

CS152
Computer Architecture and Engineering
Lecture 24

I/O and Storage Systems


April 29, 2004

John Kubiatowicz (www.cs.berkeley.edu/~kubitron)

lecture slides: <http://inst.eecs.berkeley.edu/~cs152/>

Recap: Nano-layered Disk Heads

- Special sensitivity of Disk head comes from "Giant Magneto-Resistive effect" or (GMR)
- IBM is leader in this technology
 - Same technology as TMU-RAM breakthrough we described in earlier class.




4/28/04

©UCB Spring 2004

CS152 / Kubiatowicz
Lec24.3

Recap: Disk Device Terminology



Disk Latency = Queuing Time +
Controller time +
Seek Time + Rotation Time + Xfer Time

Order of magnitude times for 4K byte transfers:

Average Seek: 8 ms or less

Rotate: 4.2 ms @ 7200 rpm

Xfer: 1 ms @ 7200 rpm

4/28/04

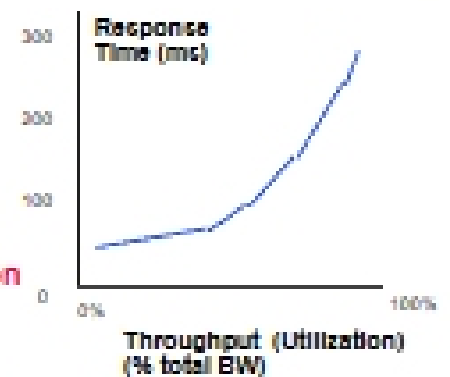
©UCB Spring 2004

CS152 / Kubiatowicz
Lec24.3

Disk I/O Performance

Metrics:
Response Time
Throughput

latency goes as $T_{over} \times u / (1-u)$
 $u = \text{utilization}$




Response time = Queue + Device Service time

4/28/04

©UCB Spring 2004

CS152 / Kubiatowicz
Lec24.4

Introduction to Queueing Theory




- Queueing Theory applies to long term, steady state behavior \Rightarrow **Arrival rate = Departure rate**
- Little's Law:**
Mean number tasks in system = arrival rate x mean reponse time
 - Observed by many, Little was first to prove
 - Simple interpretation: you should see the same number of tasks in queue when entering as when leaving.
- Applies to any system in equilibrium, as long as nothing in black box is creating or destroying tasks

4/28/04


@UCB Spring 2004

CS152 / Rubinfeld
Lec24.5

A Little Queueing Theory: Use of random distributions



- Server spends a variable amount of time with customers
 - Weighted mean $m1 = (f1 \times T1 + f2 \times T2 + \dots + fn \times Tn) / F = \sum p(T) \times T$
 - $\sigma^2 = (f1 \times T1^2 + f2 \times T2^2 + \dots + fn \times Tn^2) / F - m1^2 = \sum p(T) \times T^2 - m1^2$
 - Squared coefficient of variance: $C = \sigma^2 / m1^2$**
 - Unitless measure (100 ms² vs. 0.1 s²)
- Exponential distribution $C = 1$** : most short relative to average, few others long; 80% < 2.3 x average, 89% < average
- Hypoexponential distribution $C < 1$** : most close to average, $C=0.5 \Rightarrow$ 80% < 2.0 x average, only 57% < average
- Hyperexponential distribution $C > 1$** : further from average $C=2.0 \Rightarrow$ 80% < 2.8 x average, 89% < average




4/28/04


@UCB Spring 2004

CS152 / Rubinfeld
Lec24.5

A Little Queueing Theory: Variable Service Time



- Disk response times $C \approx 1.5$ (majority seeks < average)
- Yet usually pick $C = 1.0$ for simplicity
 - Memoryless, exponential dist
 - Many complex systems well described by memoryless distribution!
- Another useful value is average time must wait for server to complete current task: $m1(z)$
 - Called "Average Residual Wait Time"
 - Not just $1/2 \times m1$ because doesn't capture variance
 - Can derive $m1(z) = 1/2 \times m1 \times (1 + C)$
 - No variance $\Rightarrow C = 0 \Rightarrow m1(z) = 1/2 \times m1$
 - Exponential $\Rightarrow C = 1 \Rightarrow m1(z) = m1$



4/28/04

@UCB Spring 2004

CS152 / Rubinfeld
Lec24.7

A Little Queueing Theory: Average Wait Time

- Calculating average wait time in queue T_q :
 - All customers in line must complete; avg time: $m1 \times T_{ser} = 1/\omega$
 - If something at server, it takes to complete on average $m1(z)$
 - Chance server is busy = $u = \lambda/\omega$; average delay is $u \times m1(z)$
- $$T_q = u \times m1(z) + L_q \times T_{ser}$$
- $$T_q = u \times m1(z) + \lambda \times T_q \times T_{ser}$$
- $$T_q = u \times m1(z) + u \times T_q$$
- $$T_q \times (1 - u) = m1(z) \times u$$
- $$T_q = m1(z) \times u / (1 - u) = T_{ser} \times \{1/2 \times (1 + C)\} \times u / (1 - u)$$

Little's Law

Defn of utilization (u)

Notation:

- λ average number of arriving customers/second
- T_{ser} average time to service a customer
- u server utilization (0..1): $u = \lambda \times T_{ser}$
- T_q average time/customer in queue
- L_q average length of queue: $L_q = \lambda \times T_q$
- $m1(z)$ average residual wait time = $T_{ser} \times \{1/2 \times (1 + C)\}$

4/28/04

@UCB Spring 2004

CS152 / Rubinfeld
Lec24.8

A Little Queuing Theory: M/G/1 and M/M/1

- Assumptions so far:
 - System in equilibrium
 - Time between two successive arrivals in line are random
 - Server can start on next customer immediately after prior finishes
 - No limit to the queue: works First-In-First-Out
 - Afterward, all customers in line must complete; each avg T_{ser}
- Described "memoryless" or Markovian request arrival (M for C=1 exponentially random), General service distribution (no restrictions), 1 server: **M/G/1 queue**
- When Service times have C = 1, **M/M/1 queue**

$$T_q = T_{ser} \times u / (1 - u)$$
 - T_{ser} average time to service a customer
 - u server utilization (0..1): $u = \lambda \times T_{ser}$
 - T_q average time/customer in queue

4/28/04

@UCB Spring 2004

CS152 / Rubinfeld
Lec24.9

A Little Queuing Theory: An Example

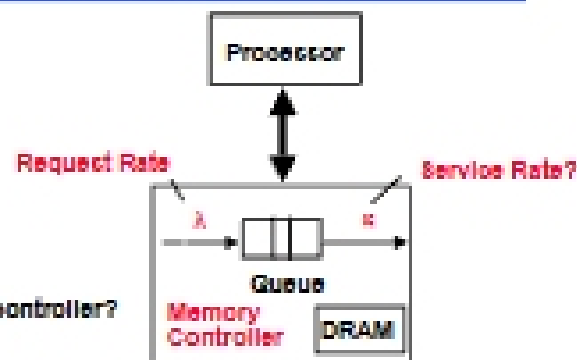
- Processor sends 10 x 8KB disk I/Os per second, requests & service exponentially distrib., avg. disk service = 20 ms
 - This number comes from disk equation:
Service time = Ave seek + ave rot delay + transfer time + otri overhead
- On average, how utilized is the disk?
 - What is the number of requests in the queue?
 - What is the average time spent in the queue?
 - What is the average response time for a disk request?
- Notation:
 - λ average number of arriving customers/second = 10
 - T_{ser} average time to service a customer = 20 ms (0.02s)
 - u server utilization (0..1): $u = \lambda \times T_{ser} = 10/s \times .02s = 0.2$
 - T_q average time/customer in queue = $T_{ser} \times u / (1 - u)$
= $20 \times 0.2 / (1 - 0.2) = 20 \times 0.25 = 5 \text{ ms (0.005s)}$
 - T_{sys} average time/customer in system: $T_{sys} = T_q + T_{ser} = 25 \text{ ms}$
 - L_q average length of queue: $L_q = \lambda \times T_q$
= $10/s \times .005s = 0.05 \text{ requests in queue}$
 - L_{sys} average # tasks in system: $L_{sys} = \lambda \times T_{sys} = 10/s \times .025s = 0.25$

4/28/04

@UCB Spring 2004

CS152 / Rubinfeld
Lec24.10

Memory System I/O Performance



- Pipelined Bus with queue at controller?
 - Time to transfer request
 - $T_{queue} = \text{Queueing Delay} + \text{service time}$
 - Time to transfer data
- DRAM has DETERMINISTIC service time
 - $T_{ser} = t_{bank} + (n-1) \times t_{bc} + t_{precharge}$
 - $T_q = m1(z) \times u / (1 - u) = T_{ser} \times (1/2 \times (1 + C)) \times u / (1 - u)$
with $C=0$

4/28/04

@UCB Spring 2004

CS152 / Rubinfeld
Lec24.11

Administrivia

- Go to the "Projects" link and describe your project (By Friday)
- Thursday: Sections in lab again (119 Cory)
 - Discuss your design document
 - Send draft to TA tonight
- Midterm II next Wednesday
 - 5:30 – 8:30 in 306 Soda Hall
 - Pizza afterwards
 - Topics
 - Pipelining
 - Caches/Memory systems
 - Busess and I/O (Disk equation)
 - Queueing theory
 - Can bring 1 page of notes and calculator
 - Handwritten, double-sided (CLOSED BOOK!)
- Oral Report
 - Powerpoint
 - 15 minute presentation, 5 minutes for questions

4/28/04

@UCB Spring 2004

CS152 / Rubinfeld
Lec24.12