

ESE680-002 (ESE534): Computer Organization

Day 25: April 16, 2007
Defect and Fault Tolerance



Today

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

Warmup Discussion

- Where do we guard against defects and faults today?

Motivation: Probabilities

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

Probabilities

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

- $N=10^7, P=0.999999$
- $P(\text{all good}) \sim 0.000045$

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
 - As devices get **faster**, failure rate **increases**

Defining Problems

Penn ESE680-002 Spring2007 -- DeHon

7

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

Penn ESE680-002 Spring2007 -- DeHon

8

Defects

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

Penn ESE680-002 Spring2007 -- DeHon

9

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)
- (2nd half of lecture)

Penn ESE680-002 Spring2007 -- DeHon

10

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?

Penn ESE680-002 Spring2007 -- DeHon

11

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 new defect every 10ms
- At 10GHz operation:
 - One new defect every 10^8 cycles
 - $P_{\text{newdefect}} = 10^{-19}$

Penn ESE680-002 Spring2007 -- DeHon

13

First Step to Recover

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

Penn ESE680-002 Spring2007 -- DeHon

14

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

Penn ESE680-002 Spring2007 -- DeHon

15

Detection

- Coding
 - space of legal values \ll 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

Penn ESE680-002 Spring2007 -- DeHon

16

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

Penn ESE680-002 Spring2007 -- DeHon

17

Defect Tolerance

Penn ESE680-002 Spring2007 -- DeHon

18

Two Models

- Disk Drives (defect map)
- Memory Chips (perfect chip)

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)
 - Never allocate a bad sector to a task or file
 - yielded capacity varies

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

Example: Memory

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

Memory Defect Tolerance

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

Memory Techniques

- Row Redundancy
- Column Redundancy
- Block Redundancy

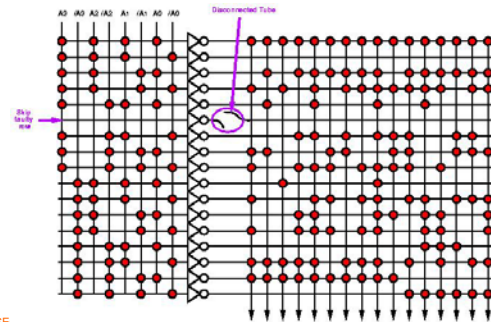
Row Redundancy

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

Penn ESE680-002 Spring2007 -- DeHon

25

Spare Row



Penn ESE

Column Redundancy

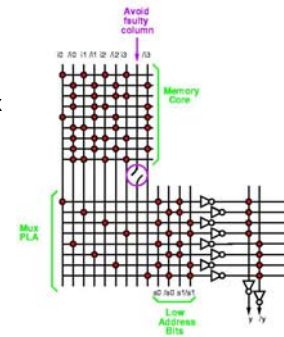
- Provide extra columns
- Program decoder/mux to use subset of columns

Penn ESE680-002 Spring2007 -- DeHon

27

Spare Memory Column

- Provide extra columns
- Program output mux to avoid



Penn ESE680-002 Spring2007 -- DeHon

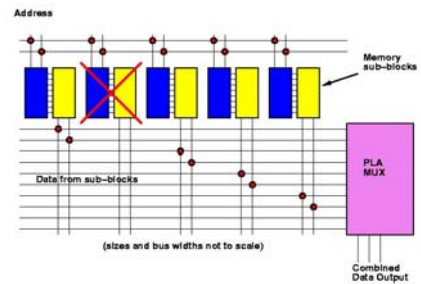
Block Redundancy

- 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)

Penn ESE680-002 Spring2007 -- DeHon

29

Spare Block



Penn ESE680-002 Spring2007 -- DeHon

30

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]

Penn ESE680-002 Spring2007 -- DeHon

31

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!³²

Penn ESE680-002 Spring2007 -- DeHon

Repairable Area

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

Penn ESE680-002 Spring2007 -- DeHon

33

Repairable Area

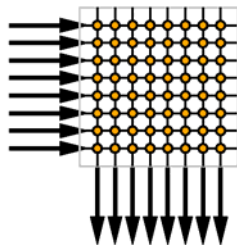
- $P(\text{yield}) = P(\text{non-repair}) \cdot P(\text{repair})$
- $P(\text{non-repair}) = P^N$
 - $N \ll N_{\text{total}}$
 - Maybe $P > P_{\text{repair}}$
 - e.g. use coarser feature size
- $P(\text{repair}) \sim P(\text{yield } M \text{ of } N)$

Penn ESE680-002 Spring2007 -- DeHon

34

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

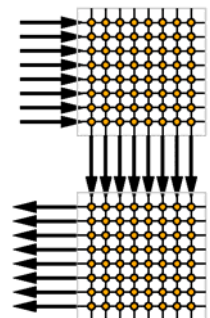


Penn ESE680-002 Spring2007 -- DeHon

35

Crossbar Buses and Defects

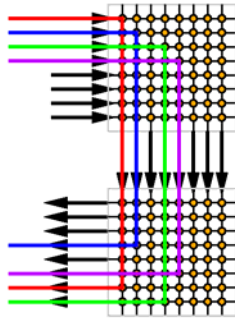
- Two crossbars
- Wires may fail
- Switches may fail
- Provide more wires
 - Any wire fault avoidable
 - M choose N



Penn ESE680-002 Spring2007 -- DeHon

Crossbar Buses and Defects

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

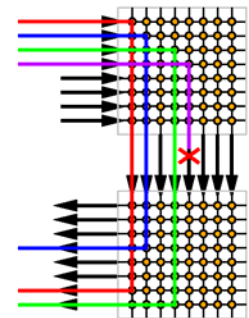


37

Penn ESE680-002 Spring2007 -- DeHon

Crossbar Buses and Faults

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

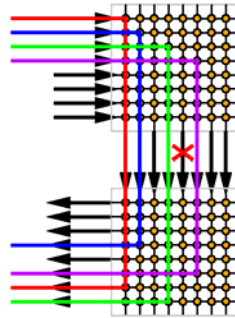


38

Penn ESE680-002 Spring2007 -- DeHon

Crossbar Buses and Faults

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

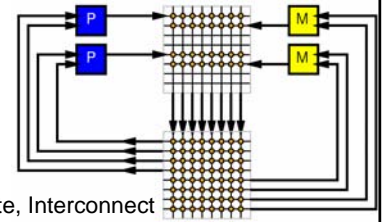


39

Penn ESE680-002 Spring2007 -- DeHon

Simple System

- P Processors
- M Memories
- Wires

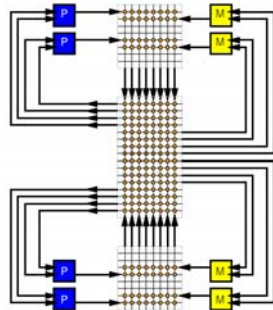


Memory, Compute, Interconnect

Penn ESE680-002 Spring2007 -- DeHon

Simple System w/ Spares

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

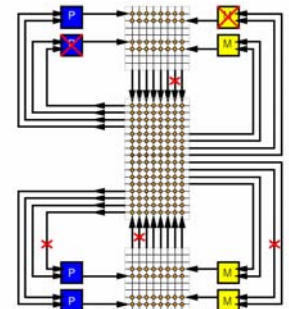


41

Penn ESE680-002 Spring2007 -- DeHon

Simple System w/ Defects

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

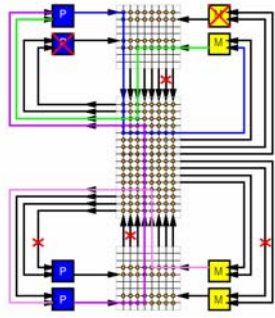


42

Penn ESE680-002 Spring2007 -- DeHon

Simple System Repaired

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

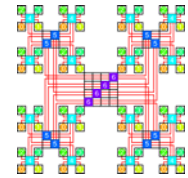


Penn ESE680-002 Spring2007 -- DeHon

43

In Practice

- Crossbars are inefficient [Day13]
- Use switching networks with
 - Locality
 - Segmentation
- ...but basic idea for sparing is the same



Penn ESE680-002 Spring2007 -- DeHon

44

Fault Tolerance

Penn ESE680-002 Spring2007 -- DeHon

45

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

Penn ESE680-002 Spring2007 -- DeHon

46

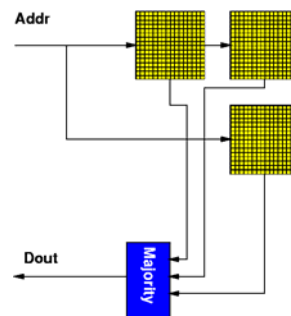
Simple Memory Example

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

Penn ESE680-002 Spring2007 -- DeHon

47

Redundant Memory

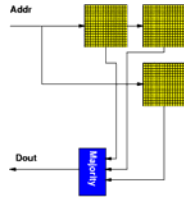


Penn ESE680-002 Spring2007 -- DeHon

48

Redundant Memory

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



49

Penn ESE680-002 Spring2007 -- DeHon

Better: Less Overhead

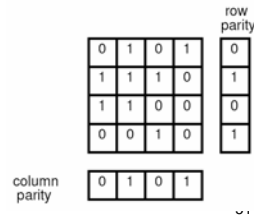
- Don't have to keep N copies
- Block data into groups
- Add a small number of bits to detect/correct errors

50

Penn ESE680-002 Spring2007 -- DeHon

Row/Column Parity

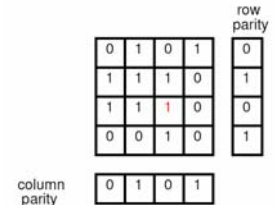
- Think of $N \times N$ bit block as array
- Compute row and column parities
 – (total of $2N$ bits)



Penn ESE680-002 Spring2007 -- DeHon

Row/Column Parity

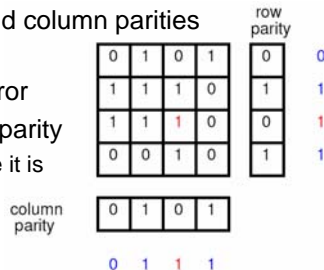
- Think of $N \times N$ bit block as array
- Compute row and column parities
 – (total of $2N$ bits)
- Any single bit error



Penn ESE680-002 Spring2007 -- DeHon

Row/Column Parity

- Think of $N \times N$ 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



Penn ESE680-002 Spring2007 -- 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

54

Penn ESE680-002 Spring2007 -- DeHon

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)

Penn ESE680-002 Spring2007 -- DeHon

55

Interconnect

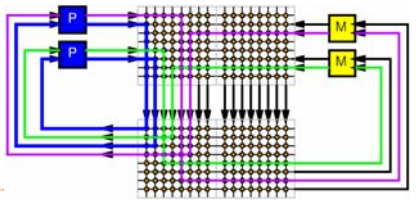
- Also guards against whole path failure
- Sender expects acknowledgement
- If no acknowledgement will retransmit
- If have multiple paths
 - ...and select well among them
 - Can route around any fault in interconnect

Penn ESE680-002 Spring2007 -- DeHon

56

Interconnect Fault Example

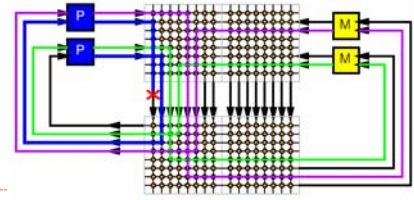
- Send message
- Expect Acknowledgement



Penn ESE680-002 Spring2007 --

Interconnect Fault Example

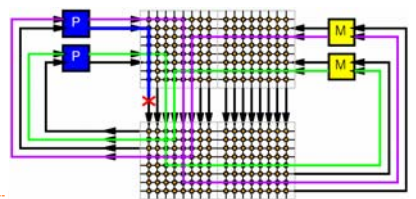
- Send message
- Expect Acknowledgement
- If Fail



Penn ESE680-002 Spring2007 --

Interconnect Fault Example

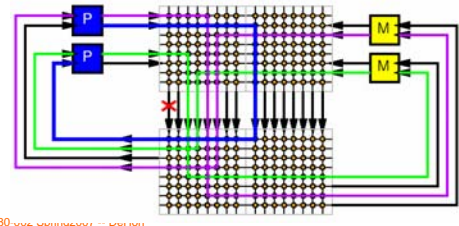
- Send message
- Expect Acknowledgement
- If Fail
 - No ack



Penn ESE680-002 Spring2007 --

Interconnect Fault Example

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

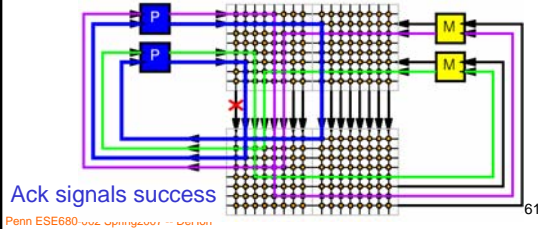


Penn ESE680-002 Spring2007 -- DeHon

30

Interconnect Fault Example

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

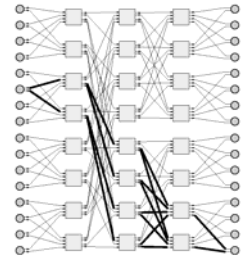


Penn ESE680-002 Spring2007 -- DeHon

61

Transit Multipath

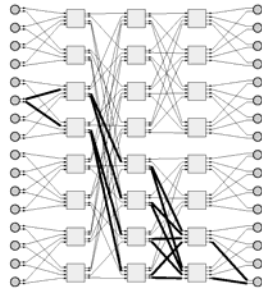
- Butterfly (or Fat-Tree) networks with multiple paths



Penn ESE680-002 Spring2007 -- DeHon

Multiple Paths

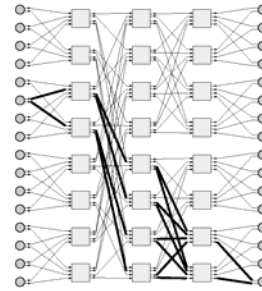
- Provide bandwidth
- Minimize congestion
- Provide **redundancy** to tolerate faults



Penn ESE680-002 Spring2007 -- DeHon

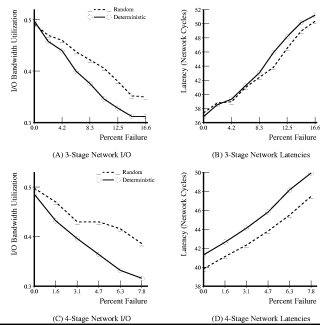
Routers May be faulty (links may be faulty)

- Dynamic
 - Corrupt data
 - Misroute
 - Send data nowhere



Penn ESE680-002 Spring2007 -- DeHon

Multibutterfly Performance w/ Faults

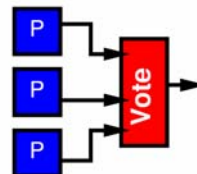


Penn ESE680-002 Spring

65

Compute Elements

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



Penn ESE680-002 Spring2007

66

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

Penn ESE680-002 Spring2007 -- DeHon

67

In Use

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

Penn ESE680-002 Spring2007 -- DeHon

68

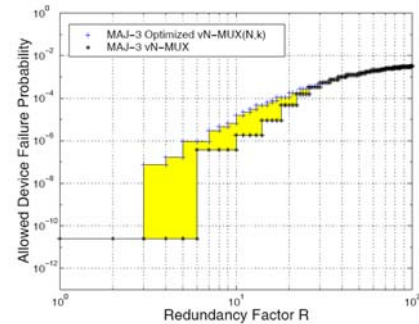
Forward Recovery

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

Penn ESE680-002 Spring2007 -- DeHon

69

Majority Multiplexing



Penn ESE680-002 Spring2007 -- DeHon

[Roy+Beiu/IEEE Nano2004]

70

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)

Penn ESE680-002 Spring2007 -- DeHon

71

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

Penn ESE680-002 Spring2007 -- DeHon

72

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

Penn ESE680-002 Spring2007 -- DeHon

73

Admin

- Final Class on Wednesday
 - Will have course feedback forms
- André traveling 18—26th
 - Won't find in office
- Final exercise
 - Due Friday May 4th
 - Proposals for Problem 3 before Friday April 27th
 - ...same goes for clarifying questions

Penn ESE680-002 Spring2007 -- DeHon

74

Big Ideas

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

Penn ESE680-002 Spring2007 -- DeHon

75

Big Ideas

- Detect failures
 - static: directed test
 - dynamic: use **redundancy** to guard
- Repair with **Redundancy**
- Model
 - establish and provide model of correctness
 - Perfect component model (memory model)
 - Defect map model (disk drive model)

Penn ESE680-002 Spring2007 -- DeHon

76