

Great Ideas in Computer Architecture

Dependability: Parity, ECC, RAID


Instructor: Suvansh Sanjeev

Agenda

- **Dependability**
- Administrivia
- Error Correcting Codes (ECC)
- RAID

Great Idea: Dependability via Redundancy

- Redundancy so that a failing piece doesn't make the whole system fail

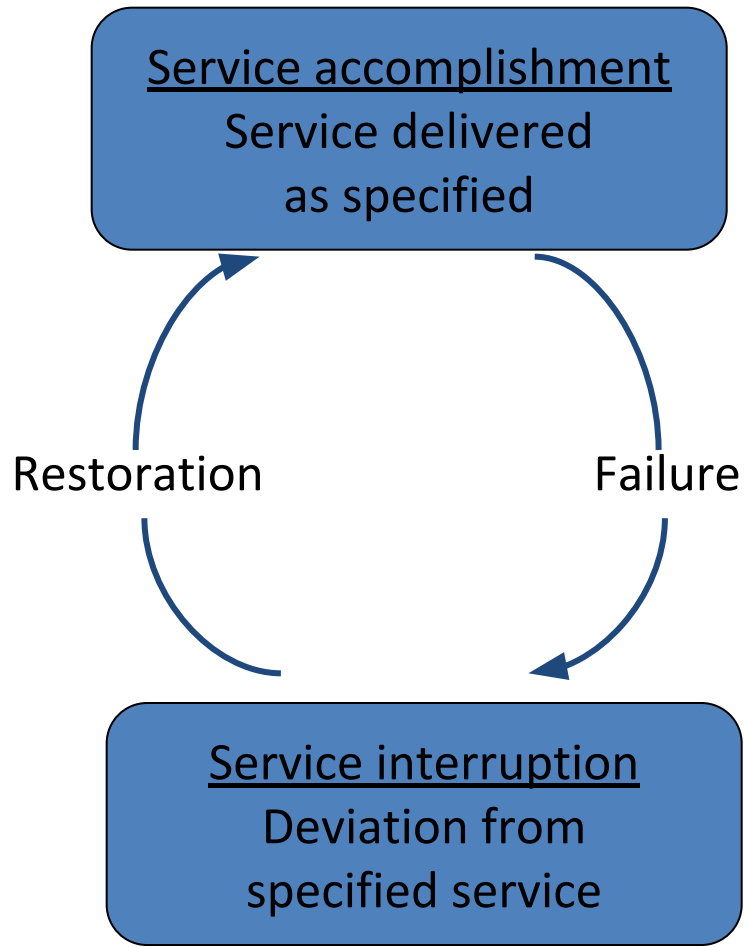


Great Idea: Dependability via Redundancy

- Applies to everything from datacenters to memory
 - Redundant datacenters so that can lose 1 datacenter but Internet service stays online
 - Redundant routes so can lose nodes but Internet doesn't fail
 - Redundant disks so that can lose 1 disk but not lose data (Redundant Arrays of Independent Disks/RAID)
 - Redundant memory bits of so that can lose 1 bit but no data (Error Correcting Code/ECC Memory)



Dependability



- **Fault:** failure of a component
 - May or may not lead to system failure
 - Applies to *any* part of the system

Dependability Measures

- *Reliability*: Mean Time To Failure (MTTF)
- *Service interruption*: Mean Time To Repair (MTTR)
- Mean Time Between Failures (MTBF)
 - $MTBF = MTTR + MTTF$
- **Availability** =
$$\frac{MTTF}{MTTF + MTTR} = \frac{MTTF}{MTBF}$$
- Improving Availability
 - Increase MTTF: more reliable HW/SW + fault tolerance
 - Reduce MTTR: improved tools and processes for diagnosis and repair

Reliability Measures

- 1) MTTF, MTBF measured in hours/failure
 - e.g. average MTTF is 100,000 hr/failure
- 2) Annualized Failure Rate (AFR)
 - Average rate of failures per year (%)

$$\text{AFR} = \underbrace{\left(\frac{\text{Disks}}{\text{MTTF}} \times 8760 \frac{\text{hr}}{\text{yr}} \right)}_{\text{Total disk failures/yr}} \times \frac{1}{\text{Disks}} = \frac{8760 \text{ hr/yr}}{\text{MTTF}}$$

Availability Measures

- Availability = $MTTF / (MTTF + MTTR)$ usually written as a percentage (%)
- Want high availability, so categorize by “number of 9s of availability per year”
 - 1 nine: 90% \Rightarrow 36 days of repair/year
 - 2 nines: 99% \Rightarrow 3.6 days of repair/year
 - 3 nines: 99.9% \Rightarrow 526 min of repair/year
 - 4 nines: 99.99% \Rightarrow 53 min of repair/year
 - 5 nines: 99.999% \Rightarrow 5 min of repair/year

Dependability Example

- 1000 disks with $MTTF = 100,000$ hr and $MTTR = 100$ hr
 - $MTBF = MTTR + MTTF = 100,100$ hr
 - $Availability = MTTF/MTBF = 0.9990 = 99.9\%$
 - $AFR = 8760/MTTF = 0.0876 = 8.76\%$

Calculating MTTR Example

- Faster repair to get 4 nines of availability?
 - $0.9999 = \text{MTTF} / (\text{MTTF} + \text{MTTR})$
 - $0.9999 \times (\text{MTTF} + \text{MTTR}) = \text{MTTF}$
 - $0.9999 \times \text{MTTF} + 0.9999 \times \text{MTTR} = \text{MTTF}$
 - $0.9999 \times \text{MTTR} = 0.0001 \times \text{MTTF}$
 - Plug in $\text{MTTF} = 100,000 \text{ hr}$
 - $\text{MTTR} = 10.001 \text{ hr}$

Dependability Design Principle

- No single points of failure
 - “Chain is only as strong as its weakest link”
- Dependability Corollary of Amdahl’s Law
 - Doesn’t matter how dependable you make *one* portion of system because dependability is limited by the part you do not improve

Question: There's a hardware glitch in our system that makes the Mean Time To Failure (MTTF) *decrease*. Are the following statements TRUE or FALSE?

- 1) Our system's Availability will *increase*.
- 2) Our system's Annualized Failure Rate (AFR) will *increase*.

	1	2
(A)	F	F
(B)	F	T
(C)	T	F
(D)	T	T

Question: There's a hardware glitch in our system that makes the Mean Time To Failure (MTTF) *decrease*. Are the following statements TRUE or FALSE?

- 1) Our system's Availability will *increase*.
- 2) Our system's Annualized Failure Rate (AFR) will *increase*.

	1	2
(A)	F	F
(B)	F	T
(C)	T	F
(D)	T	T

Administrivia

- Proj4 due on Friday (8/03)
 - Hold off on submissions for now
- HW7 due 8/06
- Guerilla Session today @Cory 540AB, 4-6p
- Regrade requests are open for MT2 until Friday
- The final will be 8/09 7-10PM @VLSB 2040/2060!
 - If you have a conflict, look out for an email with details

Agenda

- Dependability
- Administrivia
- **Error Correcting Codes (ECC)**
- RAID

Error Detection/Correction Codes

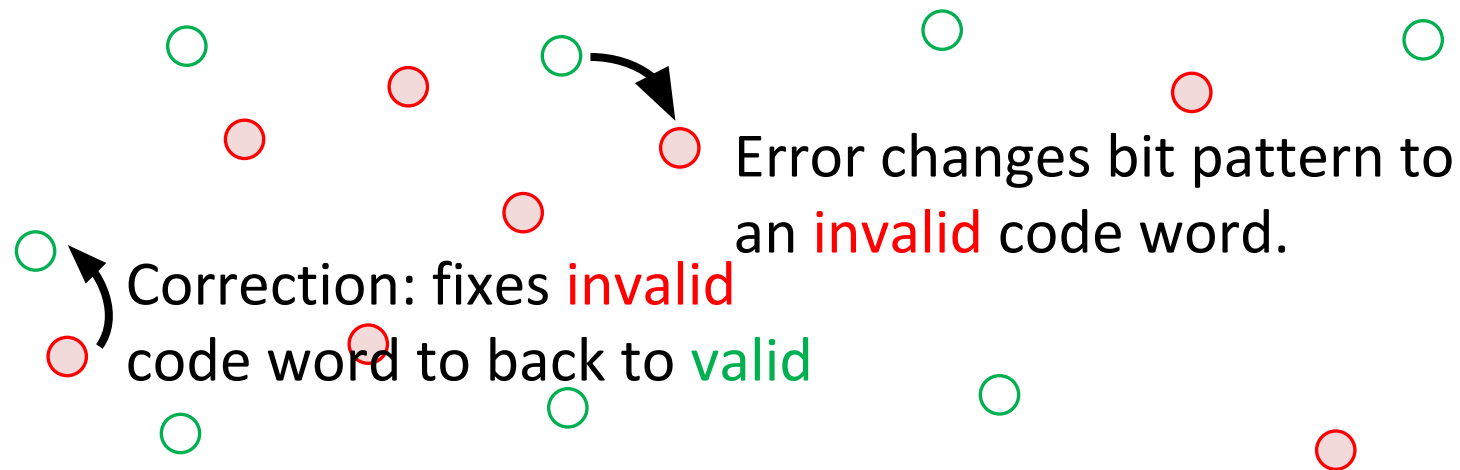
- Memory systems generate errors (accidentally flipped-bits)
 - DRAMs store very little charge per bit
 - “Soft” errors occur occasionally when cells are struck by alpha particles or other environmental upsets
 - “Hard” errors occur when chips permanently fail
 - Problem gets worse as memory systems get denser and larger

Error Detection/Correction Codes

- Protect against errors with **EDC/ECC**
- Extra bits are added to each M-bit data chunk to produce an N-bit “**code word**”
 - Extra bits are a function of the data
 - Each data word value is mapped to a valid code word
 - Certain errors change *valid* code words to *invalid* ones (i.e. you can tell something is wrong)

Detecting/Correcting Code Concept

Space of all possible bit patterns:



2^N patterns, but only 2^M are **valid** code words

- *Detection*: fails code word validity check
- *Correction*: can map to nearest valid code word

Hamming Distance

- Hamming distance = # of bit changes to get from one code word to another

- $p = 0\underline{1}1\underline{0}11$,
 $q = 0\underline{0}1\underline{1}11$, $H_{\text{dist}}(p,q) = 2$

- $p = 011011$,
 $q = 110001$, $H_{\text{dist}}(p,q) = 3$

- If all code words are valid, then

min Hdist between valid code words is 1


– Change one bit, at another valid code word




Richard Hamming
(1915-98)
Turing Award Winner

3-Bit Visualization Aid

- Want to be able to *see* Hamming distances
 - Show code words as *nodes*, Hdist of 1 as *edges*
- For 3 bits, show each bit in a different dimension:




Minimum Hamming Distance 2



- If 1-bit error, is code word still valid?
 - No! So can *detect*
- If 1-bit error, know which code word we came from?
 - No! Equidistant, so cannot *correct*

Minimum Hamming Distance 3




- How many bit errors can we detect?
 - Two! Takes 3 errors to reach another valid code word
- If 1-bit error, know which code word we came from?
 - Yes!

Parity Bit


- Describes whether a group of bits contains an even or odd number of 1's
 - Define 1 = odd and 0 = even
 - Can use XOR to compute parity bit!
- Adding the parity bit to a group will always result in an even number of 1's (“even parity”)
 - 100 Parity: 1, 101 Parity: 0
- If we know number of 1's must be even, can we figure out what a single missing bit should be?
 - 10?11 → missing bit is 1

Parity: Simple Error Detection Coding

- Add parity bit when writing block of data:



- Check parity on block read:
 - Error if odd number of 1s
 - Valid otherwise



- Minimum Hamming distance of parity code is 2
- Parity of code word = 1 indicates an error occurred:
 - 2-bit errors not detected (nor any even # of errors)
 - Detects an odd # of errors

Parity Examples

- 1) Data 0101 0101
- 4 ones, even parity now
 - Write to memory
0101 0101 **0**
to *keep* parity even

- 2) Data 0101 0111
- 5 ones, odd parity now
 - Write to memory:
0101 0111 **1**
to *make* parity even

- 3) Read from memory
0101 0101 0
- 4 ones → even parity, so no error

- 4) Read from memory
1101 0101 0
- 5 ones → odd parity, so error
 - What if error in parity bit?
 - Can detect!

How to Correct 1-bit Error?

- **Recall:** Minimum distance for correction?
 - Three
- Richard Hamming came up with a mapping to allow Error Correction at min distance of 3
 - Called Hamming ECC for Error Correction Code

Hamming ECC (1/2)

- Use *extra parity bits* to allow the position identification of a single error
 - Interleave parity bits *within* bits of data to form code word
 - **Note:** Number bits starting at 1 from the left
- 1) Use *all* bit positions in the code word that are **powers of 2** for parity bits (1, 2, 4, 8, 16, ...)
 - 2) **All other bit positions** are for the data bits (3, 5, 6, 7, 9, 10, ...)

Hamming ECC (2/2)

- 3) Set each parity bit to create even parity for *a group* of the bits in the code word
- The **position** of each parity bit determines the group of bits that it checks
 - Parity bit p checks every bit whose position number in binary has a 1 in the bit position corresponding to p
 - Bit 1 (0001_2) checks bits 1,3,5,7, ... ($XXX1_2$)
 - Bit 2 (0010_2) checks bits 2,3,6,7, ... ($XX1X_2$)
 - Bit 4 (0100_2) checks bits 4-7, 12-15, ... ($X1XX_2$)
 - Bit 8 (1000_2) checks bits 8-15, 24-31, ... ($1XXX_2$)

Graphic of Hamming Code

Bit position	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Encoded data bits	p1	p2	d1	p4	d2	d3	d4	p8	d5	d6	d7	d8	d9	d10	d11
Parity bit coverage	p1	X		X		X		X		X		X		X	
	p2		X	X			X	X			X	X			X
	p4				X	X	X	X					X	X	X
	p8								X	X	X	X	X	X	X

- http://en.wikipedia.org/wiki/Hamming_code

Hamming ECC Example (1/4)

- A byte of data: 10011010
- Create the code word, leaving spaces for the parity bits:

$_1 _2 \mathbf{1} _4 \mathbf{0} \mathbf{0} \mathbf{1} _8 \mathbf{1} \mathbf{0} \mathbf{1} \mathbf{0}$

Hamming ECC Example (2/4)

- Calculate the parity bits:
 - Parity bit 1 group (1, 3, 5, 7, 9, 11):
? _ 1 _ 0 0 1 _ 1 0 1 0 → **0**
 - Parity bit 2 group (2, 3, 6, 7, 10, 11):
? ? 1 _ 0 0 1 _ 1 0 1 0 → **1**
 - Parity bit 4 group (4, 5, 6, 7, 12):
? ? 1 ? 0 0 1 _ 1 0 1 0 → **1**
 - Parity bit 8 group (8, 9, 10, 11, 12):
? ? 1 ? 0 0 1 ? 1 0 1 0 → **0**

Hamming ECC Example (3/4)

- Valid code word: 011100101010
- Recover original data: 1 001 1010

Suppose we see 0₁1₂1₃1₄0₅0₆1₇0₈1₉**1**₁₀1₁₁0₁₂ instead –
fix the error!

But how would we figure out where the error is if we *just* see the code word? Hmm....

- Let's examine the parity bits that are dependent on bit 10
 - Maybe we can figure out a pattern?

Graphic of Hamming Code

0 1 1 1 0 0 1 0 1 **1** 1 0

Bit position	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Encoded data bits	p1	p2	d1	p4	d2	d3	d4	p8	d5	d6	d7	d8	d9	d10	d11
Parity bit coverage	p1	X	X		X		X		X		X	X		X	X
	p2		X	X			X	X		X	X			X	X
	p4				X	X	X	X				X	X	X	X
	p8								X	X	X	X	X	X	X

- Looks like p2 and p8 were responsible
 - p2 & p8 will be the only incorrect parity bits IFF bit 10 is incorrect
 - Notice that $2 + 8 = 10$
 - Seems like incorrect parity bits tell us where the error is... 0_3

Hamming ECC Example (3/4)

- Valid code word: 011100101010
- Recover original data: 1 001 1010

Suppose we see 0₁1₂1₃1₄0₅0₆1₇0₈1₉**1**₁₀1₁₁0₁₂ instead –
fix the error!

How to figure out where the error is? Hmm....

- Check the parity bits responsible for bit 10
 - Parity bits 2 and 8 are incorrect
 - As $2+8=10$, bit position 10 is the bad bit, so flip it!
- Corrected value: 011100101**0**10

Hold on...

Replace all the X's with 1's

Everywhere else, put a 0

Bit position	1	2	3	4	5	6	7	8	9
Encoded data bits	p1	p2	p3	p4	d1	d2	d3	d4	p8
OMFG Parity bit	p1	X	X	X	X	X	X	X	X
	p2	X	X			X	X		
	p4	X	X	X	X	X	X		
coverage	p8							X	X

code word!

- This tells you which bit of the code word is incorrect!

Encoded data bits	Bit position	Parity bit	coverage
p1	1	p1	p8
p2	2	p2	p4
d1	3	d1	
p4	4	p4	
d2	5	d2	
d3	6	d3	
d4	7	d4	
p8	8	p8	
d5	9	d5	
d6	10	d6	
d7	11	d7	
d8	12	d8	
d9	13	d9	
d10	14	d10	

Hamming ECC Example (4/4)

Going in reverse:

We see $0_1 1_2 1_3 1_4 0_5 0_6 1_7 0_8 1_9 1_{10} 1_{11} 0_{12}$

Figure out which bit is wrong:

$$p_1: 0_1 1_3 1_5 1_7 1_9 1_{11} = \text{even number of 1's} : 0$$

$$p_2: 1_2 1_3 0_6 1_7 1_{10} 1_{11} = \text{odd number of 1's} : 1$$

$$p_4: 1_4 0_5 0_6 1_7 0_{12} = \text{even number of 1's} : 0$$

$$p_8: 0_8 1_9 1_{10} 1_{11} 0_{12} = \text{odd number of 1's} : 1$$

Incorrect code bit:

$$0b p_8 p_4 p_2 p_1 = 0b 1010 = 10! \text{ So flip bit 10}$$

Hamming ECC “Cost”

- Space overhead in single error correction code
 - Form $p + d$ bit code word, where p = # parity bits and d = # data bits
- Want the p parity bits to indicate either “no error” or 1-bit error in one of the $p + d$ places
 - Need $2^p \geq p + d + 1$, thus $p \geq \log_2(p + d + 1)$
 - For large d , p approaches $\log_2(d)$

Hamming Single Error Correction, Double Error Detection (SEC/DED)

- Adding extra parity bit covering the entire SEC code word provides *double error detection* as well!

1 2 3 4 5 6 7 8

p_1 p_2 d_3 p_4 d_5 d_6 d_7 p_8

- Let H be the position of the incorrect bit we would find from checking p_1 , p_2 , and p_4 (0 means no error) and let P be parity of complete code word (p 's & d 's)
 - $H=0$ $P=0$, no error
 - $H \neq 0$ $P=1$, correctable single error ($P=1 \rightarrow$ odd # errors)
 - $H \neq 0$ $P=0$, double error detected ($P=0 \rightarrow$ even # errors)
 - $H=0$ $P=1$, an error occurred in p_8 bit, not in rest of word

Question: We saw that when the minimum hamming distance of codewords was 3, we get single error detection and correction. If we had a hamming distance of 4, we'd have

- 1) single error detection
- 2) single error correction
- 3) double error detection
- 4) double error correction

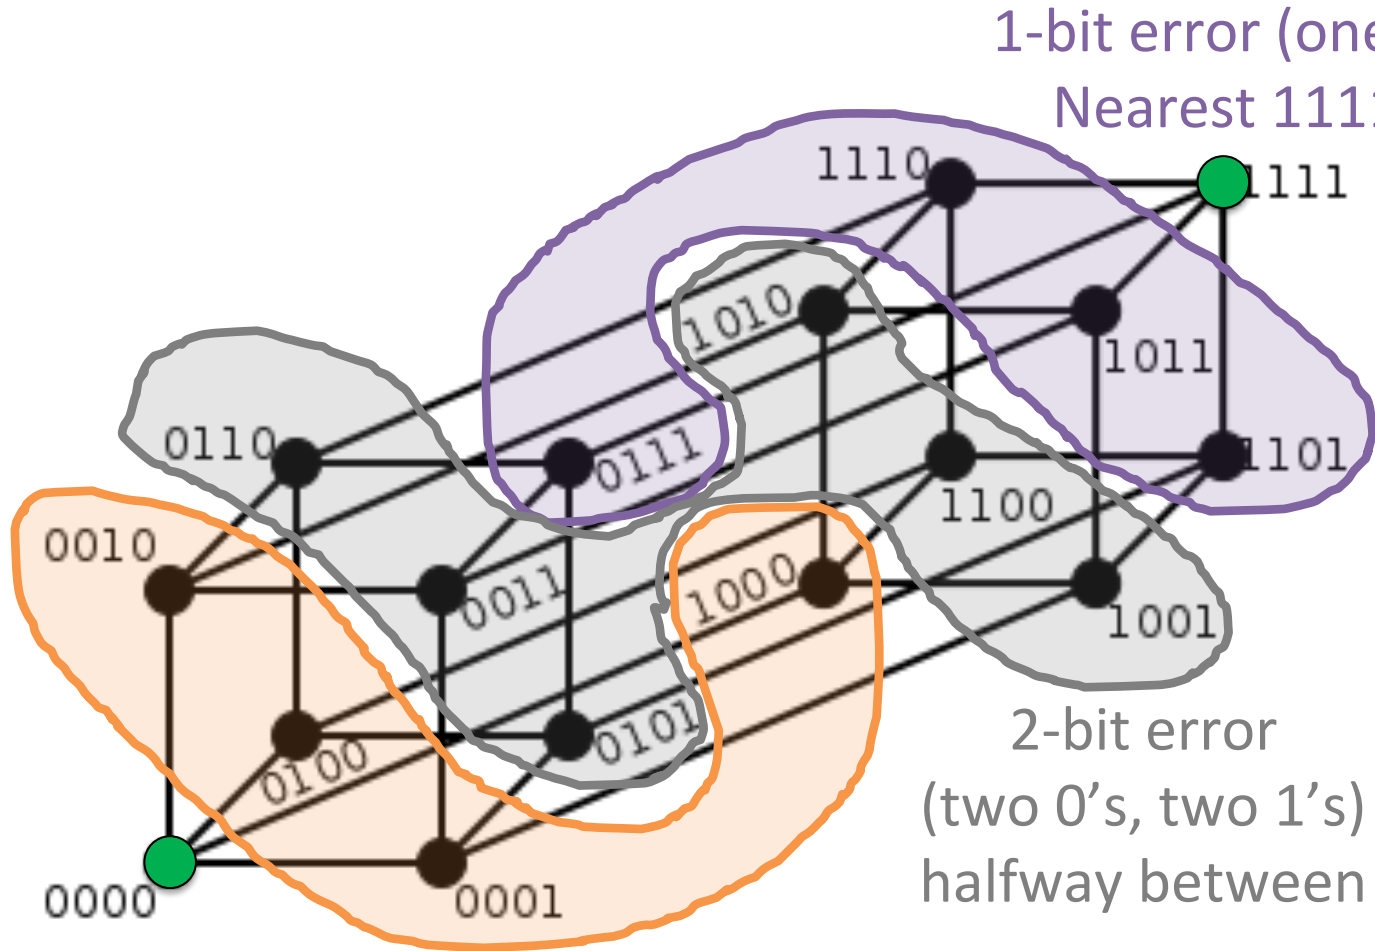
A) 1, 2

B) 1, 3

C) 1, 2, 3

D) 1, 2, 3, 4

SEC/DED: Hamming Distance 4



Summary

- Great Idea: Dependability via Redundancy
 - Reliability: MTTF & Annual Failure Rate
 - Availability: % uptime = $MTTF/MTBF$
- Memory Errors:
 - Hamming distance 2: Parity for Single Error Detect
 - Hamming distance 3: Single Error Correction Code + encode bit position of error
 - Hamming distance 4: SEC/Double Error Detection
- We'll talk about RAID next lecture!