Mommy, Daddy, Why is My New Laptop So Dog Slow?!

or

Recent Trends in Systems Research

Erez Zadok

Stony Brook University

Computer Science Department

`http://www.cs.sunysb.edu/~ezk`

Outline

- Problem motivation
- How computers work today
- Some history
- What can you do?
- What are computer scientists doing?
- The future

Motivation

- Computers “become” slower over time
- Money spent ineffectively
- Power consumption
- E-waste

How a Computer System Works

0. User clicks on "paper.doc"

Word Processor

Short term Memory (fast)

Main Memory

4.Present result to user

The brain (fastest)

3. Copy part of file into CPU

2. Copy file into memory

Central Processing Unit (CPU)

Long term Memory (slow)

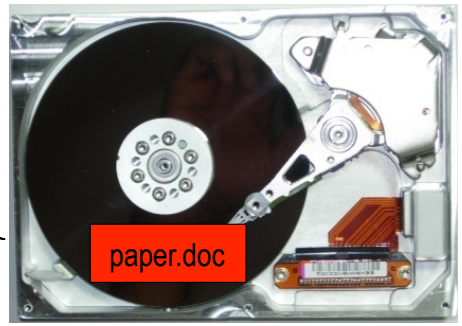

Hard Disk

1. Find file on disk

Show-n-Tell: CPU

TOP VIEW

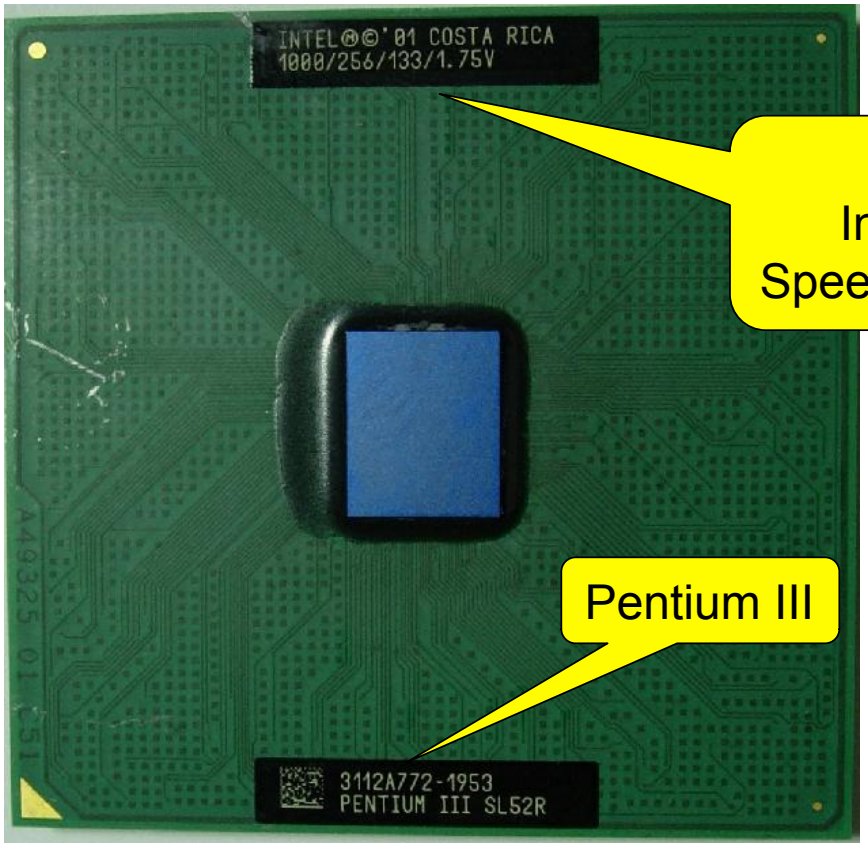

BOTTOM VIEW

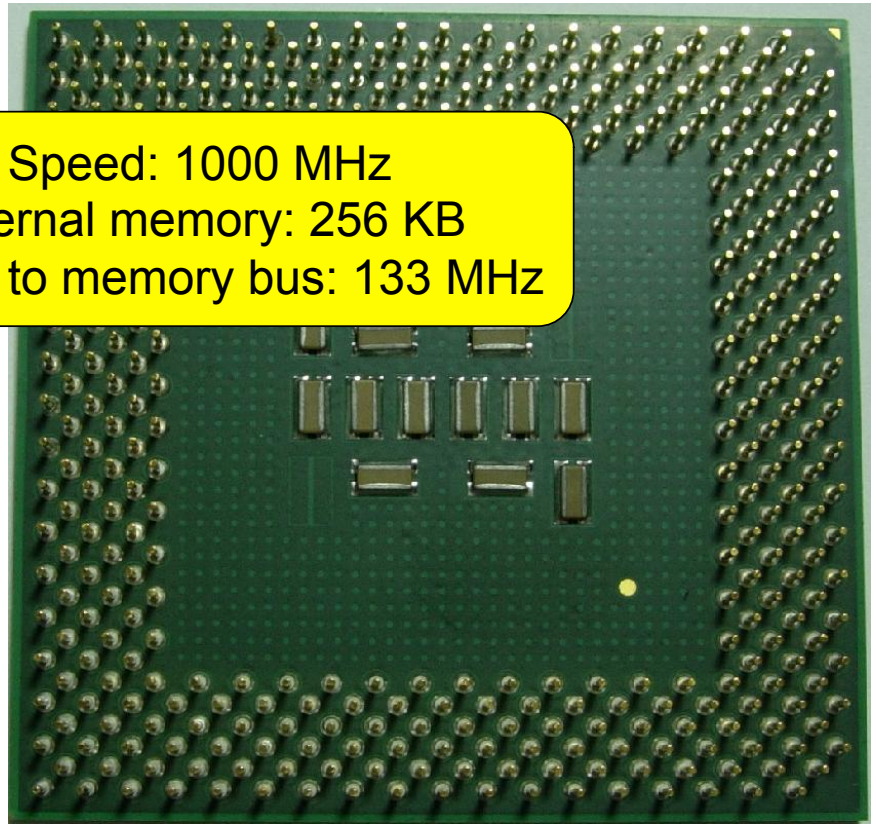

1. **CPU speeds are in the nano-seconds (billions of operations per second)**
2. **CPU's internal speed is 7.5 times faster than its access to main memory.**
3. **Typical CPU's internal memory: 1MB-512MB**

Show-n-Tell: CPU Heat Sink

Note: modern faster CPUs require an ever larger CPU fan

Show-n-Tell: Main Memory

Common name: Dynamic Random Access Memory (DRAM)

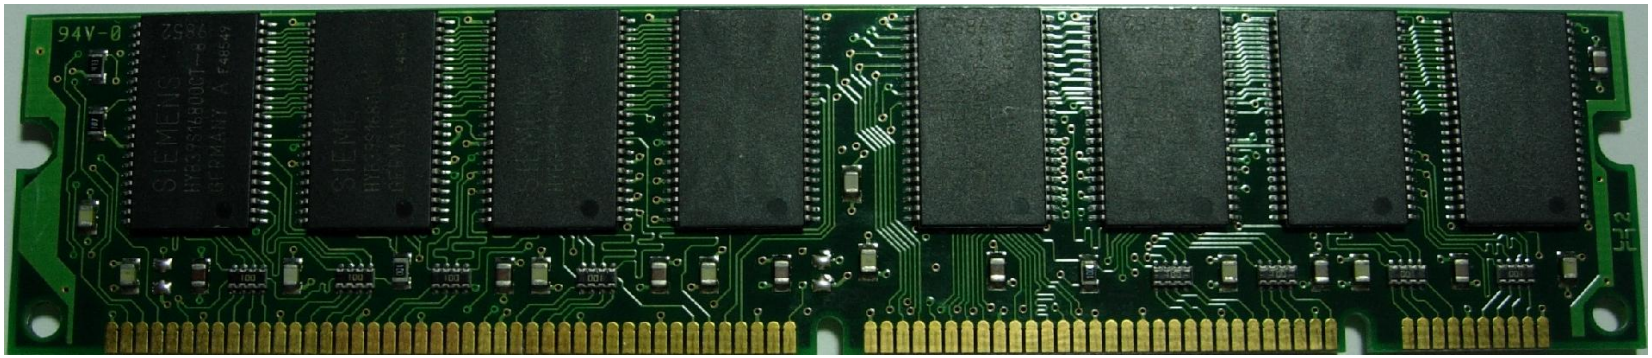

Typical memory sizes today: 2-16 GB (Giga Bytes), or billion characters

Typical memory speeds: micro-seconds (millions of operations per second)

Show-n-Tell: Flash Memory

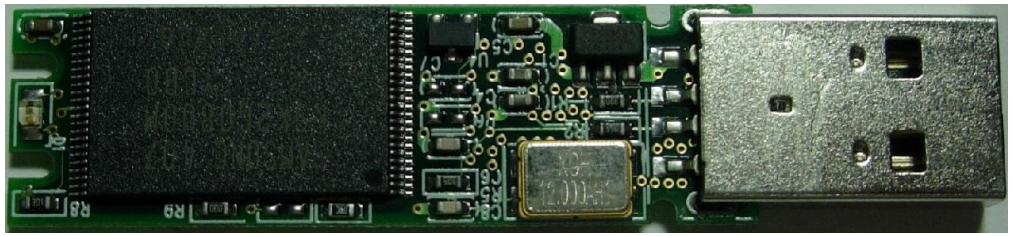

USB port

Discuss later…

Show-n-Tell: Hard Disk

Disk platter rotates
(4-18 KRPM)

Disk head moves
From side to side
(milli-seconds)

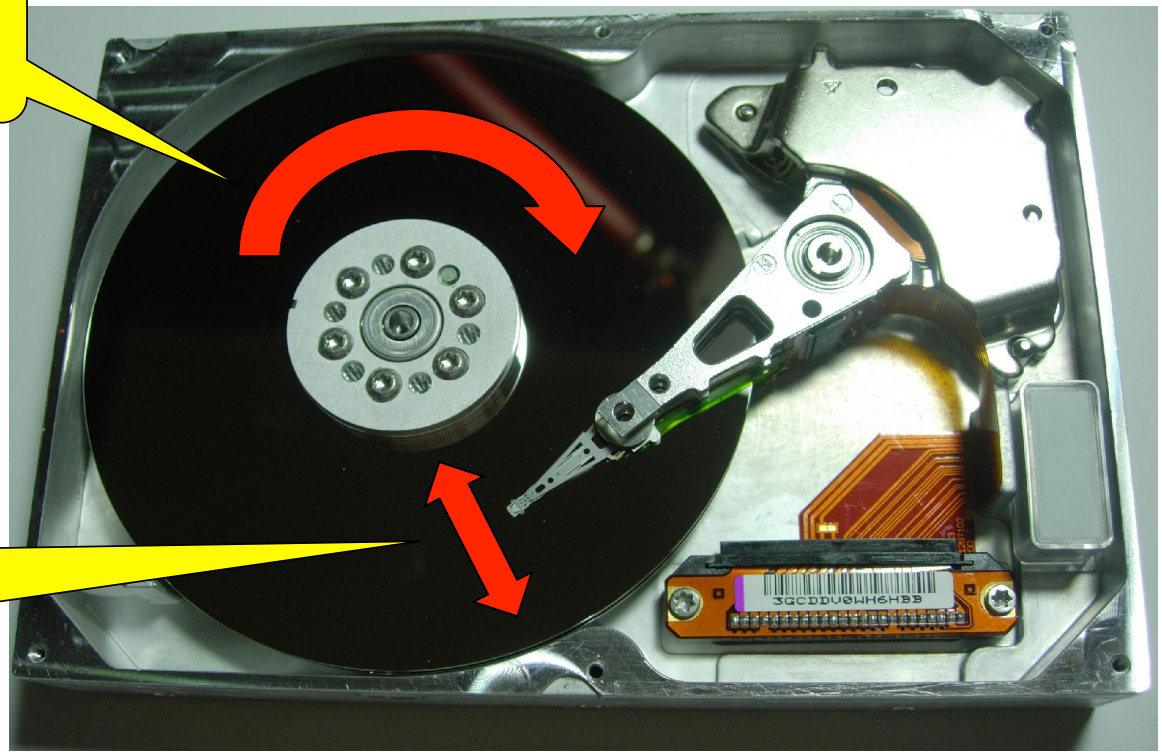

Typical disk sizes today: 250–4000 GB (Giga Bytes), or billion characters

Typical disk speeds: milli-seconds (thousands of operations per second)

The Storage Hierarchy Pyramid

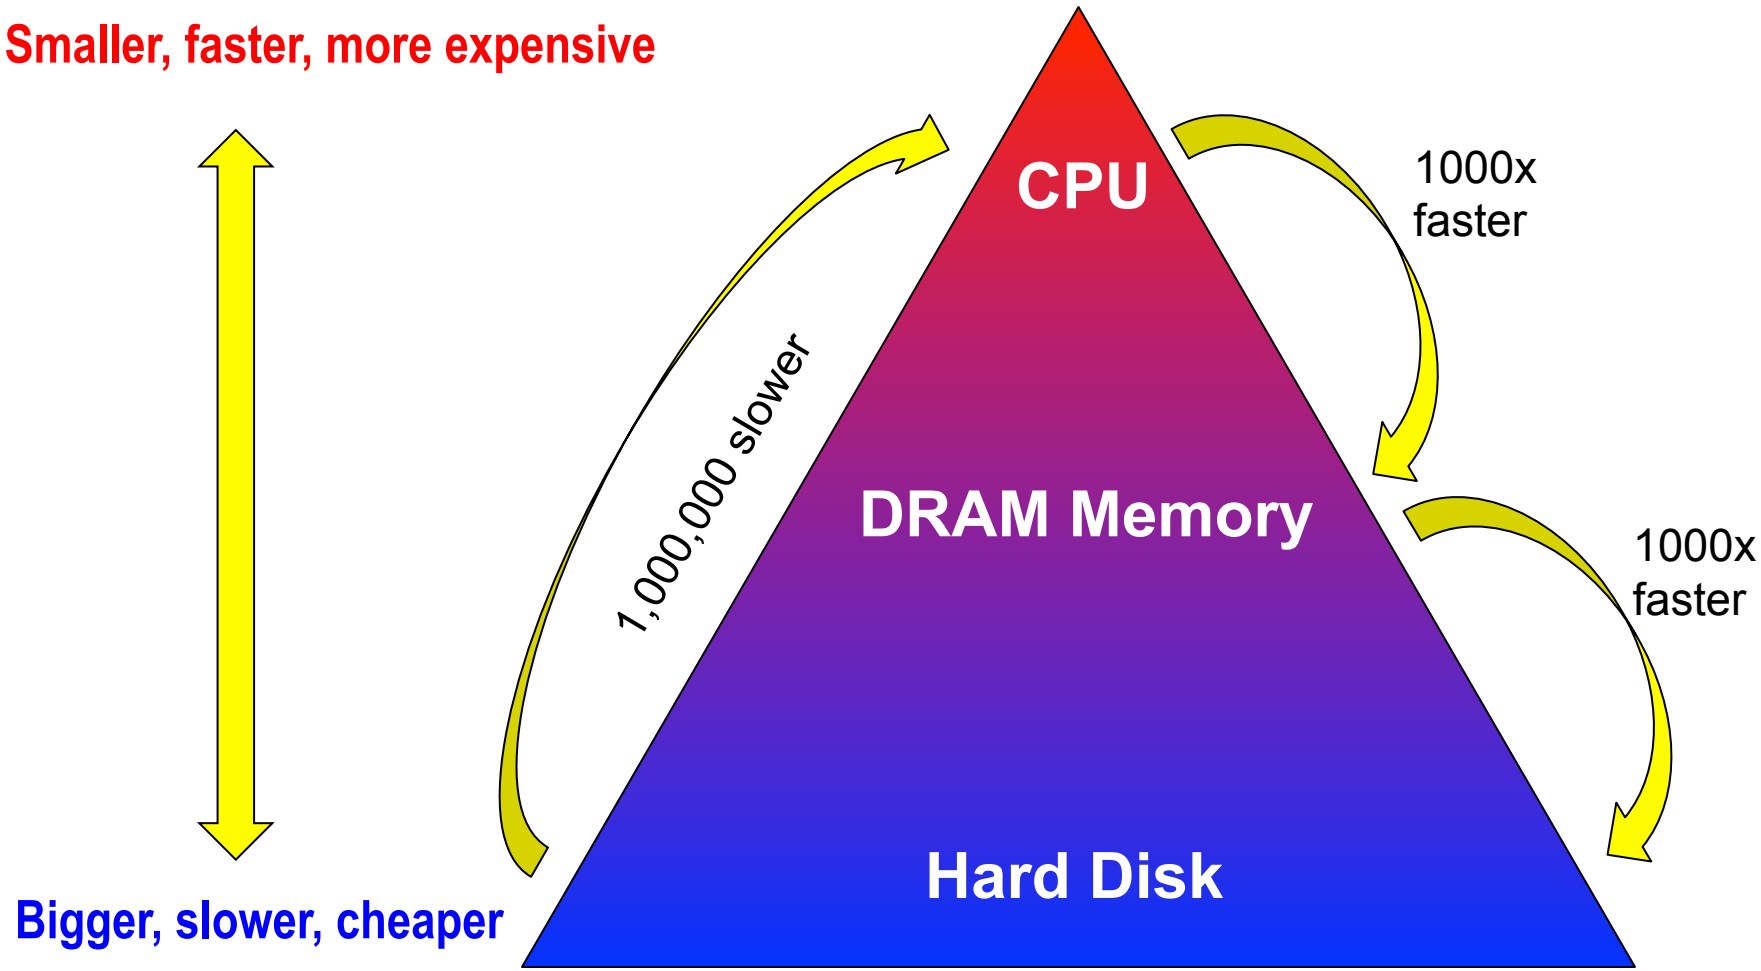

Moore's Law

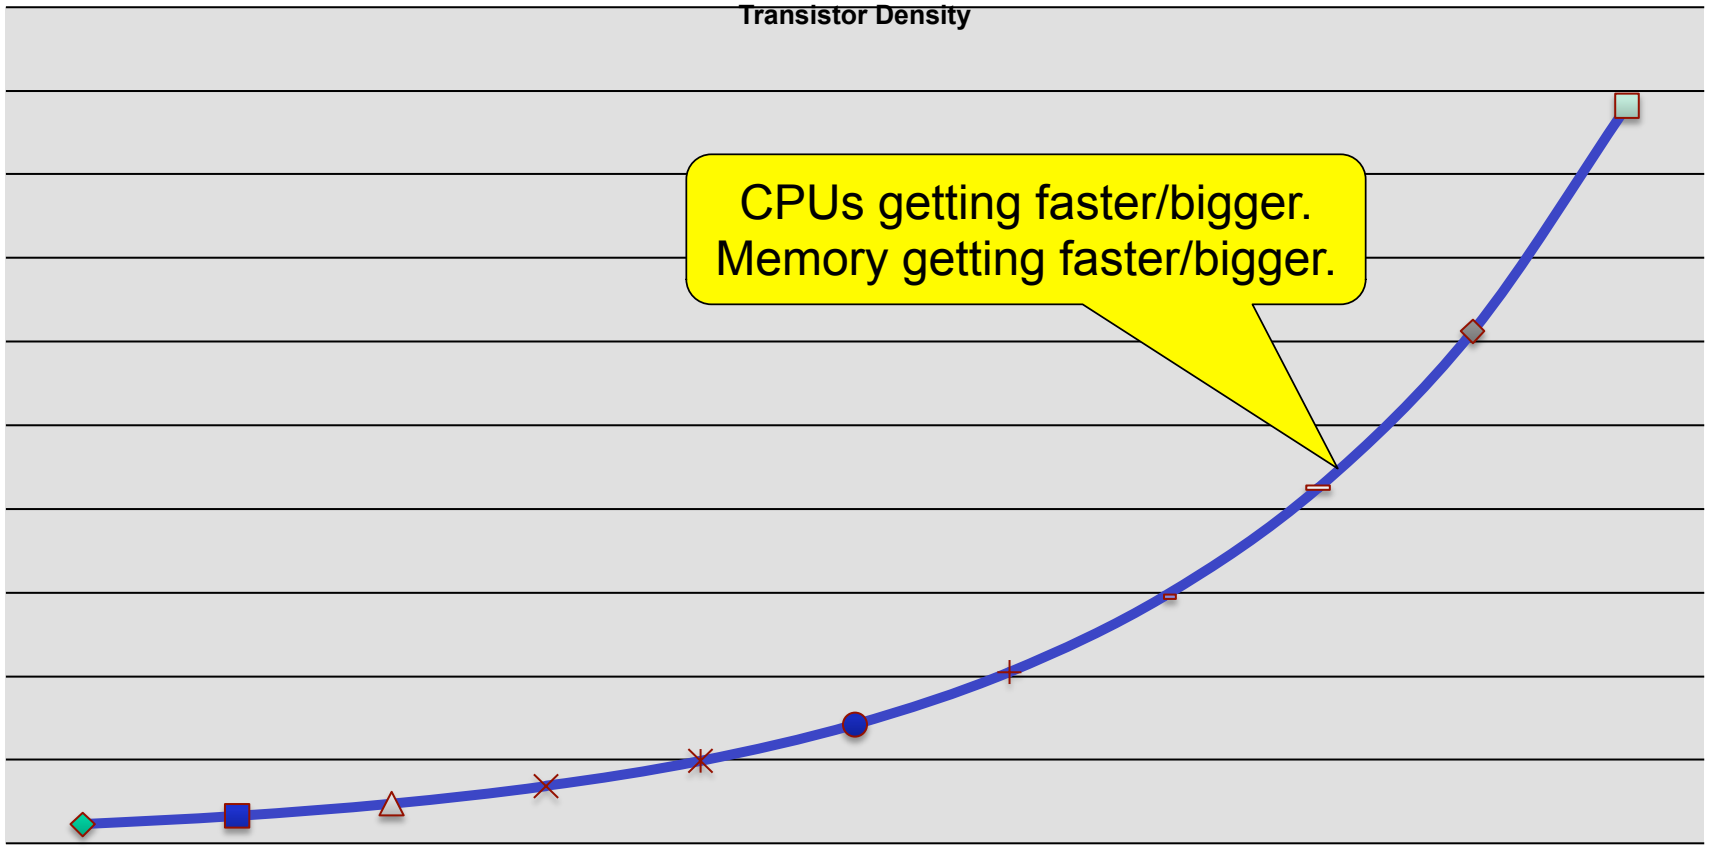

"Transistor density doubles every 18–24 months" (c. 1965)

Kryder's Law

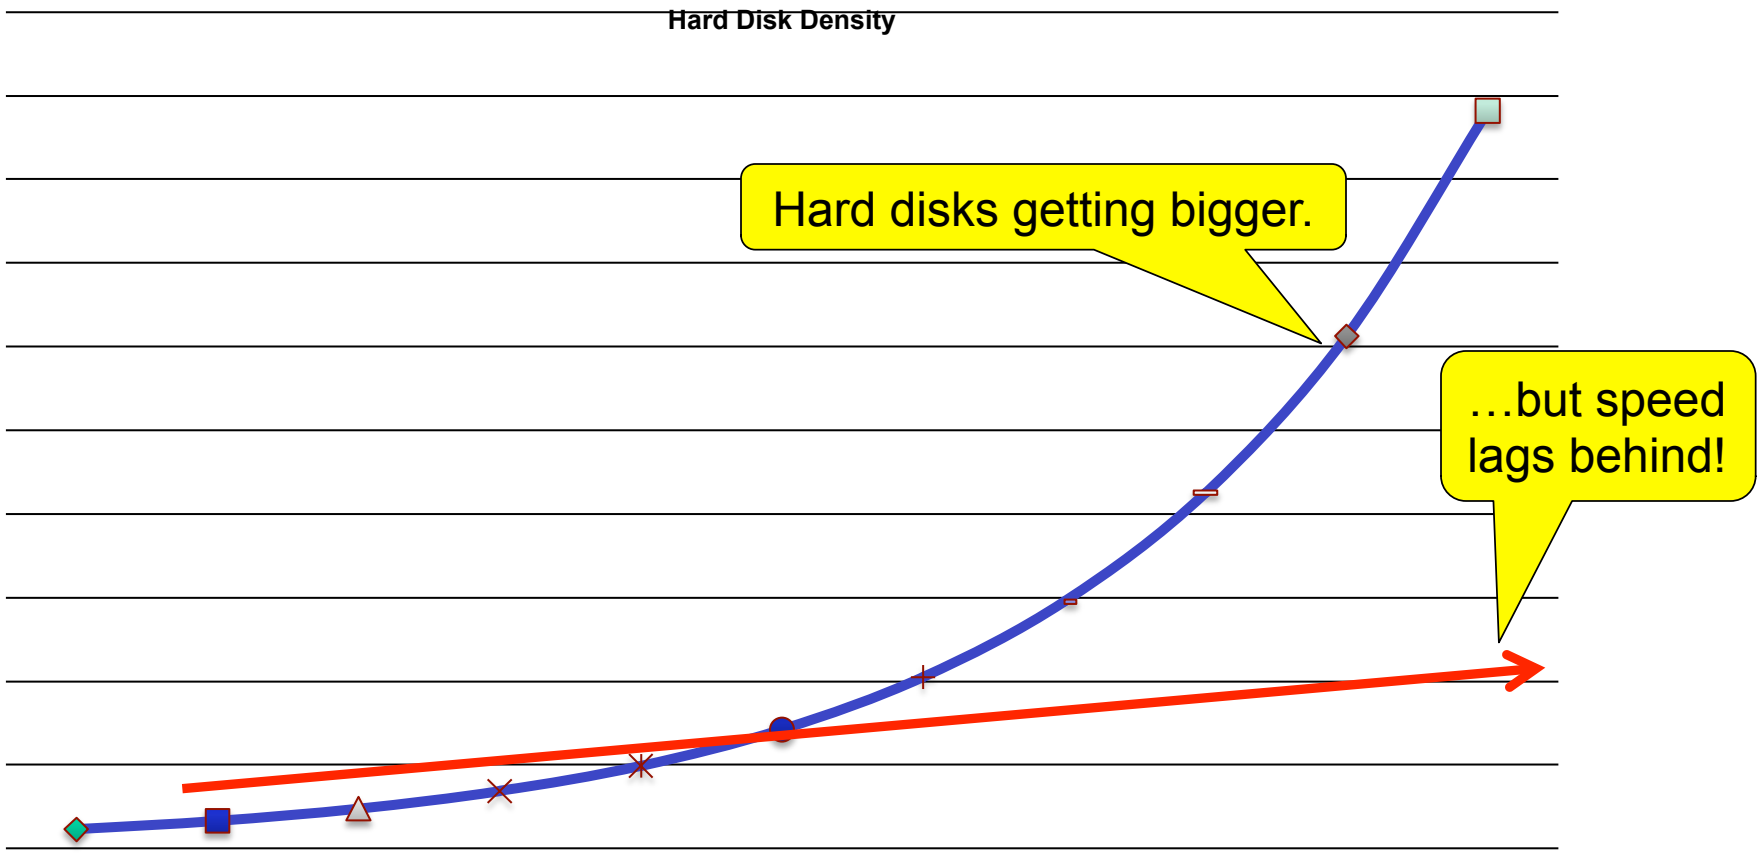

"Magnetic storage density doubles annually" (c. 1995)

Moore's vs. Kryder's Law

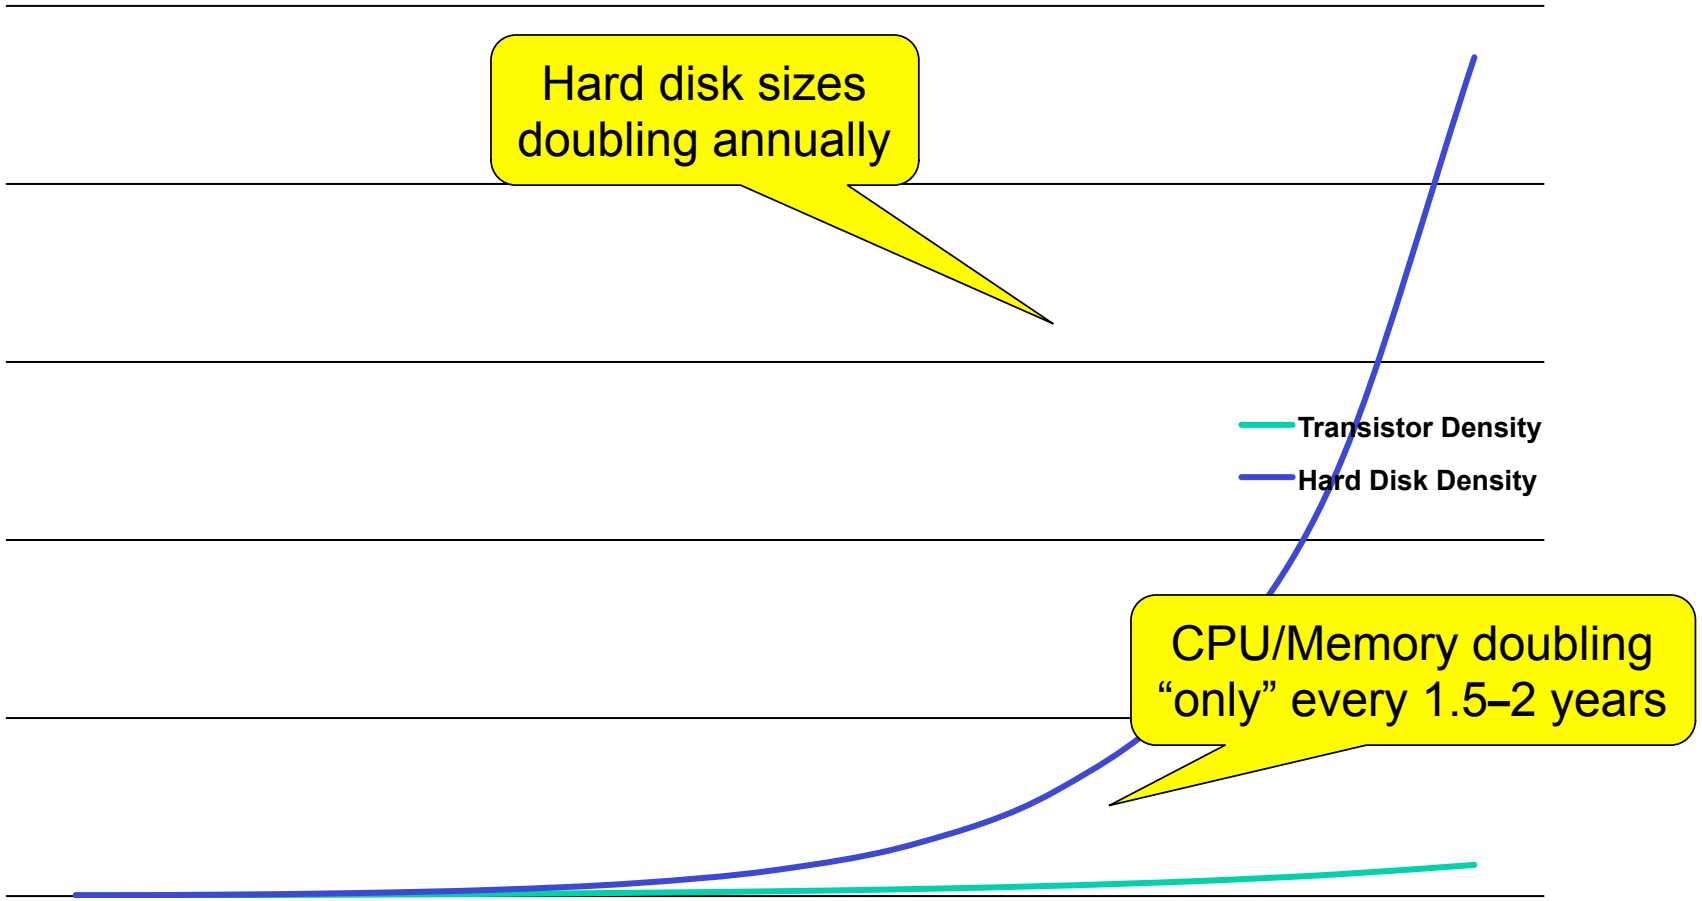

How Computers REALLY Work (1)

Scenario 1: CPU and memory are large and fast enough.

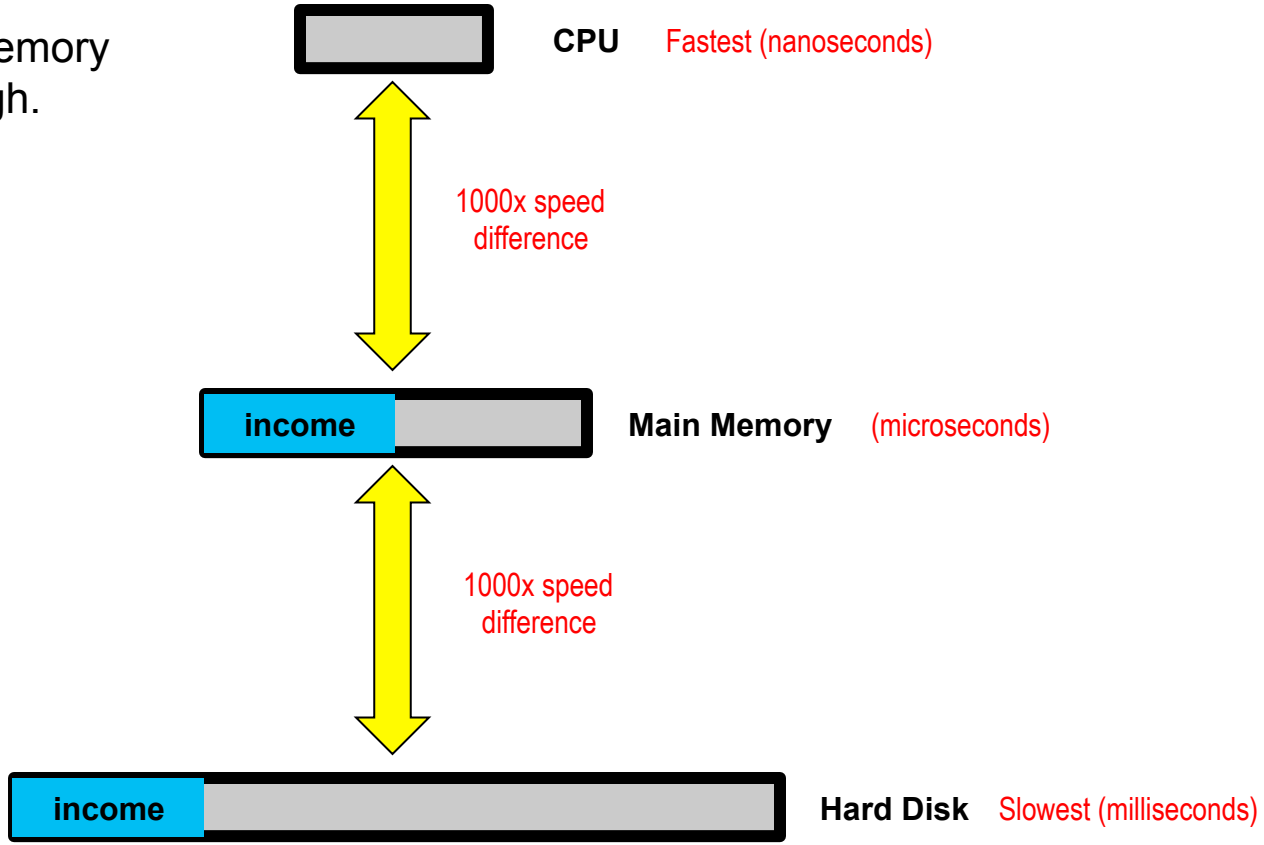

How Computers REALLY Work (2)

Scenario 2: CPU **too** small.
Memory is large enough.

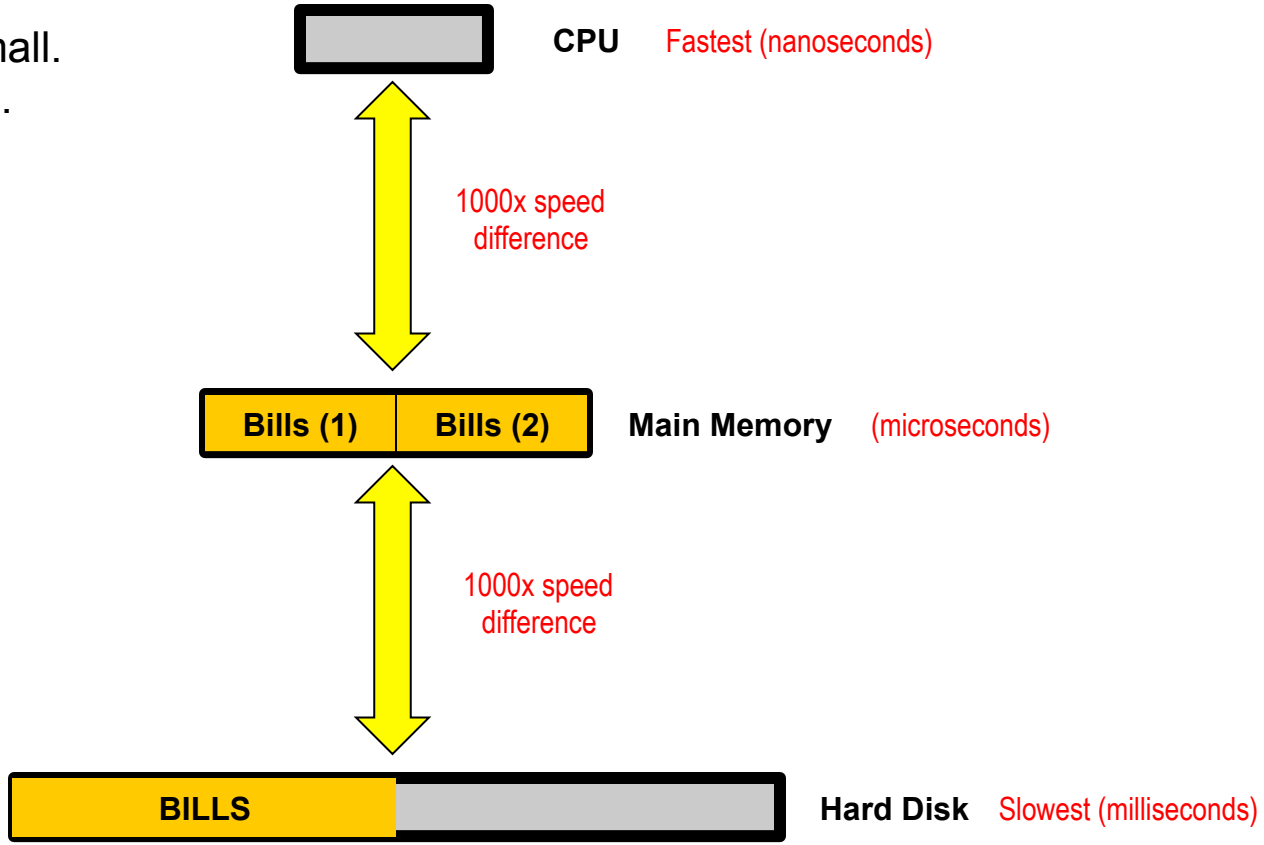

How Computers REALLY Work (3)

Scenario 3: CPU and memory are **BOTH** too small.

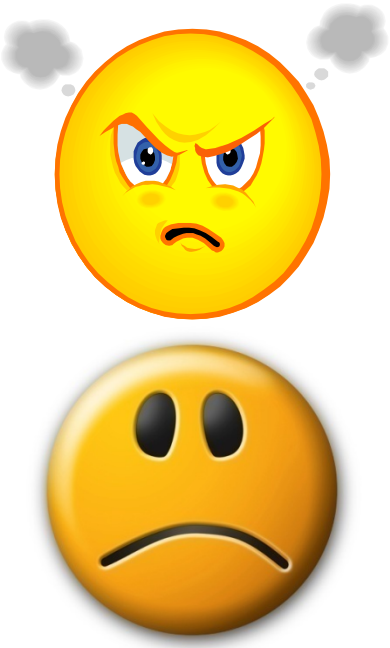

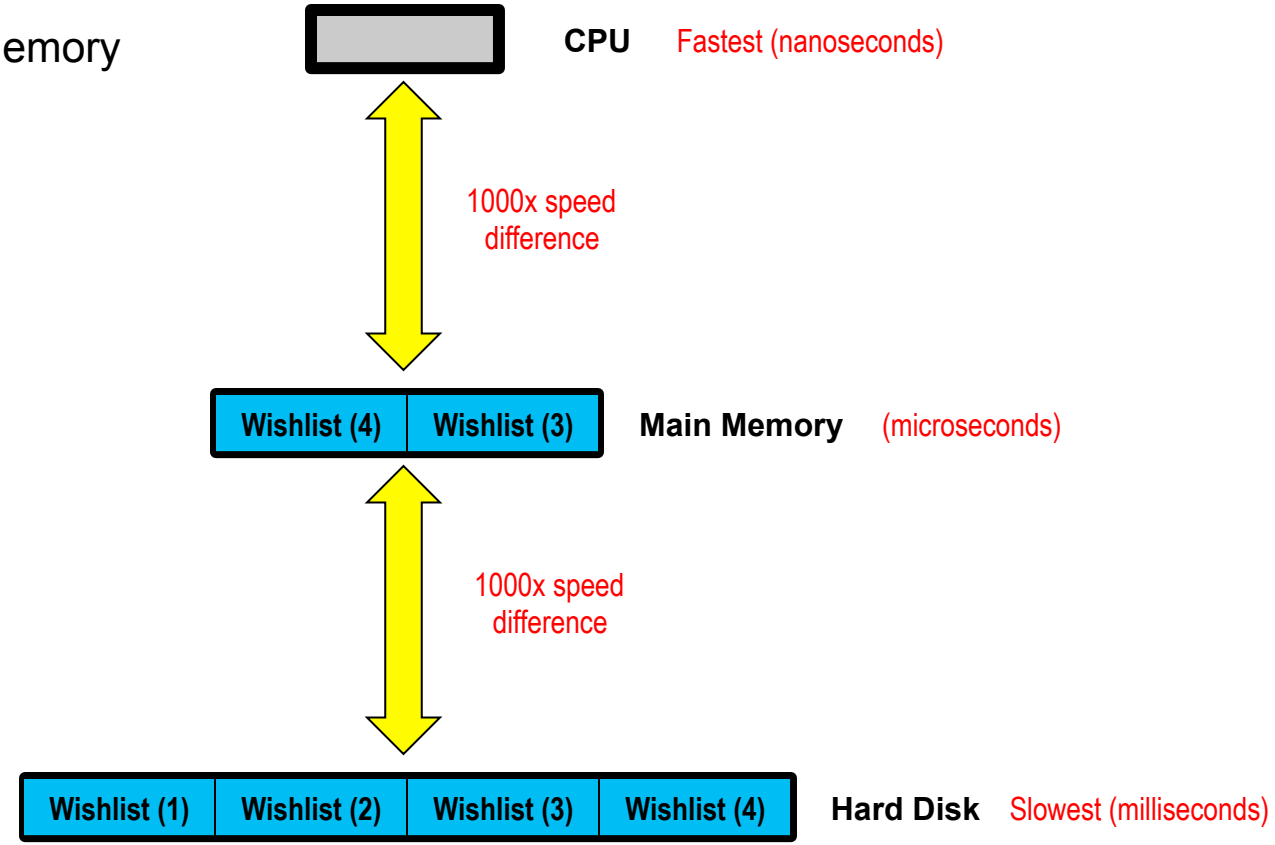

The Perfect Storm

- Disk sizes getting larger the fastest
- CPU/Memory getting larger/faster
  - But lagging behind disk-size growth rates
- Disk speeds lagging far behind the rest
- Software getting larger
  - Fills disks to capacity
  - Software updates & system service packs
- E.g., anti-virus/spyware scanners
  - How long to scan a 4TB disk?

Buying a New Computer (1)

- Given a fixed budget, spend your $$$ on the following, by importance:

1. More main memory (4-8 GB+)
   - Try to keep a free memory upgrade slot
2. Larger CPU caches
3. Faster disks

Caveat: may increase power consumption

Buying a New Computer (2)

- Where to cut back on (if fixed budget):

1. Don't get the fastest CPU (in MHz)
   - E.g., get a 2.8GHz instead of a 3.2GHz
2. You may skip on multi-core CPUs
   - Extra "brains" but share same CPU cache
3. Consider a smaller disk

Recommend: avoid "canned" solutions

Ongoing Comp. Sci. Research

- How to use multi-core CPUs effectively?
  - Trade CPU for reduced I/O?
- How to reduce the need to go to disk?
  - Compressing data (spare CPU cores)
  - Prediction/AI algorithms, prefetch disk data
- Use new storage devices
  - E.g., FLASH memory

Show-n-Tell: Flash Memory

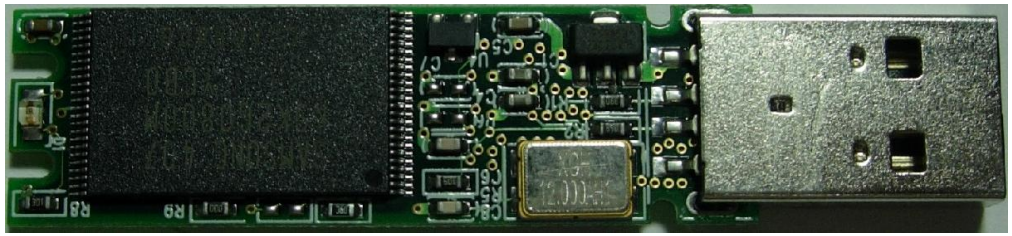

USB port

Major advantage: persistent (non-volatile) memory

Typical flash memory sizes today: 4-256GB (Giga Bytes), or billion characters

Typical flash speeds: between DRAM and hard disks

Lifetime: only millions of writes

The Future of FLASH Memory

- Intermediate cache between main memory and hard disks
- Reliability
  - Distribute writes evenly across device
- Staging device inside hard disks
  - Hybrid disks
- Replace spinning disks
  - Solid State Disks (SSDs)
    - MacBook Air (2008)

Worldwide Cost to Power and Cool Server Installed Base, 1996–2010

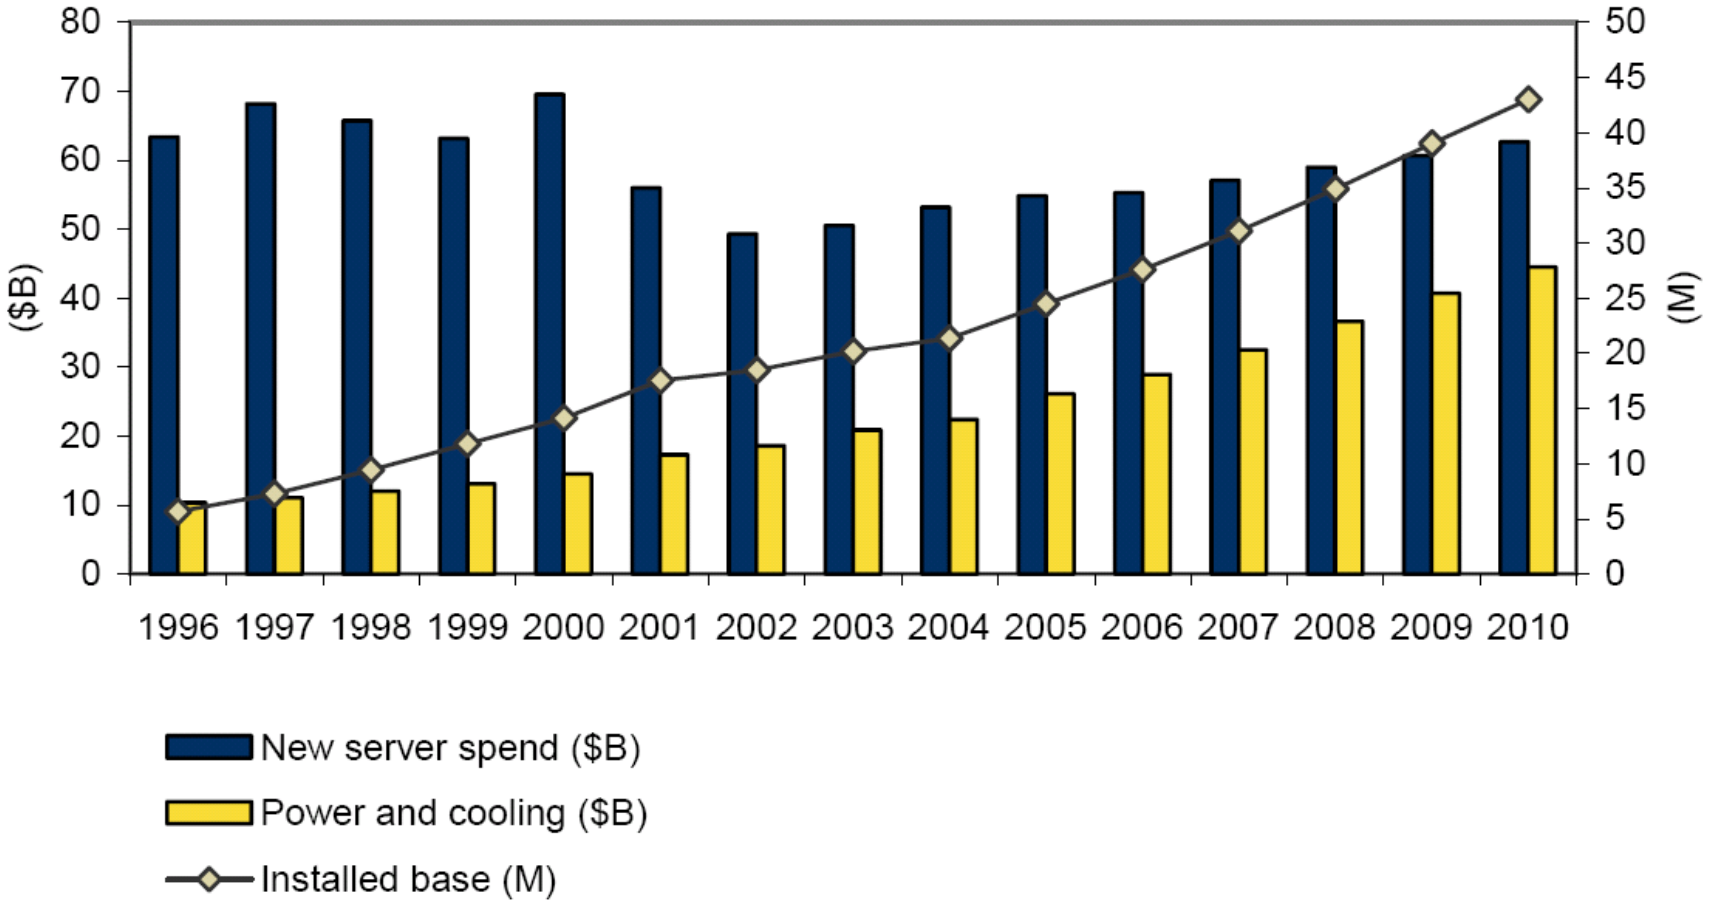

Source: IDC, 2007

Computer Power Use

- Energy savings trails off everything
  - Small percentage improvements annually
- Energy consumption secondary design consideration:
  - Idle laptop: 20–30 Watts
  - Idle workstation: 60–80 Watts
  - Idle server: 200+ Watts

“Green” Research

- Turn off components automatically
  - Without annoying users
- Better hardware components
  - Without slowing down the system
- Use FLASH to store data
  - Lower costs, improve reliability
- Lightweight software
  - Without reduced functionality

Mommy, Daddy, Why is My New Laptop So Dog Slow?!

or

Recent Trends in Systems Research

Q&A

Erez Zadok

Stony Brook University

Computer Science Department

http://www.cs.sunysb.edu/~ezk

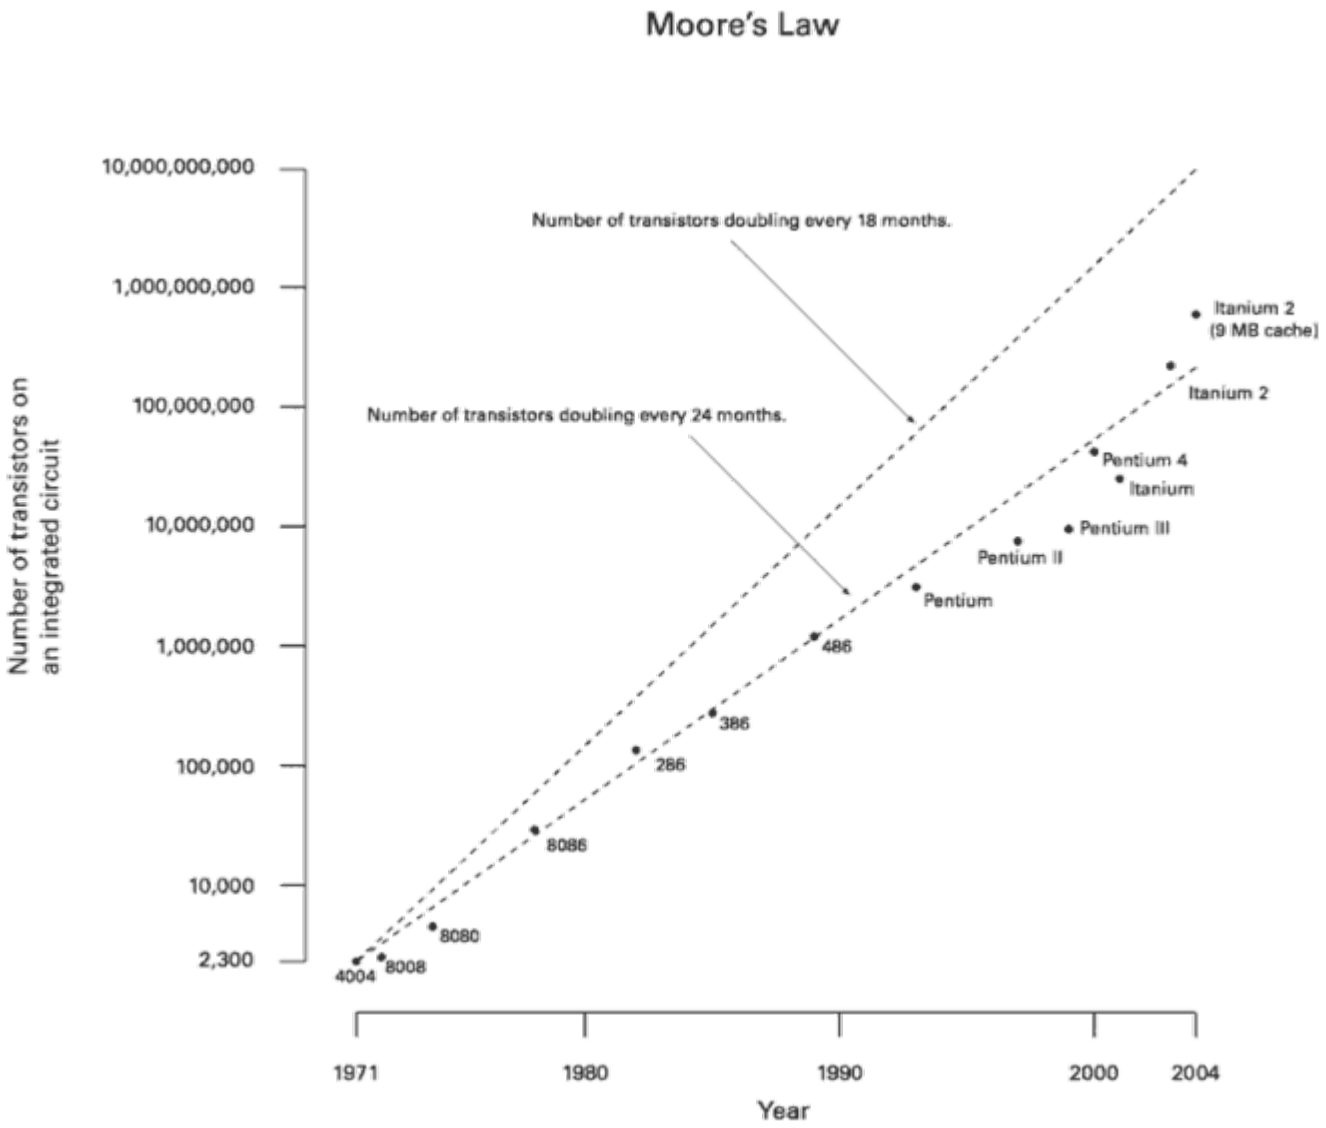
Moore's Law
Number of transistors on an integrated circuit
10,000,000,000
1,000,000,000
100,000,000
10,000,000
1,000,000
100,000
10,000
2,300
Number of transistors doubling every 18 months.
Number of transistors doubling every 24 months.
Itanium 2 (9 MB cache)
Itanium 2
Pentium 4
Itanium
Pentium III
Pentium II
Pentium
486
386
286
8086
8080
8008
4004
1971
1980
1990
2000
2004
Year

Today's Storage Hierarchy Pyramid