91.77	93.67	88.97	91.51
9	BTB#1	89.29	96.19	89.33	91.85	93.83	89.13	91.60
9	BTB#2	89.28	96.19	89.32	91.86	93.81	89.11	91.59
10	BTB#1	89.37	96.22	89.41	91.94	93.98	89.32	91.71
10	BTB#2	89.35	96.22	89.40	91.93	93.97	89.30	91.69
11	BTB#1	89.44	96.24	89.46	92.03	94.13	89.58	91.81
11	BTB#2	89.43	96.23	89.46	92.04	94.11	89.54	91.80
12	BTB#1	89.51	96.25	89.51	92.13	94.34	89.96	91.95
12	BTB#2	89.50	96.24	89.50	92.13	94.32	89.93	91.94
13	BTB#1	89.59	96.26	89.55	92.29	94.57	90.31	92.10
13	BTB#2	89.58	96.26	89.54	92.29	94.55	90.27	92.08
14	BTB#1	89.72	96.30	89.64	92.44	94.80	90.72	92.27
14	BTB#2	89.72	96.30	89.63	92.43	94.78	90.67	92.25
15	BTB#1	89.90	96.33	89.73	92.60	95.06	91.03	92.44
15	BTB#2	89.89	96.32	89.72	92.59	95.04	90.99	92.43
16	BTB#1	90.14	96.40	89.90	92.93	95.25	91.48	92.68
16	BTB#2	90.13	96.39	89.89	92.91	95.22	91.45	92.67

Table #A-4: Prediction rates for infinite size BTBs, BTB#1 and BTB#2, as a function of prediction bits and workload.

<b>Table #A-5: Effect of Context Switching on Prediction Rates</b>							
<b>Thousands of Instructions per Context Switch</b>							
<b>Prediction Bits</b>	<b>1K</b>	<b>3.16K</b>	<b>10K</b>	<b>31.6K</b>	<b>100K</b>	<b>316K</b>	<b>1000K</b>
0	71.40	77.28	80.63	82.09	81.93	81.12	79.96
1	71.78	78.43	83.01	85.76	87.12	87.75	88.04
2	72.44	79.36	84.11	86.94	88.35	89.00	89.30
3	72.50	79.53	84.36	87.27	88.71	89.39	89.69
4	72.65	79.92	84.99	88.06	89.60	90.33	90.65
5	72.68	79.98	85.08	88.19	89.74	90.47	90.80
6	72.81	80.20	85.36	88.51	90.09	90.83	91.15
7	72.82	80.23	85.43	88.60	90.19	90.93	91.27
8	72.83	80.25	85.48	88.68	90.29	91.04	91.37
9	72.84	80.28	85.54	88.75	90.37	91.12	91.46
10	72.87	80.32	85.58	88.82	90.45	91.22	91.56
11	72.88	80.34	85.63	88.89	90.55	91.32	91.66
12	72.89	80.36	85.70	89.00	90.67	91.45	91.80
13	72.91	80.39	85.75	89.10	90.80	91.59	91.94
14	72.92	80.40	85.81	89.22	90.94	91.75	92.11
15	72.93	80.42	85.86	89.31	91.07	91.91	92.28
16	72.94	80.44	85.92	89.44	91.26	92.13	92.51

**Table #A-5:** Table presents the effect of context switching on the prediction rate of infinite size BTB#2. The values are averages of the workload average prediction rates.

Table #A-6: Prediction Rates for Designs BTB#3, BTB#4, and BTB#5										
Prediction Bits	BTB Design	Set Size	Number of BTB Entries							
			16	32	64	128	256	512	1024	2048
0	BTB#3	1	46.26	50.10	53.90	56.91	59.46	61.41	62.66	63.51
0	BTB#4	1	61.00	67.70	71.98	75.10	76.77	78.00	78.90	78.97
0	BTB#5	1	61.00	67.70	71.98	75.10	76.77	78.00	78.90	78.97
0	BTB#3	2	46.76	50.64	54.95	58.13	60.80	62.66	63.62	64.06
0	BTB#4	2	62.42	68.76	73.70	76.38	78.14	78.86	79.02	79.03
0	BTB#5	2	62.42	68.76	73.70	76.38	78.14	78.86	79.02	79.03
0	BTB#3	4	47.60	50.63	55.58	58.97	61.46	63.18	63.92	64.18
0	BTB#4	4	62.72	69.47	74.68	77.10	78.32	78.96	79.00	79.05
0	BTB#5	4	62.72	69.47	74.68	77.10	78.32	78.96	79.00	79.05
0	BTB#3	8	46.96	51.16	55.95	59.27	61.72	63.44	64.00	64.23
0	BTB#4	8	62.56	69.80	75.33	77.19	78.54	78.88	79.04	79.06
0	BTB#5	8	62.56	69.80	75.33	77.19	78.54	78.88	79.04	79.06
0	BTB#3	16	46.79	51.22	56.04	59.41	61.81	63.55	64.04	64.25
0	BTB#4	16	63.43	69.64	75.42	77.18	78.60	78.91	79.06	79.09
0	BTB#5	16	63.43	69.64	75.42	77.18	78.60	78.91	79.06	79.09
1	BTB#3	1	58.11	65.84	72.01	77.14	81.02	84.07	85.92	87.04
1	BTB#4	1	61.99	69.11	74.69	79.02	82.36	84.85	86.34	87.27
1	BTB#5	1	61.99	69.11	74.69	79.02	82.36	84.85	86.34	87.27
1	BTB#3	2	58.54	67.34	74.28	79.39	83.22	85.85	87.21	87.85
1	BTB#4	2	63.32	70.51	76.75	81.05	84.41	86.46	87.51	87.96
1	BTB#5	2	63.47	70.69	76.90	81.05	84.41	86.46	87.51	87.96
1	BTB#3	4	59.39	67.68	75.31	80.67	84.20	86.58	87.63	88.01
1	BTB#4	4	63.61	71.27	77.92	82.11	85.28	87.10	87.83	88.08
1	BTB#5	4	63.92	71.50	78.14	82.28	85.35	87.13	87.84	88.08
1	BTB#3	8	59.72	68.32	75.71	81.13	84.57	86.94	87.75	88.08
1	BTB#4	8	63.79	71.73	78.63	82.56	85.74	87.36	87.92	88.12
1	BTB#5	8	64.39	72.07	78.93	82.72	85.91	87.41	87.92	88.12
1	BTB#3	16	61.49	68.60	75.97	81.37	84.71	87.12	87.81	88.10
1	BTB#4	16	64.72	71.53	78.89	82.63	85.98	87.46	87.96	88.14
1	BTB#5	16	64.90	71.92	79.23	82.80	86.17	87.49	87.97	88.14
2	BTB#3	1	58.41	66.43	72.75	77.99	82.08	85.22	87.13	88.29
2	BTB#4	1	62.34	69.79	75.52	79.95	83.43	86.01	87.55	88.50
2	BTB#5	1	62.34	69.79	75.52	79.95	83.43	86.01	87.55	88.50
2	BTB#3	2	58.87	67.90	75.12	80.37	84.33	87.05	88.45	89.10
2	BTB#4	2	63.72	71.22	77.70	82.11	85.56	87.68	88.76	89.21
2	BTB#5	2	63.72	71.23	77.72	82.15	85.59	87.68	88.76	89.21
2	BTB#3	4	59.74	68.20	76.22	81.72	85.36	87.81	88.88	89.27
2	BTB#4	4	64.01	72.03	78.88	83.21	86.47	88.33	89.09	89.33
2	BTB#5	4	64.17	72.11	79.02	83.26	86.52	88.35	89.09	89.34
2	BTB#3	8	59.98	68.92	76.60	82.18	85.73	88.17	89.00	89.34
2	BTB#4	8	64.18	72.47	79.58	83.67	86.94	88.61	89.18	89.38
2	BTB#5	8	64.46	72.59	79.77	83.77	87.05	88.64	89.18	89.38
2	BTB#3	16	61.71	69.22	76.88	82.43	85.87	88.35	89.07	89.37
2	BTB#4	16	65.19	72.29	79.86	83.73	87.18	88.71	89.22	89.40
2	BTB#5	16	65.40	72.72	80.25	83.88	87.33	88.74	89.23	89.40

Table #A-6 (cont.): Prediction Rates for Designs BTB#3, BTB#4, and BTB#5										
Prediction Bits	BTB Design	Set Size	Number of BTB Entries							
			16	32	64	128	256	512	1024	2048
3	BTB#3	1	58.45	66.51	72.89	78.19	82.35	85.54	87.47	88.66
3	BTB#4	1	62.37	69.88	75.68	80.18	83.72	86.34	87.92	88.89
3	BTB#5	1	62.37	69.88	75.68	80.18	83.72	86.34	87.92	88.89
3	BTB#3	2	58.92	67.98	75.28	80.60	84.64	87.41	88.83	89.49
3	BTB#4	2	63.76	71.32	77.89	82.36	85.89	88.05	89.14	89.61
3	BTB#5	2	63.78	71.32	77.90	82.39	85.91	88.06	89.14	89.61
3	BTB#3	4	59.78	68.28	76.38	81.97	85.68	88.18	89.27	89.67
3	BTB#4	4	64.05	72.13	79.06	83.49	86.81	88.71	89.48	89.73
3	BTB#5	4	64.12	72.22	79.20	83.57	86.88	88.74	89.48	89.73
3	BTB#3	8	60.02	69.02	76.75	82.43	86.07	88.55	89.39	89.74
3	BTB#4	8	64.23	72.58	79.76	83.94	87.29	89.00	89.57	89.78
3	BTB#5	8	64.41	72.71	79.96	84.08	87.41	89.02	89.58	89.78
3	BTB#3	16	61.78	69.31	77.04	82.69	86.21	88.74	89.46	89.76
3	BTB#4	16	65.26	72.38	80.06	84.00	87.53	89.10	89.61	89.80
3	BTB#5	16	65.47	72.84	80.46	84.26	87.68	89.12	89.63	89.80
4	BTB#3	1	58.65	66.79	73.31	78.78	83.08	86.37	88.36	89.59
4	BTB#4	1	62.56	70.18	76.13	80.78	84.45	87.17	88.80	89.83
4	BTB#5	1	62.56	70.18	76.13	80.78	84.45	87.17	88.80	89.83
4	BTB#3	2	59.12	68.30	75.73	81.24	85.44	88.30	89.77	90.45
4	BTB#4	2	63.97	71.66	78.37	83.03	86.70	88.94	90.08	90.56
4	BTB#5	2	63.98	71.67	78.38	83.04	86.70	88.95	90.08	90.56
4	BTB#3	4	59.97	68.59	76.89	82.65	86.51	89.11	90.23	90.64
4	BTB#4	4	64.25	72.50	79.57	84.19	87.65	89.64	90.43	90.70
4	BTB#5	4	64.50	72.56	79.68	84.25	87.71	89.65	90.44	90.70
4	BTB#3	8	60.22	69.34	77.25	83.13	86.91	89.48	90.36	90.70
4	BTB#4	8	64.45	72.93	80.29	84.65	88.15	89.93	90.53	90.74
4	BTB#5	8	64.82	73.05	80.45	84.75	88.24	89.96	90.53	90.74
4	BTB#3	16	62.01	69.64	77.56	83.41	87.07	89.68	90.42	90.73
4	BTB#4	16	65.52	72.75	80.60	84.71	88.41	90.03	90.57	90.76
4	BTB#5	16	65.69	73.17	80.96	84.92	88.52	90.06	90.58	90.76
5	BTB#3	1	58.68	66.83	73.37	78.87	83.19	86.50	88.50	89.74
5	BTB#4	1	62.59	70.22	76.19	80.86	84.56	87.30	88.94	89.98
5	BTB#5	1	62.59	70.22	76.19	80.86	84.56	87.30	88.94	89.98
5	BTB#3	2	59.14	68.34	75.78	81.33	85.56	88.44	89.92	90.60
5	BTB#4	2	64.00	71.70	78.42	83.12	86.81	89.08	90.23	90.71
5	BTB#5	2	64.01	71.71	78.43	83.13	86.82	89.08	90.23	90.71
5	BTB#3	4	59.99	68.63	76.95	82.74	86.64	89.24	90.37	90.79
5	BTB#4	4	64.28	72.55	79.63	84.28	87.78	89.78	90.58	90.85
5	BTB#5	4	64.53	72.61	79.74	84.34	87.83	89.80	90.58	90.85
5	BTB#3	8	60.24	69.38	77.32	83.23	87.04	89.63	90.51	90.86
5	BTB#4	8	64.47	72.98	80.36	84.75	88.28	90.08	90.67	90.89
5	BTB#5	8	64.85	73.10	80.52	84.84	88.36	90.10	90.68	90.89
5	BTB#3	16	62.03	69.68	77.62	83.50	87.20	89.82	90.58	90.89
5	BTB#4	16	65.55	72.80	80.67	84.82	88.54	90.18	90.72	90.91
5	BTB#5	16	65.72	73.22	81.02	85.01	88.65	90.20	90.73	90.91

Table #A-6 (cont.): Prediction Rates for Designs BTB#3, BTB#4, and BTB#5										
Prediction Bits	BTB Design	Set Size	Number of BTB Entries							
			16	32	64	128	256	512	1024	2048
6	BTB#3	1	58.81	67.08	73.63	79.14	83.50	86.83	88.83	90.08
6	BTB#4	1	62.73	70.47	76.46	81.15	84.88	87.64	89.28	90.33
6	BTB#5	1	62.73	70.47	76.46	81.15	84.88	87.64	89.28	90.33
6	BTB#3	2	59.27	68.57	76.04	81.62	85.88	88.77	90.26	90.95
6	BTB#4	2	64.13	71.94	78.70	83.41	87.14	89.42	90.58	91.06
6	BTB#5	2	64.14	71.95	78.72	83.42	87.14	89.43	90.58	91.06
6	BTB#3	4	60.11	68.77	77.22	83.04	86.96	89.58	90.73	91.14
6	BTB#4	4	64.41	72.79	79.92	84.58	88.11	90.13	90.93	91.20
6	BTB#5	4	64.66	72.86	80.02	84.64	88.16	90.14	90.94	91.20
6	BTB#3	8	60.27	69.62	77.60	83.53	87.36	89.98	90.86	91.22
6	BTB#4	8	64.61	73.23	80.65	85.05	88.62	90.43	91.03	91.25
6	BTB#5	8	64.97	73.35	80.81	85.14	88.70	90.45	91.04	91.25
6	BTB#3	16	62.06	69.91	77.90	83.81	87.52	90.17	90.93	91.24
6	BTB#4	16	65.77	73.04	80.96	85.12	88.88	90.53	91.08	91.27
6	BTB#5	16	65.93	73.46	81.31	85.31	88.99	90.55	91.09	91.27
7	BTB#3	1	58.83	67.11	73.66	79.18	83.57	86.92	88.93	90.19
7	BTB#4	1	62.75	70.51	76.50	81.21	84.96	87.73	89.38	90.43
7	BTB#5	1	62.75	70.51	76.50	81.21	84.96	87.73	89.38	90.43
7	BTB#3	2	59.27	68.59	76.08	81.67	85.96	88.87	90.37	91.06
7	BTB#4	2	64.16	71.98	78.74	83.48	87.23	89.53	90.69	91.18
7	BTB#5	2	64.16	71.98	78.76	83.49	87.24	89.54	90.69	91.18
7	BTB#3	4	60.13	68.80	77.27	83.10	87.04	89.69	90.84	91.26
7	BTB#4	4	64.42	72.84	79.97	84.65	88.21	90.24	91.05	91.31
7	BTB#5	4	64.68	72.89	80.07	84.70	88.26	90.25	91.05	91.32
7	BTB#3	8	60.28	69.64	77.64	83.59	87.45	90.09	90.97	91.33
7	BTB#4	8	64.62	73.27	80.70	85.12	88.72	90.54	91.14	91.36
7	BTB#5	8	64.99	73.40	80.86	85.21	88.79	90.56	91.15	91.36
7	BTB#3	16	62.07	69.93	77.95	83.86	87.62	90.28	91.04	91.36
7	BTB#4	16	65.79	73.08	81.01	85.19	88.99	90.64	91.19	91.38
7	BTB#5	16	65.95	73.50	81.36	85.38	89.08	90.67	91.20	91.39
8	BTB#3	1	58.83	67.13	73.68	79.23	83.64	87.00	89.02	90.30
8	BTB#4	1	62.75	70.52	76.53	81.26	85.02	87.81	89.47	90.54
8	BTB#5	1	62.75	70.52	76.53	81.26	85.02	87.81	89.47	90.54
8	BTB#3	2	59.28	68.61	76.10	81.72	86.04	88.97	90.47	91.17
8	BTB#4	2	64.16	72.00	78.77	83.53	87.31	89.63	90.80	91.28
8	BTB#5	2	64.17	72.00	78.78	83.54	87.32	89.63	90.80	91.28
8	BTB#3	4	60.13	68.81	77.31	83.15	87.13	89.79	90.95	91.37
8	BTB#4	4	64.43	72.85	80.01	84.70	88.30	90.34	91.16	91.42
8	BTB#5	4	64.69	72.91	80.11	84.77	88.34	90.35	91.16	91.43
8	BTB#3	8	60.28	69.65	77.67	83.64	87.54	90.19	91.09	91.44
8	BTB#4	8	64.63	73.29	80.74	85.18	88.82	90.65	91.25	91.47
8	BTB#5	8	64.99	73.41	80.89	85.27	88.89	90.67	91.26	91.47
8	BTB#3	16	62.08	69.94	77.99	83.92	87.71	90.39	91.15	91.47
8	BTB#4	16	65.79	73.10	81.05	85.26	89.08	90.75	91.30	91.49
8	BTB#5	16	65.97	73.51	81.39	85.44	89.16	90.77	91.31	91.49

Table #A-6 (cont.): Prediction Rates for Designs BTB#3, BTB#4, and BTB#5										
Prediction Bits	BTB Design	Set Size	Number of BTB Entries							
			16	32	64	128	256	512	1024	2048
9	BTB#3	1	58.85	67.14	73.72	79.28	83.69	87.07	89.10	90.38
9	BTB#4	1	62.77	70.54	76.56	81.31	85.08	87.88	89.55	90.62
9	BTB#5	1	62.77	70.54	76.56	81.31	85.08	87.88	89.55	90.62
9	BTB#3	2	59.30	68.63	76.13	81.76	86.10	89.05	90.56	91.26
9	BTB#4	2	64.18	72.03	78.81	83.58	87.38	89.71	90.88	91.37
9	BTB#5	2	64.19	72.03	78.82	83.59	87.38	89.71	90.88	91.37
9	BTB#3	4	60.15	68.83	77.33	83.20	87.19	89.87	91.03	91.45
9	BTB#4	4	64.45	72.87	80.05	84.77	88.37	90.42	91.24	91.51
9	BTB#5	4	64.71	72.93	80.14	84.82	88.41	90.43	91.24	91.51
9	BTB#3	8	60.30	69.67	77.70	83.70	87.61	90.27	91.17	91.52
9	BTB#4	8	64.65	73.31	80.78	85.25	88.89	90.73	91.34	91.55
9	BTB#5	8	65.01	73.43	80.92	85.33	88.96	90.74	91.34	91.56
9	BTB#3	16	62.09	69.96	78.01	83.98	87.78	90.47	91.24	91.55
9	BTB#4	16	65.81	73.12	81.10	85.33	89.16	90.83	91.39	91.58
9	BTB#5	16	65.99	73.54	81.42	85.49	89.24	90.85	91.39	91.58
10	BTB#3	1	58.86	67.16	73.75	79.33	83.75	87.14	89.19	90.48
10	BTB#4	1	62.79	70.57	76.60	81.36	85.15	87.96	89.64	90.71
10	BTB#5	1	62.79	70.57	76.60	81.36	85.15	87.96	89.64	90.71
10	BTB#3	2	59.31	68.65	76.17	81.81	86.17	89.13	90.65	91.36
10	BTB#4	2	64.20	72.06	78.86	83.64	87.45	89.80	90.98	91.46
10	BTB#5	2	64.20	72.07	78.87	83.65	87.46	89.80	90.98	91.46
10	BTB#3	4	60.16	68.85	77.36	83.25	87.26	89.96	91.13	91.55
10	BTB#4	4	64.47	72.91	80.09	84.82	88.45	90.51	91.34	91.61
10	BTB#5	4	64.73	72.97	80.18	84.87	88.48	90.52	91.34	91.61
10	BTB#3	8	60.32	69.70	77.74	83.74	87.68	90.36	91.27	91.63
10	BTB#4	8	64.67	73.35	80.82	85.31	88.97	90.82	91.44	91.66
10	BTB#5	8	65.03	73.47	80.97	85.38	89.04	90.84	91.44	91.66
10	BTB#3	16	62.11	69.99	78.05	84.02	87.85	90.57	91.33	91.65
10	BTB#4	16	65.84	73.16	81.14	85.39	89.24	90.93	91.48	91.68
10	BTB#5	16	65.98	73.57	81.47	85.54	89.31	90.95	91.49	91.68
11	BTB#3	1	58.88	67.19	73.79	79.38	83.81	87.21	89.28	90.57
11	BTB#4	1	62.81	70.60	76.64	81.42	85.23	88.05	89.73	90.81
11	BTB#5	1	62.81	70.60	76.64	81.42	85.23	88.05	89.73	90.81
11	BTB#3	2	59.32	68.68	76.21	81.86	86.23	89.22	90.75	91.46
11	BTB#4	2	64.22	72.09	78.90	83.70	87.54	89.89	91.08	91.57
11	BTB#5	2	64.22	72.10	78.91	83.70	87.54	89.89	91.08	91.57
11	BTB#3	4	60.18	68.88	77.41	83.31	87.33	90.05	91.23	91.66
11	BTB#4	4	64.50	72.94	80.14	84.89	88.53	90.61	91.44	91.72
11	BTB#5	4	64.75	73.00	80.23	84.93	88.56	90.62	91.45	91.72
11	BTB#3	8	60.34	69.73	77.78	83.81	87.77	90.46	91.37	91.73
11	BTB#4	8	64.70	73.38	80.87	85.38	89.06	90.93	91.54	91.77
11	BTB#5	8	65.06	73.50	81.00	85.44	89.12	90.94	91.54	91.77
11	BTB#3	16	62.13	70.02	78.10	84.08	87.93	90.66	91.44	91.76
11	BTB#4	16	65.86	73.19	81.19	85.45	89.33	91.03	91.59	91.78
11	BTB#5	16	66.01	73.60	81.50	85.61	89.40	91.05	91.60	91.79

Table #A-6 (cont.): Prediction Rates for Designs BTB#3, BTB#4, and BTB#5										
Prediction Bits	BTB Design	Set Size	Number of BTB Entries							
			16	32	64	128	256	512	1024	2048
12	BTB#3	1	58.90	67.22	73.83	79.45	83.90	87.33	89.39	90.70
12	BTB#4	1	62.83	70.63	76.69	81.49	85.33	88.16	89.86	90.94
12	BTB#5	1	62.83	70.63	76.69	81.49	85.33	88.16	89.86	90.94
12	BTB#3	2	59.34	68.71	76.25	81.93	86.34	89.34	90.88	91.60
12	BTB#4	2	64.24	72.11	78.95	83.77	87.64	90.02	91.21	91.71
12	BTB#5	2	64.24	72.12	78.96	83.78	87.64	90.02	91.21	91.70
12	BTB#3	4	60.21	68.90	77.45	83.39	87.44	90.18	91.36	91.80
12	BTB#4	4	64.52	72.97	80.19	84.97	88.65	90.74	91.58	91.85
12	BTB#5	4	64.77	73.03	80.28	85.01	88.67	90.75	91.58	91.85
12	BTB#3	8	60.37	69.75	77.82	83.89	87.87	90.59	91.51	91.87
12	BTB#4	8	64.72	73.41	80.92	85.47	89.18	91.06	91.67	91.90
12	BTB#5	8	65.08	73.52	81.04	85.53	89.23	91.07	91.68	91.90
12	BTB#3	16	62.15	70.04	78.14	84.17	88.05	90.79	91.57	91.90
12	BTB#4	16	65.89	73.22	81.25	85.55	89.46	91.16	91.72	91.92
12	BTB#5	16	66.03	73.62	81.54	85.69	89.52	91.18	91.73	91.92
13	BTB#3	1	58.93	67.26	73.88	79.51	83.98	87.43	89.52	90.83
13	BTB#4	1	62.86	70.66	76.73	81.56	85.42	88.26	89.98	91.07
13	BTB#5	1	62.86	70.66	76.73	81.56	85.42	88.26	89.98	91.07
13	BTB#3	2	59.37	68.74	76.29	81.99	86.44	89.46	91.01	91.74
13	BTB#4	2	64.26	72.15	79.00	83.85	87.75	90.14	91.35	91.85
13	BTB#5	2	64.26	72.15	79.00	83.85	87.75	90.14	91.35	91.85
13	BTB#3	4	60.23	68.93	77.49	83.46	87.55	90.31	91.51	91.94
13	BTB#4	4	64.54	73.00	80.24	85.06	88.76	90.87	91.72	92.00
13	BTB#5	4	64.80	73.06	80.32	85.08	88.79	90.88	91.73	92.00
13	BTB#3	8	60.39	69.78	77.87	83.97	87.99	90.73	91.65	92.01
13	BTB#4	8	64.75	73.44	80.97	85.56	89.31	91.19	91.82	92.04
13	BTB#5	8	65.11	73.56	81.09	85.60	89.35	91.21	91.82	92.05
13	BTB#3	16	62.18	70.07	78.18	84.25	88.17	90.93	91.72	92.04
13	BTB#4	16	65.92	73.26	81.30	85.64	89.59	91.30	91.87	92.06
13	BTB#5	16	66.05	73.65	81.58	85.76	89.63	91.31	91.88	92.07
14	BTB#3	1	58.95	67.29	73.91	79.58	84.08	87.55	89.66	90.99
14	BTB#4	1	62.88	70.69	76.78	81.64	85.53	88.40	90.13	91.23
14	BTB#5	1	62.88	70.69	76.78	81.64	85.53	88.40	90.13	91.23
14	BTB#3	2	59.39	68.77	76.33	82.06	86.55	89.61	91.18	91.91
14	BTB#4	2	64.28	72.18	79.04	83.93	87.89	90.30	91.51	92.02
14	BTB#5	2	64.29	72.18	79.04	83.93	87.88	90.29	91.51	92.02
14	BTB#3	4	60.26	68.96	77.53	83.53	87.67	90.47	91.67	92.11
14	BTB#4	4	64.56	73.03	80.30	85.15	88.90	91.03	91.89	92.17
14	BTB#5	4	64.82	73.09	80.38	85.17	88.92	91.04	91.89	92.17
14	BTB#3	8	60.41	69.81	77.91	84.06	88.12	90.89	91.82	92.19
14	BTB#4	8	64.77	73.47	81.04	85.66	89.46	91.36	91.99	92.22
14	BTB#5	8	65.12	73.59	81.14	85.68	89.49	91.37	91.99	92.22
14	BTB#3	16	62.20	70.10	78.23	84.35	88.30	91.09	91.89	92.22
14	BTB#4	16	65.94	73.29	81.37	85.74	89.73	91.47	92.04	92.24
14	BTB#5	16	66.07	73.68	81.62	85.84	89.76	91.48	92.05	92.24



Table #A-6 (cont.): Prediction Rates for Designs BTB#3, BTB#4, and BTB#5										
Prediction Bits	BTB Design	Set Size	Number of BTB Entries							
			16	32	64	128	256	512	1024	2048
15	BTB#3	1	58.96	67.31	73.94	79.64	84.17	87.67	89.80	91.15
15	BTB#4	1	62.89	70.72	76.82	81.71	85.63	88.53	90.28	91.39
15	BTB#5	1	62.89	70.72	76.82	81.71	85.63	88.53	90.28	91.39
15	BTB#3	2	59.40	68.80	76.37	82.13	86.66	89.74	91.33	92.08
15	BTB#4	2	64.30	72.21	79.10	84.01	88.00	90.44	91.67	92.18
15	BTB#5	2	64.30	72.20	79.09	84.00	87.99	90.44	91.67	92.18
15	BTB#3	4	60.27	68.99	77.56	83.61	87.79	90.61	91.84	92.28
15	BTB#4	4	64.58	73.06	80.35	85.24	89.02	91.19	92.06	92.34
15	BTB#5	4	64.84	73.11	80.41	85.23	89.03	91.19	92.06	92.34
15	BTB#3	8	60.43	69.83	77.94	84.14	88.25	91.04	91.99	92.36
15	BTB#4	8	64.79	73.50	81.09	85.76	89.59	91.52	92.16	92.39
15	BTB#5	8	65.13	73.61	81.17	85.77	89.60	91.52	92.16	92.39
15	BTB#3	16	62.22	70.12	78.26	84.43	88.43	91.25	92.05	92.39
15	BTB#4	16	65.96	73.32	81.42	85.84	89.87	91.63	92.21	92.41
15	BTB#5	16	66.08	73.69	81.66	85.90	89.87	91.63	92.21	92.41
16	BTB#3	1	58.97	67.34	73.98	79.71	84.30	87.83	89.99	91.36
16	BTB#4	1	62.90	70.75	76.88	81.80	85.77	88.70	90.48	91.61
16	BTB#5	1	62.90	70.75	76.88	81.80	85.77	88.70	90.48	91.61
16	BTB#3	2	59.42	68.83	76.42	82.21	86.79	89.94	91.56	92.31
16	BTB#4	2	64.32	72.24	79.16	84.11	88.15	90.64	91.90	92.42
16	BTB#5	2	64.32	72.23	79.15	84.10	88.14	90.63	91.89	92.42
16	BTB#3	4	60.29	69.02	77.61	83.71	87.94	90.82	92.07	92.52
16	BTB#4	4	64.60	73.09	80.41	85.35	89.19	91.40	92.29	92.58
16	BTB#5	4	64.85	73.14	80.46	85.33	89.19	91.39	92.29	92.57
16	BTB#3	8	60.45	69.86	77.99	84.24	88.42	91.25	92.22	92.60
16	BTB#4	8	64.81	73.53	81.15	85.88	89.78	91.73	92.39	92.63
16	BTB#5	8	65.15	73.63	81.22	85.86	89.76	91.73	92.39	92.63
16	BTB#3	16	62.24	70.15	78.30	84.54	88.60	91.46	92.29	92.63
16	BTB#4	16	65.98	73.35	81.49	85.96	90.06	91.85	92.45	92.65
16	BTB#5	16	66.10	73.71	81.70	85.99	90.02	91.84	92.44	92.65

**Table #A-6:** Table presents the prediction rates of BTB#3, BTB#4, and BTB#5 as a function of number of prediction bits, set size, and BTB size. BTB#3 uses LRU and enters branches on taken and not taken executions. BTB#4 uses LRU and enters branches on taken executions only. BTB#5 uses MPP and enters branches on taken executions only.

Table #A-7: Miss Rates for Instruction Cache and BTB#3								
Size	Set Size	Design	COMP	TEXT	FP	SPARC	VAX	68k
64	1	BTB#3	43.23	20.97	11.18	31.88	17.23	25.69
64*BB	1	Cache	40.61	19.84	20.35	31.48	10.47	19.66
64	2	BTB#3	36.79	17.14	8.21	30.08	9.80	23.42
64*BB	2	Cache	35.35	17.60	12.97	29.06	6.20	15.98
64	4	BTB#3	33.81	15.07	6.77	28.71	6.00	18.16
64*BB	4	Cache	33.82	16.59	5.08	28.97	5.06	12.85
64	8	BTB#3	34.77	15.21	6.43	27.86	5.44	17.52
64*BB	8	Cache	33.46	16.37	5.18	29.09	4.59	11.82
128	1	BTB#3	31.11	14.13	8.11	22.36	7.81	17.72
128*BB	1	Cache	28.02	14.21	9.91	21.21	5.21	13.42
128	2	BTB#3	23.05	10.08	4.72	20.04	3.32	12.51
128*BB	2	Cache	20.72	9.15	8.48	20.19	2.96	10.68
128	4	BTB#3	17.68	8.32	3.75	18.61	1.88	9.64
128*BB	4	Cache	14.98	7.31	2.92	18.17	1.68	7.04
128	8	BTB#3	14.18	6.94	3.56	17.77	1.43	8.74
128*BB	8	Cache	13.71	6.30	2.85	17.83	1.49	6.67
256	1	BTB#3	20.85	8.45	5.61	14.05	4.17	10.67
256*BB	1	Cache	15.70	5.79	2.32	14.20	2.99	6.01
256	2	BTB#3	13.85	4.66	2.82	10.60	1.58	5.51
256*BB	2	Cache	9.98	4.92	1.72	8.94	1.60	4.29
256	4	BTB#3	8.94	3.23	2.29	7.87	0.93	4.39
256*BB	4	Cache	7.66	3.01	1.62	7.42	0.82	3.32
256	8	BTB#3	7.06	3.00	2.05	6.67	0.79	3.58
256*BB	8	Cache	7.11	2.92	1.49	6.28	0.79	2.89
512	1	BTB#3	12.55	5.36	2.62	7.96	2.03	5.54
512*BB	1	Cache	10.59	2.35	1.35	6.18	1.23	3.02
512	2	BTB#3	7.58	2.11	1.54	4.26	0.88	2.66
512*BB	2	Cache	5.21	1.64	1.03	3.60	0.65	1.72
512	4	BTB#3	4.51	1.46	1.28	2.19	0.59	1.84
512*BB	4	Cache	4.49	1.36	0.90	1.10	0.62	1.16
512	8	BTB#3	4.16	1.24	1.13	1.02	0.53	1.36
512*BB	8	Cache	4.23	1.20	0.85	0.65	0.57	0.82
1024	1	BTB#3	7.55	3.06	1.39	4.25	1.37	3.53
1024*BB	1	Cache	6.08	1.49	0.95	3.49	0.62	1.94
1024	2	BTB#3	3.84	1.18	0.92	1.34	0.62	1.26
1024*BB	2	Cache	2.73	0.87	0.75	0.65	0.49	0.80
1024	4	BTB#3	2.41	0.76	0.76	0.53	0.46	0.81
1024*BB	4	Cache	2.21	0.69	0.70	0.33	0.45	0.54
1024	8	BTB#3	2.08	0.61	0.67	0.30	0.41	0.72
1024*BB	8	Cache	1.90	0.60	0.70	0.26	0.42	0.53

Table #A-7: Table compares the miss rate of BTB#3 and a comparable instruction cache, as a function of size, set size, and workload. Each workload miss rate is the geometric mean of the workload trace's miss rates.

Table #A-8: Miss Rates for BTB#3 and BTB#4									
Size	Set Size	Design	COMP	TEXT	FP	SPARC	VAX	68k	ave
64	1	BTB#3	43.23	20.97	11.18	31.88	17.23	25.69	22.87
64	1	BTB#4	27.15	11.23	6.48	19.35	7.50	16.19	12.91
64	2	BTB#3	36.79	17.14	8.21	30.08	9.80	23.42	18.15
64	2	BTB#4	21.25	9.58	4.81	18.09	5.42	13.57	10.45
64	4	BTB#3	33.81	15.07	6.77	28.71	6.00	18.16	14.87
64	4	BTB#4	18.12	7.93	4.02	17.67	3.13	10.56	8.34
64	8	BTB#3	34.77	15.21	6.43	27.86	5.44	17.52	14.43
64	8	BTB#4	17.01	6.00	3.41	17.82	2.60	9.79	7.35
128	1	BTB#3	31.11	14.13	8.11	22.36	7.81	17.72	14.92
128	1	BTB#4	20.16	7.48	4.39	13.78	4.47	10.79	8.72
128	2	BTB#3	23.05	10.08	4.72	20.04	3.32	12.51	9.85
128	2	BTB#4	13.68	4.91	2.94	11.89	1.99	6.85	5.64
128	4	BTB#3	17.68	8.32	3.75	18.61	1.88	9.64	7.55
128	4	BTB#4	9.89	3.94	2.37	10.49	1.31	5.90	4.42
128	8	BTB#3	14.18	6.94	3.56	17.77	1.43	8.74	6.53
128	8	BTB#4	7.89	3.19	2.15	9.66	0.93	5.43	3.72
256	1	BTB#3	20.85	8.45	5.61	14.05	4.17	10.67	9.23
256	1	BTB#4	13.27	3.68	2.72	8.71	2.43	6.47	5.13
256	2	BTB#3	13.85	4.66	2.82	10.60	1.58	5.51	5.06
256	2	BTB#4	8.34	2.43	1.83	6.10	0.97	3.36	3.00
256	4	BTB#3	8.94	3.23	2.29	7.87	0.93	4.39	3.58
256	4	BTB#4	5.42	1.87	1.48	3.79	0.65	2.61	2.14
256	8	BTB#3	7.06	3.00	2.05	6.67	0.79	3.58	3.06
256	8	BTB#4	4.39	1.63	1.35	2.89	0.55	2.07	1.78
512	1	BTB#3	12.55	5.36	2.62	7.96	2.03	5.54	5.01
512	1	BTB#4	8.12	2.36	1.64	4.80	1.21	4.06	3.01
512	2	BTB#3	7.58	2.11	1.54	4.26	0.88	2.66	2.50
512	2	BTB#4	4.63	1.24	1.01	2.33	0.60	1.65	1.54
512	4	BTB#3	4.51	1.46	1.28	2.19	0.59	1.84	1.65
512	4	BTB#4	2.76	0.85	0.84	0.90	0.43	1.09	0.97
512	8	BTB#3	4.16	1.24	1.13	1.02	0.53	1.36	1.28
512	8	BTB#4	2.49	0.74	0.72	0.41	0.38	0.84	0.75
1024	1	BTB#3	7.55	3.06	1.39	4.25	1.37	3.53	2.95
1024	1	BTB#4	4.84	1.41	1.01	2.48	0.84	2.67	1.84
1024	2	BTB#3	3.84	1.18	0.92	1.34	0.62	1.26	1.28
1024	2	BTB#4	2.48	0.67	0.66	0.71	0.41	0.87	0.81
1024	4	BTB#3	2.41	0.76	0.76	0.53	0.46	0.81	0.81
1024	4	BTB#4	1.50	0.47	0.55	0.26	0.35	0.58	0.52
1024	8	BTB#3	2.08	0.61	0.67	0.30	0.41	0.72	0.65
1024	8	BTB#4	1.27	0.41	0.50	0.20	0.33	0.50	0.45

Table #A-8: Table compares the miss rates of BTB#3 and BTB#4 as a function of size, set size, and workload. BTB#3 uses LRU and enters branches on taken and not taken executions. BTB#4 uses LRU and enters branches on taken executions only. Workload miss rates are geometric means of the workload trace miss rates. Ave is the geometric mean of the workload geometric means.

Table #A-9: Average Cycle Savings Per Branch, 4096 bit BTB, costx										
Design	lev1	lev2	lev3	COMP	TEXT	FP	SPARC	VAX	68k	ave
BTB#6	0	128	0	0.832	1.060	0.937	0.660	1.078	0.852	0.903
BTB#6	4	112	0	1.121	1.583	1.244	0.888	1.763	1.190	1.298
BTB#6	8	96	0	1.226	1.633	1.496	1.081	1.829	1.202	1.411
BTB#6	12	80	0	1.259	1.749	1.673	1.106	1.880	1.426	1.515
BTB#6	16	64	0	1.255	1.810	1.850	1.134	1.892	1.426	1.561
BTB#6	20	48	0	1.300	1.837	1.864	1.163	1.958	1.414	1.589
BTB#6	24	32	0	1.310	1.849	1.845	1.151	2.012	1.439	1.601
BTB#6	28	16	0	1.313	1.828	1.848	1.116	2.035	1.452	1.599
BTB#6	32	0	0	1.337	2.000	1.834	1.202	2.068	1.547	1.665
BTB#7	0	64	0	0.846	1.073	0.956	0.610	1.111	0.815	0.902
BTB#7	4	56	0	1.132	1.579	1.281	0.882	1.789	1.192	1.309
BTB#7	8	48	0	1.245	1.657	1.500	1.082	1.856	1.204	1.424
BTB#7	12	40	0	1.265	1.758	1.693	1.114	1.923	1.386	1.523
BTB#7	16	32	0	1.238	1.846	1.862	1.149	1.934	1.415	1.574
BTB#7	20	24	0	1.299	1.876	1.872	1.175	1.975	1.434	1.605
BTB#7	24	16	0	1.320	1.890	1.868	1.184	2.025	1.480	1.628
BTB#7	28	8	0	1.337	1.881	1.868	1.192	2.054	1.498	1.638
BTB#7	32	0	0	1.337	2.000	1.834	1.202	2.068	1.547	1.665
BTB#8	0	32	1024	0.807	1.071	0.920	0.693	1.064	0.860	0.903
BTB#8	2	28	1024	1.057	1.511	1.160	0.789	1.543	1.107	1.195
BTB#8	4	24	1024	1.107	1.519	1.251	0.958	1.741	1.223	1.300
BTB#8	6	20	1024	1.187	1.565	1.367	1.093	1.773	1.231	1.369
BTB#8	8	16	1024	1.228	1.594	1.477	1.150	1.794	1.228	1.412
BTB#8	10	12	1024	1.233	1.640	1.637	1.172	1.815	1.404	1.484
BTB#8	12	8	1024	1.267	1.686	1.675	1.176	1.843	1.405	1.509
BTB#8	14	4	1024	1.276	1.767	1.682	1.176	1.840	1.440	1.530
BTB#8	16	0	1024	1.253	1.765	1.847	1.194	1.831	1.445	1.556
BTB#9	0	0	2048	0.529	0.589	0.544	0.409	0.561	0.502	0.522
BTB#9	4	0	1792	0.942	1.341	1.000	0.788	1.575	1.044	1.115
BTB#9	8	0	1536	1.142	1.453	1.308	1.090	1.689	1.080	1.294
BTB#9	12	0	1280	1.187	1.598	1.581	1.104	1.780	1.357	1.435
BTB#9	16	0	1024	1.253	1.765	1.847	1.194	1.831	1.445	1.556
BTB#9	20	0	768	1.356	1.808	1.853	1.236	1.952	1.457	1.610
BTB#9	24	0	512	1.421	1.856	1.873	1.276	2.036	1.550	1.668
BTB#9	28	0	256	1.442	1.854	1.880	1.255	2.085	1.561	1.679
BTB#9	32	0	0	1.337	2.000	1.834	1.202	2.068	1.547	1.665

Table #A-10: Average Cycle Savings Per Branch, 4096 bit BTB, costly										
Design	lev1	lev2	lev3	COMP	TEXT	FP	SPARC	VAX	68k	ave
BTB#6	0	128	0	1.394	1.689	1.511	1.071	1.740	1.387	1.465
BTB#6	5	112	0	1.593	1.960	1.696	1.194	2.067	1.535	1.674
BTB#6	10	96	0	1.587	1.987	1.888	1.281	2.087	1.649	1.747
BTB#6	16	80	0	1.582	2.062	1.979	1.300	2.111	1.673	1.785
BTB#6	21	64	0	1.629	2.051	1.960	1.317	2.160	1.643	1.793
BTB#6	26	48	0	1.608	2.056	1.978	1.314	2.151	1.655	1.794
BTB#6	32	32	0	1.595	2.124	1.966	1.275	2.186	1.667	1.802
BTB#6	37	16	0	1.564	2.123	1.992	1.251	2.184	1.665	1.796
BTB#6	42	0	0	1.488	2.133	2.004	1.265	2.161	1.611	1.777
BTB#7	0	64	0	1.345	1.643	1.521	0.980	1.752	1.277	1.420
BTB#7	5	56	0	1.532	1.913	1.723	1.136	2.091	1.476	1.645
BTB#7	10	48	0	1.547	1.962	1.889	1.233	2.113	1.568	1.719
BTB#7	16	40	0	1.555	2.035	1.973	1.254	2.143	1.597	1.759
BTB#7	21	32	0	1.562	2.051	1.969	1.274	2.195	1.594	1.774
BTB#7	26	24	0	1.553	2.060	1.967	1.288	2.202	1.605	1.779
BTB#7	32	16	0	1.520	2.126	1.973	1.298	2.218	1.619	1.792
BTB#7	37	8	0	1.496	2.131	1.993	1.288	2.163	1.618	1.781
BTB#7	42	0	0	1.488	2.133	2.004	1.265	2.161	1.611	1.777
BTB#8	0	32	1024	1.433	1.712	1.521	1.195	1.732	1.441	1.506
BTB#8	2	28	1024	1.574	1.922	1.667	1.248	1.971	1.577	1.660
BTB#8	5	24	1024	1.640	1.948	1.735	1.351	2.082	1.640	1.733
BTB#8	8	20	1024	1.672	1.952	1.834	1.434	2.095	1.640	1.771
BTB#8	10	16	1024	1.673	1.978	1.909	1.446	2.094	1.732	1.805
BTB#8	13	12	1024	1.689	2.047	1.926	1.451	2.116	1.748	1.829
BTB#8	16	8	1024	1.693	2.050	2.001	1.465	2.119	1.756	1.847
BTB#8	18	4	1024	1.719	2.047	2.004	1.473	2.134	1.753	1.855
BTB#8	21	0	1024	1.728	2.059	2.009	1.482	2.171	1.754	1.867
BTB#9	0	0	2048	1.175	1.230	1.180	0.920	1.201	1.111	1.136
BTB#9	5	0	1792	1.521	1.760	1.513	1.193	1.903	1.472	1.560
BTB#9	10	0	1536	1.582	1.854	1.786	1.389	1.975	1.682	1.711
BTB#9	16	0	1280	1.592	1.983	1.968	1.388	2.027	1.702	1.777
BTB#9	21	0	1024	1.728	2.059	2.009	1.482	2.171	1.754	1.867
BTB#9	26	0	768	1.751	2.080	2.006	1.479	2.199	1.792	1.884
BTB#9	32	0	512	1.787	2.224	2.035	1.507	2.260	1.856	1.945
BTB#9	37	0	256	1.761	2.226	2.061	1.464	2.281	1.828	1.937
BTB#9	42	0	0	1.488	2.133	2.004	1.265	2.161	1.611	1.777

Table #A-11: Average Cycle Savings Per Branch, 8192 bit BTB, costx										
Design	lev1	lev2	lev3	COMP	TEXT	FP	SPARC	VAX	68k	ave
BTB#6	0	256	0	0.920	1.122	0.972	0.748	1.124	0.931	0.969
BTB#6	8	224	0	1.357	1.720	1.549	1.191	1.893	1.280	1.498
BTB#6	16	192	0	1.417	1.923	1.924	1.266	1.986	1.564	1.680
BTB#6	24	160	0	1.550	1.987	1.942	1.307	2.131	1.636	1.758
BTB#6	32	128	0	1.601	2.131	1.970	1.343	2.203	1.711	1.826
BTB#6	40	96	0	1.627	2.179	2.051	1.328	2.247	1.724	1.859
BTB#6	48	64	0	1.689	2.207	2.059	1.360	2.306	1.724	1.891
BTB#6	56	32	0	1.780	2.209	2.073	1.318	2.325	1.755	1.910
BTB#6	64	0	0	1.868	2.230	2.110	1.362	2.363	1.767	1.950
BTB#7	0	128	0	0.946	1.113	0.976	0.715	1.143	0.912	0.968
BTB#7	8	112	0	1.385	1.707	1.546	1.181	1.907	1.297	1.504
BTB#7	16	96	0	1.459	1.910	1.926	1.244	2.006	1.536	1.680
BTB#7	24	80	0	1.593	1.981	1.951	1.308	2.153	1.607	1.765
BTB#7	32	64	0	1.655	2.136	1.986	1.335	2.229	1.684	1.837
BTB#7	40	48	0	1.705	2.182	2.060	1.347	2.270	1.705	1.878
BTB#7	48	32	0	1.760	2.212	2.075	1.362	2.331	1.727	1.911
BTB#7	56	16	0	1.840	2.226	2.090	1.359	2.351	1.774	1.940
BTB#7	64	0	0	1.868	2.230	2.110	1.362	2.363	1.767	1.950
BTB#8	0	96	1024	0.966	1.144	0.995	0.764	1.147	0.939	0.992
BTB#8	6	84	1024	1.359	1.704	1.450	1.169	1.884	1.334	1.483
BTB#8	12	72	1024	1.464	1.844	1.767	1.256	1.987	1.539	1.643
BTB#8	18	60	1024	1.562	1.955	1.950	1.311	2.052	1.595	1.737
BTB#8	24	48	1024	1.622	2.013	1.967	1.354	2.160	1.670	1.797
BTB#8	30	36	1024	1.665	2.088	1.992	1.379	2.228	1.718	1.845
BTB#8	36	24	1024	1.701	2.190	2.045	1.399	2.252	1.756	1.890
BTB#8	42	12	1024	1.726	2.230	2.069	1.416	2.270	1.768	1.913
BTB#8	48	0	1024	1.740	2.242	2.072	1.448	2.317	1.774	1.932
BTB#9	0	0	4096	0.542	0.597	0.546	0.415	0.563	0.512	0.529
BTB#9	8	0	3584	1.164	1.463	1.311	1.099	1.693	1.092	1.304
BTB#9	16	0	3072	1.276	1.775	1.851	1.206	1.836	1.459	1.567
BTB#9	24	0	2560	1.454	1.869	1.880	1.293	2.045	1.570	1.685
BTB#9	32	0	2048	1.580	2.118	1.937	1.364	2.165	1.701	1.811
BTB#9	40	0	1536	1.644	2.188	2.040	1.387	2.222	1.735	1.869
BTB#9	48	0	1024	1.740	2.242	2.072	1.448	2.317	1.774	1.932
BTB#9	56	0	512	1.853	2.260	2.092	1.439	2.351	1.832	1.971
BTB#9	64	0	0	1.868	2.230	2.110	1.362	2.363	1.767	1.950

Table #A-12: Average Cycle Savings Per Branch, 8192 bit BTB, costly										
Design	lev1	lev2	lev3	COMP	TEXT	FP	SPARC	VAX	68k	ave
BTB#6	0	256	0	1.491	1.756	1.554	1.200	1.793	1.478	1.545
BTB#6	10	224	0	1.740	2.087	1.957	1.437	2.168	1.757	1.858
BTB#6	21	192	0	1.805	2.181	2.062	1.481	2.257	1.807	1.932
BTB#6	32	160	0	1.843	2.257	2.069	1.469	2.315	1.863	1.969
BTB#6	42	128	0	1.866	2.279	2.108	1.464	2.318	1.855	1.982
BTB#6	53	96	0	1.920	2.288	2.122	1.440	2.364	1.870	2.001
BTB#6	64	64	0	1.962	2.284	2.139	1.428	2.369	1.862	2.007
BTB#6	74	32	0	1.966	2.274	2.140	1.377	2.373	1.842	1.995
BTB#6	85	0	0	1.973	2.264	2.143	1.443	2.400	1.826	2.008
BTB#7	0	128	0	1.500	1.704	1.551	1.142	1.798	1.421	1.520
BTB#7	10	112	0	1.742	2.036	1.946	1.390	2.186	1.710	1.835
BTB#7	21	96	0	1.818	2.132	2.052	1.428	2.279	1.746	1.909
BTB#7	32	80	0	1.864	2.225	2.071	1.448	2.327	1.804	1.956
BTB#7	42	64	0	1.889	2.253	2.113	1.455	2.348	1.811	1.978
BTB#7	53	48	0	1.938	2.267	2.120	1.466	2.384	1.815	1.998
BTB#7	64	32	0	1.969	2.270	2.136	1.461	2.393	1.824	2.009
BTB#7	74	16	0	1.973	2.266	2.140	1.451	2.398	1.829	2.009
BTB#7	85	0	0	1.973	2.264	2.143	1.443	2.400	1.826	2.008
BTB#8	0	96	1024	1.592	1.786	1.595	1.268	1.816	1.513	1.595
BTB#8	8	84	1024	1.840	2.092	1.913	1.512	2.196	1.730	1.881
BTB#8	16	72	1024	1.878	2.195	2.090	1.545	2.241	1.855	1.967
BTB#8	24	60	1024	1.947	2.231	2.103	1.574	2.314	1.895	2.011
BTB#8	32	48	1024	1.978	2.313	2.117	1.582	2.353	1.939	2.047
BTB#8	40	36	1024	2.004	2.338	2.152	1.588	2.373	1.952	2.068
BTB#8	48	24	1024	2.032	2.353	2.161	1.608	2.403	1.964	2.087
BTB#8	56	12	1024	2.076	2.362	2.168	1.611	2.415	1.989	2.103
BTB#8	64	0	1024	2.095	2.365	2.181	1.614	2.423	1.988	2.111
BTB#9	0	0	4096	1.193	1.240	1.182	0.928	1.203	1.124	1.145
BTB#9	10	0	3584	1.626	1.883	1.798	1.417	1.980	1.699	1.734
BTB#9	21	0	3072	1.773	2.076	2.016	1.506	2.176	1.777	1.887
BTB#9	32	0	2560	1.864	2.262	2.051	1.557	2.274	1.899	1.984
BTB#9	42	0	2048	1.934	2.323	2.133	1.585	2.322	1.917	2.035
BTB#9	53	0	1536	2.035	2.354	2.150	1.614	2.396	1.969	2.086
BTB#9	64	0	1024	2.095	2.365	2.181	1.614	2.423	1.988	2.111
BTB#9	74	0	512	2.088	2.361	2.182	1.595	2.431	1.994	2.109
BTB#9	85	0	0	1.973	2.264	2.143	1.443	2.400	1.826	2.008

Table #A-13: Average Cycle Savings Per Branch, 16384 bit BTB, costx										
Design	lev1	lev2	lev3	COMP	TEXT	FP	SPARC	VAX	68k	ave
BTB#6	0	512	0	0.988	1.159	1.004	0.813	1.148	0.989	1.017
BTB#6	16	448	0	1.503	1.970	1.951	1.332	2.017	1.641	1.736
BTB#6	32	384	0	1.718	2.201	2.012	1.456	2.248	1.806	1.907
BTB#6	48	320	0	1.841	2.286	2.107	1.510	2.356	1.855	1.993
BTB#6	64	256	0	1.988	2.313	2.153	1.516	2.404	1.917	2.048
BTB#6	80	192	0	2.031	2.321	2.165	1.517	2.421	1.928	2.064
BTB#6	96	128	0	2.041	2.318	2.162	1.541	2.420	1.923	2.067
BTB#6	112	64	0	2.043	2.309	2.155	1.529	2.423	1.931	2.065
BTB#6	128	0	0	2.082	2.312	2.151	1.578	2.427	1.960	2.085
BTB#7	0	256	0	1.002	1.194	1.014	0.856	1.165	1.031	1.044
BTB#7	16	224	0	1.528	1.988	1.961	1.388	2.030	1.673	1.761
BTB#7	32	192	0	1.737	2.219	2.020	1.470	2.257	1.829	1.922
BTB#7	48	160	0	1.858	2.300	2.105	1.522	2.361	1.876	2.004
BTB#7	64	128	0	2.002	2.324	2.147	1.515	2.404	1.937	2.055
BTB#7	80	96	0	2.042	2.303	2.156	1.523	2.417	1.954	2.066
BTB#7	96	64	0	2.060	2.310	2.154	1.563	2.419	1.950	2.076
BTB#7	112	32	0	2.074	2.315	2.152	1.570	2.425	1.963	2.083
BTB#7	128	0	0	2.082	2.312	2.151	1.578	2.427	1.960	2.085
BTB#8	0	224	1024	1.010	1.191	1.013	0.845	1.163	1.017	1.040
BTB#8	14	196	1024	1.528	1.987	1.791	1.350	2.021	1.661	1.723
BTB#8	28	168	1024	1.727	2.081	2.004	1.459	2.221	1.784	1.879
BTB#8	42	140	1024	1.829	2.278	2.105	1.502	2.303	1.871	1.981
BTB#8	56	112	1024	1.972	2.305	2.130	1.538	2.386	1.937	2.045
BTB#8	70	84	1024	2.038	2.321	2.156	1.547	2.413	1.960	2.073
BTB#8	84	56	1024	2.069	2.329	2.168	1.572	2.425	1.961	2.087
BTB#8	98	28	1024	2.085	2.339	2.170	1.600	2.428	1.958	2.096
BTB#8	112	0	1024	2.097	2.345	2.171	1.608	2.432	1.971	2.104
BTB#9	0	0	8192	0.547	0.599	0.547	0.419	0.564	0.514	0.531
BTB#9	16	0	7168	1.299	1.781	1.854	1.217	1.836	1.466	1.576
BTB#9	32	0	6144	1.598	2.127	1.940	1.374	2.167	1.710	1.819
BTB#9	48	0	5120	1.777	2.255	2.078	1.469	2.321	1.791	1.949
BTB#9	64	0	4096	1.991	2.301	2.146	1.490	2.394	1.880	2.034
BTB#9	80	0	3072	2.053	2.320	2.164	1.519	2.419	1.917	2.065
BTB#9	96	0	2048	2.084	2.338	2.171	1.588	2.423	1.934	2.090
BTB#9	112	0	1024	2.097	2.345	2.171	1.608	2.432	1.971	2.104
BTB#9	128	0	0	2.082	2.312	2.151	1.578	2.427	1.960	2.085



Table #A-14: Average Cycle Savings Per Branch, 16384 bit BTB, costly										
Design	lev1	lev2	lev3	COMP	TEXT	FP	SPARC	VAX	68k	ave
BTB#6	0	512	0	1.586	1.800	1.591	1.293	1.813	1.541	1.604
BTB#6	21	448	0	1.908	2.230	2.093	1.577	2.297	1.905	2.002
BTB#6	42	384	0	2.002	2.363	2.157	1.632	2.373	1.976	2.084
BTB#6	64	320	0	2.099	2.382	2.188	1.632	2.431	2.014	2.124
BTB#6	85	256	0	2.116	2.379	2.188	1.639	2.437	2.025	2.131
BTB#6	106	192	0	2.122	2.385	2.189	1.642	2.439	2.022	2.133
BTB#6	128	128	0	2.123	2.370	2.177	1.627	2.437	2.027	2.127
BTB#6	149	64	0	2.111	2.357	2.171	1.616	2.435	2.042	2.122
BTB#6	170	0	0	2.136	2.344	2.176	1.652	2.436	2.083	2.138
BTB#7	0	256	0	1.586	1.826	1.608	1.350	1.832	1.598	1.633
BTB#7	21	224	0	1.920	2.252	2.108	1.646	2.315	1.956	2.033
BTB#7	42	192	0	2.005	2.374	2.164	1.668	2.387	2.017	2.102
BTB#7	64	160	0	2.097	2.396	2.187	1.666	2.436	2.042	2.138
BTB#7	85	128	0	2.118	2.399	2.184	1.670	2.445	2.051	2.145
BTB#7	106	96	0	2.126	2.396	2.181	1.673	2.445	2.061	2.147
BTB#7	128	64	0	2.135	2.395	2.177	1.661	2.442	2.075	2.147
BTB#7	149	32	0	2.138	2.350	2.175	1.658	2.438	2.089	2.141
BTB#7	170	0	0	2.136	2.344	2.176	1.652	2.436	2.083	2.138
BTB#8	0	224	1024	1.639	1.835	1.613	1.353	1.832	1.591	1.644
BTB#8	18	196	1024	1.969	2.253	2.108	1.644	2.277	1.937	2.031
BTB#8	37	168	1024	2.048	2.376	2.158	1.685	2.385	2.027	2.113
BTB#8	56	140	1024	2.141	2.412	2.186	1.707	2.433	2.073	2.159
BTB#8	74	112	1024	2.181	2.421	2.204	1.708	2.450	2.089	2.175
BTB#8	93	84	1024	2.198	2.408	2.207	1.730	2.452	2.095	2.182
BTB#8	112	56	1024	2.209	2.412	2.207	1.739	2.455	2.104	2.188
BTB#8	130	28	1024	2.218	2.413	2.205	1.746	2.458	2.115	2.192
BTB#8	149	0	1024	2.223	2.416	2.205	1.752	2.459	2.129	2.197
BTB#9	0	0	8192	1.199	1.243	1.183	0.935	1.204	1.127	1.148
BTB#9	21	0	7168	1.788	2.079	2.017	1.512	2.177	1.781	1.892
BTB#9	42	0	6144	1.956	2.331	2.136	1.595	2.324	1.927	2.045
BTB#9	64	0	5120	2.139	2.380	2.188	1.639	2.428	2.007	2.130
BTB#9	85	0	4096	2.188	2.397	2.203	1.679	2.445	2.032	2.157
BTB#9	106	0	3072	2.211	2.408	2.206	1.716	2.453	2.059	2.175
BTB#9	128	0	2048	2.228	2.415	2.205	1.741	2.458	2.093	2.190
BTB#9	149	0	1024	2.223	2.416	2.205	1.752	2.459	2.129	2.197
BTB#9	170	0	0	2.136	2.344	2.176	1.652	2.436	2.083	2.138

**Table #A-9 - Table #A-14:** These tables give the average cycle savings per branch for the multi-level BTBs BTB#6 - BTB#9. Each table gives the average cycle savings for a constant total number of storage bits in the BTB design (4096 to 16384) and a set of cycle costs (costx or costly). Each BTB design has nine different configurations in each table, varying the number of bits allocated to the various levels in the BTB, allowing optimal configurations (highest average cycle savings per branch) to be spotted.

## 11. Bibliography

- [AMD88] Advanced Micro Devices, "Am29000 Streamlined Instruction Processor User Manual," Advanced Micro Devices, 1988
- [Amdahl76] Amdahl, "Amdahl 470 V/6 Machine Reference Manual," Amdahl, Sunnyvale, Calif., 1976
- [Anderson67] D.W. Anderson, F.J. Sparacio, and R.M. Tomasulo, "The IBM System/360 Model 91: Machine Philosophy and Instruction-Handling," IBM Journal of Research and Development, Jan. 1967, pp. 8-24
- [DeRosa87] J.A. DeRosa and H.M. Levy, "An Evaluation of Branch Architectures," Proc. 14th Ann. Symp. Computer Architecture, 1987, pp. 10-16
- [DEC77] Digital Equipment Corporation, "VAX 11/780 Architecture Handbook," Digital Equipment Corporation, 1977
- [Ditzel87] D.R. Ditzel and H.R. McLellan, "Branch Folding in the CRISP Microprocessor: Reducing Branch Delay to Zero," Proc. 14th Ann. Symp. Computer Architecture, 1987, pp. 2-9.
- [Doduc89] N. Doduc, "Fortran Execution Time Benchmark," unpublished report, V.20, March 1989
- [Elrod70] T.H. Elrod, "The CDC 7600 and SCOPE 76," Datamation, Apr. 1970, pp. 80-85
- [Folger83] D. Folger and E. Basart, "Computer Architecture - Designing for Speed," Spring COMPCON 1983, pp. 25-31
- [Fujitsu87] Fujitsu Microelectronics, Inc., "MB86900 RISC Processor Architecture Manual," Fujitsu Microelectronics, Inc., 1987
- [Gaulding75] S.N. Gaulding and D.P. Madison, Jr., "Optimization of Scalar Instructions for the Advanced Scientific Computer," Spring COMPCON 1975, pp. 189-193
- [Grochowski86] E.T. Grochowski, "An Instruction Tracer for the Motorola 68010," UC Berkeley Masters Report, 1986
- [Gross82] T.R. Gross and J.L. Hennessy, "Optimizing Delayed Branches," Proc. 15th Ann. Workshop on Microprogramming, Oct. 1982, pp. 114-120
- [Henry83] R.R. Henry, "VAX Address and Instruction Traces," unpublished report, 1983
- [Hill87] M.D. Hill, "Aspects of Cache Memory and Instruction Buffer Performance," UC Berkeley Technical Report UCB/CSD 87/381, November 1987
- [Hill89] M. Hill and A.J. Smith, "Evaluating Associativity in CPU Caches," Univ. of Wisconsin Computer Sciences Technical Report #823, Feb. 1989, to appear in IEEE TOC, Dec. 1989
- [Hinton89] G. Hinton, "80960 - Next Generation," Spring COMPCON 1989, pp. 13-17
- [Hollingsworth89] W. Hollingsworth, H. Sachs, and A.J. Smith, "The CLIPPER Processor: Instruction Set Architecture and Implementation," Communications of the ACM, Feb. 1989, pp. 200-219
- [Hwang84] K. Hwang and F.A. Briggs, "Computer Architecture and Parallel Processing," McGraw-Hill, 1984, pp. 714-729
- [IBM73] IBM, "IBM Maintenance Library System/370 Model 168 Theory of Operation/Diagrams Manual," Vol. 2, 1973, IBM, Poughkeepsie, N.Y.

- [IBM78] IBM, "IBM Maintenance Library 3033 Processor Complex Theory of Operation/Diagrams Manual," Vols. 1-3, Jan. 1978, IBM, Poughkeepsie, N.Y.
- [Jordan83] H.F. Jordan, "Performance Measurements on HEP - A Pipelined MIMD Computer," Proc. 10th Ann. Symp. Computer Architecture, 1983, pp. 207-212
- [Kane89] G. Kane, "MIPS RISC Architecture," Prentice Hall, 1989
- [Kartashev90] S.P. Kartashev and S.I. Kartashev, "Supercomputing Systems," Van Nostrand Reinhold, 1990, pp. 106-153
- [Kogge81] P.M. Kogge, "The Architecture of Pipelined Computers," McGraw-Hill, 1981
- [Lee84] J.K.F. Lee and A.J. Smith, "Branch Prediction Strategies and Branch Target Buffer Design," Computer, Jan. 1984, pp. 6-22.
- [Lewis88] D.K. Lewis, J.P. Costello, and D.M. O'Connor, "Design Tradeoffs for a 40 MIPS (Peak) CMOS 32-bit Microprocessor," Proc. IEEE Int'l Conf. Computer Design: VLSI in Computers & Processors, Oct. 1988, pp. 110-113
- [Lilja88] D.J. Lilja, "Reducing the Branch Penalty in Pipelined Processors," Computer, July 1988, pp. 47-55.
- [Manuel87] T. Manuel, "Getting Mainframe Power out of a CISC Supermicro," Electronics, Sept. 3, 1987, pp. 66-69
- [McFarling86] S. McFarling and J. Hennessy, "Reducing the Cost of Branches," Proc. 13th Ann. Symp. Computer Architecture, 1986, pp. 396-403.
- [Motorola89] Motorola, "M68000 8-/16-/32-Bit Microprocessors User's Manual," Prentice Hall, 1989
- [Murphey70] J.O. Murphey and R.M. Wade, "The IBM 360/195," Datamation, Apr. 1970, pp. 72-79
- [Patterson81] D.A. Patterson and H. Sequin, "RISC-I: A Reduced Instruction Set VLSI Computer," Proc. Eighth Symp. Computer Architecture, May 1981, pp. 443-458
- [Peuto77] B.L. Peuto and L.J. Shustek, "An Instruction Timing Model of CPU Performance," Proc. 4th Ann. Symp. Computer Architecture, Mar. 1977, pp.165-178
- [Ramamoorthy77] C.V. Ramamoorthy and H.F. Li, "Pipeline Architectures," Computing Surveys, March 1977, pp. 61-102.
- [Rau77] B.R. Rau and G.E. Rossman, "The Effect of Instruction Fetch Strategies upon the Performance of Pipelined Instruction Units," 4th Ann. Symp. Computer Architecture, 1977, pp. 80-87
- [Russell78] R.M. Russell, "The CRAY-1 Computer System," Comm. of the ACM, Jan. 1978, pp. 63-72
- [Sakamura87] K. Sakamura, "TRON Project 1987," Springer-Verlag, 1987
- [Shar74] L.E. Shar and E.S. Davidson, "A Multiminiprocessor System Implemented Through Pipelining," IEEE Computer, Feb. 1974, pp. 42-51
- [SmithJ81] J.E. Smith, "A Study of Branch Prediction Strategies," Proc. Eighth Symp. Computer Architecture, May 1981, pp. 135-148

- [SmithJ83] J.E. Smith and J.R. Goodman, "A Study of Instruction Cache Organizations and Replacement Policies," Proc. 10th Symp. Computer Architecture, June 1983, pp. 132-137
- [SmithAJ85] A.J. Smith, "Cache Evaluation and the Impact of Workload Choice," Proc. 12th Symp. Computer Architecture, June 1985, pp. 64-74.
- [SmithAJ87] A.J. Smith, "Line (Block) Size Choice for CPU Cache Memories," IEEE Transactions on Computing, Sept. 1987, pp. 1063-1075.
- [SmithAJ89] A.J. Smith, conversation on branch target buffers, Spring 1989
- [Stiles89a] D.R. Stiles and H.L. McFarland, "Pipeline Control for a Single Cycle VLSI Implementation of a Complex Instruction Set Computer," Spring COMP-CON 1989, pp. 504-508
- [Stiles89b] D.R. Stiles (of NexGen Microsystems), personal interview concerning branch prediction cache of NexGen processor, Sept. 25, 1989
- [Thornton64] J.E. Thornton, "Parallel Operation in the Control Data 6600," AFIPS Fall Joint Comp. Conf., 1964, pp. 33-40
- [UCB87] U.C. Berkeley CAD/IC Group, "SPICE2G.6," March 1987
- [Walter89] S. Walter (of Edgecore), personal interview concerning branch cache in Edge 2000, Sept 20, 1989
- [Widdoes77] L.C. Widdoes, Jr., "Jump Prediction," Stanford University (unpublished draft), February 1977
- [Yoshida87] T. Yoshida and T. Enomoto, "The Mitsubishi VLSI CPU in the TRON Project," IEEE Micro, Apr. 1987, p. 24