

Queuing Theory A Quick View



Ray Wicks
561-236-5846

RayWicks@us.ibm.com

Bibliography

Ray has spent most of his career at IBM in the performance analysis and capacity planning end of the business in Poughkeepsie, London, and now at the Washington Systems Center. He is the major contributor to IBM's internal PA & CP tool zCP3000. This tool is used extensively by the IBM services and technical support staff world wide to analyze existing zSeries configurations (Processor, storage, and I/O) and make projections for capacity expectations.

Ray has given classes and lectures worldwide. He was a visiting scholar at the University of Maryland where he taught part time at the Honors College.

He won the prestigious Computer Measurement Group's A.A. Michelson award in 2000. His recent virtual sessions "Getting Started in Performance Analysis & Capacity Planning" workshop held for attendees in China and India was well accepted.

Queuing Theory

This session reviews some of the basics of queuing theory – the terminology, the assumptions, some statistics and some simple model implementations.

Although one may not do queuing theory, in Performance Analysis and Capacity planning discussions, it is very important to know what the terms mean.

Included will be Little's Law and M/M/1 and M/M/c equations. And then begins the slippery slope: once the basic equations are understood, the reality of non Markovian distributions make the match with reality a bear. Enter M/M/c/k models.

Excel graphics will be used to see what the equations are telling us. This will then be followed by a taste of the real implementation in larger queuing models: Mean Value Analysis.

Trade Marks, Copyrights & Stuff


This presentation is copyright by Ray Wicks 2008-2010.

Selected material reproduced by permission of  Corporation.

Many terms are trademarks of different companies and are owned by them.

- **On foils that appear in this presentation are not in the handout. This is to prevent you from looking ahead and spoiling my jokes and surprises.**


The Approach



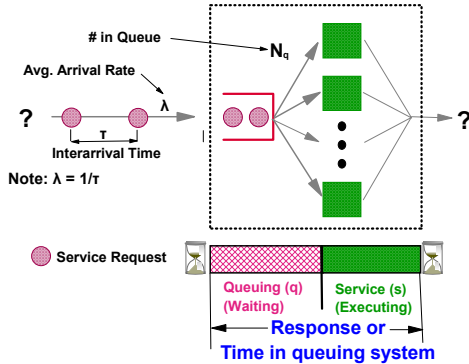
- Queuing Theory
- Shape of curve?
 - Quantification?
 - Exceptions?

As the utilization increases, response time gets worse.

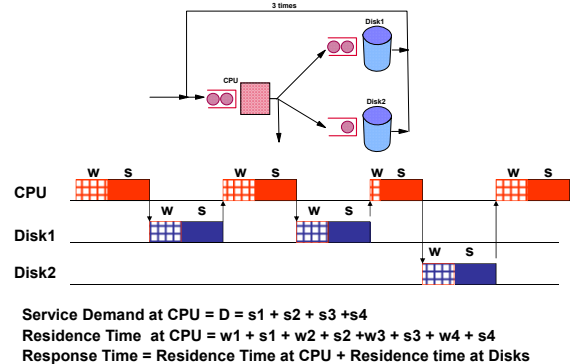
Queuing Analysis



A Service Center



Terminology



Total Response Time

$RT = V_1R_1 + V_2R_2 + V_3R_3$
 Where V_i = # visits to Service Center i
 S_i = service time at Service Center i
 R_i = Response at Service Center i

Service Center Considerations

- Arrival Rate
- Total Population (Closed or ∞ ?)
- Number of Servers
- Speed of Servers
- Queuing Discipline
- Time Distributions

- Queuing Time
- Service Time

Utilization Law


λ = Throughput (transactions/sec) 24
 t = Service time (sec/trans) 0.053 Secs
 u = utilization of server
 (Utilization = % = Secs/sec busy)
 N = number users in server "system"
 c = number of servers 2

$\lambda t = 24 * 0.053 = 1.272$ seconds/second (Traffic)
 $u = \lambda t / c = 1.272 / 2 = 63.6\%$
 $\lambda * \text{Service_Time} = \text{Traffic}$

Examples

- Workload uses 32% of a 4 way processor for 10 transactions / second. What's the service time?
 □ Total CPU seconds = #CPUs x Busy
 = $4 \times 32\%$
 = $4 \times 0.32 = 1.28$ Seconds
 □ Service time = CPU seconds/Transaction
 = $1.28 / 10 = 0.128$ seconds
- A 5 Server system is busy for 23 seconds over a minute. What's the average System utilization?
 □ $23 / 60 = 0.38$ seconds/second
 □ Utilization = total CPU seconds/second / # servers = $0.38 / 5 = 0.076$
 □ 7.6%

Little's Law



$\lambda_1 < \lambda_2?$
 $\lambda_1 > \lambda_2?$
 $\lambda_1 = \lambda_2$ (steady state)

At Steady State:
 Average Residency Time = $T = N/\lambda$
 Response Time = T
 $\lambda * \text{Response_Time} = \text{Intensity}$

Example

A server processes 630,000 request in a half hour. The average number of requests in the server is 2. What's the average Response time?

$\lambda = 630000/1800 = 350/\text{second}$
 $N = \lambda T$
 $T = N/\lambda = 2/350 = 0.0058$ seconds

→ Know ST & number in system, you can compute RT
 → $Q = RT - ST$

Philosophical Remark


We Understand a law by trying to break it

$\forall x \Phi x \approx \sim \exists y \sim \Phi y$

We Generalize on limited information


$\exists y \Phi y \approx \forall x \Phi x$

Distributions



Each circle represents an arrival. If the interval is one second, notice that the average inter-arrival time in all cases is 100Ms or 0.1 seconds. What's the impact on response time for various service times?


Exponential Function $f(x)=e^{-x}$



Full Function is: $\lambda e^{-\lambda x}$ where $E[f(x)]=1/\lambda$


Erlang $K=3 \Theta=2$

$$\frac{x^{K-1} e^{-x/\Theta}}{\Theta^K (K-1)!}$$




T-Distribution $f_n(t)$

n = number of degrees of freedom



F-Distribution $f_{m,n}(F)$

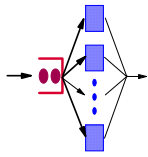
m & n = number of degrees of freedom



Kendal Notation

A/B/c/K/m/Z

A describes inter-arrival time
 B describes the service time distribution
 c the number of servers
 K maximum number of users allowed in system
 m number of users in population
 Z is the queuing discipline



Common distributions
 M exponential
 D constant
 Ek Erlang-k

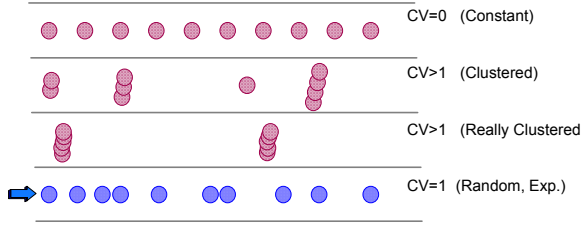
$M/M/1 = M/M/1/\infty/\infty/FCFS$

Queuing Theory Paradox

Buses pass a certain corner with an average time between them of 20 minutes. What is the average time that one would expect to wait?

Queuing Theory Paradox

Buses pass a certain corner with an average time between them of 20 minutes. What is the average time that one would expect to wait?



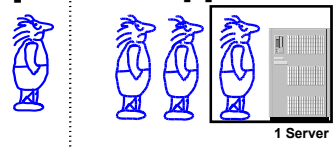
CV=0 (Constant)
 CV>1 (Clustered)
 CV>1 (Really Clustered)
 CV=1 (Random, Exp.)

Coefficient of Variation, CV = standard deviation / mean or σ/μ

Expected Response New Arrival

$E[rt]=?$

$E[s]=10$

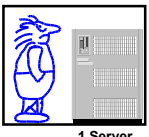


$E[rt]=E[s]*(1+N[q])$

If the service time is distributed exponentially (CV=1), the expected value of the time left of the one getting service is the expected value!

Expected Time to Completion

$E[s]=10$



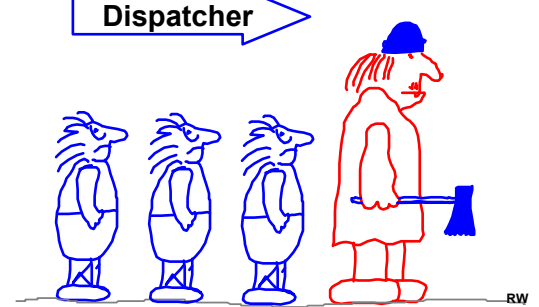
CV	Time Left
0	5
1	10
2	15

Upon arrival of a new request, how much time is remaining for someone in service?

$E[ST] \times (1+CV^2)/2$

4 Typical Tasks

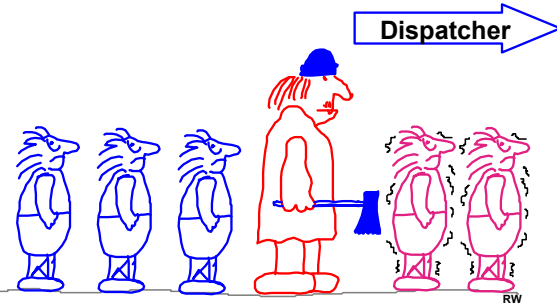
Dispatcher



Guess who gets all they want? How much do we, the little people, get if there's one server? Two servers? Four servers? What's best?

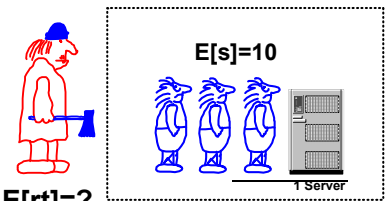
A More Complicated Situation

Dispatcher



Expected Response Big Shot?

$E[s]=10$

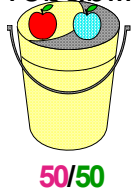


$E[rt]=?$

$E[rt]=E[s]$

If the arriving request is a "Big Shot" (high priority) and he can jump to the head of the line and preempt the one getting service, what's the expected RT?


Probability



50/50

- What's the probability of drawing a Red apple from a bucket?
- What's the probability of finding a server busy if it is averaging 50% busy?

Probability



50/50 50/50

- What's the probability of drawing a Red apple from each bucket?
- What's the probability of finding both servers busy if they are both averaging 50% busy?


An Ugly Formula

Erlang's C formula (M/M/c) for the probability of finding c Servers Busy (a Variation)


$$C(c, T) = \frac{\frac{T^c}{c!}}{\frac{T^c}{c!} + (1 - U) \sum_{n=0}^{c-1} \frac{T^n}{n!}}$$

$C(1, T) = U$
 $C(2, T) = \frac{2 U^2}{(1 + U)}$
 $C(4, T) = \frac{32 U^4}{3 + 9 U + 12 U^2 + 8 U^3}$

T=total Traffic
 U=Average Utilization = T/c



True if




Probability of finding server busy?
 Pr Server busy = busy time / clock = 0.7 secs / 1 second
 Pr(busy)=0.7?

Pr(server busy) = 0.6 / 0.9 = 0.67

Pr(server busy) = 0.1 / 0.4 = 0.25

Probability of 4 Servers busy

$$C(4,T) = (32U^4)/(3+9U+12U^2+8U^3)$$



The intuition of U^4 would be optimistic compared to Erlang's $C(4,T)$. Which would you use?

Excel Implementation

Utilization	U^4	$C(4,T)$
0.05	0.00000625	5.74548E-05
0.1	0.0001	0.000794439
0.15	0.00050625	0.00348612
0.2	0.0016	0.009580838
0.25	0.00390625	0.020408163
0.3	0.0081	0.037049743
0.35	0.01500625	0.060303906
0.4	0.0256	0.090699734
0.45	0.04100625	0.128533647
0.5	0.0625	0.173913043
0.55	0.09150625	0.226798854
0.6	0.1296	0.287043189
0.65	0.17850625	0.354420798
0.7	0.2401	0.428654318
0.75	0.31640625	0.509433962
0.8	0.4096	0.596432472
0.85	0.52200625	0.689316222
0.9	0.6561	0.787753264
0.95	0.81450625	0.891418995

Example

If a system has 2 servers. For what utilization threshold I might expect the probability of both servers being busy to be less than 0.1?

$$C(2,T) = (2U^2) / (1 + U)$$

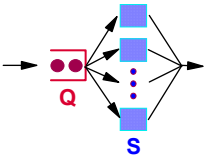
$$0.1 = (2U^2) / (1 + U)$$

$$20U^2 - U - 1 = 0$$

$$U = -0.2, +.25$$

Answer = at 25% busy the probability of finding both busy is 0.1. Or 90% of the time a request will not wait.

Erlang's M/M/c



$$E[RT] = E[S] + E[Q]$$

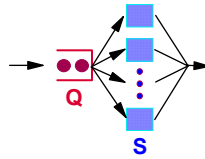
$$E[RT] = E[S] + \frac{C(c,T)E[S]}{c(1-U)}$$

$$E[RT] = E[S] \left\{ 1 + \frac{C(c,T)}{c(1-U)} \right\}$$

c = Number of CPs
U = Utilization
T = traffic or $c \cdot U$
C(c,T) is Erlang's C formula
E[s] is expected service time

From any queuing theory book: Arnold or Jain for example.

Erlang's M/M/c



$$E[RT] = E[S] + E[Q]$$

$$E[RT] = E[S] + \frac{C(c,T)E[S]}{c(1-U)}$$

$$E[RT] = E[S] \left\{ 1 + \frac{C(c,T)}{c(1-U)} \right\}$$

c = Number of CPs
U = Utilization
T = traffic or $c \cdot U$
C(c,T) is Erlang's C formula
E[s] is expected service time

Read as Contention Factor.
 When CF=0, $E[RT] = E[ST]$
 When CF=1, $E[RT] = 2 \cdot E[ST]$

M/M/1

$$E[RT] = E[S] + E[Q]$$

$$E[RT] = E[S] + \frac{C(c,T)E[S]}{c(1-U)}$$

$$E[RT] = E[S] \left\{ 1 + \frac{C(c,T)}{c(1-U)} \right\}$$

..... C=1

$$E[RT] = \frac{E[S]}{1-U}$$




IF $E[S] = 30$ Ms. And $U=80\%$
 Then $E[RT] = 30/(1-0.8) = 150$

M/M/1 Exercise

Workload	Service Time	Workload Utilization
Hi	0.05 sec	32%
Medium	0.25 sec	45%
Low	1.32 Min	12%

Assume M/M/1.

- (1) What is the expected RT for Medium?
- (2) At what effective utilization would the response time for medium exceed 1.2 seconds?



M/M/1 Exercise

Workload	Service Time	Utilization
Hi	0.05 sec	32%
Medium	0.25 sec	45%
Low	1.32 Min	12%

Assume M/M/1.

- What is the expected RT for Medium?
- At what effective utilization would the response time for medium exceed 1.2 seconds?

$$RT = ST / (1 - U)$$

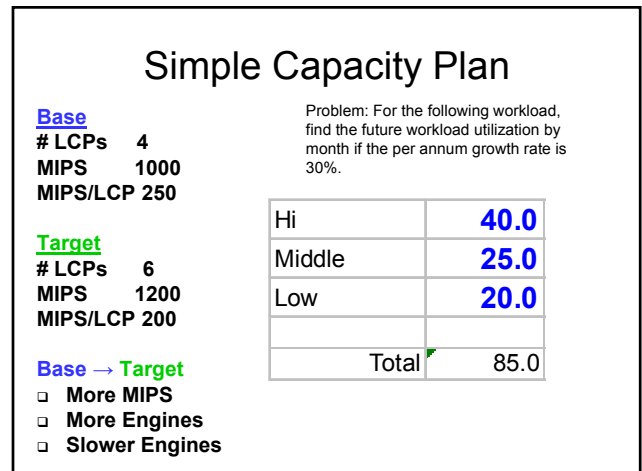
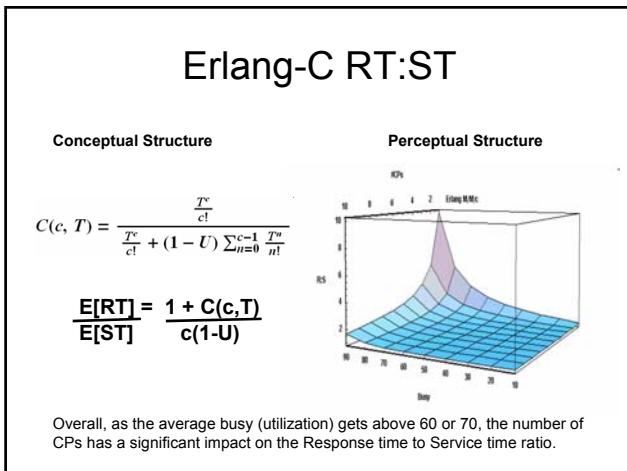


$$RT = 0.25 / (1 - .77)$$

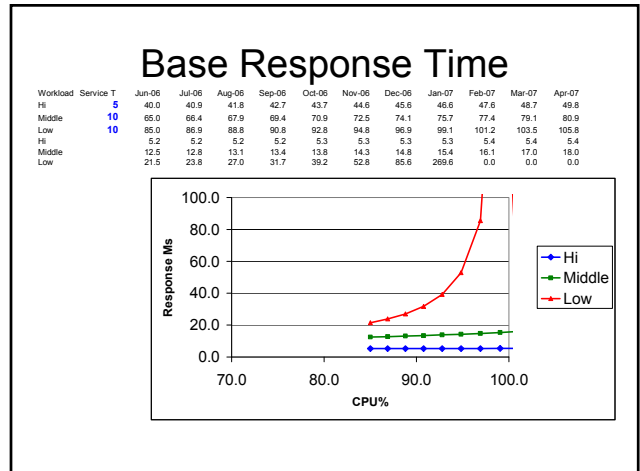

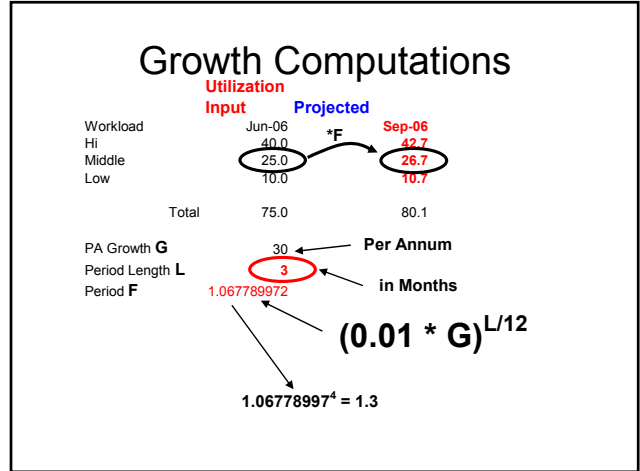
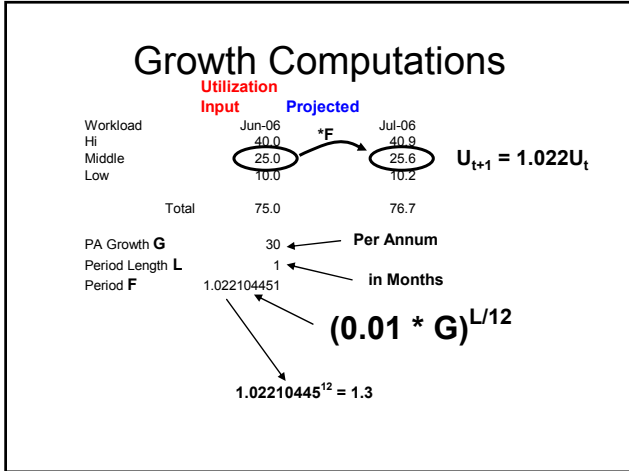
$$RT = 1.1$$

$$RT = ST / (1 - U)$$

$$1.2 = 0.25 / (1 - U)$$

$$U = 79%$$





Migration Options

- Migrate to same MIPS and more CPs.
- Migrate to same MIPS and fewer CPs.
- Migrate to more MIPS and fewer CPs.
- Migrate to more MIPS and more CPs.

Utilization?
Service Time?
Evaluation?

Computations

Base

LCPs 4
MIPS 1000
MIPS/LCP 250

$$UTILtarget = \frac{POWERbase * UTILbase}{POWERtarg} = \frac{1000}{1200} * UTILbase$$

Target


LCPs 6
MIPS 1200
MIPS/LCP 200

$$SERVtarget = \frac{PUSPEEDBase * SERVbase}{PUSPEEDtarg} = \frac{250}{200} * SERVbase$$

Target Utilization


Workload	Jun-07	Jul-07	Aug-07	Sep-07	Oct-07	Nov-07	Dec-07	Jan-08	Feb-08	Mar-08	Apr-08
Hi	33.3	34.1	34.8	35.6	36.4	37.2	38.0	38.8	39.7	40.6	41.5
Middle	20.8	21.3	21.8	22.2	22.7	23.2	23.8	24.3	24.8	25.4	25.9
Low	16.7	17.0	17.4	17.8	18.2	18.6	19.0	19.4	19.9	20.3	20.7
Total	70.8	72.4	74.0	75.6	77.3	79.0	80.8	82.5	84.4	86.2	88.1

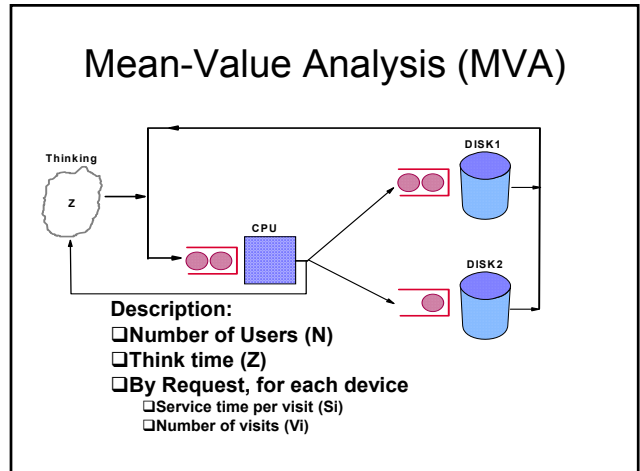
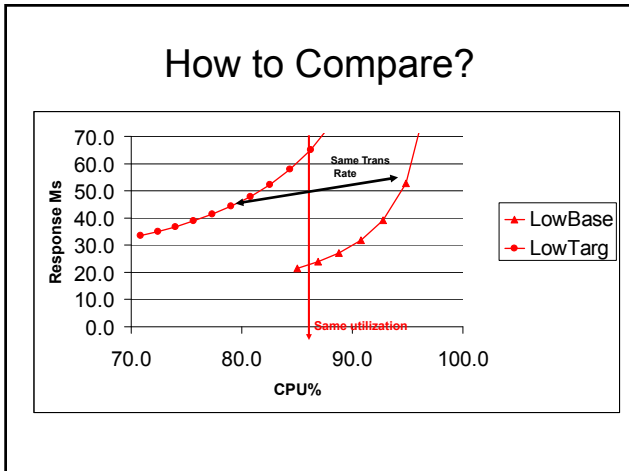

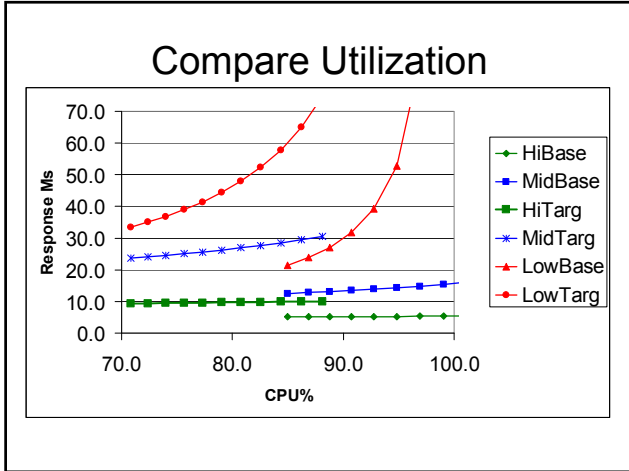
PA Growth 30
Period 1
Period F 1,022104
#CPs 2

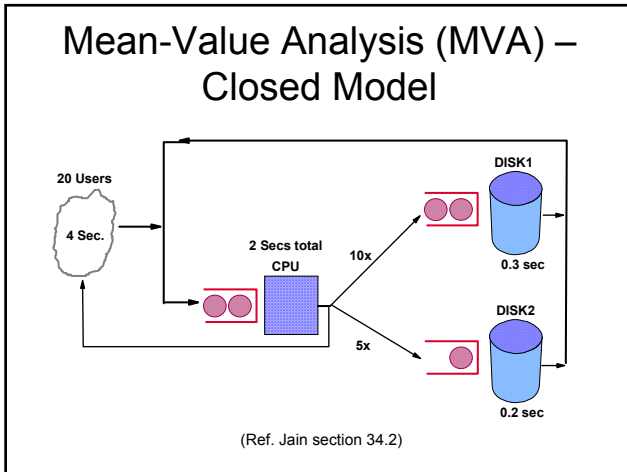


Target Response Time

Workload	Service T	Jun-07	Jul-07	Aug-07	Sep-07	Oct-07	Nov-07	Dec-07	Jan-08	Feb-08	Mar-08	Apr-08
Hi	8.333333	33.3	34.1	34.8	35.6	36.4	37.2	38.0	38.8	39.7	40.6	41.5
Middle	16.66667	54.2	55.4	56.6	57.8	59.1	60.4	61.8	63.1	64.5	65.9	67.4
Low	16.66667	70.8	72.4	74.0	75.6	77.3	79.0	80.8	82.5	84.4	86.2	88.1
Hi	8.4	8.4	8.5	8.5	8.6	8.7	8.7	8.8	8.9	9.0	9.1	9.2
Middle	23.6	24.0	24.5	25.0	25.6	26.3	26.9	27.7	28.6	29.5	30.5	
Low	33.4	35.0	36.6	38.9	41.4	44.4	47.9	52.3	57.8	65.0	74.7	







Minimum RT (1 user)

Number of Users (N)	1
Think time (Z)	4
Number of Devices (3)	2
Total CPU Service	0.1
CPU Service per Visit	0.3
Disk1 Service per Visit (S1)	10
Number of Visits (V1)	0.2
Disk2 Service per Visit (S2)	5
Number of Visits (V2)	

Minimum Residency (Response Time) is time without queuing.

Minimum RT = Total residency for CPU + Disk1 + Disk2
 $= 2 + 10 \cdot 0.3 + 5 \cdot 0.2 = 6.0$

Computations

For each server i:
 Ri = response time
 Si = service time
 Qi = queue length (included one in service)
 Vi = visits

$$R_i = S_i \cdot (1 + Q_i)$$

$$R = \sum_{i=1}^3 R_i \cdot V_i$$

$$\lambda = N / (Z + R) \quad \text{(Little's Law)}$$

Input

Number of Users (N)	20
Think time (Z)	4
Number of Devices	3
Total CPU Service	2
CPU Service per Visit	0.125
DISK1 Service per Visit (S1)	0.3
Number of Visits (V1)	10
DISK2 Service per Visit (S2)	0.2
Number of Visits (V2)	5

Minimum RT = Total residency for CPU + Disk1 + Disk2
 $= 2 + 10 \cdot 0.3 + 5 \cdot 0.2 = 6.0$

Algorithm

- Initialize $Q_i = 0$ for all i .
- Iterate by the number of users (N)


For $n = 1$ to N

- Iterate by the number of Service Centers (M)

For $i = 1$ to M

1. $R_i = S_i * (1 + Q_i)$
2. $R = \text{Sum_over_i} = R_i * V_i$
3. $\text{Thruput} = N / (Z + R)$
4. Set $Q_i = \lambda * V_i * R_i$

Results




	Iteration	CPU	Disk1	Disk2	System	Thruput	CPU QI	
	0						0	
	1	0.13	0.30	0.20	6.00	0.10	0.20	
	2	0.15	0.39	0.22	7.40	0.18	0.42	
	3	0.18	0.51	0.24	9.09	0.23	0.65	
	4	0.21	0.65	0.25	11.05	0.27	0.88	
	5	0.23	0.82	0.27	13.26	0.29	1.09	
	6	0.26	1.01	0.28	15.66	0.31	1.27	
Number of Users (N)	20	7	0.28	1.22	0.28	18.22	0.32	1.43
Think time (Z)	4	8	0.30	1.46	0.29	20.89	0.32	1.56
Number of Devices (S)		9	0.32	1.71	0.29	23.65	0.33	1.67
Total CPU Service	2	10	0.33	1.97	0.30	26.47	0.33	1.75
CPU Service per Visit	0.125	11	0.34	2.24	0.30	29.34	0.33	1.82
DASD1 Service per Visit (S1)	0.3	12	0.35	2.51	0.30	32.24	0.33	1.86
Number of Visits (V1)	10	13	0.36	2.80	0.30	35.18	0.33	1.90
DASD2 Service per Visit (S2)	0.2	14	0.36	3.08	0.30	38.13	0.33	1.93
Number of Visits (V2)	5	15	0.37	3.37	0.30	41.09	0.33	1.95
		16	0.37	3.67	0.30	44.06	0.33	1.96
		17	0.37	3.96	0.30	47.04	0.33	1.97
		18	0.37	4.26	0.30	50.03	0.33	1.98
		19	0.37	4.56	0.30	53.02	0.33	1.99
		20	0.37	4.85	0.30	56.02	0.33	1.99

Mean Value Analysis in Excel

Workload Modeling

Description	TrRate	CPU%/CR	Single CP%	MIPS	MIPS/Tr	DASD I/O	IO/Tr	MIPS/IO	IO Resp
DB2C	20.6	1.3	28.7	103.2	5.0	49.7	2.4	2.076	2.3




Flip Book Animation



Simulation an Alternate to Queuing Equations

- Initialize Model for T_0
- Schedule first event
- At each T_i for each Queue
 - Gather Stats on $T_i - T_{i-1}$
 - Busy? Population?
 - New Arrivals?
 - Start Counters
 - Departures?
 - Compute Behavior
 - Send to next queue
 - Schedule Next Event
- End Simulation
- Report



Simulation Demo(?)

A simulation is an imitation of some real thing, state of affairs, or process. The act of simulating something generally entails representing certain key characteristics or behaviors of a selected physical or abstract system.

Modeling Issues

- Analytic Queuing theory (and simulation) is difficult to apply in more than simple cases (Single server Unix).
- M/M/c can approximate (bound more complicated) cases of M/G/c/k cases. It's a good approximation at less than 100%.
- z/OS is complicated: WLM, priority, IRD, specialized PUs (zIIPs, zAAPs, IFLs).
- zSeries hardware behaves differently
- Packages & Services are available but it helps to know what's being done and what the terms mean.
- There are Single Task Multi Thread applications

Bibliography

The Art of Computer Systems Performance Analysis, by Raj Jain, Wiley. I like this one. It is thorough and complete. A very good reference. Hard to find.

Capacity Planning for Web Performance, by Daniel A. Menasce and Virgilio A.F. Almeida, Prentice Hall. A good book on queuing theory, network structure and terminology and introduction to the topic.

Probability, Statistics, and Queuing Theory, by Arnold O. Allen, Academic Press Inc. This is the classic in queuing theory.

Computer Performance Analysis with Mathematica, Arnold O. Allen, AP Press 1994. This is a clearly written textbook even if you don't have access to the Mathematica package.

Excel Data Analysis, by Jinjer Simon, Wiley Publishing. This is a step by step visual approach to data analysis with Excel.