

# branch prediction


## last time

what happens with TLB in access patterns

overlapping TLB and cache index lookup

overview of caches and page table lookups and TLB generally

# an OOO pipeline



# reorder buffer: on rename

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07
...	...

free list

%x19
%x23
...
...

# reorder buffer: on rename

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07
...	...

free list

%x19
%x23
...
...

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
14	0x1233	%rbx / %x23		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31		
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34		
19	0x1249	%rax / %x38		
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		

reorder buffer contains instructions started,  
but not fully finished new entries created on rename  
(not enough space? stall rename stage)

# reorder buffer: on rename

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07
...	...

free list

%x19
%x23
...
...

remove here  
when committed



reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
14	0x1233	%rbx / %x23		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31		
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34		
19	0x1249	%rax / %x38		
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		

add here  
on rename



place newly started instruction at end of buffer  
remember at least its destination register  
(both architectural and physical versions)

# reorder buffer: on rename

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	<del>%x07</del> %x19
...	...

free list

%x19
%x23
...
...

remove here  
when committed



reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
14	0x1233	%rbx / %x23		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31		
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34		
19	0x1249	%rax / %x38		
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		
32	0x1230	%rdx / %x19		

add here  
on rename



next renamed instruction goes in next slot, etc.

# reorder buffer: on rename

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

free list

%x19
%x23
...
...

reorder buffer (ROB)

remove here  
when committed



instr num.	PC	dest. reg	done?	mispred?
14	0x1233	%rbx / %x23		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31		
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34		
19	0x1249	%rax / %x38		
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		
32	0x1230	%rdx / %x19		

add here  
on rename





# reorder buffer: on commit

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

free list

%x19
%x13
...
...

remove here  
when committed



reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
14	0x1233	%rbx / %x24		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31		
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34		
19	0x1249	%rax / %x38		
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		

# reorder buffer: on commit

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

free list

%x19
%x13
...
...

remove here  
when committed



reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
14	0x1233	%rbx / %x24		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31	✓	
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34	✓	
19	0x1249	%rax / %x38	✓	
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		✓

instructions marked done in reorder buffer  
when result is computed  
but not removed from reorder buffer ('committed') yet

# reorder buffer: on commit

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

phys → arch. reg  
for committed

remove here  
when committed →

arch. reg	phys. reg
%rax	%x30
%rcx	%x28
%rbx	%x23
%rdx	%x21
...	...

free list

%x19
%x13
...
...

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
14	0x1233	%rbx / %x24		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31	✓	
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34	✓	
19	0x1249	%rax / %x38	✓	
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		✓

commit stage tracks architectural to physical register map  
for committed instructions

# reorder buffer: on commit

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

phys → arch. reg  
for committed

remove here  
when committed

arch. reg	phys. reg
%rax	%x30
%rcx	%x28
%rbx	%x23 %x24
%rdx	%x21
...	...

free list

%x19
%x13
...
%x23

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
14	0x1233	%rbx / %x24	✓	
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31	✓	
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34	✓	
19	0x1249	%rax / %x38	✓	
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		✓
32	0x1230	%rdx / %x19		

when next-to-commit instruction is done  
update this register map and free register list  
and remove instr. from reorder buffer

# reorder buffer: on commit

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

phys → arch. reg remove here  
for committed when committed →

arch. reg	phys. reg
%rax	%x30
%rcx	%x28
%rbx	%x23 %x24
%rdx	%x21
...	...

free list

%x19
%x13
...
%x23

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
<del>14</del>	<del>0x1233</del>	<del>%rbx / %x24</del>	<del>✓</del>	
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31	✓	
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34	✓	
19	0x1249	%rax / %x38	✓	
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		✓
32	0x1230	%rdx / %x19		

when next-to-commit instruction is done  
update this register map and free register list  
and remove instr. from reorder buffer

# reorder buffer: commit mispredict (one way)

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x19
...	...

phys → arch. reg  
for committed

arch. reg	phys. reg
%rax	%x30 %x38
%rcx	%x31 %x32
%rbx	%x23 %x24
%rdx	%x21 %x34
...	...

free list

%x19
%x13
...
...

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
<del>14</del>	<del>0x1233</del>	<del>%rbx / %x24</del>	<del>✓</del>	
<del>15</del>	<del>0x1239</del>	<del>%rax / %x30</del>	<del>✓</del>	
<del>16</del>	<del>0x1242</del>	<del>%rcx / %x31</del>	<del>✓</del>	
<del>17</del>	<del>0x1244</del>	<del>%rcx / %x32</del>	<del>✓</del>	
<del>18</del>	<del>0x1248</del>	<del>%rdx / %x34</del>	<del>✓</del>	
<del>19</del>	<del>0x1249</del>	<del>%rax / %x38</del>	<del>✓</del>	
→ 20	0x1254	PC	✓	✓
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12	✓	
32	0x1230	%rdx / %x19		

# reorder buffer: commit mispredict (one way)

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x19
...	...

phys → arch. reg  
for committed

arch. reg	phys. reg
%rax	%x30 %x38
%rcx	%x31 %x32
%rbx	%x23 %x24
%rdx	%x21 %x34
...	...

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
<del>14</del>	<del>0x1233</del>	<del>%rbx / %x24</del>	<del>✓</del>	
<del>15</del>	<del>0x1239</del>	<del>%rax / %x30</del>	<del>✓</del>	
<del>16</del>	<del>0x1242</del>	<del>%rcx / %x31</del>	<del>✓</del>	
<del>17</del>	<del>0x1244</del>	<del>%rcx / %x32</del>	<del>✓</del>	
<del>18</del>	<del>0x1248</del>	<del>%rdx / %x34</del>	<del>✓</del>	
<del>19</del>	<del>0x1249</del>	<del>%rax / %x38</del>	<del>✓</del>	
→ 20	0x1254	PC	✓	✓
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12	✓	
32	0x1230	%rdx / %x19		

free list

%x19
%x13
...
...

when committing a mispredicted instruction...  
this is where we undo mispredicted instructions

# reorder buffer: commit mispredict (one way)

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x38
%rcx	%x32
%rbx	%x24
%rdx	%x34
...	...



phys → arch. reg  
for committed

arch. reg	phys. reg
%rax	%x30 %x38
%rcx	%x31 %x32
%rbx	%x23 %x24
%rdx	%x21 %x34
...	...

free list

%x19
%x13
...
...

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
14	0x1233	%rbx / %x24	✓	
15	0x1239	%rax / %x30	✓	
16	0x1242	%rcx / %x31	✓	
17	0x1244	%rcx / %x32	✓	
18	0x1248	%rdx / %x34	✓	
19	0x1249	%rax / %x38	✓	
20	0x1254	PC	✓	✓
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12	✓	
32	0x1230	%rdx / %x19		



copy commit register map into rename register map  
so we can start fetching from the correct PC



# reorder buffer: commit mispredict (one way)

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x38
%rcx	%x32
%rbx	%x24
%rdx	%x34
...	...



phys → arch. reg  
for committed

arch. reg	phys. reg
%rax	%x30 %x38
%rcx	%x31 %x32
%rbx	%x23 %x24
%rdx	%x21 %x34
...	...

free list

%x19
%x13
...
...

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred?
14	0x1233	%rbx / %x24	✓	
15	0x1239	%rax / %x30	✓	
16	0x1242	%rcx / %x31	✓	
17	0x1244	%rcx / %x32	✓	
18	0x1248	%rdx / %x34	✓	
19	0x1249	%rax / %x38	✓	
20	0x1254	PC	✓	✓
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12	✓	
32	0x1230	%rdx / %x19		



...and discard all the mispredicted instructions  
(without committing them)

## better? alternatives

- can take snapshots of register map on each branch

  - don't need to reconstruct the table  
(but how to efficiently store them)

- can reconstruct register map before we commit the branch instruction


  - need to let reorder buffer be accessed even more?

- can track more/different information in reorder buffer

## question

how much does a branch misprediction cost?

# an OOO pipeline



# direction versus target

two parts to branch prediction:

predict branch target

pipehw2: not done — instead compute fast enough

longer pipelines: can't compute it fast enough

predict branch direction

conditional branches: taken or not taken

not relevant for unconditional branches

# direction versus target

two parts to branch prediction:

predict branch target

pipehw2: not done — instead compute fast enough

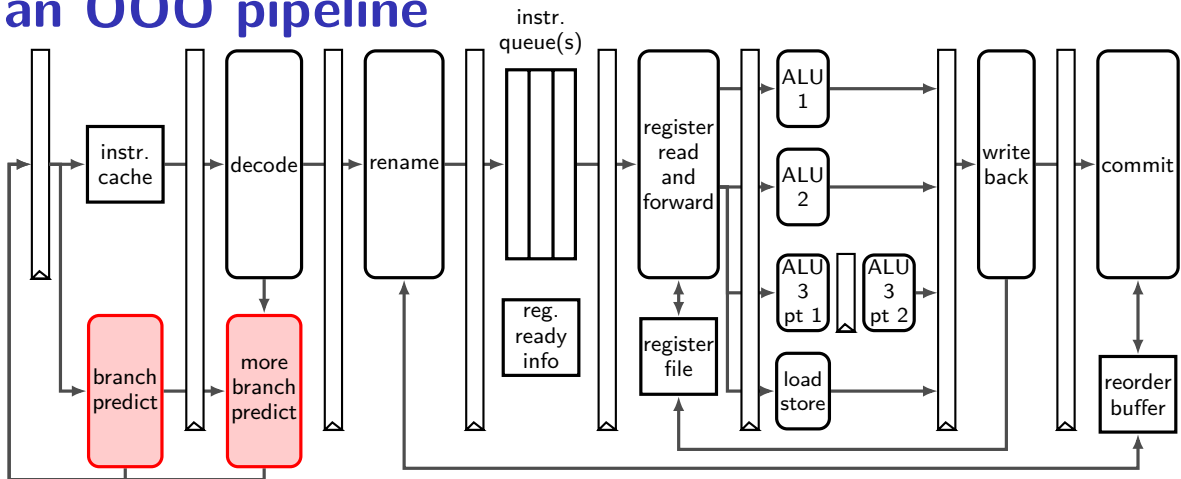
longer pipelines: can't compute it fast enough

predict branch direction

conditional branches: taken or not taken

not relevant for unconditional branches

# an OOO pipeline



# branch target buffer

can take several cycles to fetch+decode jumps, calls, returns

still want 1-cycle prediction of next thing to fetch



# BTB: cache for branches

idx	valid	tag	ofst	type	target	(more info?)	valid	...
0x00	1	0x400	5	Jxx	0x3FFFF3	...	1	...
0x01	1	0x401	C	JMP	0x401035	----	0	...
0x02	0	----	----	----	----	----	0	...
0x03	1	0x400	9	RET	----	...	0	...
...	...	...	...	...	...	...	...	...
0xFF	1	0x3FF	8	CALL	0x404033	...	0	...

```
0x3FFFF3:  movq %rax, %rsi
0x3FFFF7:  pushq %rbx
0x3FFFF8:  call 0x404033
0x400001:  popq %rbx
0x400003:  cmpq %rbx, %rax
0x400005:  jle 0x3FFFF3
...
0x400031:  ret
...
```

# BTB: cache for branches

idx	valid	tag	ofst	type	target	(more info?)	valid	...
0x00	1	0x400	5	Jxx	0x3FFFF3	...	1	...
0x01	1	0x401	C	JMP	0x401035	----	0	...
0x02	0	----	----	----	----	----	0	...
0x03	1	0x400	9	RET	----	...	0	...
...	...	...	...	...	...	...	...	...
0xFF	1	0x3FF	8	CALL	0x404033	...	0	...

```
0x3FFFF3:  movq %rax, %rsi
0x3FFFF7:  pushq %rbx
0x3FFFF8:  call 0x404033
0x400001:  popq %rbx
0x400003:  cmpq %rbx, %rax
0x400005:  jle 0x3FFFF3
...
0x400031:  ret
...
```

# BTB: cache for branches

idx	valid	tag	ofst	type	target	(more info?)	valid	...
0x00	1	0x400	5	Jxx	0x3FFFF3	...	1	...
0x01	1	0x401	C	JMP	0x401035	----	0	...
0x02	0	----	----	----	----	----	0	...
0x03	1	0x400	9	RET	----	...	0	...
...	...	...	...	...	...	...	...	...
0xFF	1	0x3FF	8	CALL	0x404033	...	0	...

```
0x3FFFF3:  movq %rax, %rsi
0x3FFFF7:  pushq %rbx
0x3FFFF8:  call 0x404033
0x400001:  popq %rbx
0x400003:  cmpq %rbx, %rax
0x400005:  jle 0x3FFFF3
...
0x400031:  ret
...
```

# predicting ret: extra copy of stack

predicting ret — ministack in processor registers

push on ministack on call; pop on ret

ministack overflows? discard oldest, mispredict it later

baz saved registers
baz return address
bar saved registers
bar return address
foo local variables
foo saved registers
foo return address
foo saved registers

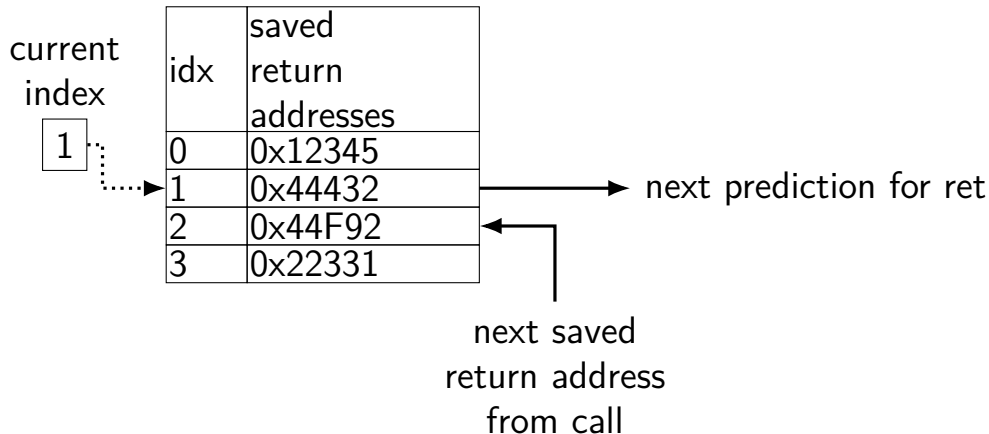
baz return address
bar return address
foo return address

(partial?) stack  
in CPU registers

stack in memory

# 4-entry return address stack

4-entry return address stack in CPU



on `call`: increment index, save return address in that slot

on `ret`: read prediction from index, decrement index

## indirect branches?

```
jmp *0x1234(%rax,8), call *(%rax)
```

e.g. switch statement (jump to case with table)

e.g. virtual method call (choose function based on which subclass)

simple solution: branch target buffer remembers last target

there are better solutions (but we probably won't have time)

# branch direction prediction

now: can predict target of most branches

for conditional branches: need to predict direction

recall: pipehw2: predict as always taken

we can do a lot better

# static branch prediction

forward (target > PC) not taken; backward taken

intuition: loops:

```
LOOP: ...  
      ...  
      je LOOP
```

```
LOOP: ...  
      jne SKIP_LOOP  
      ...  
      jmp LOOP  
SKIP_LOOP:
```



## using local history

for each branch: remember recent outcomes

prediction: based on recent pattern

# branches are consistent (1)

```
array = malloc(1024);  
if (array == NULL) ...
```

---

```
call malloc  
cmpq $0, %rax  
je handle_error
```

---

almost never taken

# branches are consistent (1)

```
array = malloc(1024);  
if (array == NULL) ...
```

---

```
call malloc  
cmpq $0, %rax  
je handle_error
```

---

almost never taken

## branches are consistent (2)

```
for (int i = 0; i != 1024; ++i)
```

---

```
    xorq %rax, %rax  
top_of_loop:  
    ...  
    incq %rax  
    cmpq $1024, %rax  
    jne top_of_loop
```

---

almost always taken

## branches are consistent (2)

```
for (int i = 0; i != 1024; ++i)
```

---

```
    xorq %rax, %rax  
top_of_loop:  
    ...  
    incq %rax  
    cmpq $1024, %rax  
    jne top_of_loop
```

---

almost always taken

# predict: repeat last

PC of branch

0x40042A

hash function

<i>index</i>	<i>prediction/ last result?</i>
0	taken (1)
1	not taken (0)
2	taken (1)
3	taken (1)
...	...
14	not taken (0)
15	taken (1)

# predict: repeat last

PC of branch

0x40042A

hash function

*index*    *prediction/  
last result?*

0

taken (1)

1

not taken (0)

2

taken (1)

3

taken (1)

...

...

14

not taken (0)

typical choice: some bits of branch address  
for our example: will use bits 4-7

# predict: repeat last

PC of branch

0x40042A

hash function

<i>index</i>	<i>prediction/ last result?</i>
0	taken (1)
1	not taken (0)
2	taken (1)
3	taken (1)
...	...
14	not taken (0)
15	taken (1)



# predict: repeat last

PC of branch

0x40042A

hash function

<i>index</i>	<i>prediction/ last result?</i>
0	taken (1)
1	not taken (0)
2	taken (1)
3	taken (1)
...	...
14	not taken (0)
15	taken (1)

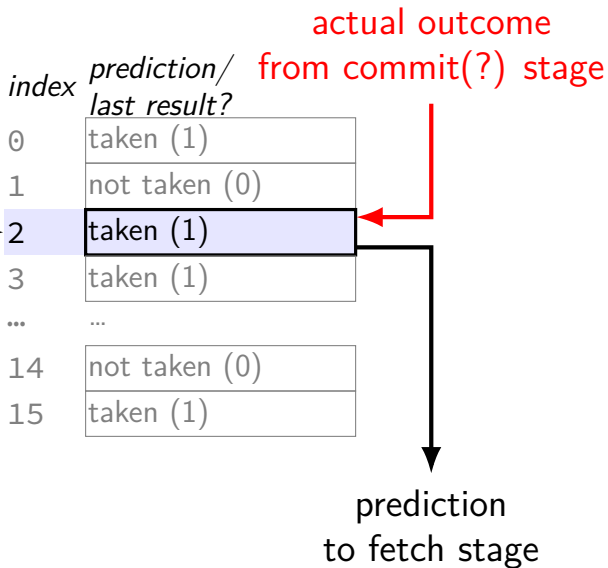
prediction  
to fetch stage

# predict: repeat last

PC of branch

0x40042A

hash function



# example

PC of branch

0x40042A

hash function

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

actual outcome  
from commit(?) stage

idx	prediction/ last result?
0	taken (1)
1	not taken (0)
2	not taken (0)
3	taken (1)
...	...
14	not taken (0)
15	taken (1)

prediction  
to fetch stage

# example

PC of branch

0x40042A

hash function

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

assembly version of:

```
i = 4; do { ...; i -= 1; } while (i)
```

idx prediction/  
last result?

0 taken (1)

1 not taken (0)

2 not taken (0)

3 taken (1)

... ..

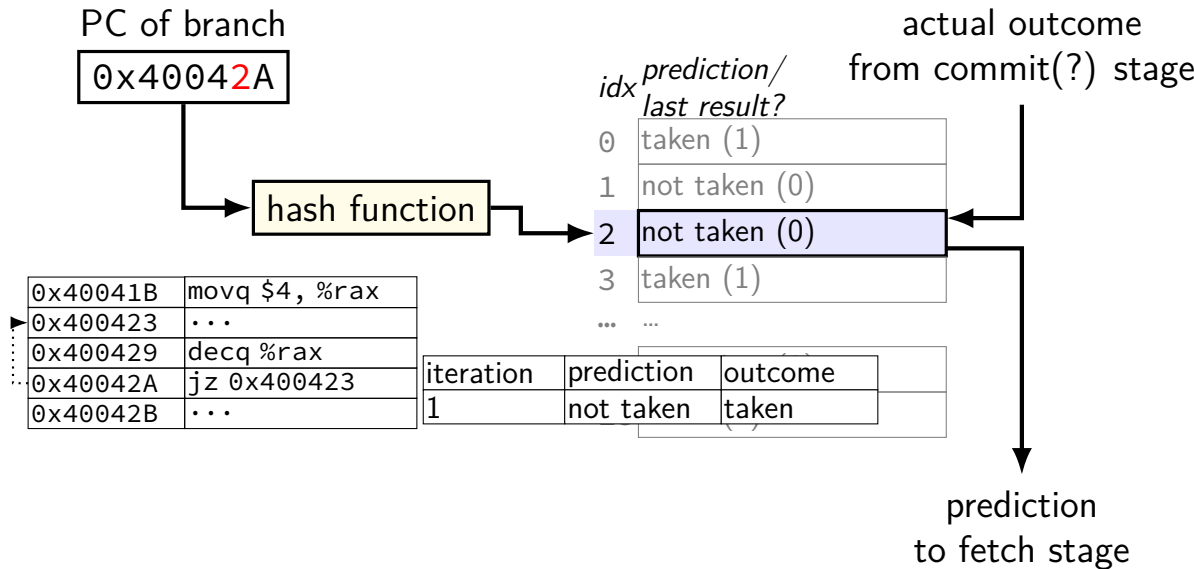
14 not taken (0)

15 taken (1)

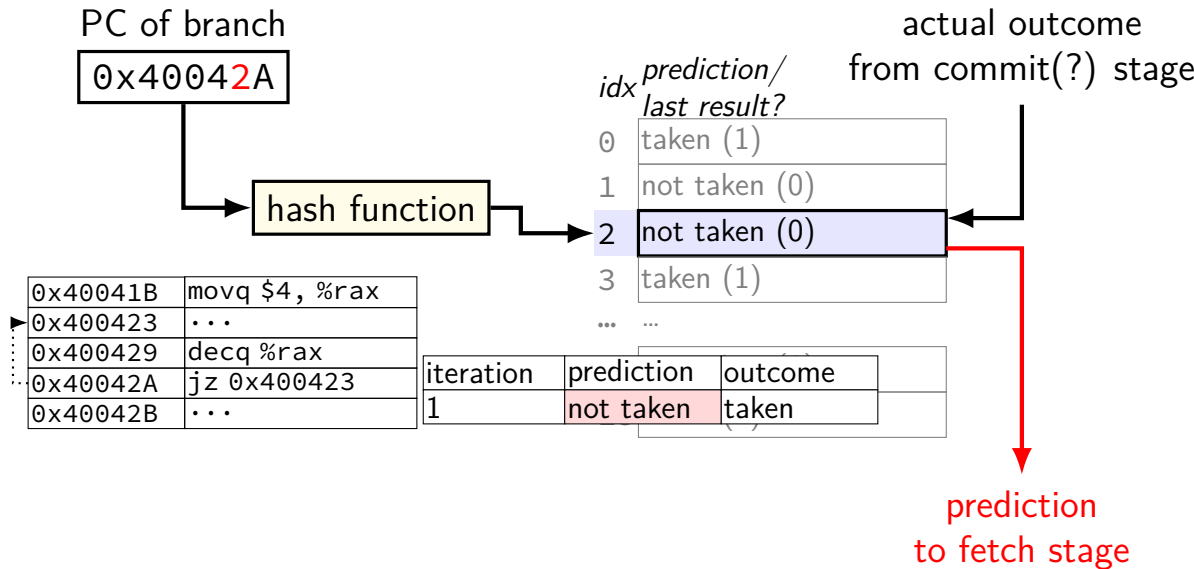
actual outcome  
from commit(?) stage

prediction  
to fetch stage

# example



# example



# example

PC of branch

0x40042A

hash function

idx prediction/  
last result?

0	taken (1)
1	not taken (0)
2	<del>not taken</del> taken (1)
3	taken (1)

actual outcome  
from commit(?) stage

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

iteration	prediction	outcome
1	not taken	taken

prediction  
to fetch stage

# example

PC of branch

0x40042A

hash function

actual outcome  
from commit(?) stage

idx	prediction/ last result?
0	taken (1)
1	not taken (0)
2	not taken taken (1)
3	taken (1)
...	...

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

iteration	prediction	outcome
1	not taken	taken
2	taken	taken

prediction  
to fetch stage



# example

PC of branch

0x40042A

hash function

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

idx prediction/  
last result?

0	taken (1)
1	not taken (0)
2	not taken taken (1)
3	taken (1)
...	...

actual outcome  
from commit(?) stage

iteration	prediction	outcome
1	not taken	taken
2	taken	taken

prediction  
to fetch stage

# example

PC of branch

0x40042A

hash function

actual outcome  
from commit(?) stage

idx	prediction/ last result?
0	taken (1)
1	not taken (0)
2	not taken taken (1)
3	taken (1)
...	...

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

iteration	prediction	outcome
1	not taken	taken
2	taken	taken
3	taken	taken
4	taken	not taken
1	not taken	taken
2	taken	taken
...	...	...

prediction  
to fetch stage

# example

PC of branch

0x40042A

hash function

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

idx prediction/  
last result?

0	taken (1)
1	not taken (0)
2	not taken taken (1)
3	taken (1)
...	...

actual outcome  
from commit(?) stage

iteration	prediction	outcome
1	not taken	taken
2	taken	taken
3	taken	taken
4	taken	not taken
1	not taken	taken
2	taken	taken
...	...	...

prediction  
to fetch stage

# example

PC of branch

0x40042A

hash function

actual outcome  
from commit(?) stage

idx	prediction/ last result?
0	taken (1)
1	not taken (0)
2	not taken taken (1)
3	taken (1)
...	...

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

iteration	prediction	outcome
1	not taken	taken
2	taken	taken
3	taken	taken
4	taken	not taken
1	not taken	taken
2	taken	taken
...	...	...

prediction  
to fetch stage

# collisions?

two branches could have same hashed PC

nothing in table tells us about this

versus direct-mapped cache: had *tag bits* to tell

is it worth it?

adding tag bits makes table *much* smaller and/or slower

but does anything go wrong when there's a collision?

## collision results

possibility 1: both branches usually taken  
no actual conflict — prediction is better(!)

possibility 2: both branches usually not taken  
no actual conflict — prediction is better(!)

possibility 3: one branch taken, one not taken  
performance probably worse

# 1-bit predictor for loops

predicts first and last iteration wrong

example: branch to beginning — but same for branch from beginning to end

everything else correct

later: we'll find a way to do better

## exercise

use 1-bit predictor on this loop

executed in outer loop (not shown) many, many times

what is the conditional branch misprediction rate?

```
int i = 0;
while (true) {
    if (i % 3 == 0) goto next;
    ...
next:
    i += 1;
    if (i == 50) break;
}
```



# beyond 1-bit predictor

devote *more space* to storing history

main goal: **rare exceptions don't immediately change prediction**

example: branch taken 99% of the time

1-bit predictor: wrong about 2% of the time


- 1% when branch not taken

- 1% of taken branches right after branch not taken

new predictor: wrong about 1% of the time

- 1% when branch not taken

## 2-bit saturating counter




PC of branch

0x40042A

hash function


<i>index</i>	<i>counter</i>
0	11
1	01
2	11
...	...
14	10
15	00

## 2-bit saturating counter



branch always taken:  
value increases to 'strongest' taken value

## 2-bit saturating counter



branch almost always taken, then not taken once:  
still predicted as taken

# example

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

iter.	table before	prediction	outcome	table after
1	01	not taken	taken	10
2	10	taken	taken	11
3	11	taken	taken	11
4	11	taken	not taken	10
1	10	taken	taken	11
2	11	taken	taken	11
3	11	taken	taken	11
4	11	taken	not taken	10
1	10	taken	taken	11
...	...	...	...	...

# generalizing saturating counters

2-bit counter: ignore one exception to taken/not taken

3-bit counter: ignore more exceptions

000 ↔ 001 ↔ 010 ↔ 011 ↔ 100 ↔ 101 ↔ 110 ↔ 111

000-011: not taken

100-111: taken

## exercise

use 2-bit predictor on this loop

executed in outer loop (not shown) many, many times

what is the conditional branch misprediction rate?

```
int i = 0;
while (true) {
    if (i % 3 == 0) goto next;
    ...
next:
    i += 1;
    if (i == 50) break;
}
```

# branch patterns

```
i = 4;  
do {  
    ...  
    i -= 1;  
} while (i != 0);
```

---

typical pattern for jump to top of do-while above:

TTTN TTTN TTTN TTTN TTTN...(T = taken, N = not taken)

goal: take advantage of recent pattern to make predictions

just saw 'NTTTNT'? predict T next

'TNTTTN'? predict T; 'TTNTTT'? predict N next

...



# local pattern predictor (incomplete)

PC of branch

0x40042A

hash function

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

4-iter loop: TTTN TTTN TTTN ...

idx recent  
pattern

0 NNNNNN

1 NNTNTT

2 TTTTNT

3 TTTTTT

... ..

14 NTNNTN

15 NNTTTT

actual outcome

from commit(?) stage

???

convert to  
prediction

???

prediction  
to fetch stage

# local pattern predictor (incomplete)

PC of branch

0x40042A

hash function

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

4-iter loop: TTTN TTTN TTTN ...

iter.	pattern tbl before	predicted	outcome	pattern tbl after
1	TTTTNT	???	taken	TTTNTT

idx recent pattern

0	NNNNNN
1	NNTNTT
2	TTTTNT
3	TTTTTT
...	...
14	NTNNTN
15	NNTTTT

actual outcome

from commit(?) stage

???

convert to prediction

???

prediction to fetch stage

# local pattern predictor (incomplete)

PC of branch

0x40042A

hash function

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

4-iter loop: TTTN TTTN TTTN ...

iter.	pattern tbl before	predicted	outcome	pattern tbl after
1	TTTTNT	???	taken	TTTNTT

idx	recent pattern
0	NNNNNN
1	NNTNTT
2	TTTNTT
3	TTTTTT
...	...
14	NTNNTN
15	NNTTTT

actual outcome from commit(?) stage

???

convert to prediction

???

prediction to fetch stage

# local pattern predictor (incomplete)

PC of branch

0x40042A

hash function

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

4-iter loop: TTTN TTTN TTTN ...

iter.	pattern tbl before	predicted	outcome	pattern tbl after
1	TTTTNT	???	taken	TTTNTT
2	TTTNTT	???	taken	TTNTTT

idx	recent pattern
0	NNNNNN
1	NNTNTT
2	TTTNTT
3	TTTTTT
...	...
14	NTNTTN
15	NNTTTT

actual outcome

from commit(?) stage

???

convert to prediction

???

prediction to fetch stage

# local pattern predictor (incomplete)

PC of branch

0x40042A

hash function

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

4-iter loop: TTTN TTTN TTTN ...

iter.	pattern tbl before	predicted	outcome	pattern tbl after
1	TTTTNT	???	taken	TTTNTT
2	TTTNTT	???	taken	TTNTTT
3	TTNTTT	???	taken	TNTTTT

idx	recent pattern
0	NNNNNN
1	NNTNTT
2	TTTNTTT
3	TTTTTT
...	...
14	NTNNTN
15	NNTTTT

actual outcome

from commit(?) stage

???

convert to prediction

???

prediction to fetch stage

# local pattern predictor (incomplete)

PC of branch

0x40042A

hash function

0x40041B	movq \$4, %rax
0x400423	...
0x400429	decq %rax
0x40042A	jz 0x400423
0x40042B	...

idx	recent pattern
0	NNNNNN
1	NNTNTT
2	TTTTNTTT
3	TTTTTT
...	...
14	NTNTTN
15	NNTTTT

actual outcome

from commit(?) stage

???

convert to prediction

???

prediction to fetch stage

4-iter loop: TTTN TTTN TTTN ...

iter.	pattern tbl before	predicted	outcome	pattern tbl after
1	TTTTNT	???	taken	TTTNTT
2	TTTNTT	???	taken	TTNTTT
3	TTNTTT	???	taken	TNTTTT
4	TNTTTN	???	not taken	NTTTTN
1	NTTTTN	???	taken	TTTTNT
2	TTTTNT	???	taken	TTTNTT

# recent pattern to prediction?

easy cases:

just saw TTTTTT: predict T

just saw NNNNNN: predict N

just saw TNTNTN: predict T

hard cases:

TTNTTTTT

predict T? loop with many iterations  
(NTTTTTTTTTTTTTTTTTTTTTTTTTTT...)

predict T? if statement mostly taken  
(TTTTNTTNTTTTTTTTTTTTTTTTTTT...)

predict N? loop with 5 iterations  
(NTTTTNTTTTTTTTTTTTTTTTTTT...)

# history of history

PC of branch

0x40042A

hash

idx	recent pattern
0	NNNN
1	TNTT
2	TTTN
3	TTTT
...	...
14	NTTN
15	TTTT

actual outcome  
from commit(?) stage

pattern  
NNNN  
NNNT  
...  
NTTT  
...  
TNTT  
...  
TTNT  
TTTN  
TTTT

counter
00
00
...
10
...
11
...
01
01
11

prediction  
to fetch sta

iter.	branch to pat. tbl before	pat. to counter before	predict	actual	pat. to counter after	branch to pat. tbl after
1	TTTN	01	not taken	taken	10	TTNT



# history of history

PC of branch

0x40042A

hash

idx	recent pattern
0	NNNN
1	TNTT
2	FTTN
3	TTTT
...	...
14	NTTN
15	TTTT

actual outcome  
from commit(?) stage

pattern  
NNNN  
NNNT  
...  
NTTT  
...  
TNTT  
...  
TTNT  
TTTT

counter
00
00
...
10
...
11
...
01
01 10
11

prediction  
to fetch sta

iter.	branch to pat. tbl before	pat. to counter before	predict	actual	pat. to counter after	branch to pat. tbl after
1	TTTN	01	not taken	taken	10	TTNT

# history of history

PC of branch

0x40042A

hash

idx	recent pattern
0	NNNN
1	TNTT
2	TTNT
3	TTTT
...	...
14	NTTN
15	TTTT

actual outcome  
from commit(?) stage

pattern
NNNN
NNNT
...
NTTT
...
TNTT
...
TTNT
TTTN
TTTT

counter
00
00
...
10
...
11
...
01
01 10
11

iter.	branch to pat. tbl before	pat. to counter before	predict	actual	pat. to counter after	branch to pat. tbl after
1	TTTN	01	not taken	taken	10	TTNT
2	TTNT	01	not taken	taken	10	TNTT

prediction  
to fetch stage

# history of history

PC of branch

0x40042A

hash

idx	recent pattern
0	NNNN
1	TNTT
2	TTNTT
3	TTTT
...	...
14	NTTN
15	TTTT

actual outcome  
from commit(?) stage

pattern
NNNN
NNNT
...
NTTT
...
TNTT
...
TTNT
TTTN
TTTT

counter
00
00
...
10
...
11
...
01 10
01 10
11

iter.	branch to pat. tbl before	pat. to counter before	predict	actual	pat. to counter after	branch to pat. tbl after
1	TTTN	01	not taken	taken	10	TTNT
2	TTNT	01	not taken	taken	10	TNTT

prediction  
to fetch stage

# history of history

PC of branch

0x40042A

hash

idx	recent pattern
0	NNNN
1	TNTT
2	TTTNTT
3	TTTT
...	...
14	NTTN
15	TTTT

actual outcome  
from commit(?) stage

pattern
NNNN
NNNT
...
NNTT
...
TNTT
...
TTNT
TTTN
TTTT

counter
00
00
...
10
...
11
...
01 10
01 10
11

iter.	branch to pat. tbl before	pat. to counter before	predict	actual	pat. to counter after	branch to pat. tbl after
1	TTTN	01	not taken	taken	10	TTNT
2	TTNT	01	not taken	taken	10	TNTT
3	TNTT	11	taken	taken	11	NNTT

prediction  
to fetch sta

# history of history

PC of branch

0x40042A

hash

idx	recent pattern
0	NNNN
1	TNTT
2	TTTNTTT
3	TTTT
...	...
14	NTTN
15	TTTT

actual outcome  
from commit(?) stage

pattern
NNNN
NNNT
...
NNTT
...
TNTT
...
TTNT
TTTN
TTTT

counter
00
00
...
10
...
11
...
01 10
01 10
11

iter.	branch to pat. tbl before	pat. to counter before	predict	actual	pat. to counter after	branch to pat. tbl after
1	TTTN	01	not taken	taken	10	TTNT
2	TTNT	01	not taken	taken	10	TNTT
3	TNTT	11	taken	taken	11	NNTT

prediction  
to fetch stage

# history of history

PC of branch

0x40042A

hash

idx	recent pattern
0	NNNN
1	TNTT
2	TTTNTTT
3	TTTT
...	...
14	NTTN
15	TTTT

actual outcome  
from commit(?) stage

pattern
NNNN
NNNT
...
NNTT
...
TTNT
...
TTTN
TTTT

counter
00
00
...
10
...
11
...
01 10
01 10
11

iter.	branch to pat. tbl before	pat. to counter before	predict	actual	pat. to counter after	branch to pat. tbl after
1	TTTN	01	not taken	taken	10	TTNT
2	TTNT	01	not taken	taken	10	TNTT
3	TNTT	11	taken	taken	11	NNTT
4	NNTT	01	not taken	taken	10	TTTT

prediction  
to fetch sta

# history of history

PC of branch

0x40042A

hash

idx	recent pattern
0	NNNN
1	TNTT
2	TTTNTTT
3	TTTT
...	...
14	NTTN
15	TTTT

actual outcome  
from commit(?) stage

pattern
NNNN
NNNT
...
NNTT
...
TNTT
...
TTNT
TTTN
TTTT

counter
00
00
...
10 11
...
11
...
01 10
01 10
11

iter.	branch to pat. tbl before	pat. to counter before	predict	actual	pat. to counter after	branch to pat. tbl after
1	TTTN	01	not taken	taken	10	TTNT
2	TTNT	01	not taken	taken	10	TNTT
3	TNTT	11	taken	taken	11	NNTT
4	NNTT	01	not taken	taken	10	TTTT

prediction  
to fetch stage

# history of history

PC of branch

0x40042A

hash

idx	recent pattern
0	NNNN
1	TNTT
2	TTTNTTT
3	TTTT
...	...
14	NTTN
15	TTTT

actual outcome  
from commit(?) stage

pattern  
NNNN  
NNNT  
...  
NTTT  
...  
TNTT  
...  
TTNT  
TTTT

counter
00
00
...
<del>10</del> 11
...
11
...
<del>01</del> 10
<b>01 10</b>
11

prediction  
to fetch sta

iter.	branch to pat. tbl before	pat. to counter before	predict	actual	pat. to counter after	branch to pat. tbl after
1	TTTN	01	not taken	taken	10	TTNT
2	TTNT	01	not taken	taken	10	TNTT
3	TNTT	11	taken	taken	11	NTTT
4	NTTT	01	not taken	taken	10	TTTT
1	TTTN	10	taken	taken	11	TTNT



# history of history

PC of branch

0x40042A

hash

idx	recent pattern
0	NNNN
1	TNTT
2	TTTNTTT
3	TTTT
...	...
14	NTTN
15	TTTT

actual outcome  
from commit(?) stage

pattern  
NNNN  
NNNT  
...  
NTTT  
...  
TNTT  
...  
TTNT  
TTTT

counter

00
00
...
10 11
...
11
...
01 10
01 10 11
11

prediction  
to fetch sta

iter.	branch to pat. tbl before	pat. to counter before	predict	actual	pat. to counter after	branch to pat. tbl after
1	TTTN	01	not taken	taken	10	TTNT
2	TTNT	01	not taken	taken	10	TNTT
3	TNTT	11	taken	taken	11	NTTT
4	NTTT	01	not taken	taken	10	TTTT
1	TTTN	10	taken	taken	11	TTNT

# history of history

PC of branch

0x40042A

hash

idx	recent pattern
0	NNNN
1	TNTT
2	TTTNTTT
3	TTTT
...	...
14	NNTN
15	TTTT

actual outcome  
from commit(?) stage

pattern  
NNNN  
NNNT  
...  
NTTT  
...  
TNTT  
...  
TTNT  
TTTN  
TTTT

counter

00
00
...
<del>10</del> 11
...
11
...
<del>01</del> 10
<del>01</del> <del>10</del> 11
11

iter.	branch to pat. tbl before	pat. to counter before	predict	actual	pat. to counter after	branch to pat. tbl after
1	TTTN	01	not taken	taken	10	TTNT
2	TTNT	01	not taken	taken	10	TNTT
3	TNTT	11	taken	taken	11	NTTT
4	NTTT	01	not taken	taken	10	TTTT
1	TTTN	10	taken	taken	11	TTNT

prediction  
to fetch sta

# local patterns and collisions (1)

```
i = 10000;  
do {  
    p = malloc(...);  
    if (p == NULL) goto error; // BRANCH 1  
    ...  
} while (i-- != 0); // BRANCH 2
```

---

what if branch 1 and branch 2 hash to same table entry?

# local patterns and collisions (1)

```
i = 10000;  
do {  
    p = malloc(...);  
    if (p == NULL) goto error; // BRANCH 1  
    ...  
} while (i-- != 0); // BRANCH 2
```

---

what if branch 1 and branch 2 hash to same table entry?

pattern: TNTNTNTNTNTNTNTNT...

actually no problem to predict!

## local patterns and collisions (2)

```
i = 10000;  
do {  
    if (i % 2 == 0) goto skip; // BRANCH 1  
    ...  
    p = malloc(...);  
    if (p == NULL) goto error; // BRANCH 2  
skip: ...  
} while (i-- != 0); // BRANCH 3
```

---

what if branch 1 and branch 2 and branch 3 hash to same table entry?

## local patterns and collisions (2)

```
i = 10000;  
do {  
    if (i % 2 == 0) goto skip; // BRANCH 1  
    ...  
    p = malloc(...);  
    if (p == NULL) goto error; // BRANCH 2  
skip: ...  
} while (i-- != 0); // BRANCH 3
```

---

what if branch 1 and branch 2 and branch 3 hash to same table entry?

pattern: TTNNTTNNTTNNTTNNTT

also no problem to predict!

## local patterns and collisions (2)

```
i = 10000;  
do {  
    if (A) goto one // BRANCH 1  
    ...  
one:  
    if (B) goto two // BRANCH 2  
    ...  
two:  
    if (A or B) goto three // BRANCH 3  
    ...  
    if (A and B) goto three // BRANCH 4  
    ...  
three:  
    ... // changes A, B  
} while (i-- != 0);
```

---

what if branch 1-4 hash to same table entry?

better for prediction of branch 3 and 4

# global history predictor: idea

one predictor idea: ignore the PC

just record taken/not-taken pattern for all branches

lookup in big table like for local patterns



# global history predictor (1)

branch history register

NTTT

*pat*      *counter*

NNNN	00
NNNT	00
...	...
NTTT	10
TNNN	01
TNNT	10
TNTN	11
...	...
TTTN	10
TTTT	11

outcome  
from  
commit(?)

prediction  
to fetch stage

# global history predictor (1)

```

i = 10000;
do {
  if (i % 2 == 0) goto skip;
  ...
  if (p == NULL) goto error;
skip:
  ...
} while (i-- != 0);
  
```

branch history register

NTTT

<i>pat</i>	<i>counter</i>
NNNN	00
NNNT	00
...	...
<b>NTTT</b>	<b>10</b>
TNNN	01
TNNT	10
TNTN	11
...	...
TTTN	10
TTTT	11

outcome  
from  
commit(?)

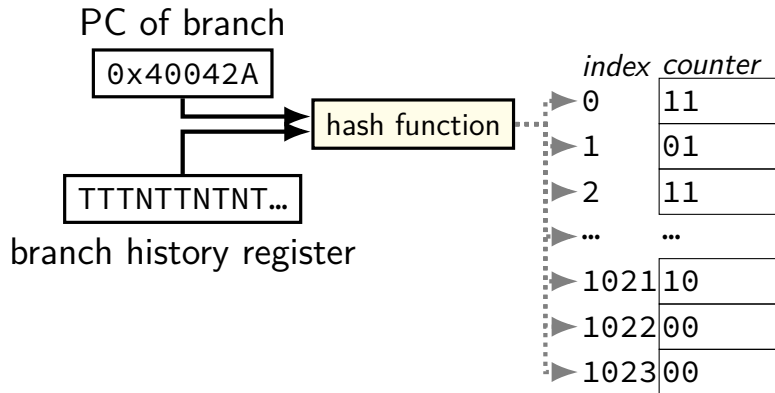
prediction  
to fetch stage

iter./branch	history before	counter before	predict	outcome	counter after	history after
0/mod 2	NTTT	10	taken	taken	11	TTTT
0/loop	TTTT			taken		TTTT
1/mod 2	TTTT			not taken		TTTN
1/error	TTTN			not taken		TTNN
1/loop	TNNT			taken		NNTT
2/mod 2	NNTT			taken		NTTT
2/loop	TTTT			taken		TTTT

# correlating predictor

global history *and* local info good together

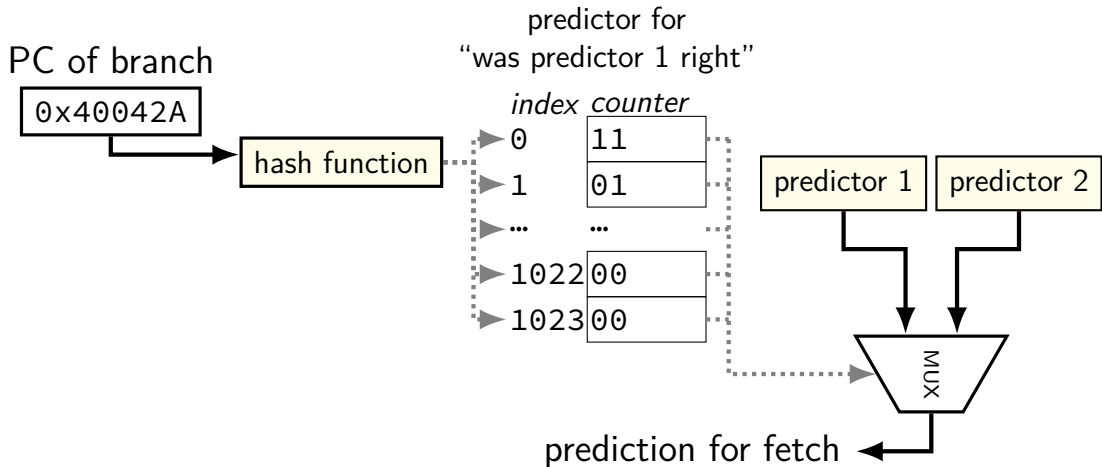
one idea: **combine history register + PC** (“gshare”)



# mixing predictors

different predictors good at different times

one idea: have two predictors, + predictor to predict which is right



# loop count predictors (1)

```
for (int i = 0; i < 64; ++i)  
    ...
```

can we predict this perfectly with predictors we've seen

yes — local or global history with 64 entries

but this is very important — more efficient way?

## loop count predictors (2)

loop count predictor idea: look for NNNNNNT+repeat (or TTTTTTN+repeat)

track for each possible loop branch:

- how many repeated Ns (or Ts) so far

- how many repeated Ns (or Ts) last time before one T (or N)

known to be used on Intel

# benchmark results

from 1993 paper


(not representative of modern workloads?)

rate for conditional branches on benchmark

variable table sizes

## 2-bit ctr + local history


from McFarling, "Combining Branch Predictors" (1993)






# 2-bit (bimodal) + local + global hist

from McFarling, "Combining Branch Predictors" (1993)



# global + hash(global+PC) (gshare/gselect)

from McFarling, "Combining Branch Predictors" (1993)



# real BP?

details of modern CPU's branch predictors often not public  
but...

Google Project Zero blog post with reverse engineered details

`https://googleprojectzero.blogspot.com/2018/01/reading-privileged-memory-with-side.html`

for RE'd BTB size:

`https://xania.org/201602/haswell-and-ivy-btb`

# reverse engineering Haswell BPs

## branch target buffer

- 4-way, 4096 entries

- ignores bottom 4 bits of PC?

- hashes PC to index by shifting + XOR

- seems to store 32 bit offset from PC (not all 48+ bits of virtual addr)

## indirect branch predictor

- like the global history + PC predictor we showed, but...

- uses history of recent branch addresses instead of taken/not taken

- keeps some info about last 29 branches

## what about conditional branches??? loops???

- couldn't find a reasonable source