

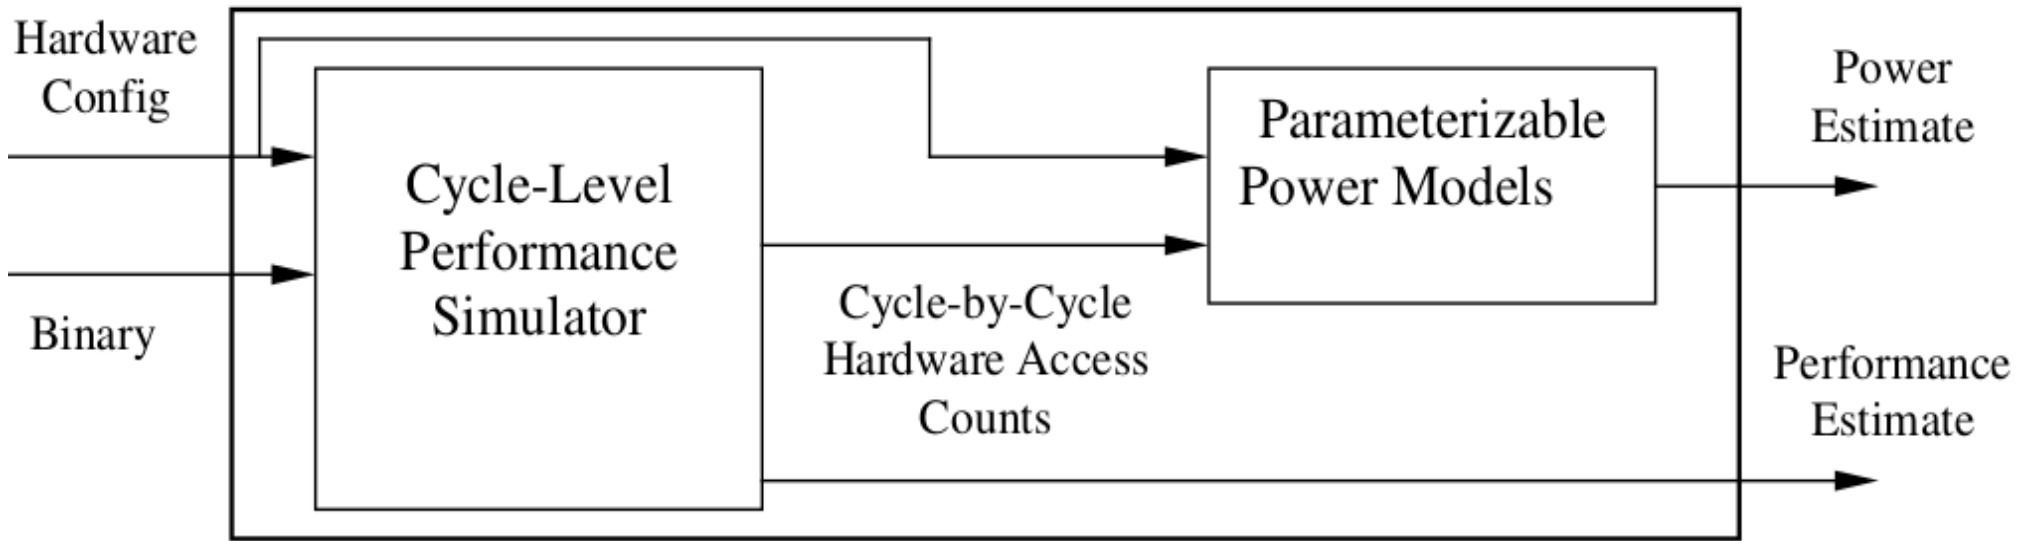
Wattch: A Framework for Architectural-Level Power Analysis and Optimization, 2000

- Power dissipation of processor is significant (power wall)
 - Conventional air-cooling techniques limits power
- Existing power analysis tools are highly accurate but require complete circuit design or HDL
- Simulator Framework to estimate power early in the design process based on high-level description of processor.
 - Accuracy within 10%

Wattch: A Framework for Architectural-Level Power Analysis and Optimization, 2000

- Parameterizable power models for different hardware structures
 - Array Structures (e.g. Cache)
 - # of rows, # of columns, # of read/write ports
 - CAM Structures (e.g. TLB)
 - # of tags, # of bits per tag, # of ports
 - Complex Logic Blocks (e.g. ALU)
 - Clocking

Wattch: A Framework for Architectural-Level Power Analysis and Optimization, 2000



- SimpleScalar (Performance Simulator) keeps track of which units are accessed per cycle and how (# of ports)
- Different models for clock gating: disable unused resources within a unit

Questions

- Other (affordable) techniques instead of air-cooling?
- The paper mentions that 14% of the Pentium Pros power is consumed by instruction decode. How does this number look like for an RISC machine?
 - At a high-level: How much does the ISA affect the power consumption?
- Which architectural improvements were implemented in the last decade to reduce power consumption?

Optimizing Pipelines for Power and Performance, 2002

- Determine optimal pipeline depth and target frequency.
 - Optimal for performance
 - Optimal for power
- 18 FO4 delay per stage for power-performance optimality
 - 15 FO4 for logic, 3 FO4 for latch insertion delay
- 10 FO4 delay per stage for performance optimality

Optimizing Pipelines for Power and Performance, 2002

- Analytical Pipeline Model

- Time per pipeline stage:

- t_i : completion time (without latches)
 - s_i : # of stages
 - c_i : latch delay per stage

$$T_i = ((\tilde{t}_i / s_i) + c_i),$$

- Throughput

$$G = \frac{u_1}{T_{fxu}} + \frac{u_2}{T_{fpu}} + \frac{u_3}{T_{lsu}} + \frac{u_4}{T_{bru}}$$

Optimizing Pipelines for Power and Performance, 2002

- Simulation Model
 - Start at a design of FO4 19 and scale power usage with FO4-depth

Questions

- How many pipeline stages would we have on a modern processor, given a FO4 of 18?
- We saw two factors that influence pipeline depth: performance and power. How hard is it to engineer deep pipelines? What about validation?