CS10
The Beauty and Joy of Computing

Lecture #8 : Concurrency

UC Berkeley EECS
Summer Instructor
Ben Chun

2012-07-02

LEAP SECOND BUG

Adding a single second to the official UTC time on Saturday (June 30) caused certain web sites to go down, servers to crash, and even some airline flight delays.

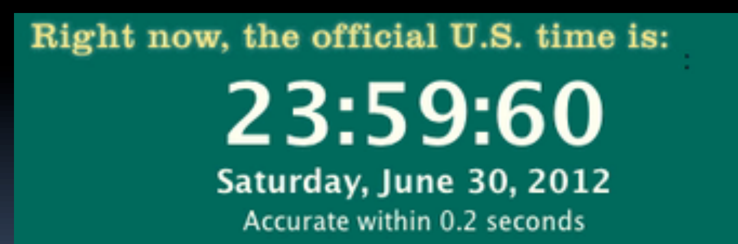

http://bit.ly/OWqNKN

Concurrency & Parallelism, 10 mi up…

Intra-computer

- Today!
- Computation split between cores <u>within one machine</u>
- Aka "multi-core"
  - Although GPU parallism is also "intra-computer"

Inter-computer

- Coming in Week 6!
- Computation split between <u>different machines</u>
- Aka "distributed computing"
  - Grid & cluster computing

Anatomy: 5 components of any Computer

Anatomy: 5 components of any Computer

John von Neumann invented this architecture

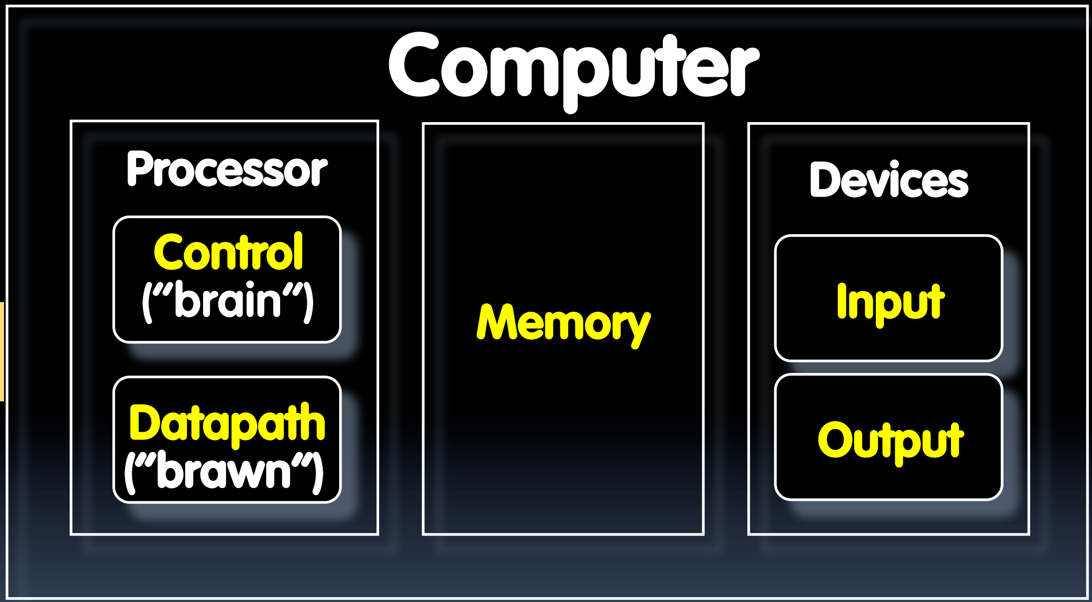

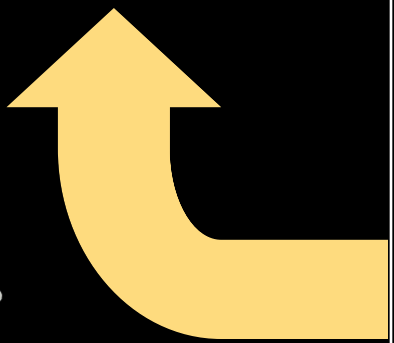

a) Control
b) Datapath
c) Memory
d) Input
e) Output

What causes the most headaches for SW and HW designers with multi-core computing?

But what is INSIDE a Processor?

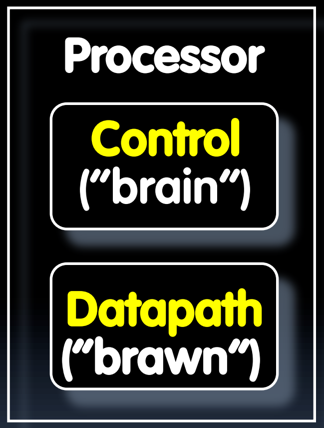

But what is INSIDE a Processor?

Processor
- **Control** ("brain")
- **Datapath** ("brawn")

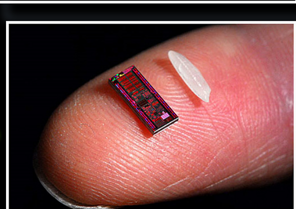

Bare Processor Die

Chip in Package

- Primarily Crystalline Silicon
- 1 mm – 25 mm on a side
- 2010 "feature size" (aka process) ~ 32 nm = 32 x 10^{-9} m (22nm in 2012, 14nm coming in 2013)
- 200 - 2000M transistors
- 3 - 10 conductive layers
- "CMOS" (complementary metal oxide semiconductor) - most common

- Package provides:
  - spreading of chip-level signal paths to board-level
  - heat dissipation.
- Ceramic or plastic with gold wires.

en.wikipedia.org/wiki/Moore's_law

Moore's Law

Predicts: 2X Transistors / chip every 2 years

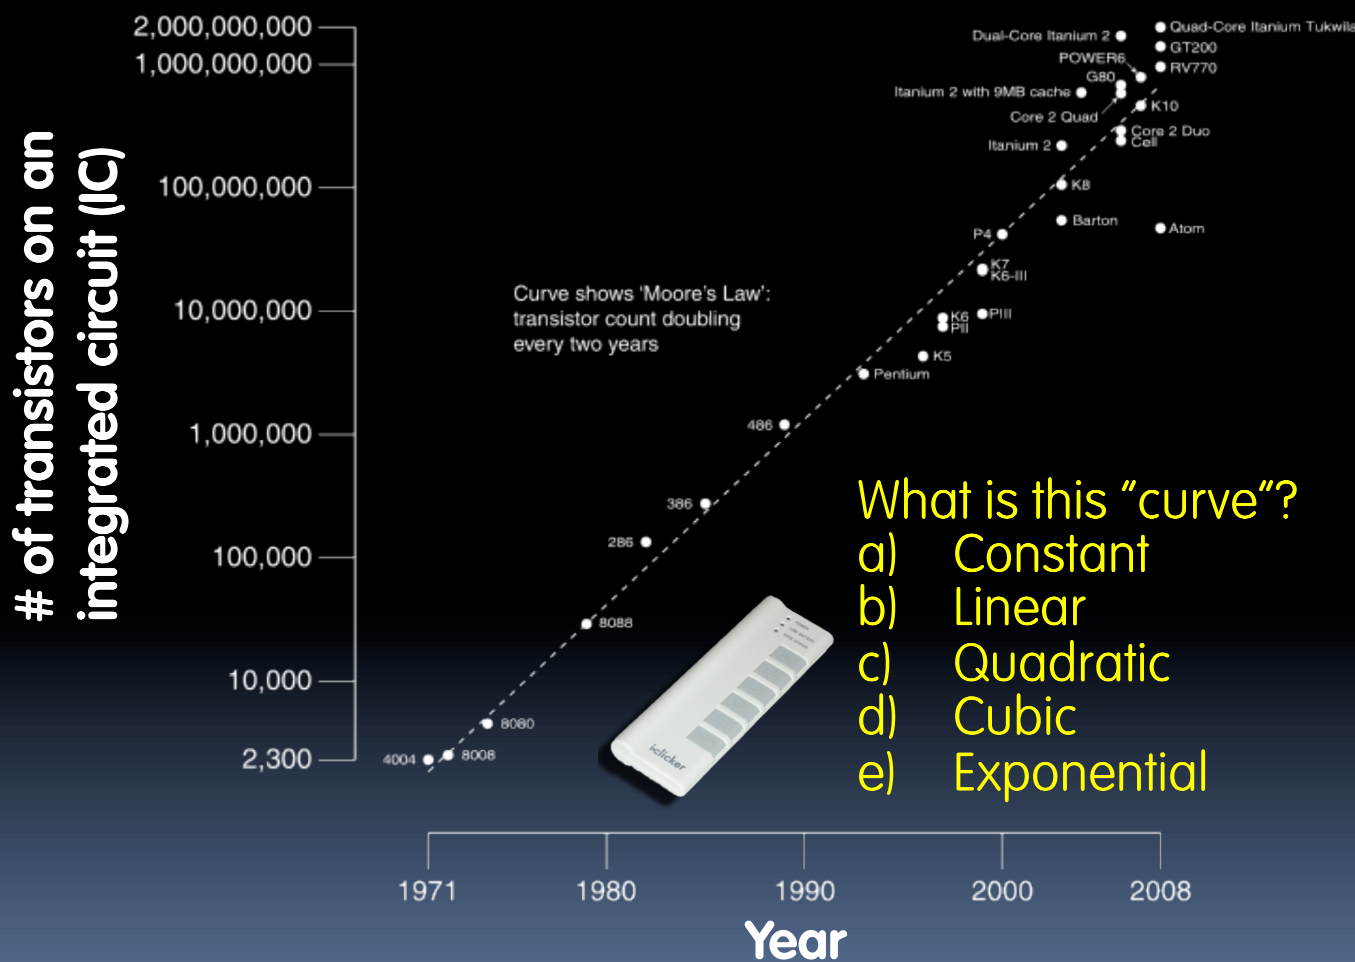

What is this "curve"?

a) Constant
b) Linear
c) Quadratic
d) Cubic
e) Exponential

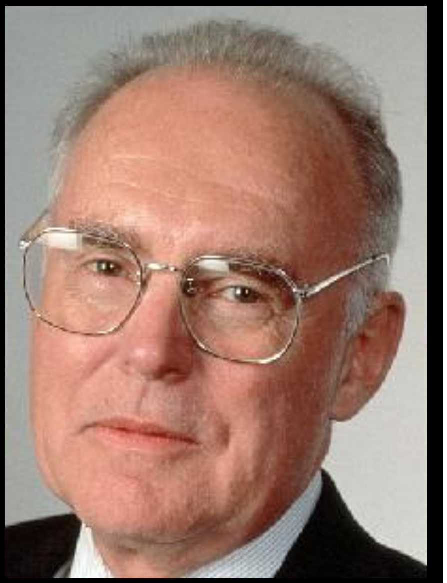

Gordon Moore
Intel Cofounder
B.S. Cal 1950!

Moore's Law and related curves

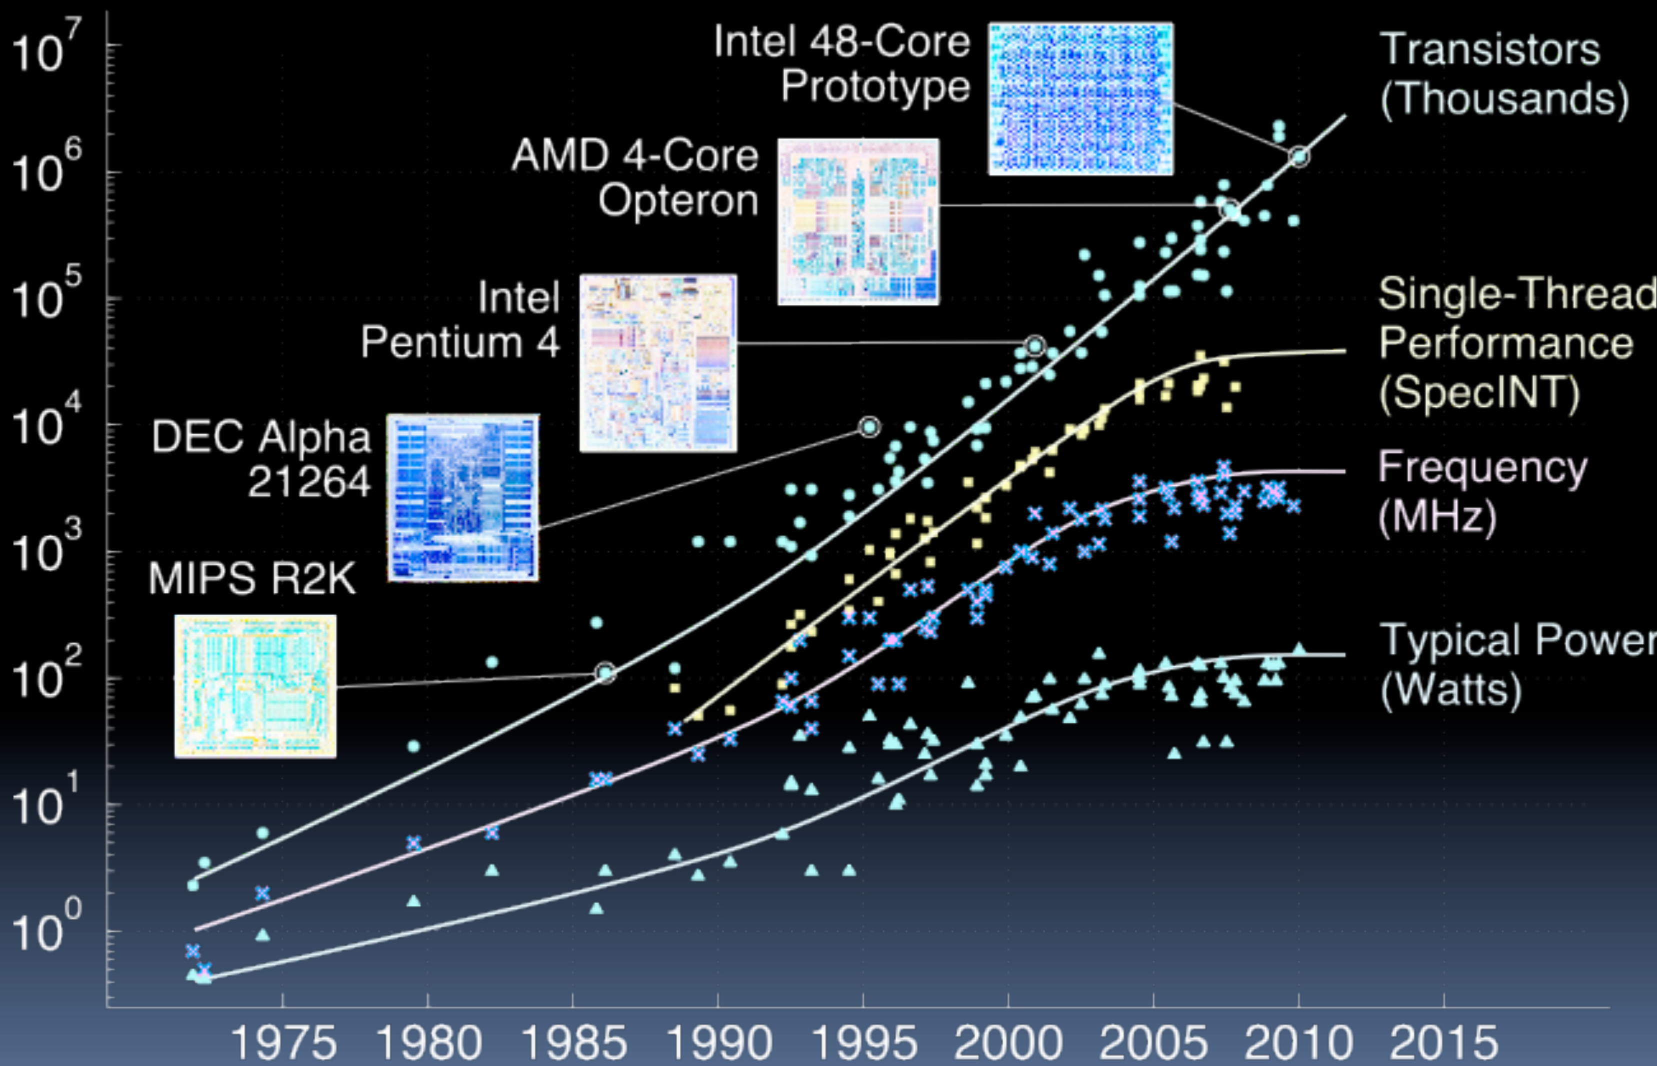

Data partially collected by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond

Moore's Law and related curves

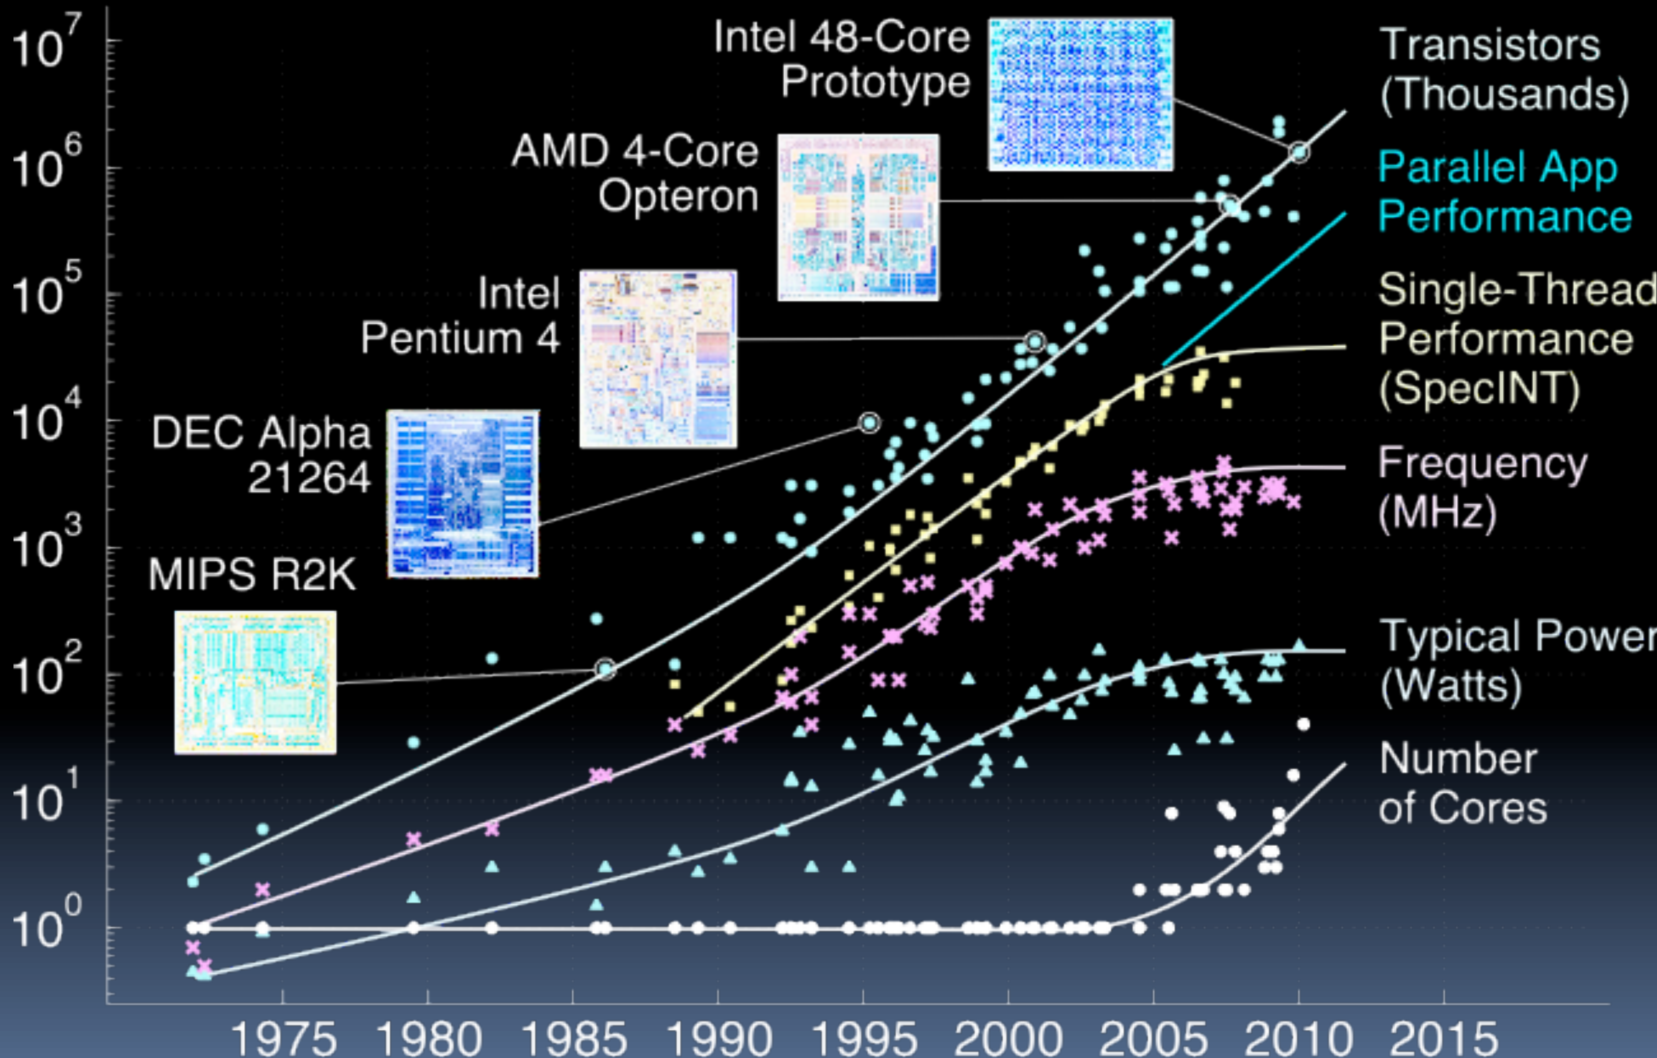

Data partially collected by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond

Chun, Summer 2012

Power Density Prediction circa 2000

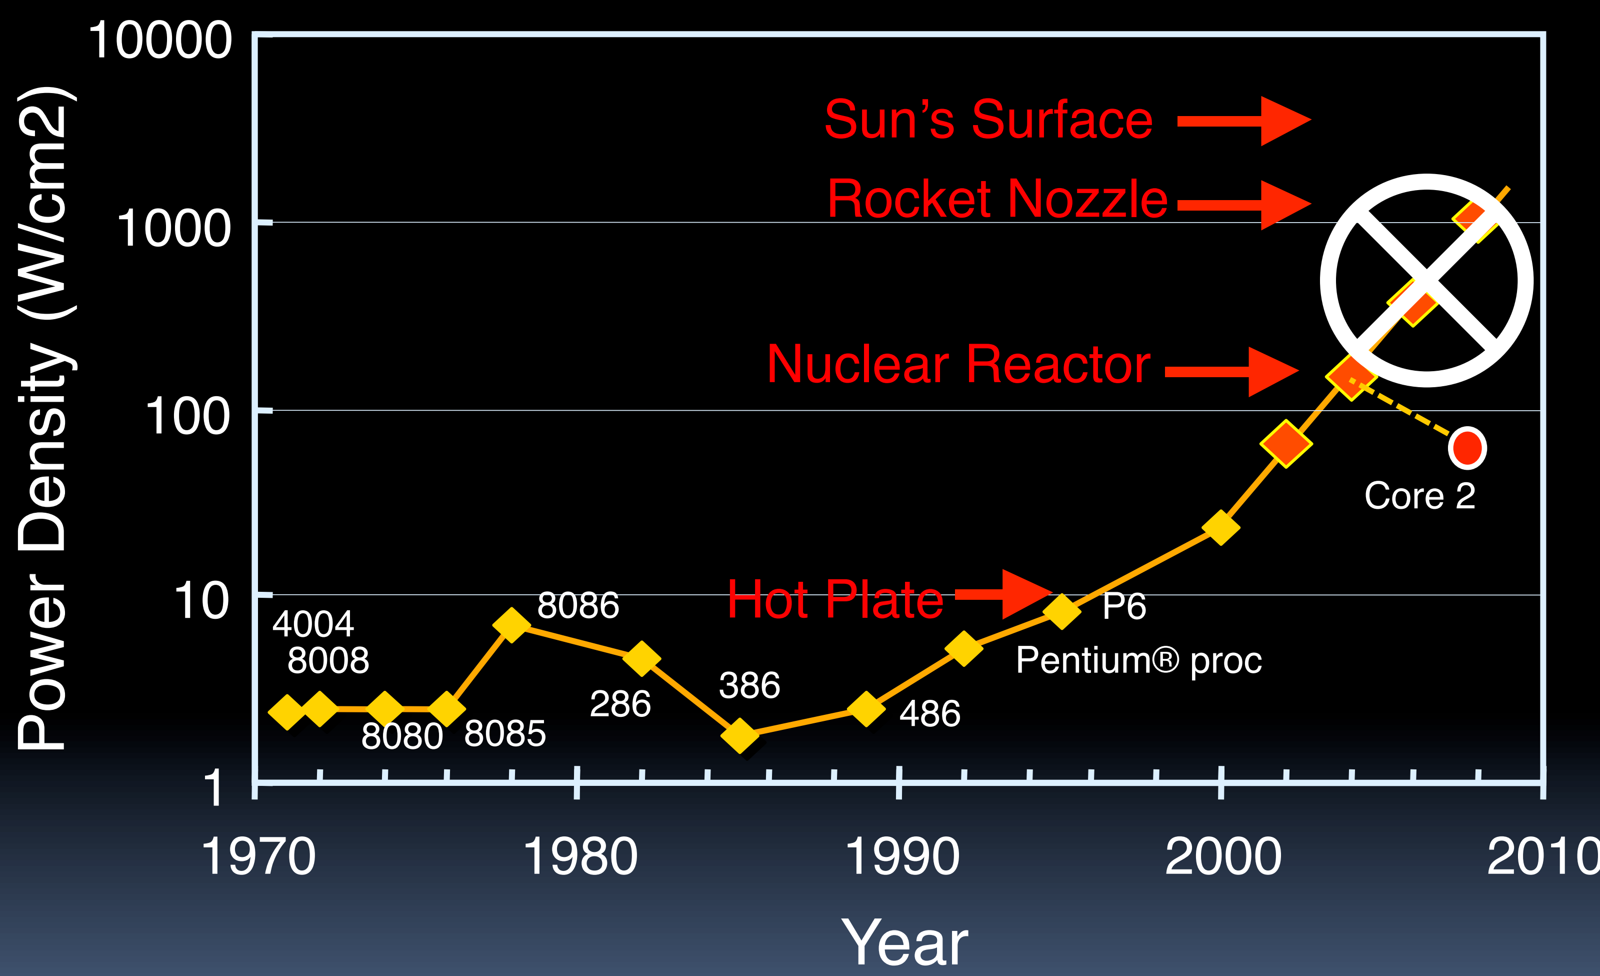

Source: S. Borkar (Intel)

Going Multi-core Helps Energy Efficiency

- **Power of typical integrated circuit ~ C V² *f***
  - C = Capacitance, how well it "stores" a charge
  - V = Voltage
  - *f* = frequency, i.e., how fast clock is (e.g., 3 GHz)

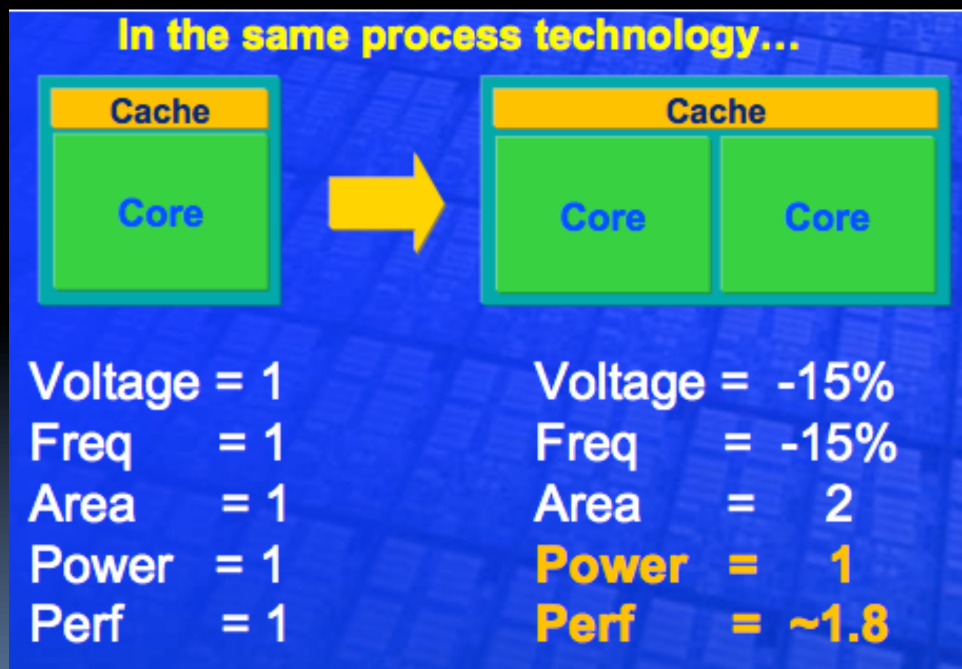

William Holt, HOT Chips 2005

Activity Monitor (on the lab Macs) shows how active your cores are

Energy & Power Considerations

$$\text{Power} = \frac{\text{Energy}}{\text{Second}} = \frac{\text{Energy}}{\text{Op}} \times \frac{\text{Ops}}{\text{Second}}$$

Power	Energy
Chip Packaging	Battery Life
Chip Cooling	Electricity Bill
System Noise	Mobile Device Weight
Case Temperature	
Data-Center Air Conditioning	

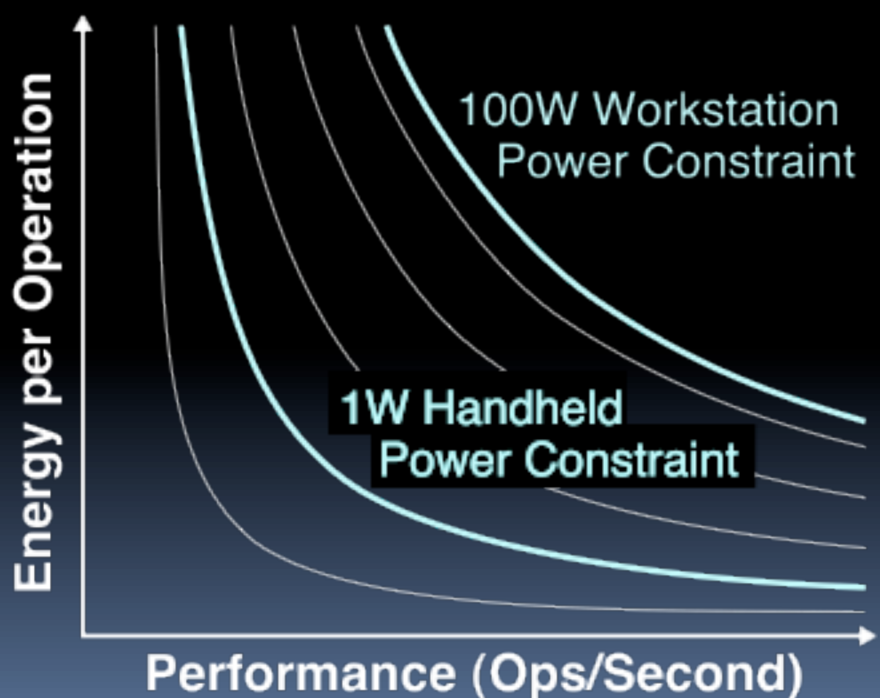

Courtesy: Chris Batten

Parallelism again? What's different this time?

"This shift toward increasing parallelism is not a triumphant stride forward based on breakthroughs in novel software and architectures for parallelism; instead, this plunge into parallelism is actually a retreat from even greater challenges that thwart efficient silicon implementation of traditional uniprocessor architectures."

– Berkeley View, December 2006

- HW/SW Industry bet its future that breakthroughs will appear before it's too late

Background: Threads

- A ***Thread*** **stands for "thread of execution", is a single stream of instructions**
  - A program / process can split, or fork itself into separate threads, which can (in theory) execute simultaneously.
  - An easy way to describe/think about parallelism
- **A single CPU can execute many threads by *Time Division Multipexing***

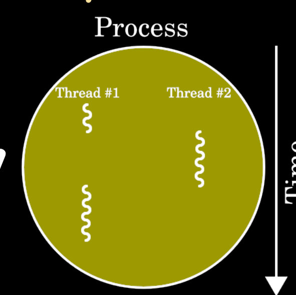

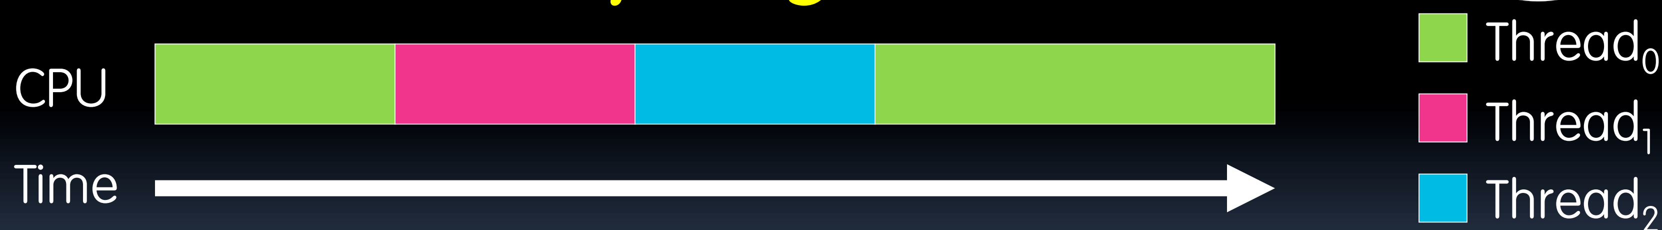

- ***Multithreading*** **is running multiple threads through the same hardware** en.wikipedia.org/wiki/Amdahl's_law

Speedup Issues : Amdahl's Law

- Applications can almost never be completely parallelized; some serial code remains

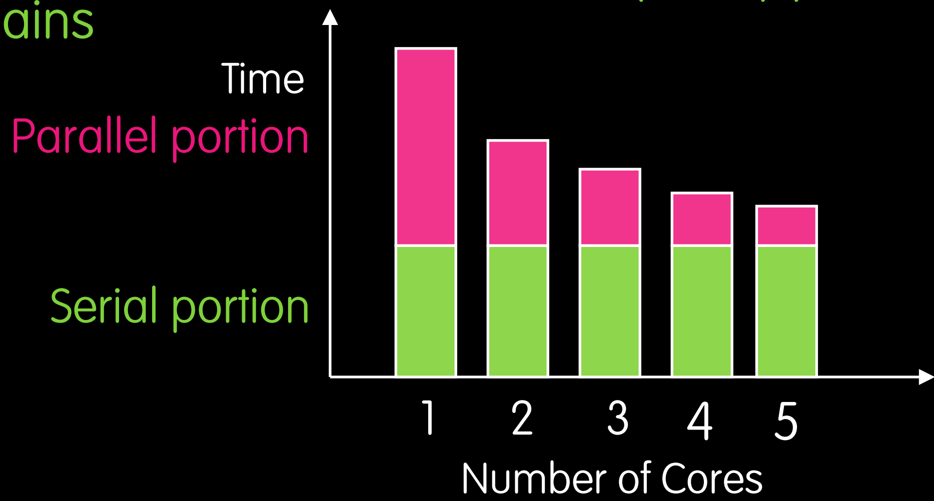

- s is serial fraction of program, P is # of cores (was processors)
- **Amdahl's law:**

Speedup(P) = Time(1) / Time(P)

$\leq 1 / (s + [(1-s) / P)]$, and as $P \rightarrow \infty$

$\leq 1 / s$

- Even if the parallel portion of your application speeds up perfectly, your performance may be limited by the sequential portion

Speedup Issues : Overhead

- **Even assuming no sequential portion, there's…**
  - Time to think how to divide the problem up
  - Time to hand out small "work units" to workers
  - All workers may not work equally fast
  - Some workers may fail
  - There may be contention for shared resources
  - Workers could overwriting each others' answers
  - You may have to wait until the last worker returns to proceed (the slowest / weakest link problem)
  - Time to put the data back together in a way that looks as if it were done by one

Life in a multi-core world…

- **This "sea change" to multi-core parallelism means that the computing community has to rethink:**
  - a) Languages
  - b) Architectures
  - c) Algorithms
  - d) Data Structures
  - e) All of the above

en.wikipedia.org/wiki/Concurrent_computing

But parallel programming is hard!

- What if two people were calling withdraw at the same time?
  - E.g., balance=100 and two withdraw 75 each
  - Can anyone see what the problem *could* be?
  - Called a race condition
- In most languages, this is a problem.
  - In Scratch, the system doesn't let two of these run at once.

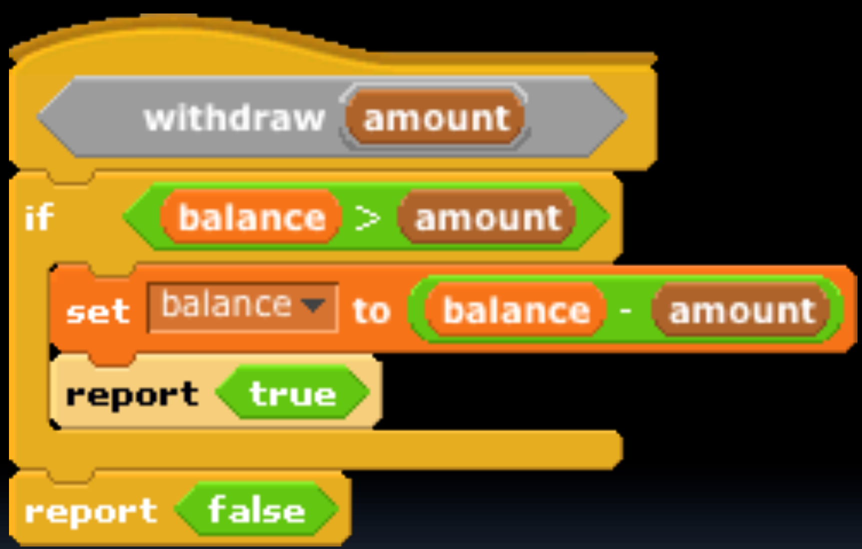

Chun, Summer 2012

en.wikipedia.org/wiki/Deadlock

Another concurrency problem … deadlock!

- Two people need to draw a graph but there is only one pencil and one ruler.
  - One grabs the pencil
  - One grabs the ruler
  - Neither release what they hold, waiting for the other to release

- Livelock also possible
  - Movement, no progress

Summary

- A sea change in computing because of inability to cool CPUs means we're now in multi-core world
- Lots of potential for innovation by computing professionals, but challenges persist