ECSE 425 Lecture 4: Dependability; Quantitative Principles of Design

H&P Chapter 1

Last Time

- Trends in Power
- Trends in IC Cost
  - Non-recurring expenses (NREs)
  - Recurring expenses
    - Manufacturing
    - Testing
    - Losses

Today

- Trends in Dependability
- Quantitative Principles of Computer Design
- Looking ahead …
  - On Friday, Pipelining
    - Read Appendix A!
  - Homework 1 due Monday
    - OH Today: 2-3 PM, Thursday: 11 AM-12 PM

IC Scaling and Resilience

Power

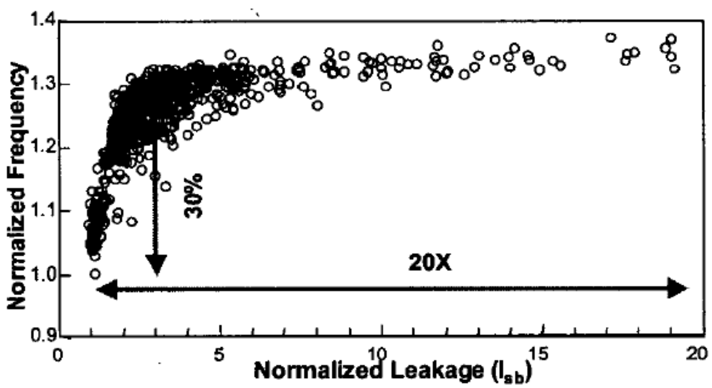

Variability

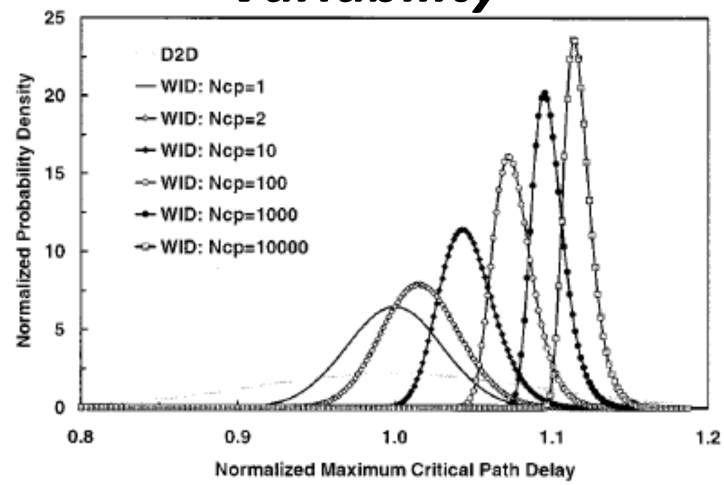

Reliability

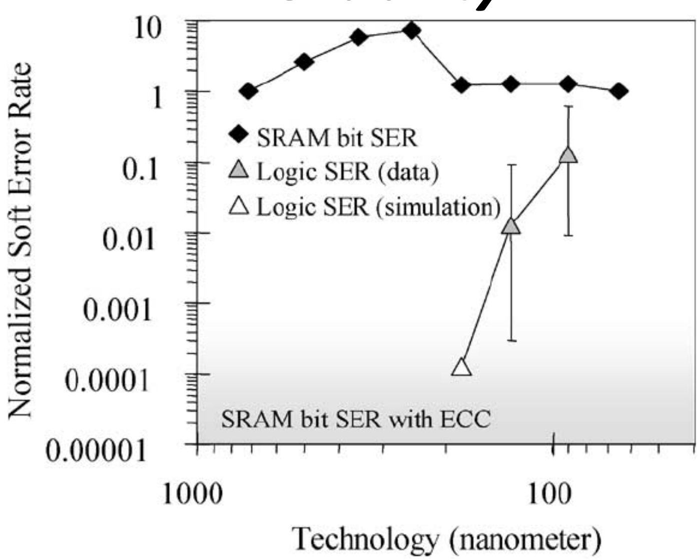

Degradation

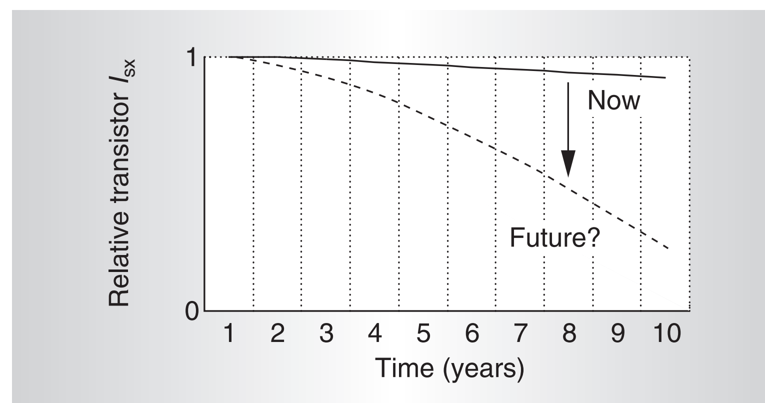

[Sources: Borkar, 2003; Bowman, 2005; Baumann, 2005; Borkar, 2005]

Service Level Agreements (SLAs)

- Provider pays a penalty to the customer when *downtime* > *threshold* over a given period
- Two states
    1. Service accomplishment (service delivered as spec.)
    2. Service interruption (service deviates from spec.)
- Service failures: state 1->2
- Service restoration: state 2->1

Quantifying Reliability

- Mean Time To Failure (MTTF)
  - On average, how long until the first failure?
- Failures in Time (FIT)
  - On average, how many failures per 10^9 hours?
  - MTTF = 1,000,000 hours ⇒ $10^9/10^6$ = 1000 FIT

Quantifying Availability

- Mean Time To Repair (MTTR)
  - Service time
- Mean Time Between Failures (MTBF)
  - MTBF = MTTF + MTTR
- Availability = MTTF / MTBF

Dependability Examples (1)

- Example: Disk subsystem
  - 10 disks, each rated at 1,000,000 hour MTTF
  - 1 SCSI controller, 500,000 hour MTTF
  - 1 power supply, 200,000 hour MTTF
  - 1 fan, 200,000 hour MTTF
  - 1 SCSI cable, 1,000,000 hour MTTF
- Assume
  - Failures are independent
  - Exponentially distributed lifetimes
- Compute the MTTF of the disk subsystem

Dependability Examples (2)

- Example: Disk subsystem with redundant power supplies

$$MTTF_{psp} = \frac{MTTF_{ps}/2}{\frac{MTTR_{ps}}{MTTF_{ps}}}$$

- What is the MTTF of the power supply pair, compared with a single power supply?

- What is the new MTTF of the disk subsystem?

Quantitative Principles of Design

- Take advantage of parallelism
- Take advantage of locality
- Make the common case fast

Exploit Parallelism at All Levels of Design

- Improve performance by performing many tasks simultaneously
- Hardware level: Multiple memory banks
  - Set-associative caches search the bank in parallel
- Instruction level: Pipelining and ILP exploitation
  - Overlapping instruction execution stages
- Thread and data level: Multiprocessors
  - Dividing the workload among multiple processors
- What are the trade-offs? Is parallelism “free?”

Exploit Spatial and Temporal Locality

- Locality and the 90%/10% rule
  - 90% of execution time spent on only 10% of the code
- Spatial locality
  - Items with nearby addresses tend to be referenced nearby in time (code without branching, array accesses)
- Temporal locality
  - Recently accessed items are likely to be accessed again soon (loop, reuse)
- Exploiting Locality
  - Branch or trace prediction to guess next instructions
  - Hierarchical memory with multi-word cache lines

Make the common case fast

- Focus optimization effort on the common case
- E.g., fetch and decode vs. multiplication
  - Fetch and decode are performed on every instruction
  - Multiply instructions occur only occasionally
- Small improvements in the common case trump massive improvements in the uncommon case
- The common case is often simpler
  - *E.g.*, overflow is rare when adding 2 numbers, so optimize the more common case of no overflow

Amdahl's Law

- When optimizing a fraction of total execution, what is the resulting speedup?
- *F* is the fraction of computation affected by the improvement
- This fraction is sped up by factor *S*

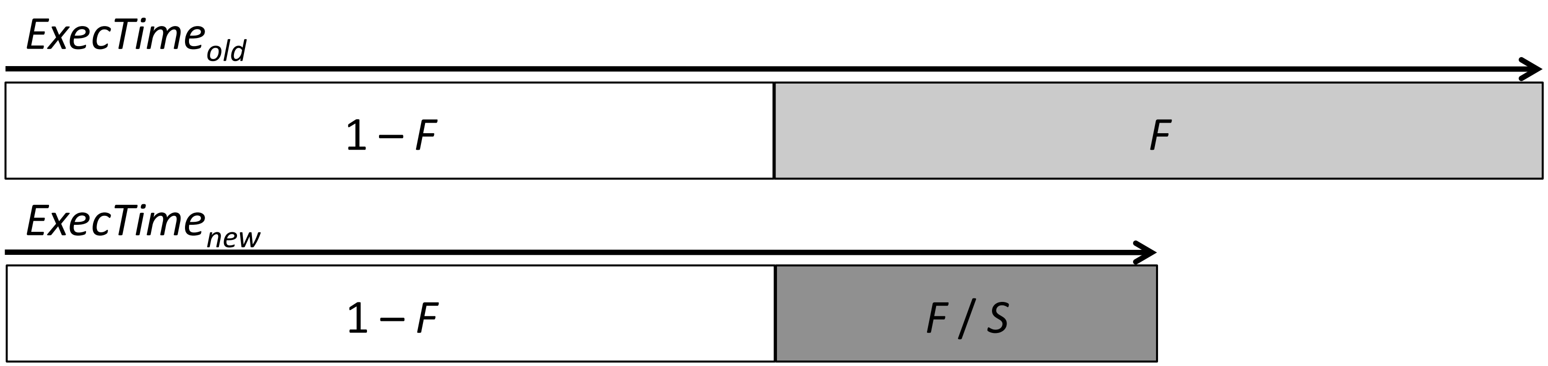

$$Speedup_{overall} = \frac{ExecTime_{old}}{ExecTime_{new}} = \frac{1}{(1-F)+\frac{F}{S}}$$

Amdah's Law: the Limits of Parallelism

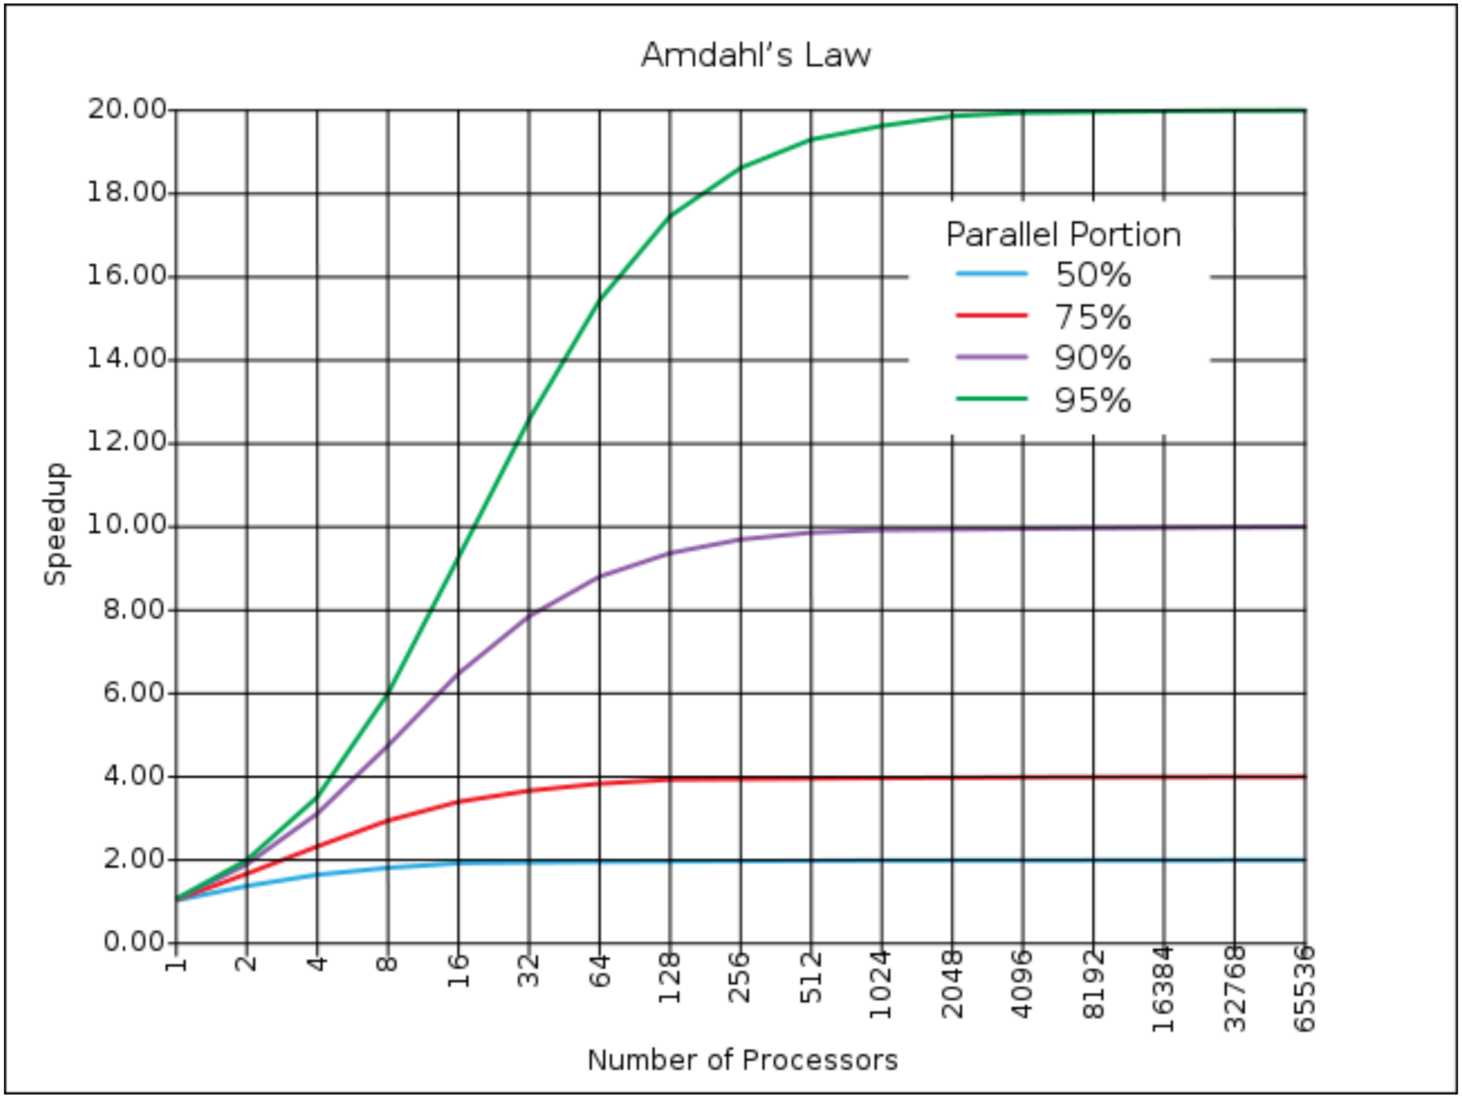

[Credit: Daniel, wikipedia.com]

Amdahl's Law and Reliability

- Example: Disk subsystem
  - 10 disks, each rated at 1,000,000 hour MTTF
  - 1 SCSI controller, 500,000 hour MTTF
  - 1 power supply, 200,000 hour MTTF
  - 1 fan, 200,000 hour MTTF
  - 1 SCSI cable, 1,000,000 hour MTTF
- Assume
  - Failures are independent
  - Exponentially distributed lifetimes
- The MTTF of the disk subsystem = 43,500 hours
- Power supply MTTF: 200K to 830M hours, 4150x better
  - What is the improvement in system MTTF?

Summary

- Trends in Dependability
  - Emerging resilience challenges
    - Manufacturing defects
    - Failure in the field (transient and permanent)
  - Availability = Time to Failure / Time Between Failures
- Principles of Computer Design
  - Exploit locality
  - Exploit parallelism
  - Make the common case fast

Next Time

- Quantifying Processor Performance