Reduction of Control Hazards (Branch) Stalls
with Dynamic Branch Prediction

- So far we have dealt with <u>control hazards</u> in instruction pipelines by:

  ISA Support Needed

  1 – Assuming that the branch is not taken (i.e stall when branch is taken).

  2 – Reducing the branch penalty by resolving the branch early in the pipeline
    - Branch penalty if branch is taken = stage resolved - 1

  3 – Branch delay slot and canceling branch delay slot. (ISA support needed)

  4 – Compiler-based static branch prediction encoded in branch instructions
    - Prediction is based on program profile or branch direction
    - ISA support needed.

How to further reduce the impact of branches on pipeline processor performance ?

- **<u>Dynamic Branch Prediction</u>:**

  | Why? Better branch prediction accuracy than static prediction and thus fewer branch stalls | **Why?** |

  – Hardware-based schemes that utilize run-time behavior of branches to make dynamic predictions:

    How?
    - Information about outcomes of previous occurrences of branches are used to dynamically predict the outcome of the current branch.

- **<u>Branch Target Buffer (BTB)</u>:** *(Goal: zero stall taken branches)*

  – A hardware mechanism that aims at reducing the stall cycles resulting from <u>correctly predicted taken branches to zero cycles.</u>

4th Edition: Static and Dynamic Prediction in ch. 2.3, BTB in Ch. 2.9
(3rd Edition: Static Pred. in Ch. 4.2 Dynamic Pred. in Ch. 3.4, BTB in Ch. 3.5)

EECC551 - Shaaban

Static Conditional Branch Prediction

- **Branch prediction schemes can be classified <u>into static (at compilation</u> time) and <u>dynamic (at runtime)</u> schemes.**

Branch Encoding	X

  X = Static
  Prediction bit
  X = 0 Not Taken
  X = 1 Taken

- **Static methods are carried out by the compiler. They are static because the prediction is already known before the program is executed.**

- **Static Branch prediction is encoded in branch instructions using one prediction (<u>or branch direction hint</u>) bit = 0 = Not Taken, = 1 = Taken**
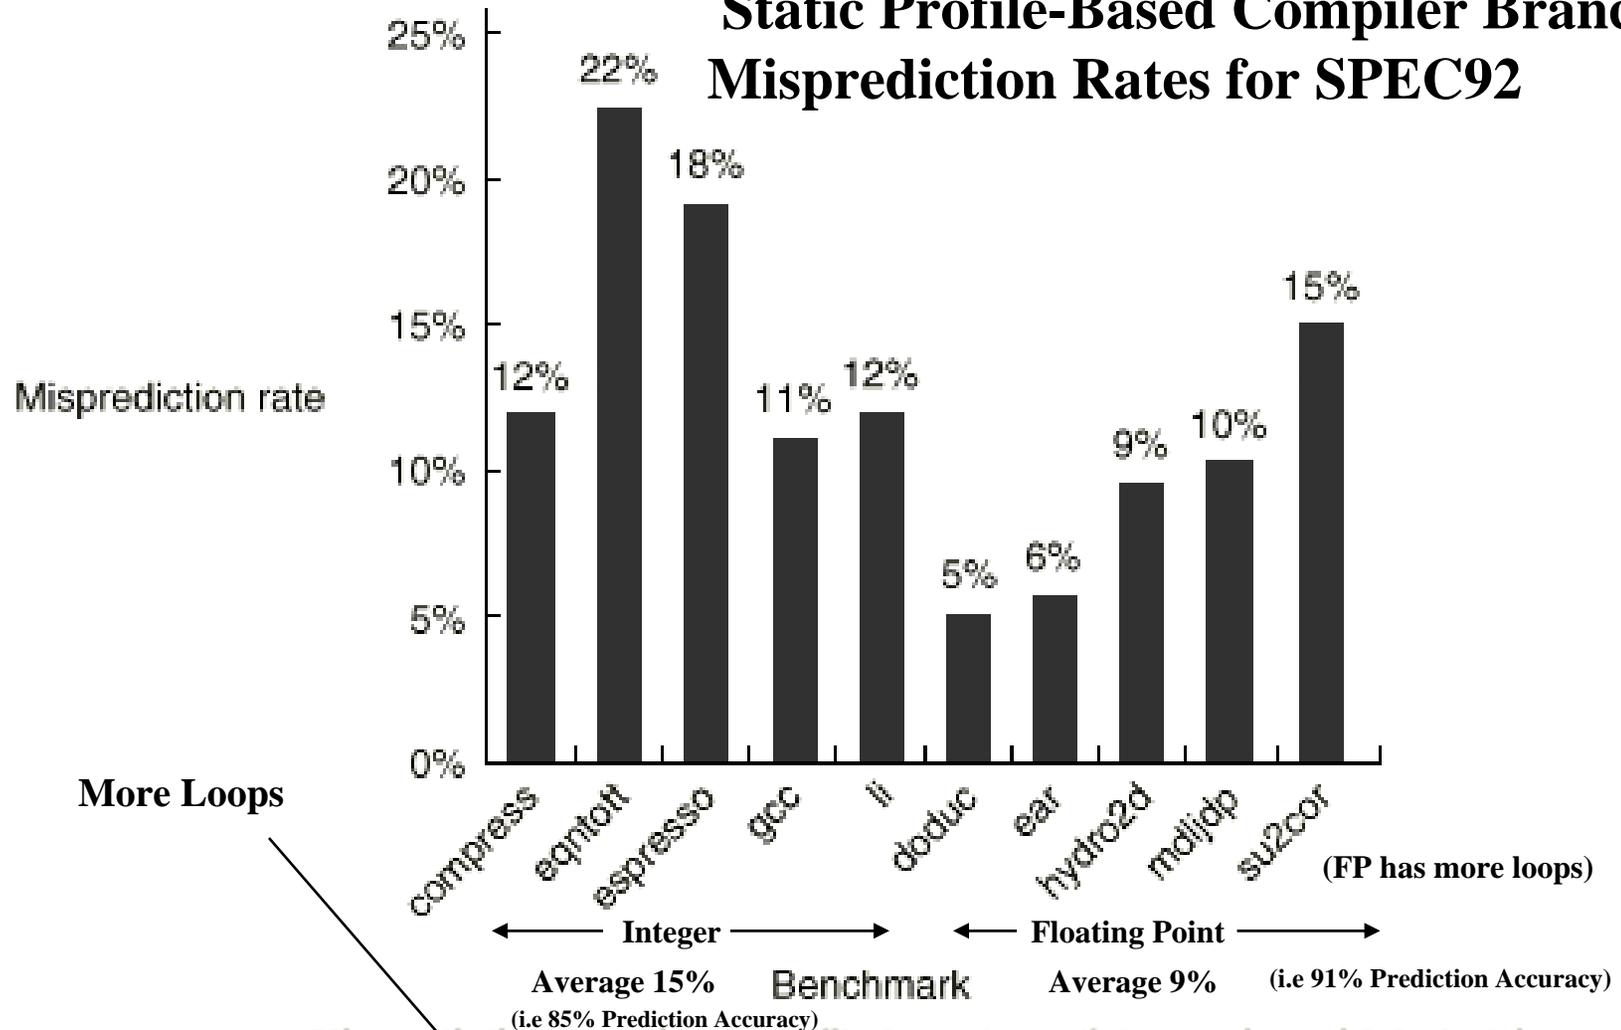 - **Must be supported by ISA, Ex: HP PA-RISC, PowerPC, UltraSPARC**

How? **Two basic methods to statically predict branches at compile time:**

1 – **<u>Use the direction of a branch to base the prediction on.</u> Predict backward branches (branches which decrease the PC) to be taken (e.g. loops) and forward branches (branches which increase the PC) not to be taken.**

2 – **<u>Profiling</u> can also be used to predict the outcome of a branch.**
  - **A number runs of the program are used to collect program behavior information (i.e. if a given branch is likely to be taken or not)**
  - **This information is included in the opcode of the branch (one bit branch direction hint) as the static prediction.**

<u>Static prediction was previously discussed in lecture 2</u>
4th edition: in Chapter 2.3, 3rd Edition: In Chapter 4.2

EECC551 - Shaaban

Static Profile-Based Compiler Branch Misprediction Rates for SPEC92

Misprediction rate

25%

22%

20% 18%

15% 15%

12% 12% 11%

10% 9% 10%

6%

5% 5%

0%

compress eqntott espresso gcc li doduc ear hydro2d mdljdp su2cor

More Loops

(FP has more loops)

←――― **Integer** ―――→ ←――― **Floating Point** ―――→

Average 15% Benchmark **Average 9%** **(i.e 91% Prediction Accuracy)**

(i.e 85% Prediction Accuracy)

Misprediction rate for a profile-based predictor varies widely but is generally better for the FP programs, which have an average misprediction rate of 9% with a standard deviation of 4%, than for the integer programs, which have an average misprediction rate of 15% with a standard deviation of 5%.

EECC551 - Shaaban

Dynamic Conditional Branch Prediction

- Dynamic branch prediction schemes are different from static mechanisms because they utilize hardware-based mechanisms that use the <u>run-time behavior of branches</u> to make <u>more accurate predictions</u> than possible using static prediction.

How?

Why?

- Usually information about outcomes of previous occurrences of branches (<u>branching history</u>) is used to dynamically predict the outcome of the current branch. The two main types of dynamic branch prediction are:

  1 – <u>One-level or Bimodal:</u> Usually implemented as a Pattern History Table (PHT), a table of usually two-bit saturating counters which is indexed by a portion of the branch address (low bits of address). (First proposed mid 1980s)
    - Also called *non-correlating* dynamic branch predictors.

  2 – <u>Two-Level Adaptive Branch Prediction.</u> (First proposed early 1990s).
    - Also called *correlating* dynamic branch predictors.

- To reduce the stall cycles resulting from correctly predicted taken branches to <u>zero cycles</u>, a <u>Branch Target Buffer (BTB)</u> that includes the addresses of conditional branches that were taken along with their targets is added to the <u>fetch</u> stage.

BTB

BTB discussed next →

4th Edition: Dynamic Prediction in Chapter 2.3, BTB in Chapter 2.9
(3rd Dynamic Prediction in Chapter 3.4, BTB in Chapter 3.5)

Branch Target Buffer (BTB)

- **Effective branch prediction requires the <u>target of the branch at an early pipeline stage.</u> (resolve the branch early in the pipeline)**

  – **One can use additional adders to calculate the target, as soon as the branch instruction is decoded. This would mean that one has to wait until <u>the ID stage</u> before the target of the branch can be fetched, taken branches would be fetched with a one-cycle penalty (this was done in the enhanced MIPS pipeline Fig A.24).**

- **To avoid this problem and to achieve <u>zero stall cycles for taken branches</u>, one can use <u>a Branch Target Buffer (BTB).</u>**

- **A typical BTB is an associative memory where the <u>addresses of taken branch instructions are stored together with their target addresses.</u>**

- **The <u>BTB is is accessed in Instruction Fetch (IF)</u> cycle and provides answers to the following questions while the current instruction is being fetched:**

  1 – **Is the instruction a branch?**
  2 – **If yes, is the branch predicted taken?**
  3 – **If yes, what is the branch target?**

- **Instructions are fetched from the target stored in the BTB in case the branch is predicted-taken and found in BTB.**

- **After the branch has been resolved the BTB is updated. If a branch is encountered for the first time a new entry is created once it is resolved as taken.**

Goal of BTB: Zero stall taken branches

EECC551 - Shaaban

Basic Branch Target Buffer (BTB)

PC of instruction to fetch

Fetch instruction from instruction memory (I-L1 Cache)

Is the instruction a branch?

1

Look up **(for address match)**

Predicted PC

Instruction Fetch

Number of entries in branch-target buffer

IF

Branch Address

Branch Target if predicted taken **3**

2

Branch Taken?

Branch Targets

BTB is accessed in Instruction Fetch (IF) cycle

=

No: instruction is not predicted to be branch. Proceed normally

Branch predicted taken or untaken

Yes: then instruction is branch and predicted PC should be used as the next PC

i.e target

0 = NT = Not Taken

1 = T = Taken

A branch-target buffer.

Goal of BTB: Zero stall taken branches
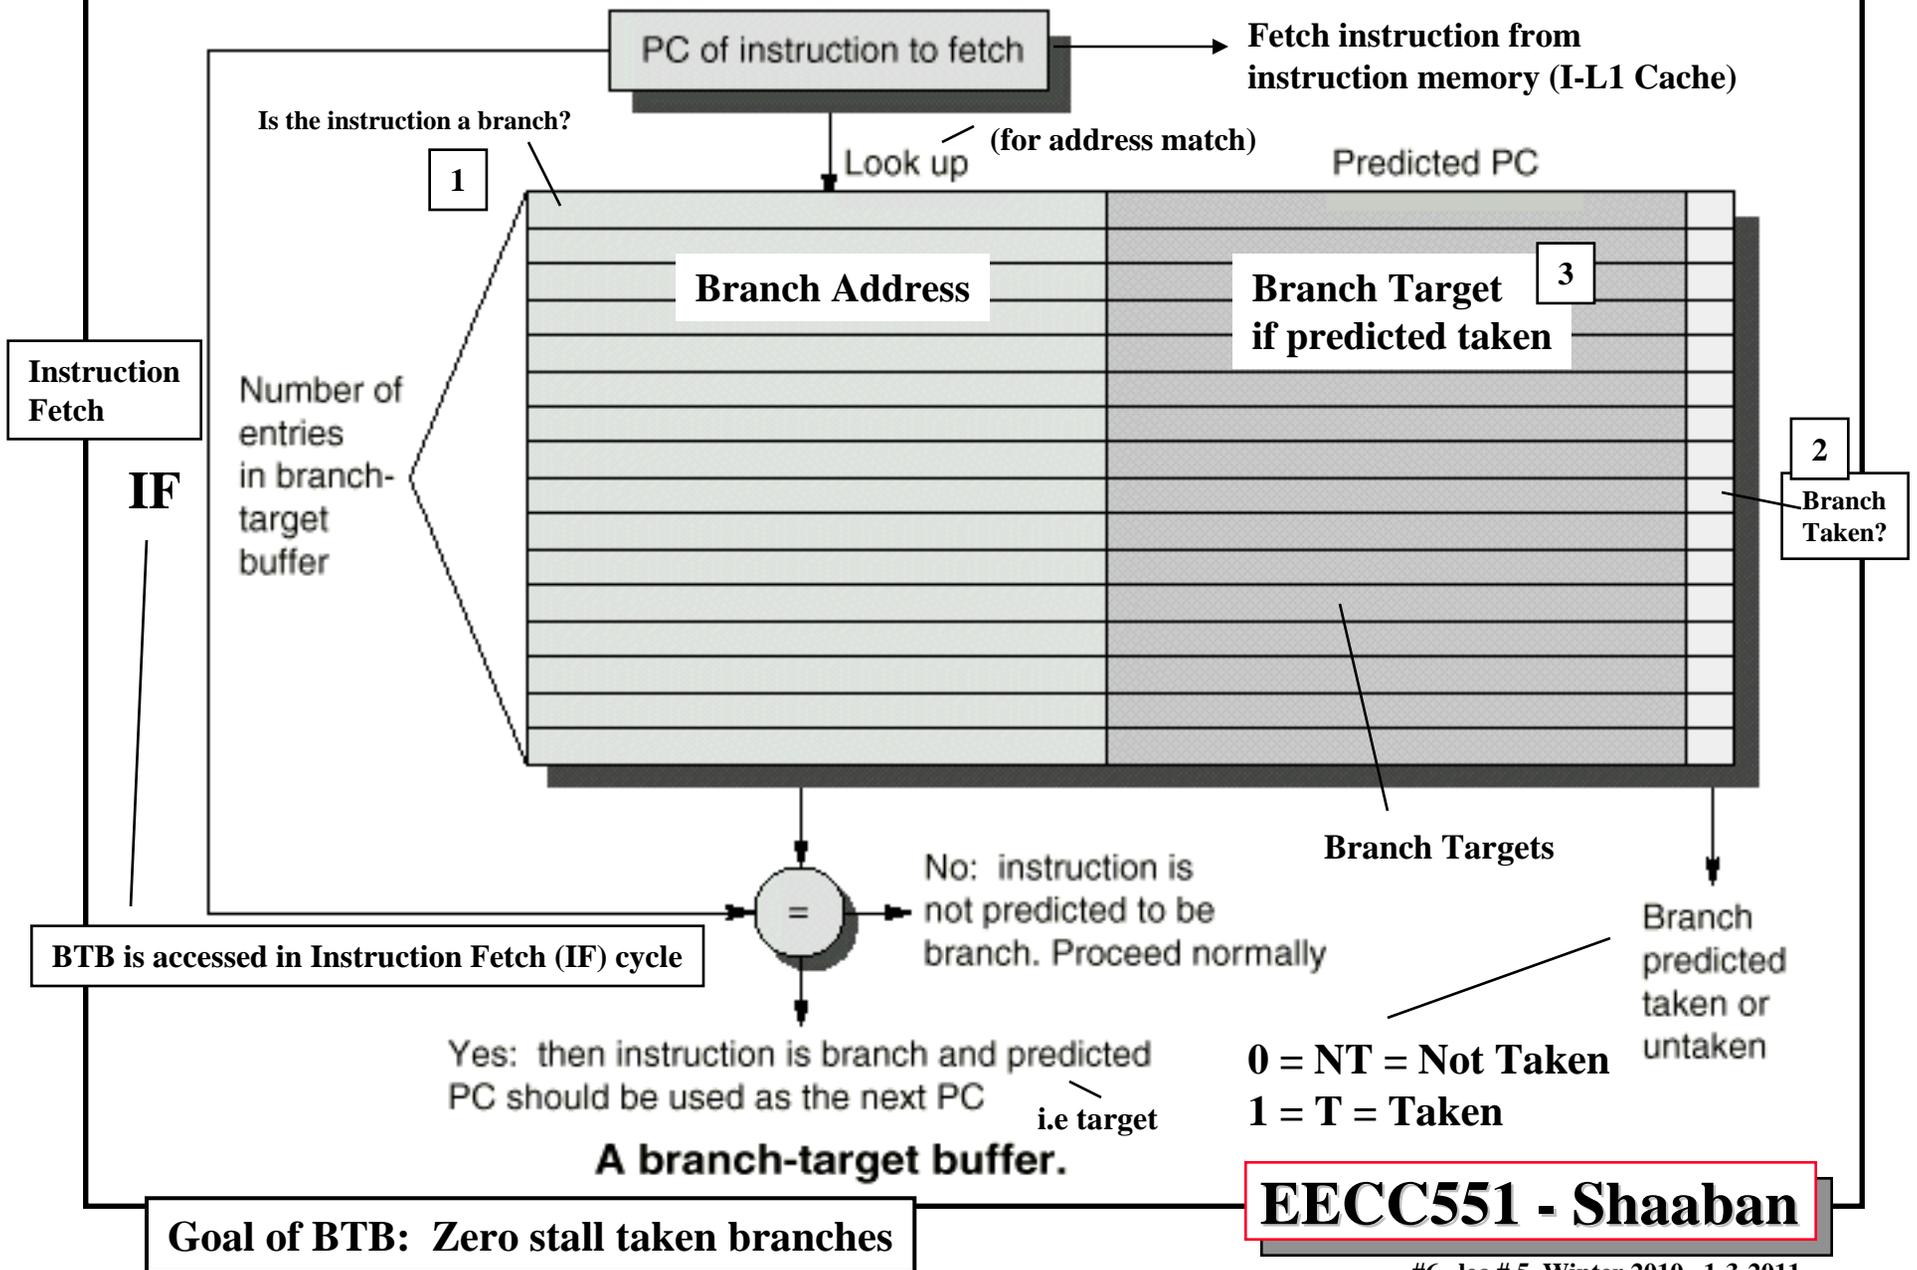
BTB Operation

Here, branches are assumed to be resolved in ID

BTB Lookup

Send PO to memory and branch-target buffer

Entry found in branch-target buffer?

Yes → Send out predicted PC → Taken branch?

Yes → Branch correctly predicted; continue execution with no stalls

No → Mispredicted branch, kill fetched instruction; restart fetch at other target; delete entry from target buffer

No → Is instruction a taken branch?

Yes → Enter branch IO and next PC into branch target buffer

No → Normal instruction execution

Update BTB

One more stall to update BTB
Penalty = 1 + 1 = 2 cycles

IF

ID

EX

Branch Penalty Cycles
Using A Branch-Target Buffer (BTB)

Base Pipeline Taken Branch Penalty = 1 cycle (*i.e. branches resolved in ID*)

i.e In BTB?

No	Not Taken	Not Taken	0
Instruction in buffer	**Prediction**	**Actual branch**	**Penalty cycles**
Yes	Taken	Taken	0
Yes	Taken	Not taken	2
No		Taken	2

BTB Goal: Taken Branches with zero stalls

Assuming one more stall cycle to update BTB
Penalty = 1 + 1 = 2 cycles

Penalties for all possible combinations of whether the branch is in the buffer and what it actually does, assuming we store only taken branches in the buffer.

Basic Dynamic Branch Prediction

- **Simplest method: (<u>One-Level or Non-Correlating</u>)**
  - – **A branch prediction buffer or <u>Pattern History Table (PHT)</u> indexed by <u>low address bits of the branch instruction</u>.**
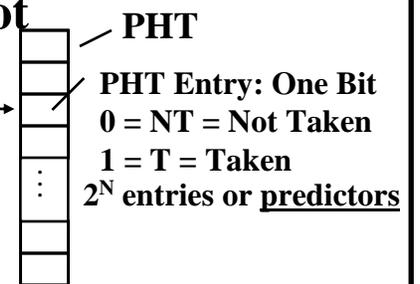
    Saturating counter

  - – **Each buffer location (or PHT entry or <u>predictor</u>) contains one bit indicating whether the branch was recently taken or not**
    - **e.g   0 = not taken , 1 =taken** ^{NT} ^{T} ^{T} **N Low Bits of Branch Address** →

      PHT

      Predictor = Saturating Counter → **0** ^{NT} **1**
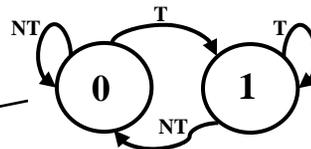
      PHT Entry: One Bit
      0 = NT = Not Taken
      1 = T = Taken
      2^N entries or <u>predictors</u>

  - – **Always mispredicts in first and last loop iterations.**

- **To improve prediction accuracy, <u>two-bit prediction is used</u>:**

  (Smith Algorithm, 1985)

  - – **A prediction must miss twice before it is changed.**

    Why 2-bit Prediction?

    - Thus, a branch involved in a loop will be mispredicted only once when encountered the next time as opposed to twice when one bit is used.

  - – **Two-bit prediction is a specific case of n-bit saturating counter incremented when the branch is taken and decremented when the branch is not taken.**   The counter (predictor) used is updated after the branch is resolved

    Smith Algorithm

  - – **<u>Two-bit saturating counters (predictors)</u> are usually always used based on observations that the performance of two-bit PHT prediction is comparable to that of n-bit predictors.**

4th Edition: In Chapter 2.3 (3rd Edition: In Chapter 3.4)

EECC551 - Shaaban

One-Level Bimodal Branch Predictors
Pattern History Table (PHT)

Most common one-level implementation

2-bit saturating counters (predictors)

Sometimes referred to as
Decode History Table (DHT)
or
Branch History Table (BHT)

Indexed by

N Low Bits of Branch Address

High bit determines branch prediction

0 = NT = Not Taken
1 = T = Taken

X   X

X   X

Prediction Bits

Table (PHT) has 2^N entries
(also called **predictors**) .

2-bit saturating counters

0	**0**	**Not Taken (NT)**
0	**1**	
1	**0**	
1	**1**	**Taken (T)**

Example:

For N =12
Table has $2^N = 2^{12}$ entries
$= 4096 = 4k$ entries

Number of bits needed = $2 \times 4k = 8k$ bits
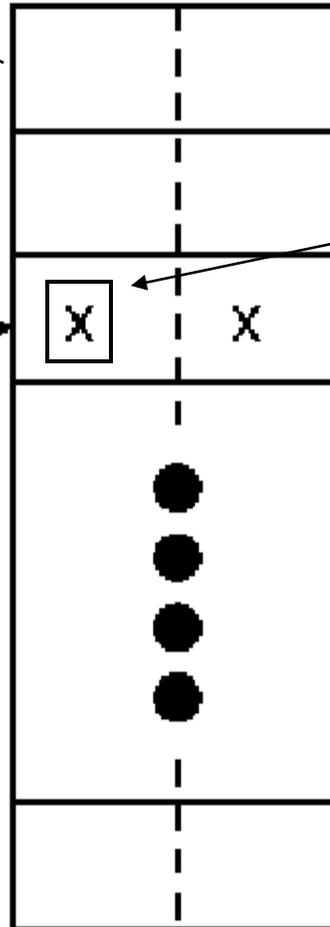
When to update

Update counter after branch is resolved:
-Increment counter used if branch is taken
- Decrement counter used if branch is not taken

What if different branches map to the same predictor (counter)?
This is called branch address aliasing and leads to interference with current branch prediction by other branches and may lower branch prediction accuracy for programs with aliasing.

EECC551 - Shaaban

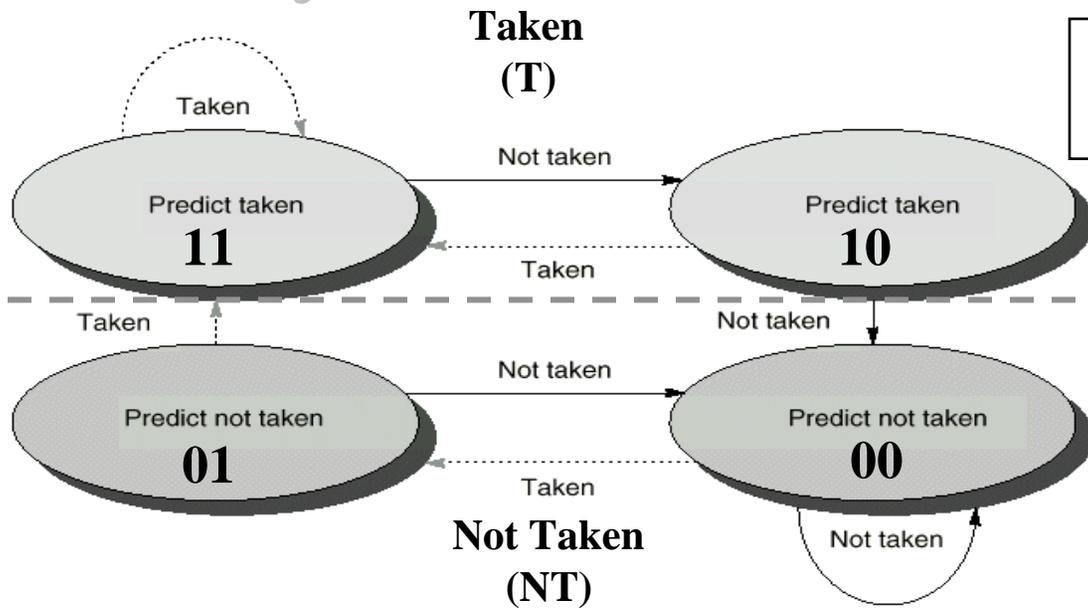# Basic Dynamic Two-Bit Branch Prediction:

Taken (T)

Taken

Predict taken
11

Not taken

Predict taken
10

Taken

Taken

Not taken

Not taken

Predict not taken
01

Not taken

Predict not taken
00

Taken

Not taken

Not Taken (NT)

0	0	**Not Taken (NT)**
0	1	
1	0	**Taken (T)**
1	1	

Or Two-bit saturating counter predictor state transition diagram (Smith Algorithm):

Not Taken (NT)

Taken (T)   **Taken (T)**   **Taken (T)**   **Taken (NT)**

_Predict Not Taken **00**

_Predict Not Taken **01**

_Predict Taken **10**

_Predict Taken **11**

Not Taken (NT)

Not Taken (NT)   Not Taken (NT)   Not Taken (NT)

Taken (T)
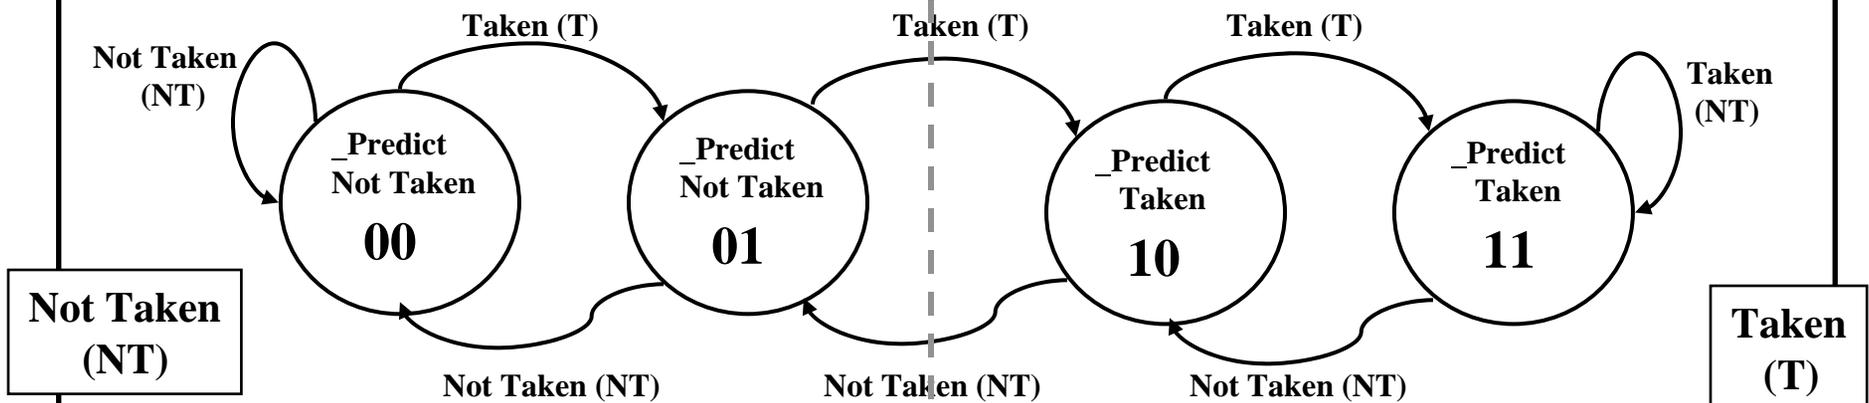
The two-bit predictor used is updated after the branch is resolved

N=12
$2^N = 4096$

Prediction Accuracy of A 4096-Entry Basic One-Level Dynamic Two-Bit Branch Predictor

FP

nasa7 — 1%
matrix300 — 0%
tomcatv — 1%
doduc — 5%

SPEC89 benchmarks

spice — 9%
fpppp — 9%

Misprediction Rate:

Integer average 11%
FP average 4%

gcc — 12%

(Lower misprediction rate due to more loops)

Integer

espresso — 5%
eqntott — 18%
li — 10%
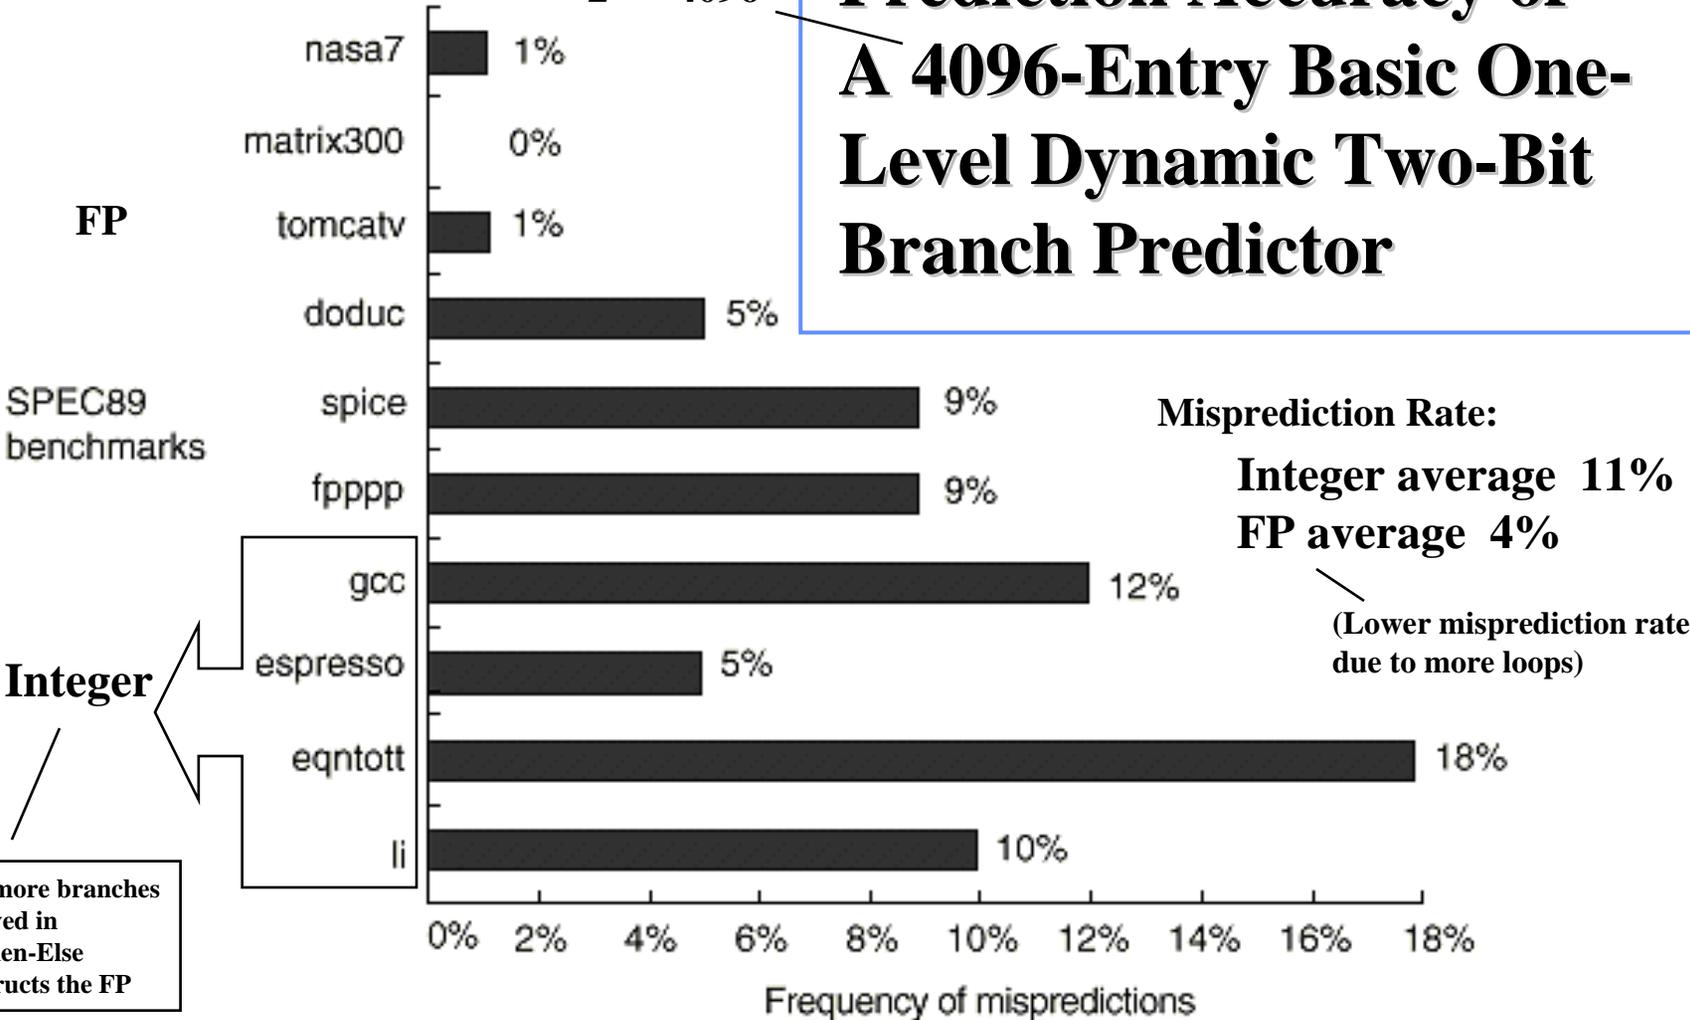
Has, more branches involved in IF-Then-Else constructs the FP

Frequency of mispredictions

Prediction accuracy of a 4096-entry two-bit prediction buffer for the SPEC89 benchmarks.

From The Analysis of Static Branch Prediction :
MIPS Performance Using Canceling Delay Branches

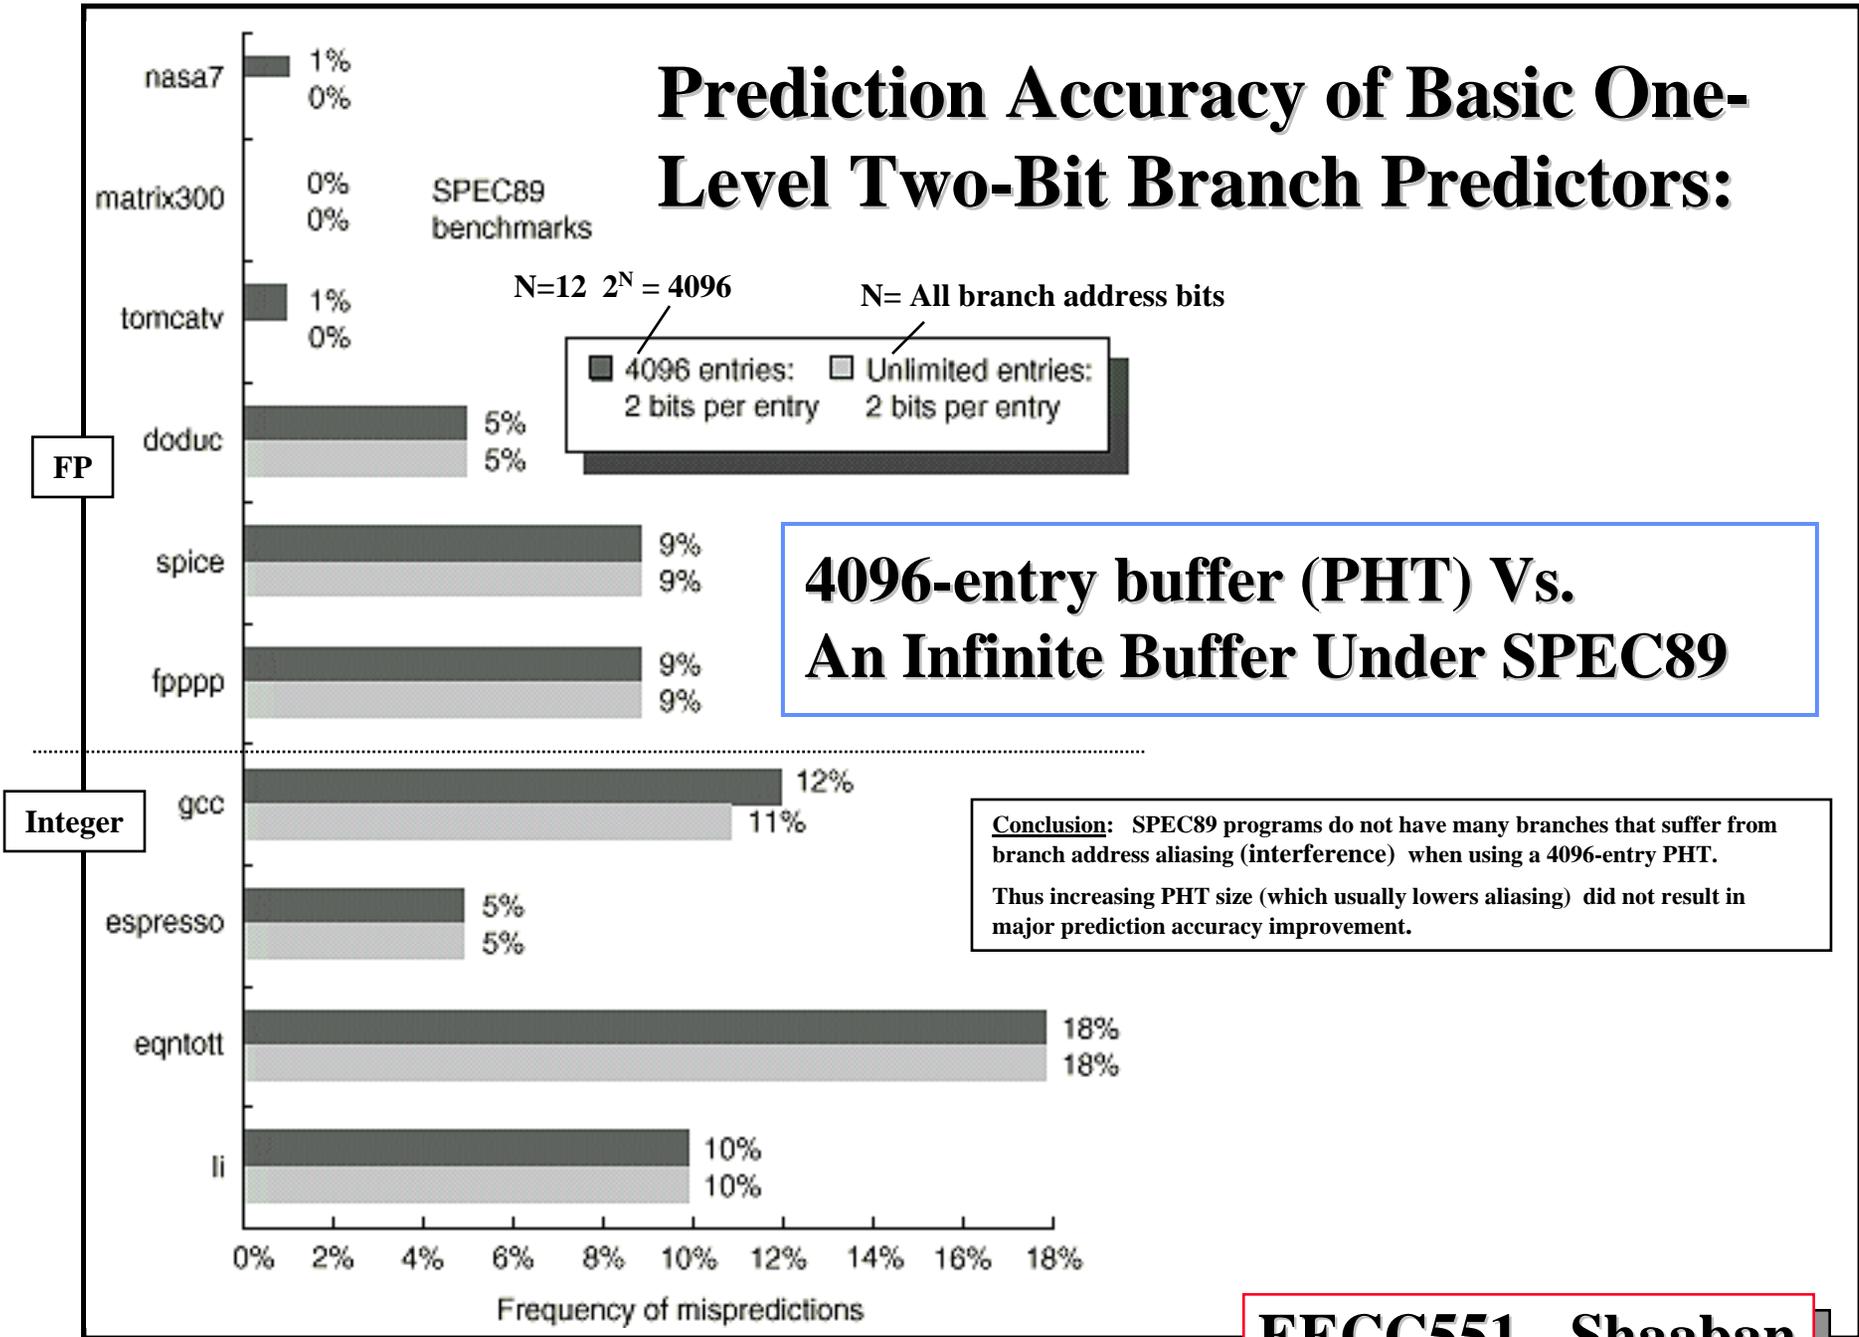| Benchmark	% conditional branches	% conditional branches with empty slots	% conditional branches that are cancelling	% cancelling branches that are cancelled	% branches with cancelled delay slots	Total % branches with empty or cancelled delay slot
compress	14%	18%	31%	43%	13%	31%
eqntott	24%	24%	50%	24%	12%	36%
espresso	15%	29%	19%	21%	4%	33%
gcc	15%	16%	33%	34%	11%	27%
li	15%	20%	55%	48%	26%	46%
Integer average	17%	21%	38%	34%	13%	35%
doduc	8%	33%	12%	62%	8%	41%
ear	10%	37%	36%	14%	5%	42%
hydro2d	12%	0%	69%	24%	16%	17%
mdljdp2	9%	0%	86%	10%	8%	8%
su2cor	3%	7%	17%	57%	10%	17%
FP average	8%	16%	44%	34%	9%	25%
Overall average	12%	18%	41%	34%	11%	30%

Delayed and cancelling delay branches for MIPS allow branch hazards to be hidden 70% of the time on average for these 10 SPEC benchmarks.

70% Static Branch Prediction Accuracy

(**repeated here from lecture2**)

Prediction Accuracy of Basic One-Level Two-Bit Branch Predictors:

nasa7 — 1% / 0%

matrix300 — 0% / 0%

tomcatv — 1% / 0%

doduc — 5% / 5%

spice — 9% / 9%

fpppp — 9% / 9%

gcc — 12% / 11%

espresso — 5% / 5%

eqntott — 18% / 18%

li — 10% / 10%

FP / **Integer**

SPEC89 benchmarks

N=12 $2^N = 4096$

N= All branch address bits

■ 4096 entries: 2 bits per entry   □ Unlimited entries: 2 bits per entry

Frequency of mispredictions — 0% 2% 4% 6% 8% 10% 12% 14% 16% 18%

4096-entry buffer (PHT) Vs. An Infinite Buffer Under SPEC89

Conclusion: SPEC89 programs do not have many branches that suffer from branch address aliasing (**interference**) when using a 4096-entry PHT.

Thus increasing PHT size (which usually lowers aliasing) did not result in major prediction accuracy improvement.

Correlating Branches

Recent branches are possibly correlated: The behavior of recently executed branches affects prediction of current branch.

Occur in branches used to implement if-then-else constructs
Which are more common in integer than floating point code

Example:

Here aa = R1   bb = R2

		DSUBUI	R3, R1, #2	; R3 = R1 - 2
B1	if (aa==2)	BNEZ	R3, L1	; B1 (aa!=2)
	aa=0; (not taken)	DADD	R1, R0, R0	; aa==0 **B1 not taken**
		L1: DSUBUI	R3, R2, #2	; R3 = R2 - 2
B2	if (bb==2)	BNEZ	R3, L2	; B2 (bb!=2) +
	bb=0; (not taken)	DADD	R2, R0, R0	; bb==0 **B2 not taken**
		L2: DSUBUI	R3, R1, R2	; R3=aa-bb
B3	if (aa!==bb){ (not taken)	BEQZ	R3, L3	; B3 (aa==bb)

B3 taken if aa=bb

Branch **B3** is **correlated** with branches B1, B2. If **B1, B2 are both not taken, then B3 will be taken.** Using only the behavior of one branch cannot detect this behavior.

aa=bb=2

B3 in this case

Both B1 and B2 Not Taken $\rightarrow$ B3 Taken

EECC551 - Shaaban

Correlating Two-Level Dynamic GAp Branch Predictors

- Improve branch prediction by looking not only at the history of the branch in question but also at that of other branches using two levels of branch history.

- Uses two levels of branch history:

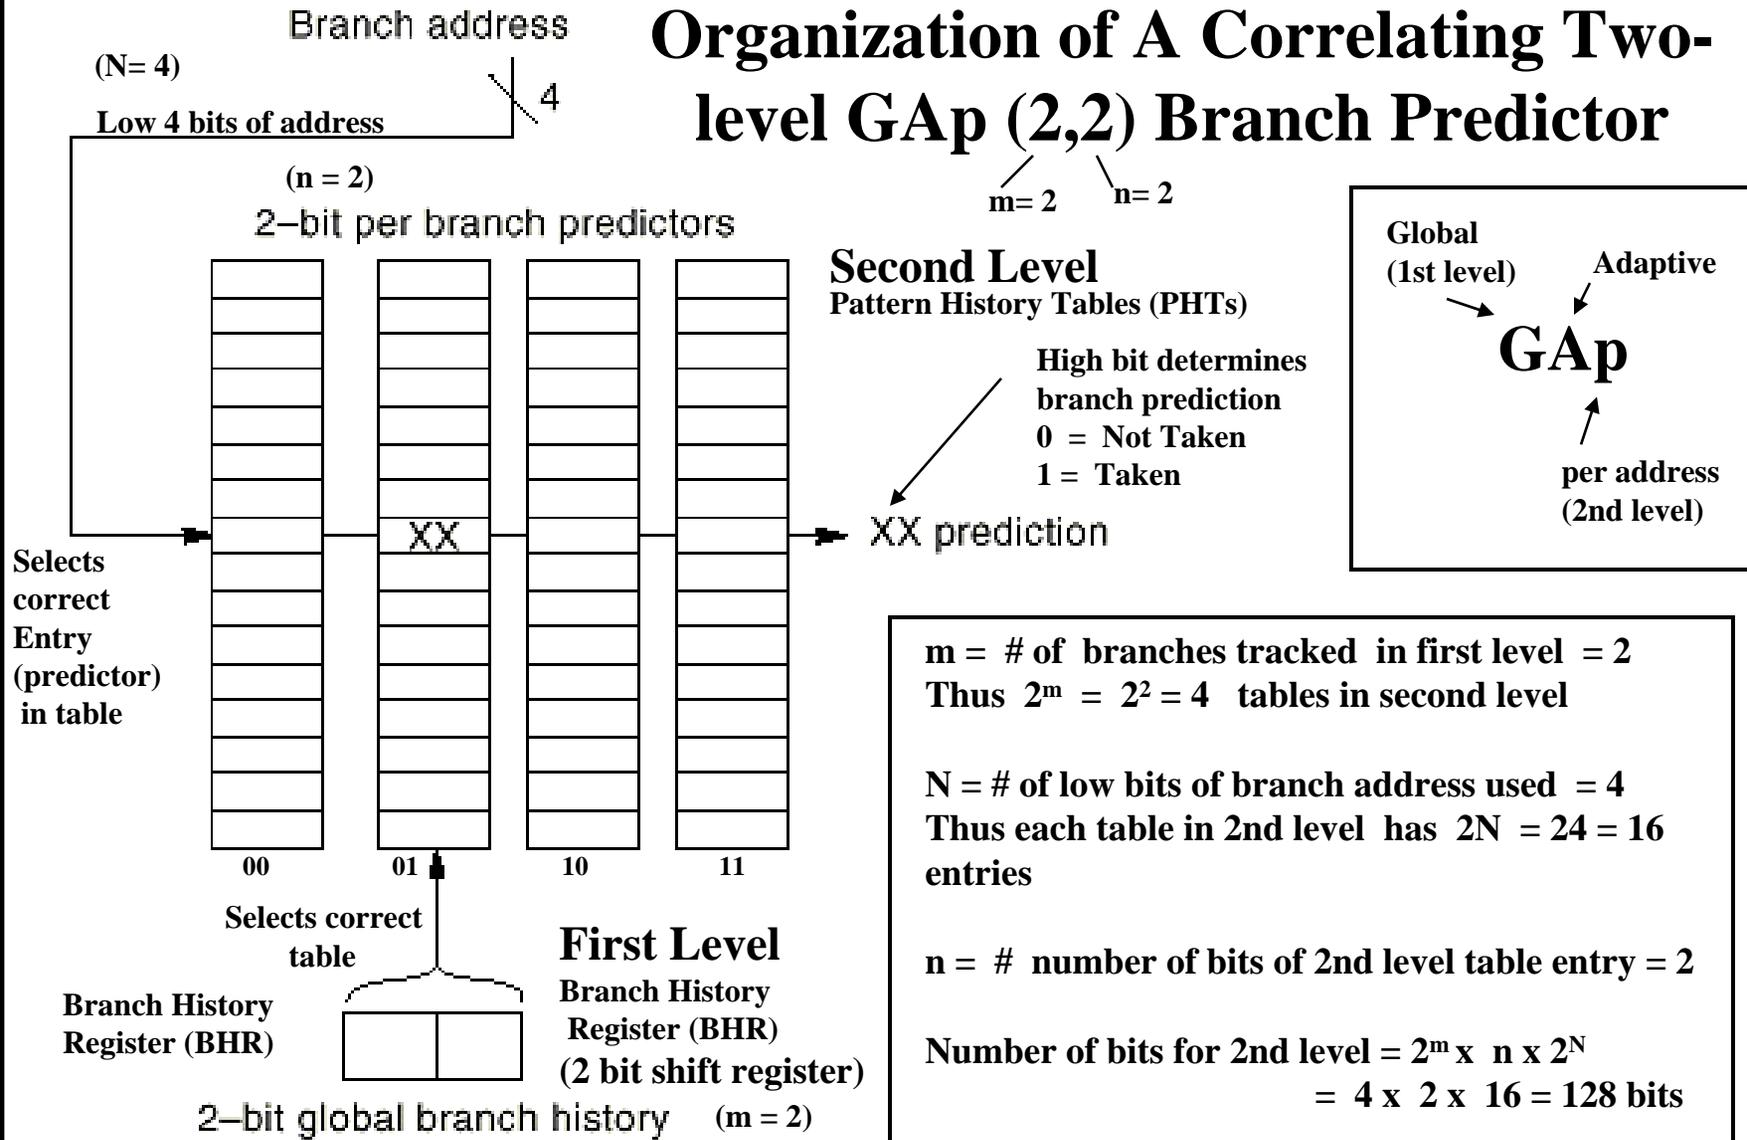| m-bit shift register	

Last Branch
0 =Not taken
1 = Taken

Branch History Register (BHR)
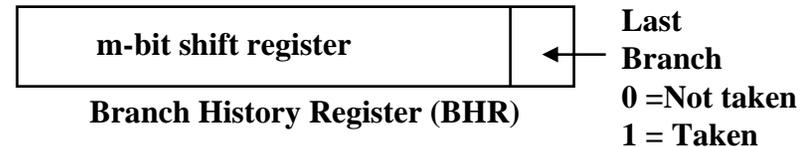
1 – **First level (global):**   BHR
  - Record the global pattern or history of the m most recently executed branches as taken or not taken.   Usually an m-bit shift register.

2 – **Second level (per branch address):** | Pattern History Tables (PHTs) |
  - 2^m prediction tables, each table entry has n bit saturating counter.
  - The branch history pattern from first level is used to select the proper branch prediction table in the second level.
  - The low N bits of the branch address are used to select the correct prediction entry (predictor)within a the selected table, thus each of the 2^m tables has 2^N entries and each entry is 2 bits counter.
  - Total number of bits needed for second level = $2^m \times n \times 2^N$ bits

- In general, the notation: **GAp (m,n) predictor means:**

GAp (m,n)

  – Record last m branches to select between 2^m history tables.
  – Each second level table uses n-bit counters (each table entry has n bits).

- Basic two-bit single-level Bimodal BHT is then a (0,2) predictor.

Organization of A Correlating Two-level GAp (2,2) Branch Predictor

Branch address

(N= 4)

Low 4 bits of address

4

(n = 2)

2–bit per branch predictors

m= 2   n= 2

Second Level
Pattern History Tables (PHTs)

**Global
(1st level)**   **Adaptive**

GAp

**per address
(2nd level)**

High bit determines
branch prediction
0 = Not Taken
1 = Taken

Selects
correct
Entry
(predictor)
in table

XX

XX prediction

| 00 | 01 | 10 | 11 |

Selects correct
table

First Level
Branch History
Register (BHR)
(2 bit shift register)

Branch History
Register (BHR)

2–bit global branch history   (m = 2)

m = # of branches tracked in first level = 2
Thus $2^m = 2^2 = 4$ tables in second level

N = # of low bits of branch address used = 4
Thus each table in 2nd level has $2N = 24 = 16$ entries

n = # number of bits of 2nd level table entry = 2

Number of bits for 2nd level = 2^m x n x 2^N
= 4 x 2 x 16 = 128 bits

A (2,2) branch-prediction buffer uses a two-bit global history to choose from among four predictors for each branch address.

Dynamic Branch Prediction: Example

if (d==0)
         d=1;
if (d==1)

BNEZ	R1, L1	; branch b1 (d!=0)
DADDIU	R1, R0, #1	; d==0, so d=1
L1: DADDIU	R3, R1, # -1	
BNEZ	R3, L2	; branch b2 (d!=1)
. . .		
L2:		

Possible execution sequences for a code fragment.

Initial value of d	d==0?	b1	Value of d before b2	d==1?	b2
0	Yes	Not taken	1	Yes	Not taken
1	No	Taken	1	Yes	Not taken
2	No	Taken	2	No	Taken

One Level

Behavior of a one-bit predictor initialized to not taken.

d=?	b1 prediction	b1 action	New b1 prediction	b2 prediction	b2 action	New b2 prediction
2	NT	T	T	NT	T	T
0	T	NT	NT	T	NT	NT
2	NT	T	T	NT	T	T
0	T	NT	NT	T	NT	NT

One Level with one-bit table entries (predictors) :   NT = 0 = Not Taken
                                                       T = 1 = Taken

Dynamic Branch Prediction: Example (continued)

if (d==0)	
	d=1;
if (d==1)	

```
            BNEZ      R1, L1        ; branch b1 (d!=0)
            DADDIU    R1, R0, #1    ;  d==0, so d=1
L1:         DADDIU    R3, R1, # -1
            BNEZ      R3, L2        ;  branch  b2  (d!=1)
. . .
L2:
```

Combinations and meaning of the taken/not taken prediction bits.

Initial value of d	d==0?	b1	Value of d before b2	d==1?	b2
0	Yes	Not taken	1	Yes	Not taken
1	No	Taken	1	Yes	Not taken
2	No	Taken	2	No	Taken

The action of the one-bit predictor with one bit of correlation, initialized to not taken/not taken.

d=?	b1 prediction	b1 action	New b1 prediction	b2 prediction	b2 action	New b2 prediction
2	NT/NT	T	T/NT	NT/NT	T	NT/T
0	T/NT	NT	T/NT	NT/T	NT	NT/T
2	T/NT	T	T/NT	NT/T	T	NT/T
0	T/NT	NT	T/NT	NT/T	NT	NT/T

Two level GAp(1,1)

m= 1   n= 1

Prediction Accuracy of Two-Bit Dynamic Predictors Under SPEC89
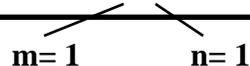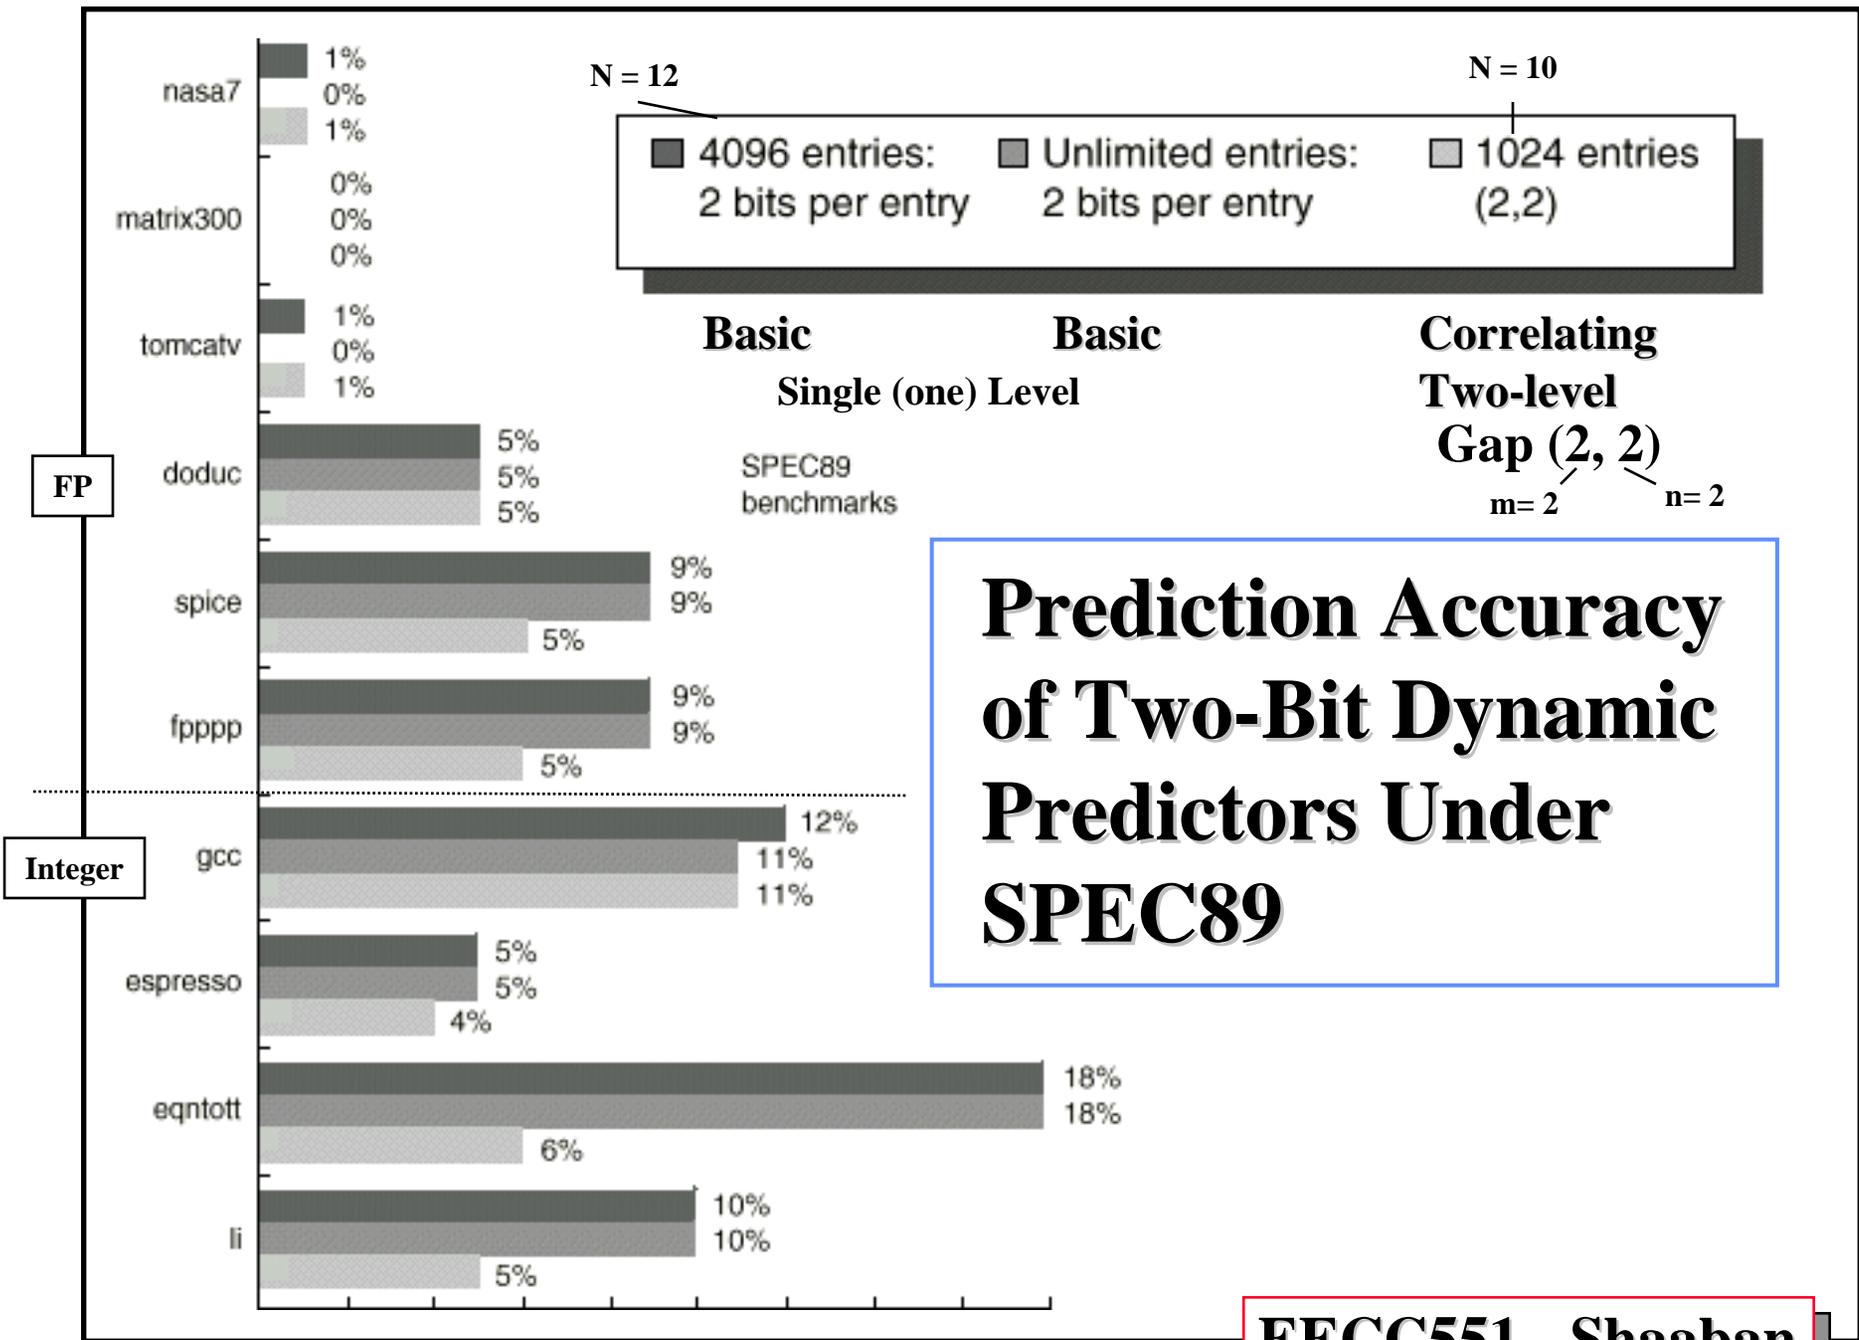
MCFarling's gshare Predictor
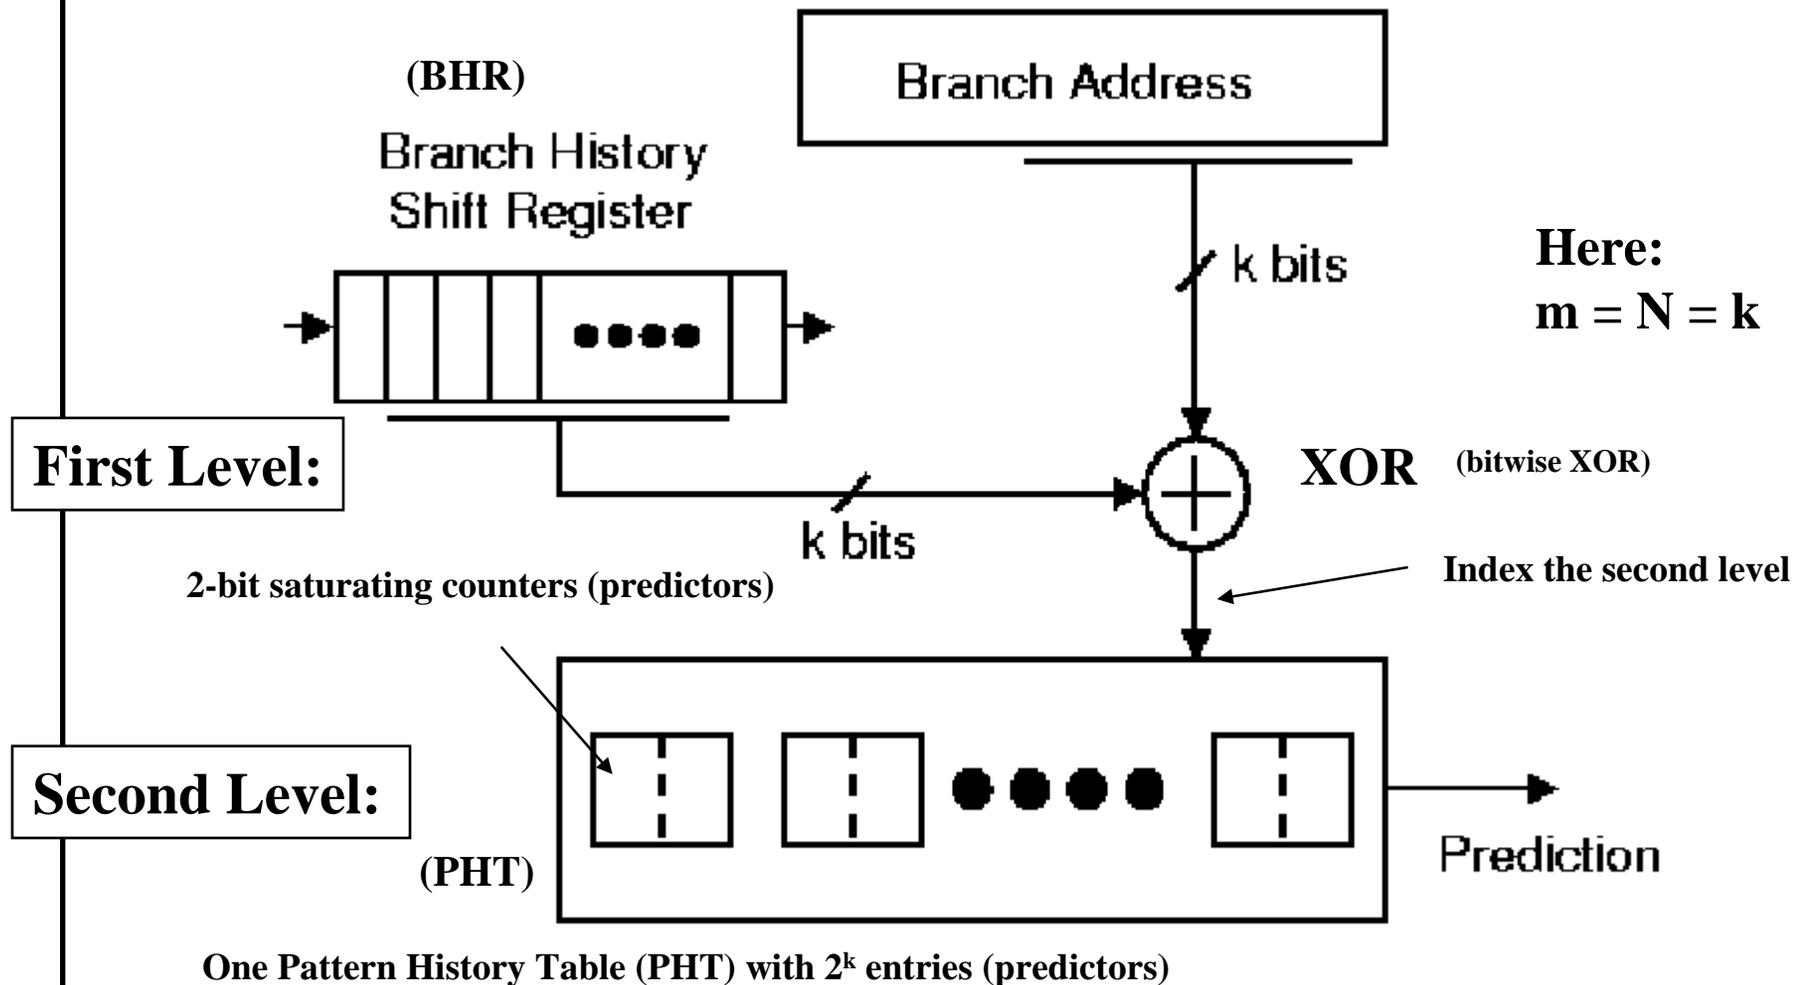
gshare = global history with index sharing

- **McFarling noted (1993) that using global history information might be less efficient than simply using the address of the branch instruction, especially for small predictors.**

- **He suggests using both <u>global history</u> (BHR) and <u>branch address</u> by hashing them together. He proposed <u>using the XOR of global branch history register (BHR) and branch address</u> since he expects that this value has more information than either one of its components. The result is that this mechanism outperforms GAp scheme by a small margin.**

- **This mechanism uses less hardware than GAp, since both branch history (first level) and pattern history (second level) are kept globally.**

- **The hardware cost for k history bits is $k + 2 \times 2^k$ bits, neglecting costs for logic.**

gshare is one one the most widely implemented two level dynamic branch prediction schemes
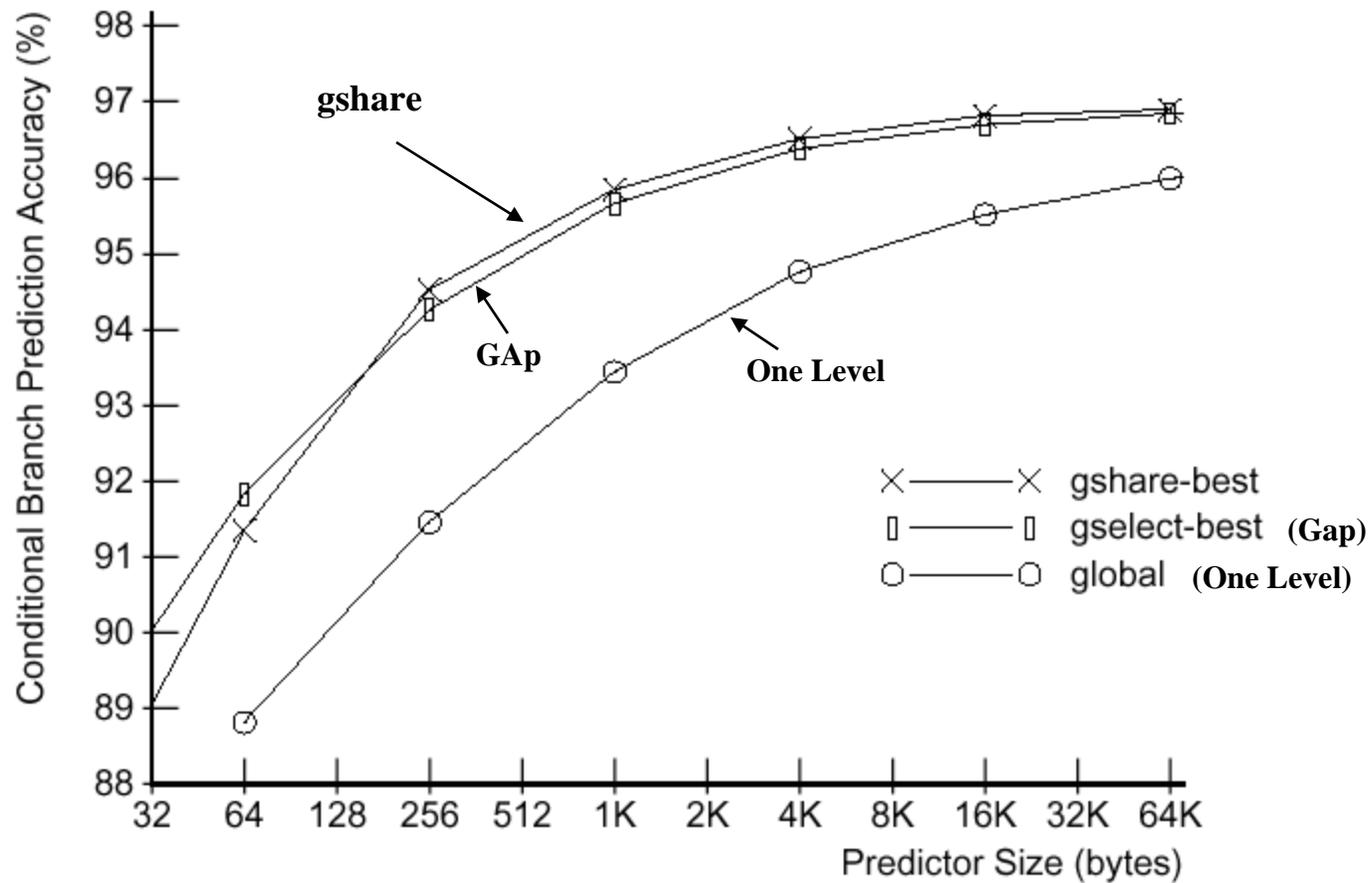
EECC551 - Shaaban

gshare Predictor

Branch and pattern history are kept globally. History and branch address are XORed and the result is used to index the pattern history table.

(BHR)

Branch History Shift Register

Branch Address

Here:
$m = N = k$

/ k bits

First Level:

k bits

XOR **(bitwise XOR)**

Index the second level

2-bit saturating counters (predictors)

Second Level:

(PHT)

Prediction

One Pattern History Table (PHT) with 2^k entries (predictors)

gshare = global history with index sharing

gshare Performance

Hybrid Predictors
(Also known as *tournament* or *combined predictors*)

- Hybrid predictors are simply combinations of two (most common) or more branch prediction mechanisms.

- This approach takes into account that different mechanisms may perform best for different branch scenarios.

- McFarling presented (1993) a number of <u>different combinations of two branch prediction mechanisms.</u>

- He proposed to use an additional 2-bit counter selector array which serves to <u>select the appropriate predictor for each branch</u>.

- One predictor is chosen for the higher two counts, the second one for the lower two counts. The selector array counter used is updated as follows:

  1. If the first predictor is wrong and the second one is right the selector counter used counter is decremented,

  2. If the first one is right and the second one is wrong, the selector counter used is incremented.

  3. No changes are carried out to selector counter used if both predictors are correct or wrong.

Predictor Selector Array Counter Update
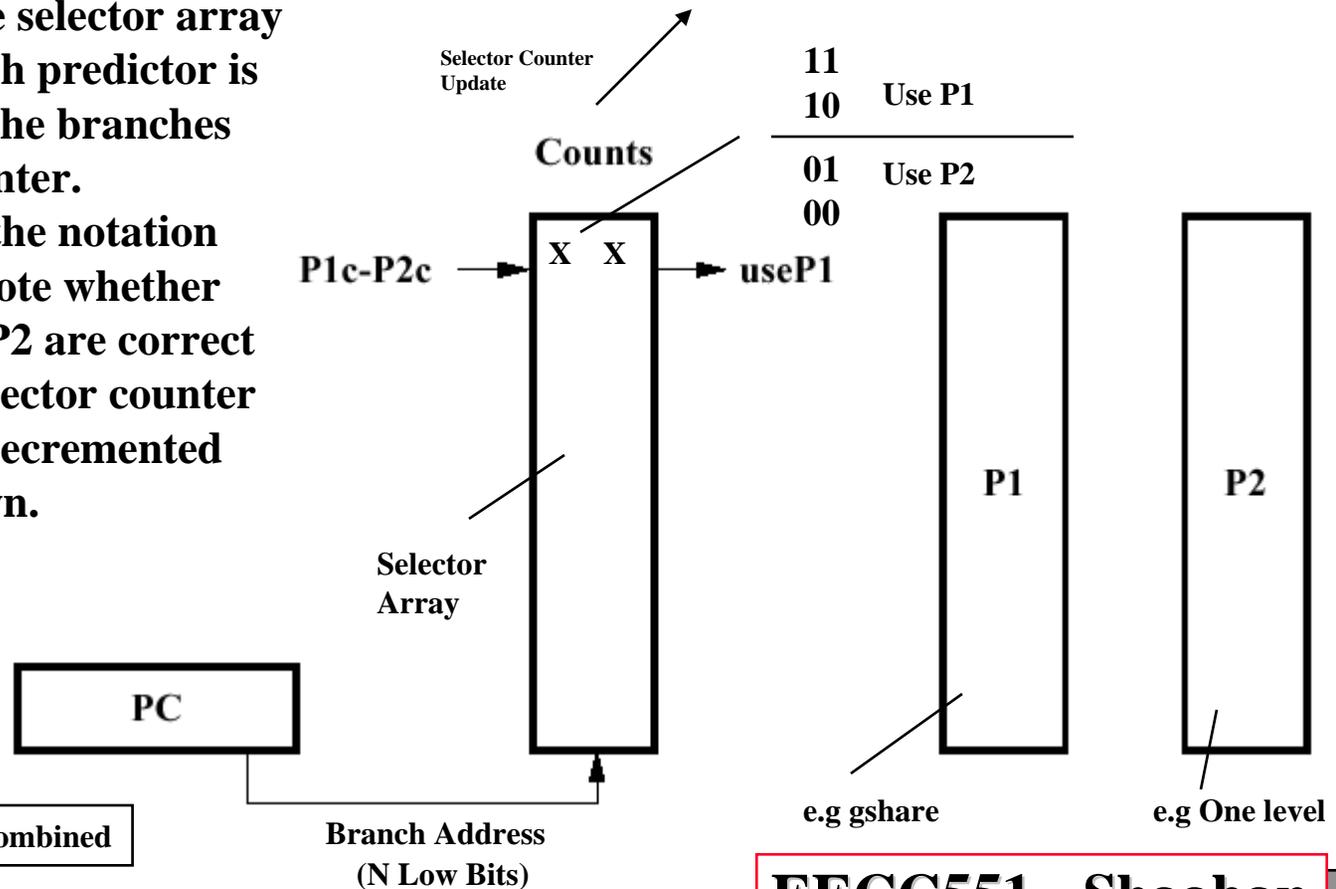
A Generic Hybrid Predictor

BHR
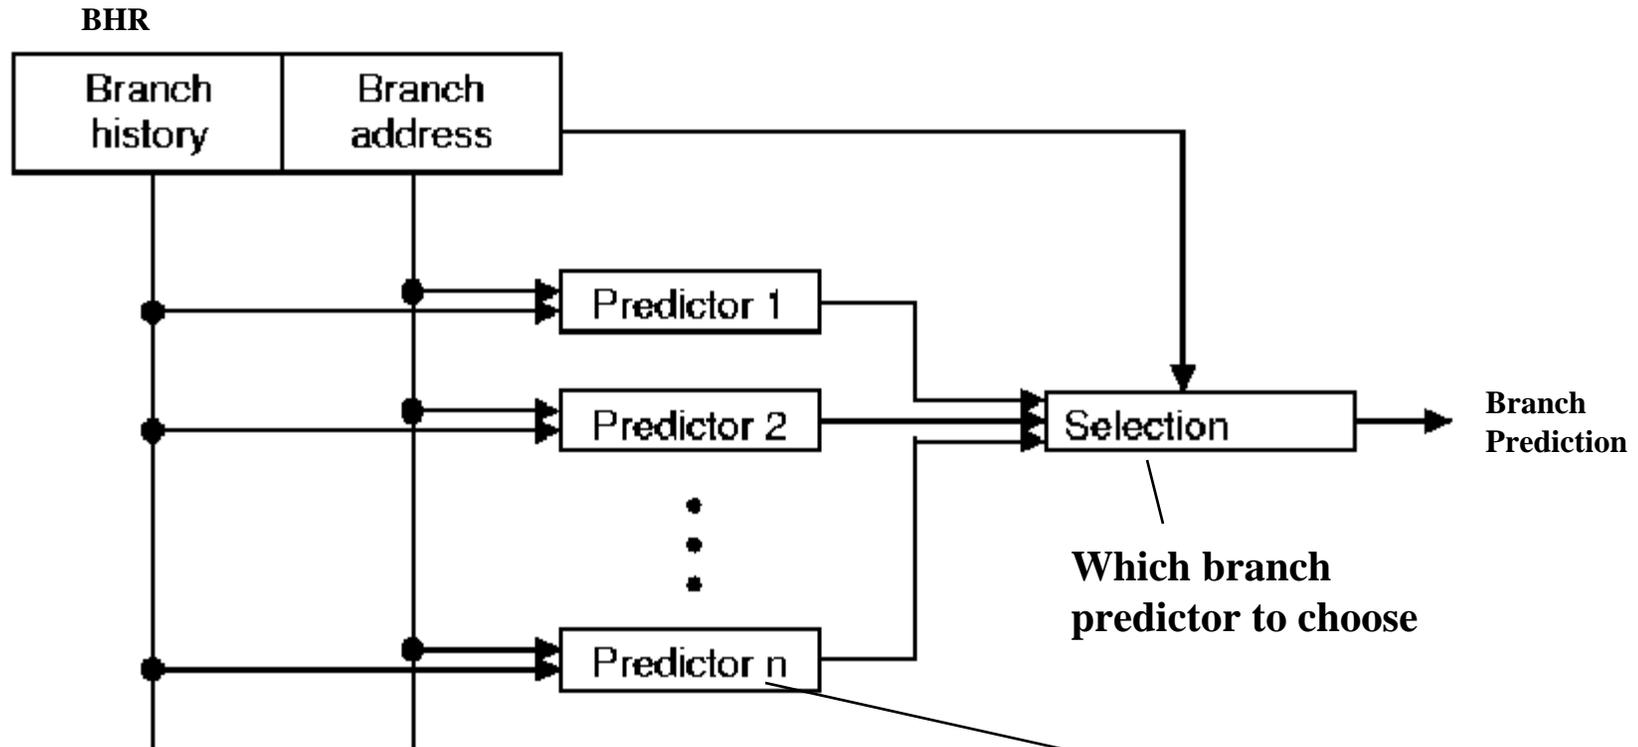
Usually only two predictors are used (i.e. n =2)
e.g. As in Alpha, IBM POWER 4-7
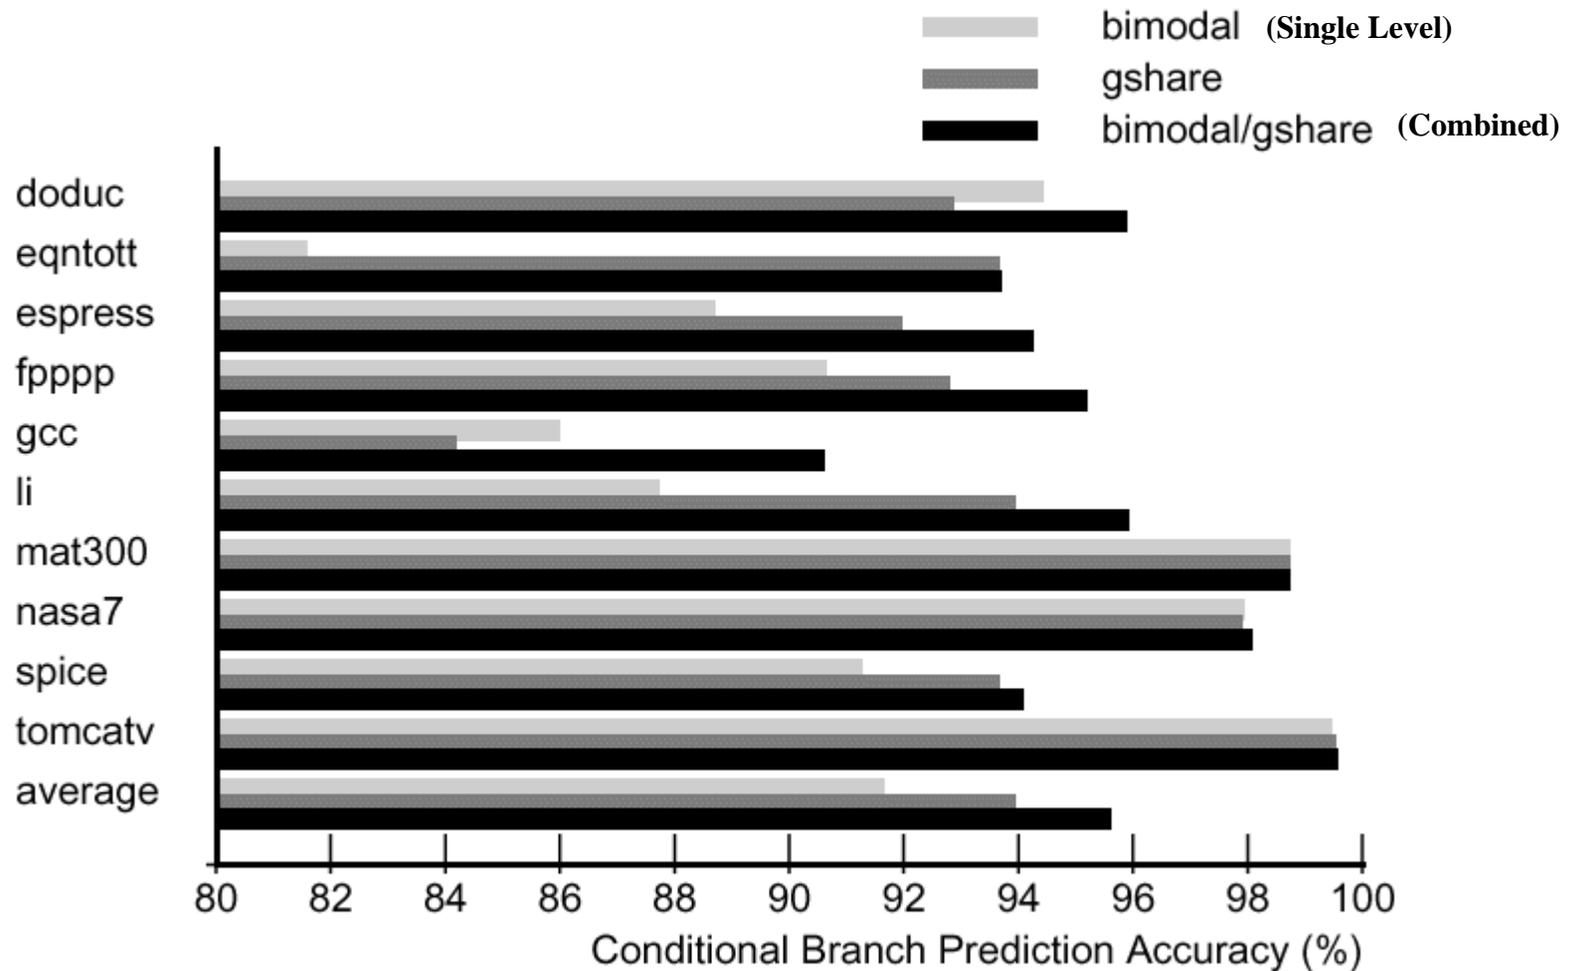
MCFarling's Hybrid Predictor Structure

The hybrid predictor contains an additional counter array (selector array) with 2-bit up/down saturating counters. Which serves to select the best predictor to use. Each counter in the selector array keeps track of which predictor is more accurate for the branches that share that counter. Specifically, using the notation P1c and P2c to denote whether predictors P1 and P2 are correct respectively, the selector counter is incremented or decremented by P1c-P2c as shown.

	P1c	P2c	P1c-P2c	
Both wrong	0	0	0	(no change)
P2 correct	0	1	-1	(decrement counter)
P1 correct	1	0	1	(increment counter)
Both correct	1	1	0	(no change)

Selector Counter Update

Counts

11
10   Use P1
—————
01   Use P2
00

P1c-P2c → X   X → useP1

Selector Array

PC

P1   P2

e.g gshare   e.g One level

Branch Address (N Low Bits)

Here two predictors are combined

(Current example implementations: IBM POWER4, 5, 6)

MCFarling's Hybrid Predictor Performance by Benchmark

Legend:
- bimodal **(Single Level)**
- gshare
- bimodal/gshare **(Combined)**

X-axis: Conditional Branch Prediction Accuracy (%)

Benchmarks: doduc, eqntott, espress, fpppp, gcc, li, mat300, nasa7, spice, tomcatv, average

Processor Branch Prediction Examples

Processor	Released	Accuracy	Prediction Mechanism
Cyrix 6x86	early '96	ca. 85%	PHT associated with BTB
Cyrix 6x86MX	May '97	ca. 90%	PHT associated with BTB
AMD K5	mid '94	80%	PHT associated with I-cache
AMD K6	early '97	95%	2-level adaptive associated with BTIC and ALU
Intel Pentium	late '93	78%	PHT associated with BTB
Intel P6	mid '96	90%	2 level adaptive with BTB
PowerPC750	mid '97	90%	PHT associated with BTIC
MC68060	mid '94	90%	PHT associated with BTIC
DEC Alpha	early '97	95%	Hybrid 2-level adaptive associated with I-cache
HP PA8000	early '96	80%	PHT associated with BTB
SUN UltraSparc	mid '95	88%int 94%FP	PHT associated with I-cache

S+D (PowerPC750)
S+D (HP PA8000)
S+D (SUN UltraSparc)

S+D : **Uses both static (ISA supported) and dynamic branch prediction**

PHT = One Level