

## Chapter 1: Fundamentals...

- Trends
  - ILP, DLP, TLP, RLP
  - Performance, Parallelism, Power, Energy, Feature size
  - Bandwidth, Latency, Capacity
  - Which factor limits which other factor?
- Classes
  - PMD, Desktop, Server, Cluster, WSC, Embedded
- Flynn Taxonomy
- Formulas
  - Cost of integrated circuit
  - MTTF, MTBF, Availability
  - Speedup, Amdahl's law, CPU time, CPI

1

## Chapter 4: Vector, SIMD, GPUs

- Vector RISC vs. GPU vs. SIMD
- CPU vs. GPU
  - Memory, ALU, FPU, branching, bandwidth, programming
- Code vectorization and transformations
  - GCD test, loop-carried dependences
  - Fixed vs. variable length loops
  - Strip-mining, conditional handling, parallel reductions

4

## Chapter 2: Memory

- Memory hierarchy, sizes, types, timing, optimizations
- Associativity, TLB mapping, misses
- SRAM, DRAM, DDR, GDDR, Flash
- Virtualization
  - Memory protection, CPU states
- Formulas
  - Misses/instruction
  - AMAT
  - Instruction mixes

2

## Chapter 5: Threads

- Coherency vs. Consistency
- Cache states and their transitions
- Coherency protocols
  - Snooping, directories, NUMA
- Forms of Consistency
- Example codes,
  - Instruction interleaving
  - Atomics
  - Synchronization primitives

5

## Chapter 3: ILP (Appendix C)

- Dependences vs. hazards
- Data hazards vs. pipeline hazards
- Scheduling, unrolling, pipelining
- RISC pipeline and its variants
  - 5-stage and multistage
  - Many-cycle instructions
  - Bypassing and forwarding
- Tomasulo, scoreboard, and speculation basics

3

## Chapter 6: Warehouse-Scale Computers

- WSC vs. Server vs. Data Center vs. Cluster
- Programming model
- Hardware
  - Memory, network, disk hierarchy, timing
- Power, energy, cost breakdown

6