Technology and Business Driven Challenges

Basil Smith
IBM T. J. Watson Research Center
Hawthorne, NY 10532

Accelerating technology

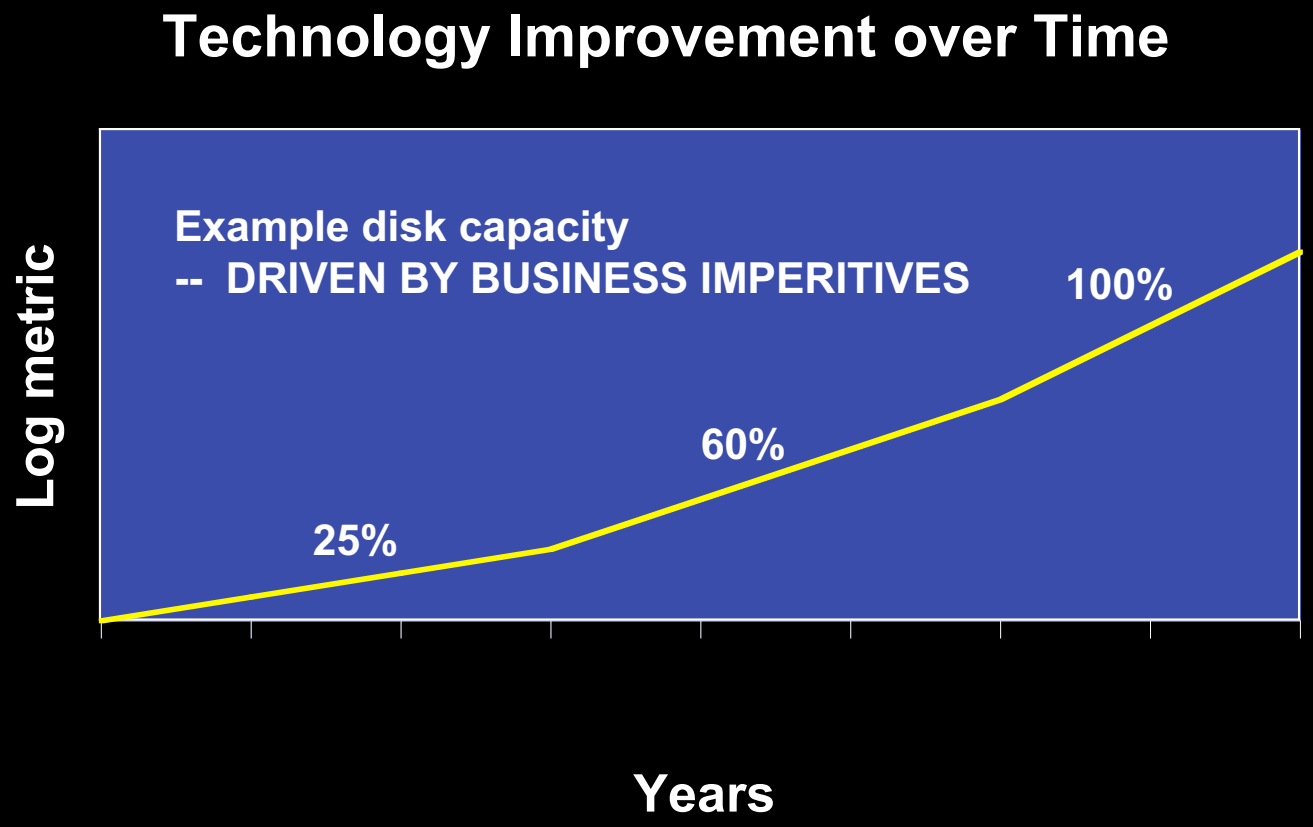

Accelerating technology – other examples

- Semiconductor Technology
  - 3 years / generation to 2 years / generation
- Processor performance
  - 25% CGR to 60% CGR (system throughput)
- Optical link bandwidth
  - Now growing faster than 100% CGR (alas telcos faster than demand)

Becoming Technology Challenged
(i.e., things are falling apart)

- Circuit Technology
  - Supply voltage (noise margins) 5V – 0.4V
  - Operation with Gausian noise upsets
  - Deterministic to statistical design methodologies
  - Thinner insulation layers and high leakage
  - New materials (CU and high K)
  - Smaller features with SEU's in logic
  - o o o
- Disk Technology (RAID 5 is now inadequate)
  - Tiny head flight distances
  - Increasing areal (bit and track) densities
  - Approaching paramagnetic limits
  - Smaller and less robust mechanical designs
  - Acoustical cross talk
  - o o o

Base Semiconductor Technology

Circuit Failure Rate

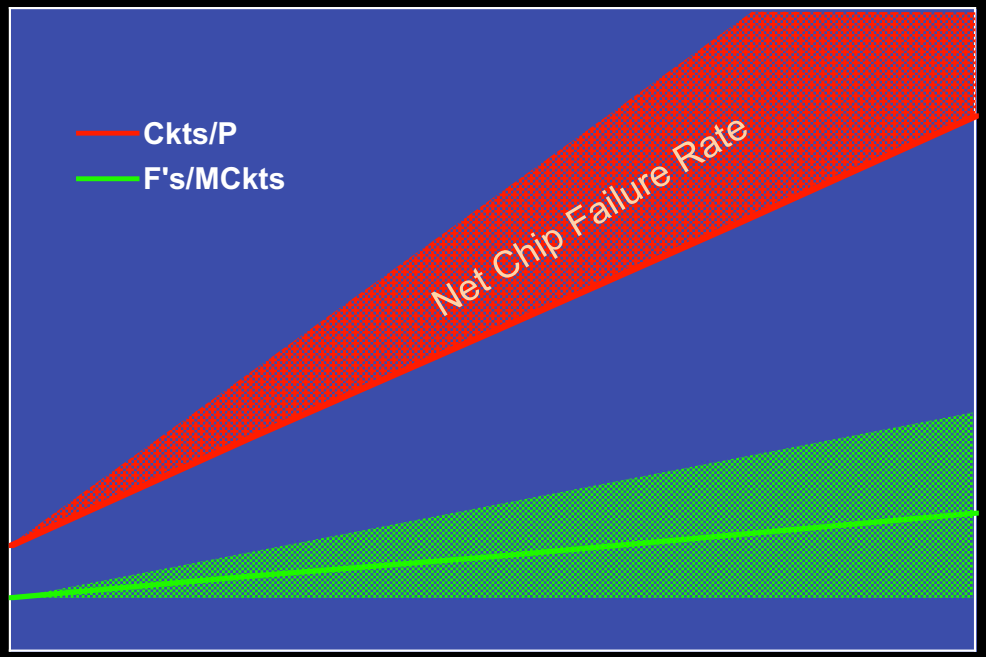

We are ignoring “learning curve” fundamentals

- “Hard” technology is raising and stretching learning curve
  - Things start out worse
  - Things get better more slowly
- Yet we are shortening technology cycles
- And introducing poorly understood technology faster

Ignoring the learning curve

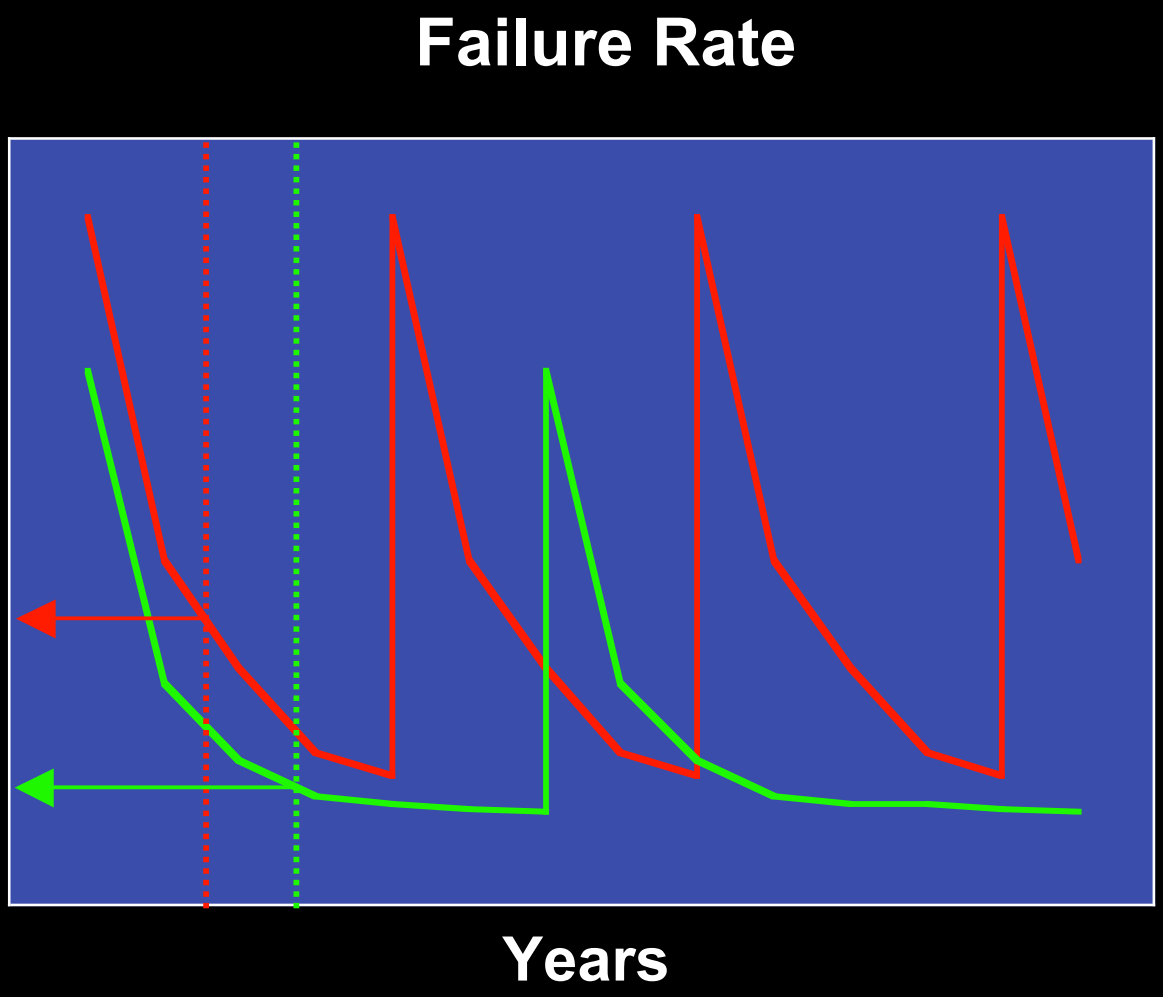

But we see a strong marketplace bias toward "CRAP" (Commodity Reliability And Practices)

- Google model – Reliable Systems from CRAP
  - Massively parallel, with resulting massive redundancy
  - Simple fault detection mechanisms (time outs)
  - Disposable logic/nodes
  - Demonstrated preference to trade reliability for other metrics
    - $'s/MB for disk vs. reliability
    - $'s for cooling vs. reliability
  - Still things broken things have to be fixed
    - Cost of repair deceptively high
    - See some signs of backing off
      - Raw boards > 1U Servers > 2U Servers

But we see a strong marketplace bias toward “CRAP” (Commodity Reliability And Practices)

- Desktop Market
  - Performance over reason
  - Arguably for most of market
    - Now performance insensitive
    - But still investment in performance vs. reliability
      (i.e., why don’t we have extremely reliable 1 GHz desktops)
  - ~Nil market for higher reliability drives
    - “SCSI” reliability vs. ATA reliability
  - Unwillingness to pay for even parity in memory
    - Given caches this could even be word or line parity (2% overhead)

So what should we do?

- Reliability cannot get much worse without becoming a maintenance/warranty headache
  - Much of the maintenance cost is labor costs of repair ops
- Fail in place could solve some maintenance / reliability problems
  - Common for memory – spare chips, chip kill ECC, …
  - Fail in place is out of the question without improving reliability
- Simple system level redundancy is going to have to work
  - Other options seem uneconomic
  - Need to be able to detect and correct multiple simultaneous fails
- Really good dependability will require better error detection
  - Undetected errors and data corruption are simply too likely with CRAP to achieve really dependable systems in many situations.
- Design mistakes all too likely, and source of common mode failures
  - A message straight from the learning curve
  - Fail in place will not work when its likely everything will fail