

Storage Systems

NPTEL Course

Jan 2012

(Lecture 28)

K. Gopinath

Indian Institute of Science

Performance Modelling

- Storage devices often the slowest component in a system
- Often determine the speed and responsiveness of the system
- Modelling useful
 - Simple models sometimes very useful!
 - Effectiveness of, say, caching may/may not be useful
 - Two zones of operation

What/How to model?

CPU utilization for processing storage requests

Compression, Encryption, Deduplication, QoS for mm

I/O Completion Processing and Interrupt handling

Network packet processing

QoS for packets

Scheduling groups of processes across nodes:
distributed search

Two main types

- Operational models: limited assumptions
- Stochastic models: mathematically tractable (eg. Queuing theory)

Operational Laws

Little's Laws: true if during observed period T ,

- arrivals (a) \sim completions (c)
- ie: $(a-c)$ small compared to c

Let J = time in system

Mean time spent in system = J/N

Mean # in system = $J/T = J/N * N/T = \text{response time} * \text{throughput}$

Or: Q_i (mean # in device i) = $X_i R_i$

Usually written as $L = \lambda W$:

Mean number in system = arrival rate * mean time spent in system

Poisson Arrivals

Assumes that in a small interval δ

of arrivals: $\lambda * \delta$

Prob of more than 1 arrival in δ : negligible

Arrivals in nonoverlapping intervals statistically indep

Expected arrival time = $1/\lambda$


Probability of an arrival in time $t = 1 - \exp(-\lambda t)$

Probability of no arrivals in time $t = \exp(-\lambda t)$

Probability of k arrivals in time $t = \exp(-\lambda t)(\lambda t)^k/k!$


Similarly, Poisson departures

M/M/4!



M/M/1 Analysis

- Arrival rate Poisson λ
- Service rate exponential μ
- Model as birth-death process
- Steady state ($\lambda < \mu$): Probability stable that i customers in system ($\pi(i)$). Given by
 - $\pi(i) * \mu = \pi(i-1) * \lambda, i = 1, \dots$
 - ie. $\pi(i) = \rho \pi(i-1), i = 1, \dots$



M/M/1

- $\pi(i) = \rho^i \pi(0), i = 1, \dots$
- Since sum of all $\pi(i), i = 1, \dots$ should be 1
 - $\pi(i) = (1-\rho) \rho^i, i = 1, \dots$
- Mean number of customers $E[N] = \rho / (1-\rho)$
- Expected Waiting time $E[W] = E[N]/\mu$
- Expected service time $E[T] = E[W] + 1/\mu = 1/(\mu-\lambda)$

M/M/1 vs M/M/2

Response time for separate Qs (M/M/1) = $1/\mu(1-\rho)$

- with $\rho = (\lambda/2)/\mu$,
 - $E[W] = 2/(2\mu - \lambda) = (4\mu + 2\lambda)/(4\mu^2 - \lambda^2)$

Response time for combined Qs (M/M/2):

$$E[N] = 2\rho/(1-\rho^2) \text{ where } \rho = \lambda/2\mu$$

$$E[W] = E[N]/\lambda \text{ (Little's Law)} = 4\mu/(4\mu^2 - \lambda^2)$$

Less than that for separate Qs !

Utilization Law

Utilization = Busy Time/T = (completions/T) * (Busy Time/completions) = throughput * service time

Or, at device i , $U_i = X_i S_i$

Forced Flow Law: Device throughput of device i =
completions(i)/T = completions(i)/completions of jobs *
completions of jobs/T = visit ratio * system throughput

Or, $X_i = X v_i$

Utilization of device U_i = throughput $_i$ * service time $_i$ = visit ratio $_i$ *
system throughput * service time $_i$ = system throughput * total
service demand on device $_i$

Or, $U_i = X D_i$

Example (from Jain'91)

Each prog requires 5 secs of CPU time & 80 I/O reqs to disk A and 100 I/O reqs to disk B. Disk A takes 50ms; disk B takes 30ms. Total of 17 terminals with disk A's thruput 15.7 I/O reqs/s. Ave think time: 18s. What is the system thruput? Device utilization?

$$D_{\text{diskA}} = V_A S_A = 80 * 1/20 = 4s$$

$$D_{\text{diskB}} = V_B S_B = 100 * 3/100 = 3s$$

$$X_A = 15.7 = X V_A \Rightarrow X = 15.7/80$$

$$D_{\text{cpu}} = 5s; V_{\text{cpu}} = V_A + V_B + 1 = 181;$$

$$U_{\text{cpu}} = X D_{\text{cpu}} = (15.7/80) * 5 = 98\%; U_A = X D_A; U_B = X D_B$$

If $Q_{\text{cpu}} = 8.88$, $Q_A = 3.2$; $Q_B = 1.4$, what is response time?

Use Little's law: $R_i = Q_i / X_i$

General Response Time Law: $Q = Q_1 + \dots + Q_n$; $Q_i = X_i R_i$
 (Little's Law); $Q = XR = \sum Q_i = \sum X_i R_i$ $R = \sum V_i R_i$

Interactive Response Time Law: Z (think time) +
 response (R). In time T , $T/(Z+R)$ requests.

With N users, system thruput $X = (N \cdot T / (Z+R)) / T = N / (Z+R)$

$$\Rightarrow R = N/X - Z$$

Bottleneck Analysis: $X(N) \leq \min(1/D_{\max}, N/(D+Z))$ where
 $D = \sum(D_i)$ (the sum of total service demands on all
 devices)

$$U_{\text{bottleneck device}} = X D_{\max} \leq 1 \Rightarrow X \leq 1/D_{\max}$$

With 1 job: $R(1) = D_1 + \dots + D_M = D$

$$R(N) \geq R(1) = D; \quad X(N) = N/(R(N)+Z) \leq N/D+Z$$

Conclusion

- Simple Models often give quick insight
 - Esp if devices are of widely varying speeds!