

# FAULT TOLERANT SYSTEMS

<http://www.ecs.umass.edu/ece/koren/FaultTolerantSystems>

## Part 8 - RAID Systems


### Chapter 3 - Information Redundancy

Part.8 .1

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## RAID - Redundant Arrays of Inexpensive (Independent) Disks


- ◆ **RAID1** - two mirrored disks
- ◆ If one disk fails, the other can continue
- ◆ If both work:
  - \* speeds up read accesses - divides them among two disks
  - \* Write accesses are slowed down
- ◆ Computing reliability, availability, and MTDL (mean time to data loss) of **RAID1**



Part.8 .2

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## RAID1 - Reliability Calculation



### ◆ Assumptions:

- \* disks fail independently
- \* failure process - Poisson process with rate  $\lambda$
- \* repair time - exponential with mean time  $1/m$

### ◆ Markov chain: state - number of good disks

$$\frac{dP_2(t)}{dt} = -2\lambda P_2(t) + \mu P_1(t) \quad \frac{dP_1(t)}{dt} = -( \lambda + \mu ) P_1(t) + 2\lambda P_2(t)$$

$$P_0(t) = 1 - P_1(t) - P_2(t) \quad P_2(0) = 1; \quad P_0(0) = P_1(0) = 0$$

### ◆ Reliability at time $t$ -

$$R(t) = P_1(t) + P_2(t) = 1 - P_0(t)$$

Part.8 .3

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## RAID1 - MTDDL Calculation



- ◆ Starting in **state 2** at  $t=0$   
- mean time before entering **state 1** =  $1/(2\lambda)$
- ◆ Mean time spent in **state 1** is  $1/(\lambda + \mu)$
- ◆ Go back to **state 2** with probability  $q = \mu / (\mu + \lambda)$   
or to **state 0** with probability  $p = \lambda / (\mu + \lambda)$
- ◆ Probability of  $n$  visits  
to **state 1** before transition to **state 0** is  $q^{n-1} p$
- ◆ Mean time to enter **state 0** :

$$T_{2 \rightarrow 0}(n) = n \left( \frac{1}{2\lambda} + \frac{1}{\lambda + \mu} \right) = n \frac{3\lambda + \mu}{2\lambda(\lambda + \mu)}$$

$$MTDDL = \sum_{n=1}^{\infty} q^{n-1} p T_{2 \rightarrow 0}(n) = \sum_{n=1}^{\infty} n q^{n-1} p T_{2 \rightarrow 0}(1) = \frac{T_{2 \rightarrow 0}(1)}{p} = \frac{3\lambda + \mu}{2\lambda^2}$$


Part.8 .4

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Approximate Reliability of RAID1



- ◆ If  $m \gg 1$ , the transition rate into state 0 from the aggregate of states 1 and 2 is  $1/MTTDL$
- ◆ Approximate reliability:

$$R(t) = e^{-t/MTTDL}$$



Impact of Disk lifetime


Part.8 .5



Impact of Disk Repair time

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## RAID1 - Availability Calculation



- ◆ Markov chain: state - number of good disks
- ◆ The solution to the differential equations:

$$P_2(t) = \frac{m^2}{(1+m)^2} + \frac{2Im}{(1+m)^2} e^{-(1+m)t} + \frac{I^2}{(1+m)^2} e^{-2(1+m)t}$$

$$P_1(t) = \frac{2Im}{(1+m)^2} + \frac{2I(1-m)}{(1+m)^2} e^{-(1+m)t} - \frac{2I^2}{(1+m)^2} e^{-2(1+m)t}$$

$$P_0(t) = 1 - P_2(t) - P_1(t)$$

- ◆ Long-term availability -

$$A = P_2 + P_1 = 1 - P_0$$

$$= \frac{(m^2 + 2Im)}{(1+m)^2} = 1 - \frac{I^2}{(1+m)^2}$$

Part.8 .6

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## RAID 2

- ◆ A bank of **data disks** in parallel with **Hamming-coded disks**
- ◆ **d** data disks and **c** code disks
- ◆ **i-th** bit of each disk - bit of a **c+d**-bit word
- ◆ From Hamming codes theory - to permit the correction of one bit per word -


$$2^c \geq c + d + 1$$

- ◆ We will not spend more time on **RAID 2** because other **RAID** designs impose much less overhead

Part.8 .7

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## RAID 3



- ◆ Modification of **RAID 2**
- ◆ **Observation** - each disk has error-detection coding per sector - a bad sector can be identified
- ◆ Bank of **d** data disks together with **one** parity disk
- ◆ Data are bit-interleaved across the data disks
- ◆ The **i-th** position of the parity disk contains the parity bit associated with the bits in the **i-th** position of each of the data disks

Part.8 .8

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Error Detection and Correction in RAID3


- ◆ **i-th** bits of each disk form a **d+1**-bit word
  - **d** data and **1** parity bits
- ◆ If **j-th** bit in word is incorrect - sector error-detecting code in **j-th** disk will indicate a failure - fault will be located - remaining bits can be used to restore the faulty bit
- ◆ **Example:** word - **11100** ; data bits - **1110** ; parity bit - **0**
- ◆ If even parity is being used - a bit is in error
- ◆ Third disk indicates an error in the relevant sector and the other disks show no such errors
- ◆ The correct word is **11000**

Part.8 .9

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Reliability of RAID3

- ◆ Similar analysis to **RAID1**: **(d+1)** disks instead of **2**



- ◆ System fails (data loss) if **two** or more disks fail
- ◆ Mean time to data loss for this group is

$$MTTDL = \frac{(2d+1)l + m}{d(d+1)l^2}$$


- ◆ The reliability is given approximately by

$$R(t) = e^{-t/MTTDL}$$

Part.8 .10

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Comparing Different RAID3 Systems




- ◆ Unreliability of **RAID3** systems for different values of **d** - mean lifetime of a single disk is **500,000** hours
- ◆ The **d=1** case - identical to **RAID1**
- ◆ Reliability goes down as **d** increases

Part.8 .11

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## RAID4




- ◆ Similar to **RAID3** - but unit of interleaving block of arbitrary size - a **stripe**
- ◆ Advantage - a small read may be contained in one single data disk, rather than interleaved over all disks
- ◆ Small read operations are faster in **RAID4**
- ◆ Similarly for small write operations
- ◆ Write - affected data disk and parity disk must be updated
- ◆ Parity update simple - parity bit toggles if data bit being written is different from one being overwritten
- ◆ **Reliability** model for **RAID4** - identical to **RAID3**

Part.8 .12

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## RAID5



- ◆ **Observation** - parity disk can be system bottleneck
- ◆ In **RAID4** - parity disk accessed in each write
- ◆ In **RAID5** - parity blocks interleaved among disks
- ◆ Every disk has some data blocks and some parity blocks
- ◆ Reliability model for **RAID5** same as for **RAID4**
- ◆ Only the performance model is different

## Modeling Correlated Failures

- ◆ We assumed until now that disks are independent with respect to failures
- ◆ Disk failures may be correlated - power supply and control are typically shared among multiple disks
- ◆ Disk systems are usually made up of **strings** - consisting of disks that share power supply, cabling, cooling, and a controller
- ◆ If any of these shared support items fail, the entire string can fail
- ◆ If the string constitutes the **RAID** group - data loss can occur

## Approximate Reliability of String

- ◆  $I_{str}$  - failure rate of the support elements (power, cabling, cooling, control) of a string
- ◆  $I_{indep}$  - approximate failure rate due to independent failures
- ◆ If a RAID group is controlled by a single string - the aggregate failure rate of the group is

$$I_{total} = I_{indep} + I_{str}$$


- ◆ And the reliability is

$$R_{total}(t) = e^{-I_{total}t}$$

Part.8 .15

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Impact of String Failures on RAID1




- ◆ Similar results for RAID3 and higher levels
- ◆ Figures of 150,000 hours for the mean string lifetime have been quoted in the literature
- ◆ At least one manufacturer claims mean disk lifetimes of 1,000,000 hours
- ◆ Grouping an entire RAID array as a single string increases unreliability by orders of magnitude

Part.8 .16

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Orthogonal Arrangement of Strings and RAID Groups



- ◆ Failure of a string affects only one disk in each RAID group
- ◆ Since each RAID can tolerate the failure of up to one disk, this reduces the impact of string failures
- ◆ Data loss will happen only if any RAID group has at least two disks down at the same time

Part.8 .17

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Approximate Modeling of Orthogonal Systems

- ◆ Data loss is caused by a sequence of events
- ◆ A failure can be triggered by an individual disk failure or by a string failure - very low failure rates
- ◆ We will find the (approximate) failure rate due to each -  $\Lambda_{indiv}$ ;  $\Lambda_{str}$
- ◆ Sum of these two failure rates - the approximate overall failure rate -  $\Lambda_{data\_loss}$   

$$\Lambda_{data\_loss} = \Lambda_{indiv} + \Lambda_{str}$$
- ◆ It can then be used to approximately determine MTTDL - Mean time to data loss, and reliability - probability of no data loss over any given period of time

Part.8 .18

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Orthogonal Arrangement - Notations

- ◆ **d+1** strings, **g** RAID groups - total of **(d+1)g** disks
- ◆  $f_{disk}(t)$  - density function of the disk repair time
- ◆  $I_{disk}$  - failure rate of a single disk
- ◆  $P_{indiv}$  - probability that a given individual failure triggers data loss
- ◆ Approximate rate (per disk) at which individual failures trigger data loss -  $I_{disk}P_{indiv}$
- ◆  $P_{indiv}$  = probability that a second disk fails in the affected RAID group while the first failure is not yet repaired
- ◆ This second failure has the rate  $d(I_{disk} + I_{str})$  - the second disk failure can happen either due to an individual disk or string failure

Part.8 .19

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Calculating Failure Rates - $\Lambda_{indiv}$

- ◆ Conditioning on **t** - repair time of first disk failure
 
$$\text{Prob}\{\text{Data Loss} \mid \text{repair takes } t\} = 1 - e^{-d(I_{disk} + I_{str})t}$$
- ◆ Unconditional probability of data loss -
 
$$\begin{aligned} \pi_{indiv} &= \int_0^{\infty} \text{Prob}\{\text{Data loss} \mid \text{the repair takes } \tau\} \cdot f_{disk}(\tau) d\tau \\ &= \int_0^{\infty} (1 - e^{-d(\lambda_{disk} + \lambda_{str})\tau}) f_{disk}(\tau) d\tau \\ &= \int_0^{\infty} f_{disk}(\tau) d\tau - \int_0^{\infty} e^{-d(\lambda_{disk} + \lambda_{str})\tau} f_{disk}(\tau) d\tau \\ &= 1 - F_{disk}^*(d[\lambda_{disk} + \lambda_{str}]) \end{aligned}$$
- ◆  $F_{disk}^*(\cdot)$  - the Laplace transform of  $f_{disk}(\cdot)$
- ◆ Approximate rate at which data loss is triggered by individual disk failure -
 
$$\Lambda_{indiv} \approx (d + 1)g\lambda_{disk}\{1 - F_{disk}^*(d[\lambda_{disk} + \lambda_{str}])\}$$

Part.8 .20

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Calculating Failure Rates - $\Lambda_{str}$

- ◆ Total rate of string failures:  $(d + 1)I_{str}$
- ◆ When a string fails - we repair the string, and any individual disks affected by this string failure
- ◆ **Pessimistic assumption** - a second failure can happen at any group or disk before all groups are fully restored
- ◆ **Example:** arrival of a second string failure to the same string before first failure has gone
- ◆ **Optimistic assumption:** disks affected by string failure are immune to further failures before string and affected disks are fully restored
- ◆ The difference between failure rates determines how tight bounds are

Part.8 .21

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Pessimistic Calculation

- ◆  $t$  - (random) time taken to repair the failed string and all disks affected by it
- ◆  $f_{str}()$  - probability density function of  $t$
- ◆  $F_{str}^*()$  - Laplace transform of  $f_{str}()$
- ◆ Pessimistic assumption - rate of additional failures
 
$$I_{pess} = (d + 1)I_{str} + (d + 1)gI_{disk}$$
- ◆ Conditioning upon  $t$  - the probability of data loss
 
$$p_{pess} = 1 - e^{-I_{pess}t}$$
- ◆ Integrating on  $t$  - unconditional pessimistic probability of data loss
 
$$P_{pess} = 1 - F_{str}^*(I_{pess})$$

Part.8 .22

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Optimistic Calculation

- ◆ Optimistic assumption - rate of additional failures

$$l_{opt} = d l_{str} + d g l_{disk}$$

- ◆ Conditioning upon  $t$  - the probability of data loss is

$$p_{opt} = 1 - e^{-l_{opt} t}$$

- ◆ Integrating on  $t$  - unconditional optimistic probability of data loss

$$P_{opt} = 1 - F_{str}^*(l_{opt})$$

Part.8 .23

Copyright 2007 Koren & Krishna, Morgan-Kaufman

## Reliability of Orthogonal System

- ◆ Rate of string failures triggering data loss -

$$\Lambda_{str} = (d + 1) l_{str} p; \quad (p_{pess} \quad \text{or} \quad p_{opt})$$

- ◆ Approximate rate of data loss in the system -

$$\Lambda_{data\_loss} \approx \Lambda_{indiv} + \Lambda_{str}$$

- ◆ Mean Time To Data Loss -

$$MTTDL \approx \frac{1}{\Lambda_{data\_loss}}$$

- ◆ System reliability -

$$R(t) \approx e^{-\Lambda_{data\_loss} t}$$

Part.8 .24

Copyright 2007 Koren & Krishna, Morgan-Kaufman