

STAM Center  
SECURE, TRUSTED, AND ASSURED MICROELECTRONICS

ASU Engineering  
The Fulton School of  
Arizona State University

## CSE 520 Computer Architecture II

### Branch Prediction

Prof. Michel A. Kinsy

1

---

---

---

---

---

---

---

---

STAM Center  
SECURE, TRUSTED, AND ASSURED MICROELECTRONICS

ASU Engineering  
The Fulton School of  
Arizona State University

### Instruction Interactions

- An instruction in the pipeline may need a resource being used by another instruction in the pipeline
  - Structural hazard
- An instruction may depend on something produced by an earlier instruction
  - Dependence may be for a data calculation
    - Data hazard
  - Dependence may be for calculating the next address
    - Control hazard (branches, interrupts)

2

---

---

---

---

---

---

---

---

STAM Center  
SECURE, TRUSTED, AND ASSURED MICROELECTRONICS

ASU Engineering  
The Fulton School of  
Arizona State University

### Modern Processors: Complex Pipeline Structures

The diagram illustrates a complex pipeline structure. It starts with four sequential stages: IF (Instruction Fetch), ID (Instruction Decode), Issue (Instruction Issue), and WB (Write Back). The Issue stage is highlighted in pink and is associated with 'GPR's' and 'FPR's'. From the Issue stage, the pipeline branches into several parallel functional units: ALU (Arithmetic Logic Unit), Mem (Memory Access), Fadd (Floating Point Add), Fmul (Floating Point Multiply), and Fdiv (Floating Point Divide). There are vertical ellipsis dots between Fmul and Fdiv, indicating other units. All these units feed into the WB stage. There is also a feedback loop from the WB stage back to the Issue stage.

3

---

---

---

---

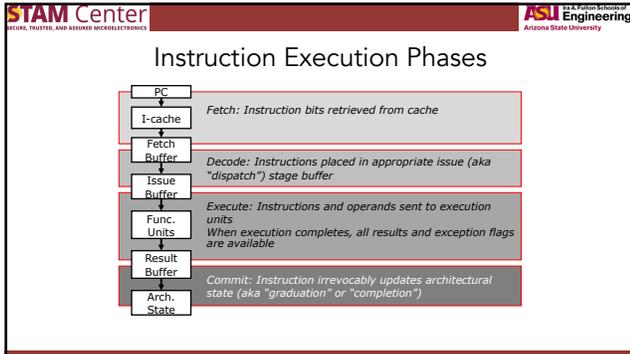
---

---

---

---





7

---

---

---

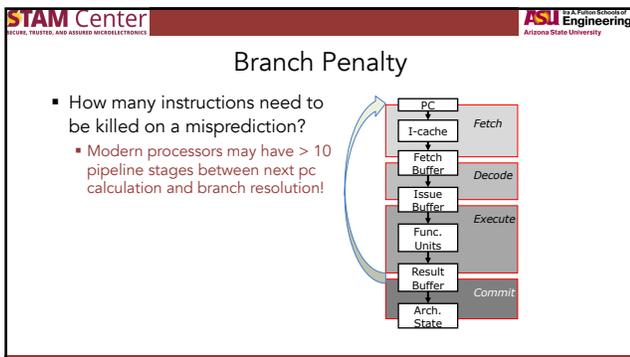
---

---

---

---

---



8

---

---

---

---

---

---

---

---

**Average Branches Distance**

- Average Run-Length between Branches
  - Average dynamic instruction mix from SPEC92:
 

	SPECint92	SPECfp92
ALU	39 %	13 %
FPU Add		20 %
FPU Mult		13 %
load	26 %	23 %
store	9 %	9 %
branch	16 %	8 %
other	10 %	12 %

SPECint92: compress, eqntott, espresso, gcc, li  
SPECfp92: doduc, ear, hydro2d, mdjdp2, su2cor

9

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Branches and Jumps

- Each instruction fetch depends on one or two pieces of information from the preceding instruction:
  - Is the preceding instruction a taken branch?
  - If so, what is the target address?

Instruction	Taken known?	Target known?
J	After Inst. Decode	After Inst. Decode
JAL/JALR	After Inst. Decode	After Reg Fetch
BEQ/BNE	After Inst. Execute	After Inst. Decode

10

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Branches and Jumps

- Each instruction fetch depends on one or two pieces of information from the preceding instruction:
  - Is the preceding instruction a taken branch?
  - If so, what is the target address?

Type	Direction at fetch time	Number of possible next fetch addresses?	When is next fetch address resolved?
Conditional	Unknown	2	Execution (register dependent)
Unconditional	Always taken	1	Decode (PC + offset)
Call	Always taken	1	Decode (PC + offset)
Return	Always taken	Many	Execution (register dependent)
Indirect	Always taken	Many	Execution (register dependent)

11

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Branch Penalties in Modern Pipelines

- UltraSPARC-III instruction fetch pipeline stages (in-order issue, 4-way superscalar, 750MHz, 2000)

12

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Control Dependencies

- Branches are very frequent
  - Approx. 20% of all instructions
- Can not wait until we know where it goes
  - Long pipelines
    - Branch outcome known after x cycles
    - No scheduling past the branch until outcome known
  - Superscalars (e.g., 4-way)
    - Branch every cycle or so!
    - One cycle of work, then bubbles for -x cycles?

13

---

---

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Control Flow Penalty

- Assume a 4-wide superscalar pipeline with 20-cycle branch resolution latency
- How long does it take to fetch 400 instructions?
  - Assume no fetch breaks and 1 out of 4 instructions is a branch
  - 100% accuracy
    - 100 cycles (all instructions fetched on the correct path)
    - No wasted work
  - 99% accuracy
    - 100 (correct path) + 20 (wrong path) = 120 cycles
    - 20% extra instructions fetched
  - 98% accuracy
    - 100 (correct path) + 20 \* 2 (wrong path) = 140 cycles
    - 40% extra instructions fetched
  - 95% accuracy
    - 100 (correct path) + 20 \* 5 (wrong path) = 200 cycles
    - 100% extra instructions fetched

14

---

---

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Reducing Control Flow Penalty

- Software solutions
  - Eliminate branches - loop unrolling increases the run length
  - Reduce resolution time - instruction scheduling compute the branch condition as early as possible (of limited value)
- Hardware solutions
  - Stall the pipeline until we know the next fetch address
  - Guess the next fetch address
    - Speculate - branch prediction speculative execution of instructions beyond the branch
  - Employ delayed branching (branch delay slot)
  - Do something else (fine-grained multithreading)
  - Eliminate control-flow instructions (predicated execution)
  - Fetch from both possible paths (if you know the addresses of both possible paths)
    - Multipath execution

15

---

---

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Branch Prediction

- Motivation:
  - Branch penalties limit performance of deeply pipelined processors
  - Modern branch predictors have high accuracy (>95%) and can reduce branch penalties significantly
- Required hardware support:
  - Prediction structures: branch history tables, branch target buffers, etc.
  - Mispredict recovery mechanisms:
    - Keep result computation separate from commit
    - Kill instructions following branch in pipeline
    - Restore state to state following branch

16

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Branch Prediction Techniques

- Compile time (static)
  - Always not taken
  - Always taken
  - BTFN (Backward taken, forward not taken)
  - Profile based (likely direction)
  - Program analysis based (likely direction)
- Run time (dynamic)
  - Last time prediction (single-bit)
  - Two-bit counter based prediction
  - Two-level prediction (global vs. local)
  - Hybrid

17

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Static Branch Prediction

- Overall probability a branch is taken is ~60-70% but:
 

*backward*  
90%

*forward*  
50%
- ISA can attach preferred direction semantics to branches, e.g., Motorola MC88110
  - `bne0` (preferred taken) `beq0` (not taken)
- ISA can allow arbitrary choice of statically predicted direction, e.g., HP PA-RISC, Intel IA-64
  - Typically reported as ~80% accurate

18

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Static Branch Prediction

- Always not-taken
  - Simple to implement: no need for BTB, no direction prediction
  - Low accuracy: ~30-40%
  - Compiler can layout code such that the likely path is the "not-taken" path
- Always taken
  - No direction prediction
  - Better accuracy: ~60-70%
    - Backward branches (i.e. loop branches) are usually taken
    - Backward branch: target address lower than branch PC
- Backward taken, forward not taken (BTFN)
  - Predict backward (loop) branches as taken, others not-taken

19

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Branch Prediction Techniques

- Compile time (static)
  - Always not taken
  - Always taken
  - BTFN (Backward taken, forward not taken)
  - Profile based (likely direction)
  - Program analysis based (likely direction)
- Run time (dynamic)
  - Last time prediction (single-bit)
  - Two-bit counter based prediction
  - Two-level prediction (global vs. local)
  - Hybrid

20

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Dynamic Branch Prediction

- Key Idea: Predict branches based on the dynamic execution behavior of the program
  - Collected at run-time
- Advantages
  - Prediction based on history of the execution of branches
  - It can adapt to dynamic changes in branch behavior
  - No need for static profiling
- Disadvantages
  - More complex architecture (requires additional hardware)

21

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Dynamic Branch Prediction

- Learning based on past behavior
- Temporal correlation
  - The way a branch resolves may be a good predictor of the way it will resolve at the next execution
- Spatial correlation
  - Several branches may resolve in a highly correlated manner (a preferred path of execution)

22

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Dynamic Branch Prediction

- Requires three things to be predicted at fetch stage:
  - Whether the fetched instruction is a branch
  - (Conditional) branch direction
  - Branch target address (if taken)
    - So we also need to predict the next fetch address (to be used in the next cycle)
- Target address remains the same for a conditional direct branch across dynamic instances
  - Store the target address from previous instance and access it with the PC
  - Branch Target Buffer (BTB) or Branch Target Address Cache

23

---

---

---

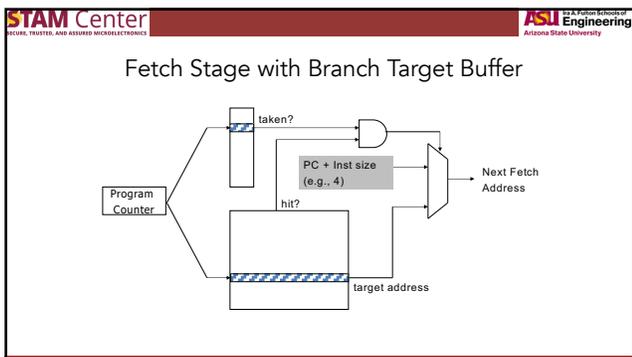
---

---

---

---

---



24

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Dynamic Prediction

Prediction as a feedback control process

- Predict
- Update

25

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Predictor Primitive

- Indexed table holding values
- Operations
  - Predict
  - Update
- Algebraic notation
  - Prediction = P[Width, Depth](Index; Update)

26

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### One-bit Predictor

- $A21064(PC; T) = P[1, 2K](PC; T)$
- What happens on loop branches?
  - At best, mispredicts twice for every use of loop

27

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** Arizona State University

### Last Time Predictor

- Last time predictor
  - Single bit per branch (stored in BTB)
  - Indicates which direction branch went last time it executed
  - TTTTTTTTTTNNNNNNNNNN → 90% accuracy
- Always mispredicts the last iteration and the first iteration of a loop branch
  - Accuracy for a loop with N iterations =  $(N-2)/N$
  - Works well for loops with a large number of iterations
  - Works poorly for loops with small number of iterations or non-correlated branches
  - TNTNTNTNTNTNTNTNTNTN → 0% accuracy

28

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** Arizona State University

### Implementation of the 1-bit Predictor

29

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** Arizona State University

### State Machine of the 1-bit Predictor

30

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Two-bit Predictor

- Counter[W,D](I; T) = P[W, D]
  - (I; if T then P+1 else P-1)
- A21164(PC; T) = MSB(Counter[2, 2K](PC; T))

31

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Branch Prediction Bits

- Assume 2 BP bits per instruction
- Use saturating counter

On ↑ -taken ↓	1	1	Strongly taken
	1	0	Weakly taken
	0	1	Weakly -taken
	0	0	Strongly -taken

32

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** The Future School of Arizona State University

### Two-bit Saturation Predictor

33

---

---

---

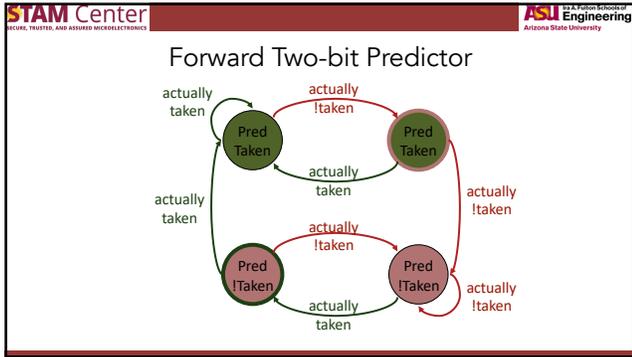
---

---

---

---

---



34

---

---

---

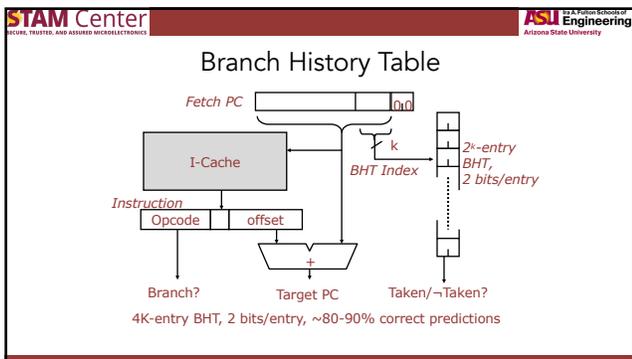
---

---

---

---

---



35

---

---

---

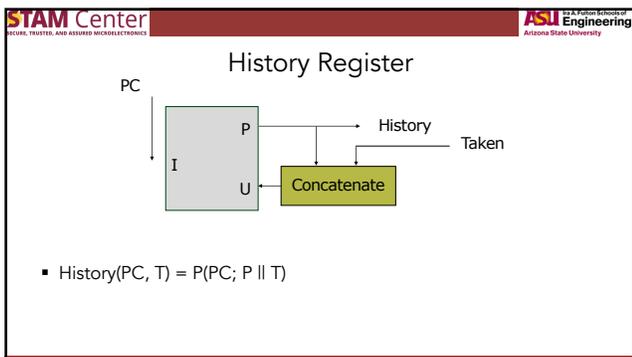
---

---

---

---

---



36

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** Arizona State University

### Exploiting Spatial Correlation

```

if (x[i] < 7) then
  y += 1;
if (x[i] < 5) then
  c -= 4;
    
```

- If first condition false, second condition also false
- History register, H, records the direction of the last N branches executed by the processor

37

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** Arizona State University

### Two-level Predictor

$$\text{GHist}(;T) = \text{Counter}(\text{History}(0, T); T)$$

$$\text{Ind-Ghist}(\text{PC};T) = \text{Counter}(\text{PC} \parallel \text{Hist}(\text{GHist}(;T);T))$$

38

---

---

---

---

---

---

---

---

**STAM Center** SECURE, TRUSTED, AND ASSURED MICROELECTRONICS **ASU Engineering** Arizona State University

### Tournament Predictor

- Alpha 21264
  - Minimum branch penalty: 7 cycles
  - Typical branch penalty: 11+ cycles
  - 48K bits of target addresses stored in I-cache
  - Predictor tables are reset on a context switch

39

---

---

---

---

---

---

---

---



**STAM Center**  
SYSTEMS, TESTED, AND ASSURED MICROELECTRONICS

**ASU** Arizona State University  
**Engineering**

### Next learning Module

- Complex Pipelining: Superscalar Architecture

---

---

---

---

---

---

---