Scheduling soft real-time jobs over a dual disk system

Shweta Agrawal (02329013)

Chetan Vaity (02329901)

1 December, 2003

System description

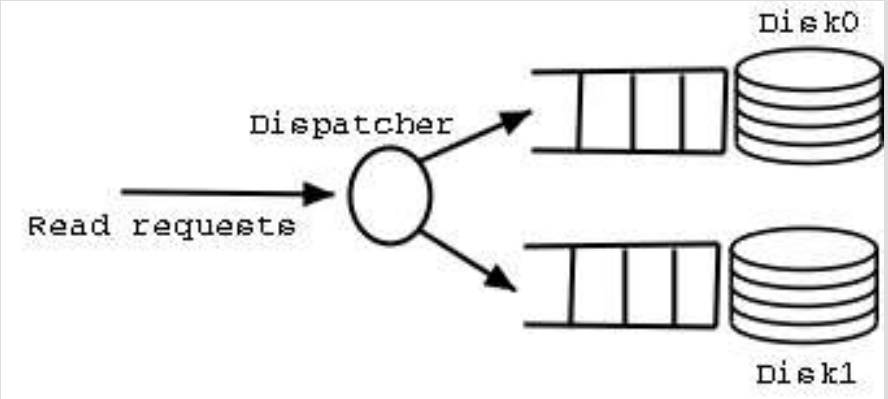

- Mirrored disks
- Only read requests
- Elevator algorithm at each disk
- Requests: Poisson arrivals
- Each request has a “slack” time within which it should be completed
- Slacks are uniformly distributed [S_{min}, S_{max}]

Elevator algorithm on disks

- Request is represented by track number
- Track numbers are uniformly distributed in the range [1, MaxTrack]
- Seek time = DiskFactor $\sqrt{n}$ + DiskConstant
  - DiskFactor: seek time scaling factor
  - DiskConstant: rotational latency + transfer time
  - n: distance (in tracks) of head from request track
- For the current simulation,
  - MaxTrack = 1000, DiskFactor = 0.6, DiskConstant = 15ms

Dispatcher strategies

- Balance
  - Dispatch a job to either disk controller with equal probability
- Chop
  - One disk is reserved for “tight” slack jobs
  - Handles jobs whose slacks are at the lower *p* quartiles of the slack distribution (i.e. Jobs with more stringent time-constraints are dispatched to one of the disks)

Results

Arrival rate = 43 jobs/s, SMIN=36ms, SMAX=326ms, T= 100s,
Number of runs = 100

P	Cutoff	Mean	Variance	95% conf. Interval
0.2	200	3.076	0.185	±0.085
0.25	200	1.996	0.083	±0.057
0.3	200	1.558	0.057	±0.047
0.35	200	1.538	0.049	±0.044
0.4	200	1.826	0.063	±0.049
0.45	200	2.182	0.083	±0.057
0.5	200	2.764	0.111	±0.066
Balance	200	2.716	0.082	±0.057

Graph comparing chop and balance