

# CS184b: Computer Architecture (Abstractions and Optimizations)

Day 17: May 9, 2005  
Defect and Fault Tolerance



Caltech CS184 Spring2005 -- DeHon

## Today

- Defect and Fault Tolerance
  - Problem
  - Defect Tolerance
  - Fault Tolerance

2

Caltech CS184 Spring2005 -- DeHon

## Motivation: Probabilities

- Given:
  - N objects
  - P yield probability
- What's the probability for yield of composite system of N items?
  - Assume iid faults
  - $P(N \text{ items good}) = P^N$

3

Caltech CS184 Spring2005 -- DeHon

## Probabilities

- $P(N \text{ items good}) = P^N$
- $N=10^6, P=0.999999$
- $P(\text{all good}) \approx 0.37$
  
- $N=10^7, P=0.999999$
- $P(\text{all good}) \approx 0.000045$

4

Caltech CS184 Spring2005 -- DeHon

## Simple Implications

- As N gets large
  - must either increase reliability
  - ...or start tolerating failures
- N
  - memory bits
  - disk sectors
  - wires
  - transmitted data bits
  - processors
  - transistors
  - molecules
- As devices get smaller, failure rates increase
- chemists think  $P=0.95$  is good

5

Caltech CS184 Spring2005 -- DeHon

## Defining Problems

6

Caltech CS184 Spring2005 -- DeHon

## Three “problems”

- Manufacturing imperfection
  - Shorts, breaks
  - wire/node X shorted to power, ground, another node
  - Doping/resistance variation too high
- Parameters vary over time
  - Electromigration
  - Resistance increases
- Incorrect operation
  - node X value flips
    - crosstalk
    - alpha particle
    - bad timing

Caltech CS184 Spring2005 -- DeHon

7

## Defects

- Shorts example of defect
- Persistent problem
  - reliably manifests
- Occurs before computation
- Can test for at fabrication / boot time and then avoid
- (1<sup>st</sup> half of lecture)

Caltech CS184 Spring2005 -- DeHon

8

## Faults

- Alpha particle bit flips is an example of a fault
- Fault occurs dynamically during execution
- At any point in time, can fail
  - (produce the wrong result)
- (2<sup>nd</sup> half of lecture)

Caltech CS184 Spring2005 -- DeHon

9

## Lifetime Variation

- Starts out fine
- Over time changes
  - *E.g.* resistance increases until out of spec.
- Persistent
  - So can use defect techniques to avoid
- But, onset is dynamic
  - Must use fault detection techniques to recognize?

Caltech CS184 Spring2005 -- DeHon

10

## In a Nut-shell...

Sherkhar Bokar  
Intel Fellow  
Micro37 (Dec.2004)



100 BT integration capacity  
20 BT unusable (variations)  
10 BT will fail over time  
Intermittent failures

Yet, deliver high performance in the power & cost envelope

44

## Defect Rate

- Device with  $10^{11}$  elements (100BT)
- 3 year lifetime =  $10^8$  seconds
- Accumulating up to 10% defects
- $10^{10}$  defects in  $10^8$  seconds
  - 1 defect every 10ms
- At 10GHz operation:
  - One new defect every  $10^8$  cycles
  - $P_{\text{newdefect}} = 10^{-19}$

Caltech CS184 Spring2005 -- DeHon

12

## First Step to Recover

Admit you have a problem  
(observe that there is a failure)

## Detection

- Determine if something wrong?
  - Some things easy
    - ...won't start
  - Others tricky
    - ...one **and** gate computes False & True → True
- Observability
  - can see effect of problem
  - some way of telling if defect/fault present

## Detection

- Coding
  - space of legal values < space of all values
  - should only see legal
  - e.g. parity, ECC (Error Correcting Codes)
- Explicit test (defects, recurring faults)
  - ATPG = Automatic Test Pattern Generation
  - Signature/BIST=Built-In Self-Test
  - POST = Power On Self-Test
- Direct/special access
  - test ports, scan paths

## Coping with defects/faults?

- **Key idea:** redundancy
- Detection:
  - Use redundancy to detect error
- Mitigating: use redundant hardware
  - Use spare elements in place of faulty elements (defects)
  - Compute multiple times so can discard faulty result (faults)
  - **Exploit Law-of-Large Numbers**

## Defect Tolerance

## Two Models

- Disk Drives
- Memory Chips

## Disk Drives

- Expose defects to software
  - software model expects faults
    - Create table of good (bad) sectors
  - manages by masking out in software
    - (at the OS level)
  - yielded capacity varies

Caltech CS184 Spring2005 -- DeHon

19

## Memory Chips

- Provide model in **hardware** of perfect chip
- Model of perfect memory at capacity X
- Use redundancy in hardware to provide perfect model
- Yielded capacity fixed
  - discard part if not achieve

Caltech CS184 Spring2005 -- DeHon

20

## Example: Memory

- Correct memory:
  - N slots
  - each slot reliably stores last value written
- Millions, billions, etc. of bits...
  - have to get them all right?

Caltech CS184 Spring2005 -- DeHon

21

## Memory Defect Tolerance

- Idea:
  - few bits may fail
  - provide more raw bits
  - configure so yield what looks like a perfect memory of specified size

Caltech CS184 Spring2005 -- DeHon

22

## Memory Techniques

- Row Redundancy
- Column Redundancy
- Block Redundancy

Caltech CS184 Spring2005 -- DeHon


23

## Row Redundancy

- Provide extra rows
- Mask faults by avoiding bad rows
- Trick:
  - have address decoder substitute spare rows in for faulty rows
  - use fuses to program

Caltech CS184 Spring2005 -- DeHon

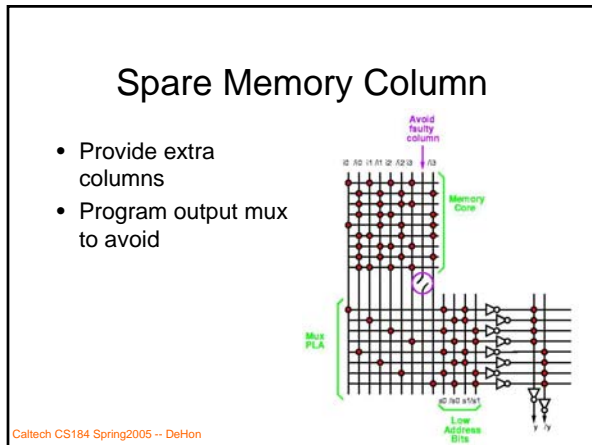
24



### Column Redundancy

- Provide extra columns
- Program decoder/mux to use subset of columns


26



### Block Redundancy

- Substitute out entire block
  - e.g. memory subarray
    - include 5 blocks
      - only need 4 to yield perfect
    - (N+1 sparing more typical for larger N)

28



### Yield M of N

- $P(M \text{ of } N) = P(\text{yield } N)$ 
  - +  $(N \text{ choose } N-1) P(\text{exactly } N-1)$
  - +  $(N \text{ choose } N-2) P(\text{exactly } N-2) \dots$
  - +  $(N \text{ choose } N-M) P(\text{exactly } N-M) \dots$

[think binomial coefficients]

30

## M of 5 example

- $1 \cdot P^5 + 5 \cdot P^4(1-P)^1 + 10P^3(1-P)^2 + 10P^2(1-P)^3 + 5P^1(1-P)^4 + 1 \cdot (1-P)^5$

- Consider  $P=0.9$

- $1 \cdot P^5$	0.59	M=5	P(sys)=0.59
- $5 \cdot P^4(1-P)^1$	0.33	M=4	P(sys)=0.92
- $10P^3(1-P)^2$	0.07	M=3	P(sys)=0.99
- $10P^2(1-P)^3$	0.008		
- $5P^1(1-P)^4$	0.00045		
- $1 \cdot (1-P)^5$	0.00001		

Can achieve higher system yield than individual components!

Caltech CS184 Spring2005 -- DeHon

31

## Repairable Area

- Not all area in a RAM is repairable
  - memory bits spare-able
  - io, power, ground, control not redundant

Caltech CS184 Spring2005 -- DeHon

32

## Repairable Area


- $P(\text{yield}) = P(\text{non-repair}) \cdot P(\text{repair})$
- $P(\text{non-repair}) = P^N$ 
  - $N \ll N_{\text{total}}$
  - Maybe  $P > P_{\text{repair}}$ 
    - e.g. use coarser feature size
- $P(\text{repair}) \sim P(\text{yield } M \text{ of } N)$

Caltech CS184 Spring2005 -- DeHon

33

## Consider a Crossbar

- Allows me to connect any of N things to each other
- E.g.
  - N processors
  - N memories
  - N/2 processors + N/2 memories




Caltech CS184 Spring2005 -- DeHon

34

## Crossbar Buses and Defects


- Two crossbars
- Wires may fail
- Switches may fail
- Provide more wires
  - Any wire fault avoidable
    - M choose N



Caltech CS184 Spring2005 -- DeHon

## Crossbar Buses and Defects

- Two crossbars
- Wires may fail
- Switches may fail
- Provide more wires
  - Any wire fault avoidable
    - M choose N




Caltech CS184 Spring2005 -- DeHon

36

## Crossbar Buses and Faults

- Two crossbars
- Wires may fail
- Switches may fail
- Provide more wires
  - Any wire fault avoidable
    - M choose N




Caltech CS184 Spring2005 -- DeHon

37

## Crossbar Buses and Faults

- Two crossbars
- Wires may fail
- Switches may fail
- Provide more wires
  - Any wire fault avoidable
    - M choose N
  - Same idea




Caltech CS184 Spring2005 -- DeHon

38

## Simple System


- P Processors
- M Memories
- Wires



Caltech CS184 Spring2005 -- DeHon

## Simple System w/ Spares

- P Processors
- M Memories
- Wires
- Provide spare
  - Processors
  - Memories
  - Wires




Caltech CS184 Spring2005 -- DeHon

40

## Simple System w/ Defects

- P Processors
- M Memories
- Wires
- Provide spare
  - Processors
  - Memories
  - Wires
- ...and defects




Caltech CS184 Spring2005 -- DeHon

41

## Simple System Repaired

- P Processors
- M Memories
- Wires
- Provide spare
  - Processors
  - Memories
  - Wires
- Use crossbar to switch together good processor and memories




Caltech CS184 Spring2005 -- DeHon

42

## In Practice

- Crossbars are inefficient [CS184A]
- Use switching networks with
  - Locality
  - Segmentation
  - CS184A
- ...but basic idea for sparing is the same



43

Caltech CS184 Spring2005 -- DeHon

## Fault Tolerance

44

Caltech CS184 Spring2005 -- DeHon

## Faults

- Bits, processors, wires
  - May fail during operation
- Basic Idea same:
  - Detect failure using redundancy
  - Correct
- Now
  - Must identify and correct **online** with the computation

45

Caltech CS184 Spring2005 -- DeHon


## Simple Memory Example

- **Problem:** bits may lose/change value
  - Alpha particle
  - Molecule spontaneously switches
- **Idea:**
  - Store multiple copies
  - Perform majority vote on result

46

Caltech CS184 Spring2005 -- DeHon

## Redundant Memory




47

Caltech CS184 Spring2005 -- DeHon

## Redundant Memory

- Like M-choose-N
- Only fail if  $>(N-1)/2$  faults
- $P=0.9$
- $P(2 \text{ of } 3)$ 
  - All good:  $(0.9)^3 = 0.729$
  - + Any 2 good:  $3(0.9)^2(0.1)=0.243$
  - = **0.971**



48

Caltech CS184 Spring2005 -- DeHon



## Better: Less Overhead


- Don't have to keep N copies
- Block data into groups
- Add a small number of bits to detect/correct errors

Caltech CS184 Spring2005 -- DeHon

49

## Row/Column Parity


- Think of NxN bit block as array
- Compute row and column parities
  - (total of 2N bits)



Caltech CS184 Spring2005 -- DeHon

## Row/Column Parity


- Think of NxN bit block as array
- Compute row and column parities
  - (total of 2N bits)
- Any single bit error



Caltech CS184 Spring2005 -- DeHon

## Row/Column Parity

- Think of NxN bit block as array
- Compute row and column parities
  - (total of 2N bits)
- Any single bit error
- By recomputing parity
  - Know which one it is
  - Can correct it



Caltech CS184 Spring2005 -- DeHon

## In Use Today

- Conventional DRAM Memory systems
  - Use 72b ECC (Error Correcting Code)
  - On 64b words
  - Correct any single bit error
  - Detect multibit errors
- CD blocks are ECC coded
  - Correct errors in storage/reading
- Learn more about ECC in EE127

Caltech CS184 Spring2005 -- DeHon

53

## Interconnect

- Also uses checksums/ECC
  - Guard against data transmission errors
  - Environmental noise, crosstalk, trouble sampling data at high rates...
- Often just detect error
- Recover by requesting retransmission
  - *E.g.* TCP/IP (Internet Protocols)

Caltech CS184 Spring2005 -- DeHon

54

## Interconnect


- Also guards against whole path failure
- Sender expects acknowledgement
- If no acknowledgement will retransmit
- If have multiple paths
  - ...and select well among them
  - Can route around any fault in interconnect

Caltech CS184 Spring2005 -- DeHon

55

## Interconnect Fault Example


- Send message
- Expect Acknowledgement



Caltech CS184 Spring2005 -- De

## Interconnect Fault Example


- Send message
- Expect Acknowledgement
- If Fail



Caltech CS184 Spring2005 -- DeH

## Interconnect Fault Example


- Send message
- Expect Acknowledgement
- If Fail
  - No ack



Caltech CS184 Spring2005 -- DeH

## Interconnect Fault Example

- If Fail → no ack
  - Retry
  - Preferably with different resource




Caltech CS184 Spring2005 -- DeHon

59

## Interconnect Fault Example

- If Fail → no ack
  - Retry
  - Preferably with different resource




Ack signals success

Caltech CS184 Spring2005 -- DeHon

60

## Transit Multipath


- Butterfly (or Fat-Tree) networks with multiple paths
  - CS184B:Day4



Caltech CS184 Spring2005 -- DeHon

## Multiple Paths


- Provide bandwidth
- Minimize congestion
- Provide **redundancy** to tolerate faults



Caltech CS184 Spring2005 -- DeHon





## Routers May be faulty (links may be faulty)

- Dynamic
  - Corrupt data
  - Misroute
  - Send data nowhere



Caltech CS184 Spring2005 -- DeHon

## Multibutterfly Performance w/ Faults




Caltech CS184 Spring200

64

## Compute Elements

- Simplest thing we can do:
  - Compute redundantly
  - Vote on answer
  - Similar to redundant memory



Caltech CS184 Spring2005

65

## Compute Elements

- Unlike Memory
  - State of computation important
  - Once a processor makes an error
    - All subsequent results may be wrong
- Response
  - “reset” processors which fail vote
  - Go to spare set to replace failing processor

Caltech CS184 Spring2005 -- DeHon

66

## In Use

- NASA Space Shuttle
  - Uses set of 4 voting processors
- Boeing 777
  - Uses voting processors
    - (different architectures, code)

Caltech CS184 Spring2005 -- DeHon

67


## Forward Recovery

- Can take this voting idea to gate level
  - VonNeuman 1956
- Basic gate is a majority gate
  - Example 3-input voter
- Number of technical details...
- High level bit:
  - Requires  $P_{\text{gate}} > 0.996$
  - Can make whole system as reliable as individual gate

Caltech CS184 Spring2005 -- DeHon

68

## Majority Multiplexing



Maybe there's  
a better  
way...  
...next time.

Caltech CS184 Spring2005 -- DeHon

[Roy+Beiu/IEEE Nano2004]

69

## Rollback Recovery

- Commit state of computation at key points
  - to memory (ECC, RAID protected...)
  - ...reduce to previously solved problem...
- On faults (lifetime defects)
  - recover state from last checkpoint
  - like going to last backup....
  - ...(snapshot)
  - [analysis next time]

Caltech CS184 Spring2005 -- DeHon

70

## Defect vs. Fault Tolerance

- Defect
  - Can tolerate large defect rates (10%)
    - Use virtually all good components
    - Small overhead beyond faulty components
- Fault
  - Require lower fault rate (e.g. VN < 0.4%)
    - Overhead to do so can be quite large

Caltech CS184 Spring2005 -- DeHon

71

## Summary

- Possible to engineer practical, reliable systems from
  - Imperfect fabrication processes (defects)
  - Unreliable elements (faults)
- We do it today for large scale systems
  - Memories (DRAMs, Hard Disks, CDs)
  - Internet
- ...and critical systems
  - Space ships, Airplanes
- Engineering Questions
  - Where invest area/effort?
    - Higher yielding components? Tolerating faulty components?
  - Where do we invoke law of large numbers?
    - Above/below the device level

Caltech CS184 Spring2005 -- DeHon

72

## Big Ideas

- Left to itself:
  - reliability of system  $\ll$  reliability of parts
- Can design
  - system reliability  $\gg$  reliability of parts [defects]
  - system reliability  $\approx$  reliability of parts [faults]
- For large systems
  - must engineer reliability of system
  - ...all systems becoming “large”

Caltech CS184 Spring2005 -- DeHon

73

## Big Ideas

- Detect failures
  - static: directed test
  - dynamic: use **redundancy** to guard
- Repair with **Redundancy**
- Model
  - establish and provide model of correctness
    - perfect model part (memory model)
    - visible defects in model (disk drive model)

Caltech CS184 Spring2005 -- DeHon

74