

CS184c: Computer Architecture [Parallel and Multithreaded]

Day 16: May 31, 2001
Defect and Fault Tolerance



CALTECH cs184c Spring2001 -- DeHon

Today

- EAS Questionnaire (10 min)
- Project Report
- Defect and Fault Tolerance
- Concepts

CALTECH cs184c Spring2001 -- DeHon

Project Report

- Option 1: Slide presentation
 - Wednesday 6th
- Option 2: Paper writeup
 - Due: Saturday 9th

CALTECH cs184c Spring2001 -- DeHon

Concept Review

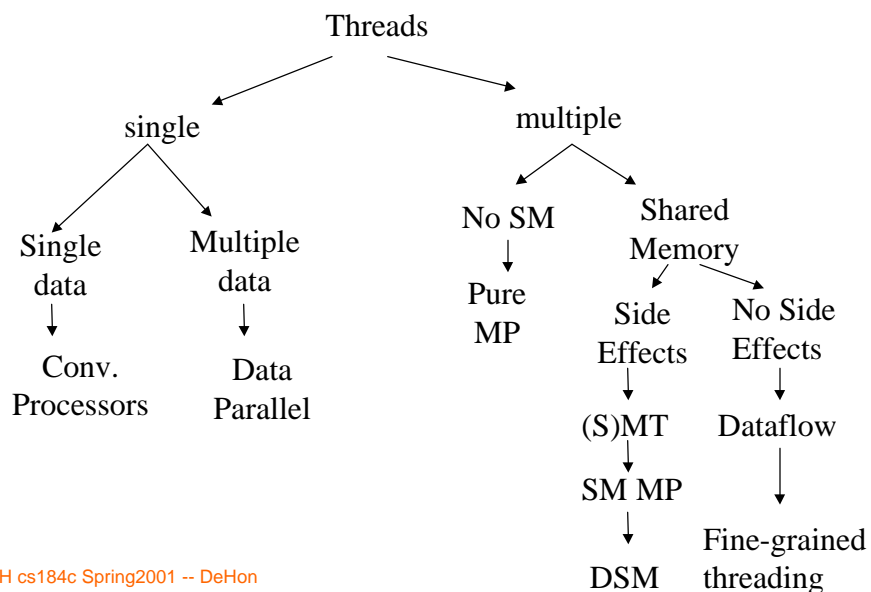
CALTECH cs184c Spring2001 -- DeHon

Models of Computation

- Single threaded, single memory
 - conventional
- Message Passing
- Multithreaded
- Shared Memory
- Dataflow
- Data Parallel
- SCORE

CALTECH cs184c Spring2001 -- DeHon

Models and Concepts



CALTECH cs184c Spring2001 -- DeHon

Mechanisms

- Communications
 - networks
 - io / interfacing
 - models for
- Synchronization
- Memory Consistency
- (defect + fault tolerance)

CALTECH cs184c Spring2001 -- DeHon

Key Issues

- Model
 - allow scaling and optimization w/ stable semantics
- Parallelism
- Latency
 - Tolerance
 - Minimization
- Bandwidth
- Overhead/Management
 - minimizing cost

CALTECH cs184c Spring2001 -- DeHon

Defect and Fault Tolerance

CALTECH cs184c Spring2001 -- DeHon

Probabilities

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

CALTECH cs184c Spring2001 -- DeHon

Probabilities

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

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

CALTECH cs184c Spring2001 -- 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

CALTECH cs184c Spring2001 -- DeHon

Increase Reliability?

- $P_{\text{sys}} = P^N$
- $P_{\text{sys}} = \text{constant}$
- $c = \ln(P_{\text{sys}}) = N \ln(P)$
- $\ln(P) = \ln(P_{\text{sys}})/N$
- $P = N\text{th root of } P_{\text{sys}}$

CALTECH cs184c Spring2001 -- DeHon

Two Models

- Disk Drives
- Memory Chips

CALTECH cs184c Spring2001 -- DeHon

Disk Drives

- Expose faults to software
 - software model expects faults
 - manages by masking out in software
 - (at the OS level)
 - yielded capacity varies

CALTECH cs184c Spring2001 -- DeHon

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 cs184c Spring2001 -- DeHon

Two “problems”

- Shorts
 - wire/node X shorted to power, ground, another node
- Noise
 - node X value flips
 - crosstalk
 - alpha particle
 - bad timing

CALTECH cs184c Spring2001 -- DeHon

Defects

- Shorts example of defect
- Persistent problem
 - reliably manifests
- Occurs before computation
- Can test for at fabrication / boot time and then avoid

CALTECH cs184c Spring2001 -- DeHon

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)

CALTECH cs184c Spring2001 -- DeHon

First Step to Recover

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

CALTECH cs184c Spring2001 -- DeHon

Detection

- Determine if something wrong?
 - Some things easy
 -won't start
 - Others tricky
 - ...one and gate computes $F * T \Rightarrow T$
- Observability
 - can see effect of problem
 - some way of telling if fault present

CALTECH cs184c Spring2001 -- DeHon

Detection

- Coding
 - space of legal values < space of all values
 - should only see legal
 - e.g. parity, redundancy, ECC
- Explicit test
 - ATPG, Signature/BIST, POST
- Direct/special access
 - test ports, scan paths

CALTECH cs184c Spring2001 -- DeHon

Coping with defects/faults?

- Key idea:
 - detection
 - redundancy
- Redundancy
 - spare elements can use in place of faulty components

CALTECH cs184c Spring2001 -- DeHon

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 cs184c Spring2001 -- DeHon

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 cs184c Spring2001 -- DeHon

Memory Techniques

- Row Redundancy
- Column Redundancy
- Block Redundancy


CALTECH cs184c Spring2001 -- DeHon

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

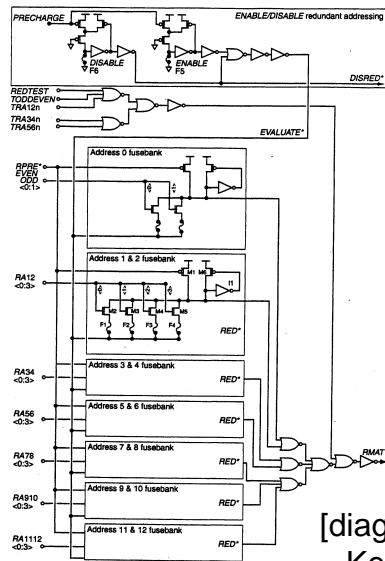
CALTECH cs184c Spring2001 -- DeHon

Spare Row



CALTECH

Row Redundancy



[diagram from
Keeth&Baker 2001]

CALTECH cs184c Spring2001 -- Def...


Column Redundancy

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

CALTECH cs184c Spring2001 -- DeHon


Spare Memory Column

- Provide extra columns
- Program output mux to avoid



CALTECH cs184c Spring2001 -- DeHon

Column Redundancy



[diagram from Keeth&Baker 2001]


CALTECH cs184c Spring2001

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)

CALTECH cs184c Spring2001 -- DeHon

Spare Block



CALTECH cs184c Spring2001 -- DeHon

Yield M of N

- $P(\text{M of N}) = P(\text{yield N})$
 - + $\binom{N}{N-1} P(\text{exactly N-1})$
 - + $\binom{N}{N-2} P(\text{exactly N-2}) \dots$
 - + $\binom{N}{N-M} P(\text{exactly N-M}) \dots$[think binomial coefficients]

CALTECH cs184c Spring2001 -- DeHon

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

CALTECH cs184c Spring2001 -- DeHon

Repairable Area

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

CALTECH cs184c Spring2001 -- DeHon

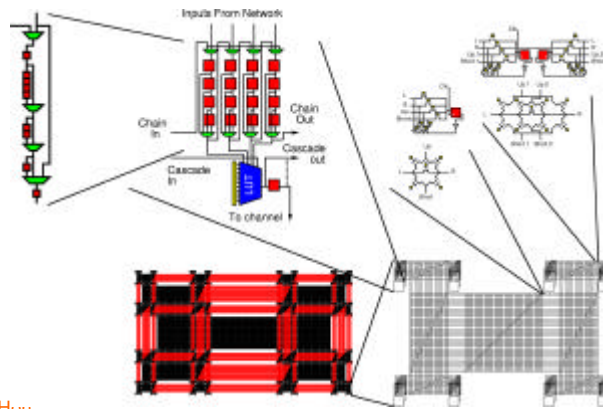
Repairable Area

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

CALTECH cs184c Spring2001 -- DeHon

Consider HSRA


- Contains
 - wires
 - luts
 - switches



CALTECH cs184c Spring2001 -- DeHon

HSRA


- Spare wires
 - most area in wires and switches
 - most wires interchangeable
- Simple model
 - just fix wires



CALTECH cs184c Spring2001 -- DeHon

HSRA “domain” model


- Like “memory” model
- spare entire domains by remapping
- still looks like perfect device



CALTECH cs184c Spring2001 -- DeHon

HSRA direct model


- Like “disk drive” model
- Route design around known faults
 - designs become device specific



CALTECH cs184c Spring2001 -- DeHon

HSRA: LUT Sparing


- All LUTs are equivalent
- In pure-tree HSRA
 - placement irrelevant
- skip faulty LUTs



CALTECH cs184c Spring2001 -- DeHon

Simple LUT Sparing

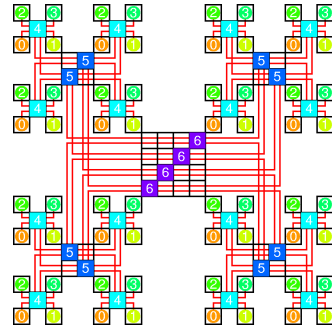
- Promise N-1 LUTs in subtree of some size
 - e.g. 63 in 64-LUT subtree
 - shift try to avoid fault LUT
 - tolerate any one fault in each subtree



CALTECH cs184c Spring2001 -- DeHon

More general LUT sparing

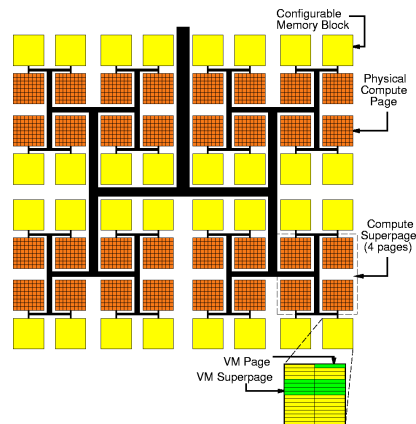
- “Disk Drive” Model
- Promise M LUTs in N-LUT subtree
 - do unique placement around faulty LUTs



CALTECH cs184c Spring2001 -- DeHon

SCORE Array

- Has memory and HSRA LUT arrays



CALTECH cs184c Spring2001 -- DeHon


SCORE Array

- ...but already know how to spare
 - LUTs
 - interconnect
 - in LUT array
 - among LUT arrays and memory blocks
 - memory blocks
- Example how can spare everything in universal computing block

CALTECH cs184c Spring2001 -- DeHon

Transit Multipath


- Butterfly (or Fat-Tree) networks with multiple paths
 - showed last time



CALTECH cs184c Spring2001 -- DeHon

Multiple Paths


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



CALTECH cs184c Spring2001 -- DeHon

Routers May be faulty (links may be faulty)

- Static
 - always corrupt message
 - not (mis) route message
- Dynamic
 - occasionally corrupt or misroute




CALTECH cs184c Spring2001 -- DeHon

Metro: Static Faults

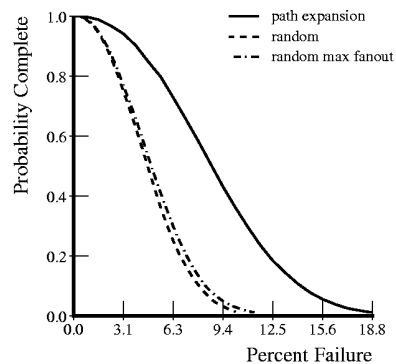
- Turn off
 - faulty ports
 - ports connected to faulty channels
 - ports connected to faulty routers
- As long as paths remain between all communication endpoints
 - still functions

CALTECH cs184c Spring2001 -- DeHon

Multibutterfly Yield



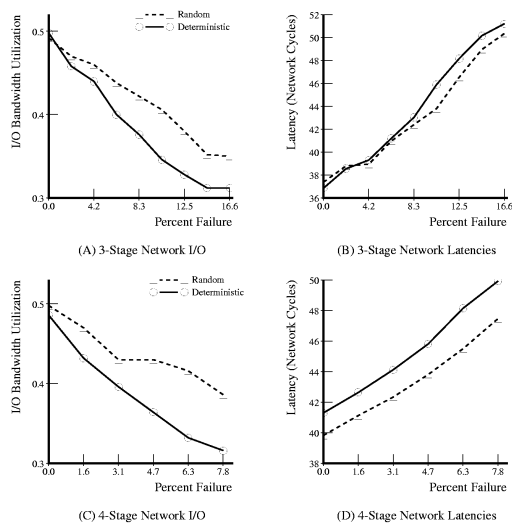
(A) 3-Stage Network Completeness



(B) 4-Stage Network Completeness

CALTECH cs184c Spring2001 -- DeHon

Multibutterfly Performance w/ Faults



CALTECH cs184c Spring

Metro: dynamic faults

- Detection: Check success
 - checksums on packets to see data intact
 - check destination (arrived at right place)
 - acknowledgement from receiver
 - know someone received correctly
- If fail
 - resend message
 - same as blocked route case

CALTECH cs184c Spring2001 -- DeHon

Metro: dynamic faults

- Consequence
 - may have faulty components
 - as long as
 - detection strong
 - there is a non-faulty path
 - will eventually deliver an intact message
 - may deliver multiple times if fault in ack
 - hence earlier concern about idempotence

CALTECH cs184c Spring2001 -- DeHon

Memory: Dynamic Faults

- Error Correcting Codes
- Provide enough redundancy to
 - detect most any errors
 - correct typical errors
- Simple scheme:
 - row and column parity
- ...better schemes in practice
 - [Caltech has whole course on this]

CALTECH cs184c Spring2001 -- DeHon

Processing Faults?

- Simplest model detection:
 - parallel checking
 - run N copies in parallel
 - compare results

CALTECH cs184c Spring2001 -- DeHon

Processor Fault Handling

- What do on fault?
 - Stop (not do anything wrong)
 - maybe just restart
 - adequate if soft error
 - maybe “reconfigure” to substitute out faulty processor
 - Vote
 - if have enough redundancy take most likely

CALTECH cs184c Spring2001 -- DeHon

Checkpoint and Rollback

- Commit state of computation at key points
 - to memory (ECC, RAID protected...)
- On faults
 - recover state from last checkpoint
 - like going to last backup....

CALTECH cs184c Spring2001 -- DeHon

Together

- Examples of handling faults in
 - processing
 - storage
 - interconnect
- All components of our system

CALTECH cs184c Spring2001 -- DeHon

Big Ideas

- Left to itself:
 - reliability of system \ll reliability of parts
- Can design
 - system reliability \gg reliability of parts
- For large systems
 - must engineer reliability of system

CALTECH cs184c Spring2001 -- DeHon

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 cs184c Spring2001 -- DeHon