Storage (RAM) in a Balanced System Session 431

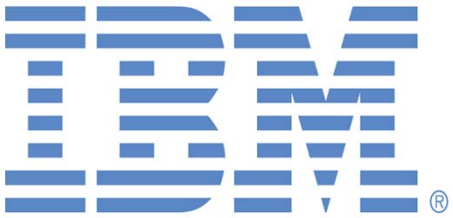

Ray Wicks
561-236-5846
RayWicks@us.ibm.com

Abstract

The forced flow law in a balanced system would lead one to think that the size of a resource, such as storage (RAM), could be a function of processor usage. This presentation looks at the balanced system set of resource ratios, storage in particular. Topics will include
*** Balanced systems metrics,**
*** Building a function Resource Usage = F(CPU Usage),**
*** New zSeries storage data in SMF 113 as a picture of storage usage,**
*** Is it possible to project storage usage for capacity planning?**

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.

Trade Marks, Copyrights & Stuff

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. Also foils added after I made handouts.**

A Plumbing Problem

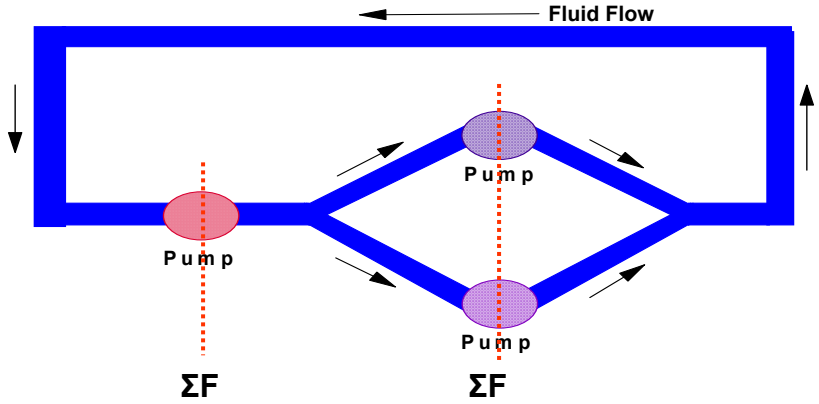

Conceptual Framework Forced Flow Law Flow Law

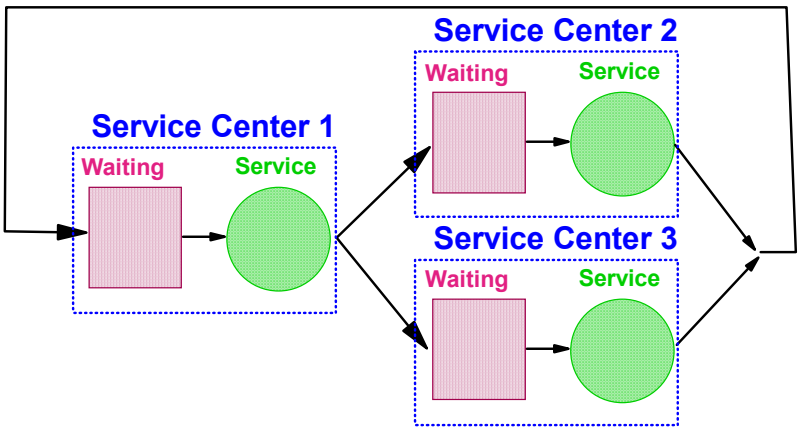

Conceptual Framework

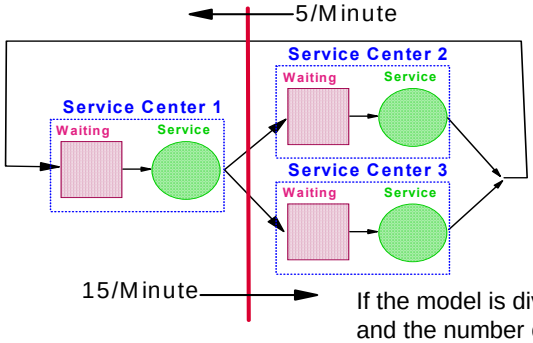

If the model is divided in two and the number of transactions crossing the boundaries is as indicated, what would the forced flow law say?

B.S. Conceptual Framework

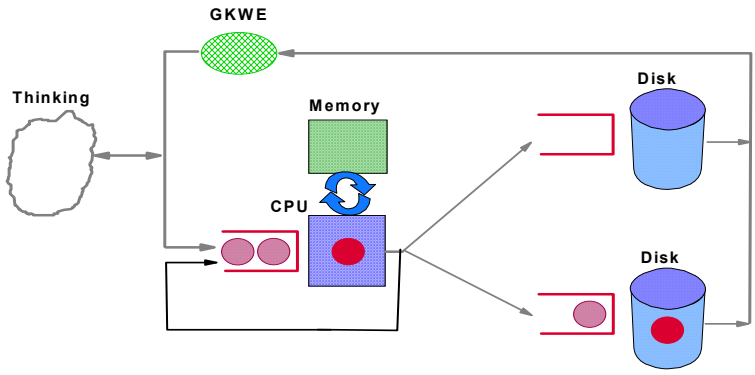

2010 z/OS B.S. Metrics

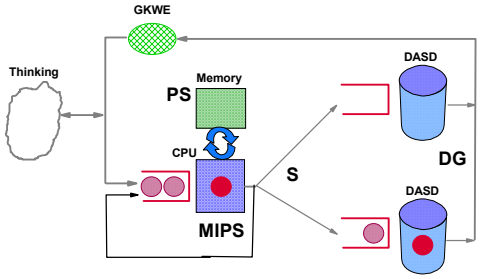

MIPS used
S = SSCH rate
DG = DASD gigabytes. The computation is nominal in that it is 2.83/act
PS = Central Storage configured

	10%	50%	90%
S/MIPS	1.201	2.349	3.707
DG/MIPS	2.236	6.593	52.539
PS/MIPS	6.766	14.055	37.306

S/MIPS Samples

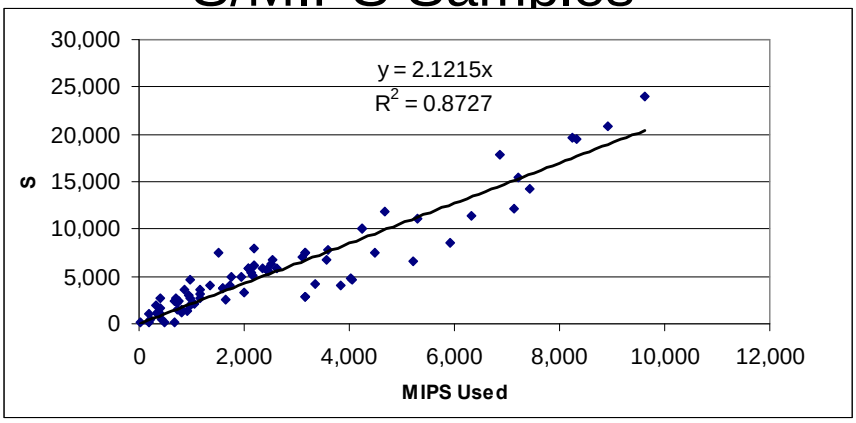

MIPS Used and SSCH are in proportion (correlated).

Good linear relationship.

DG/MIPS

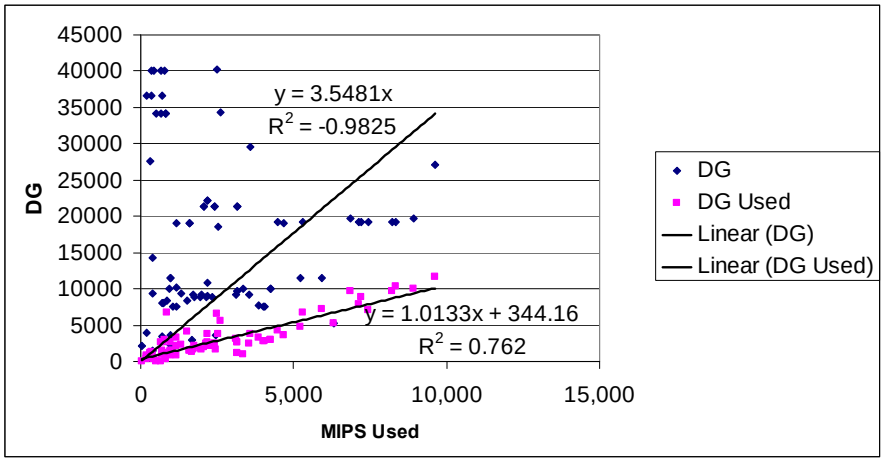

PS/MIPS

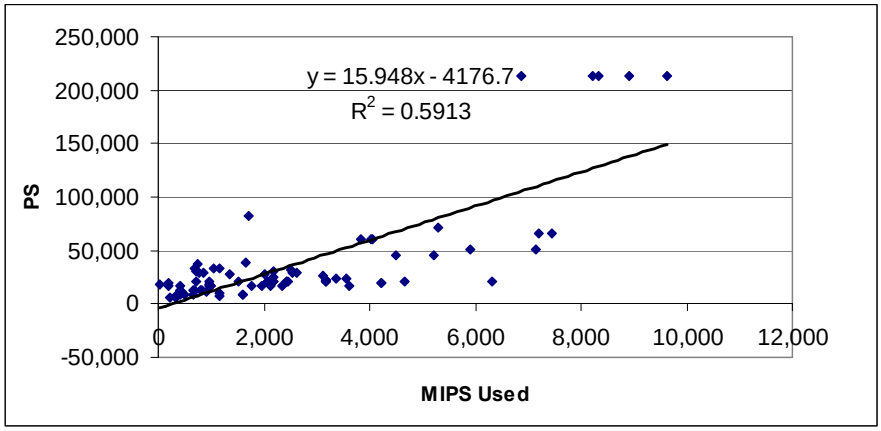

MIPS Used and PS are not in proportion (correlated).

More Metrics Defined

MIPS used

S = SSCH rate

DG = DASD gigabytes. The computation is nominal in that it is 2.83/act

DG = Nacts * 2.83

CS = Central Storage configured

D as in CS = 4000 + 0.04(MIPS)^D

where CS (or PS) is configured Central Storage (or RAM)

NNTacts = acts with rate >=2

AD = Access Density = S/DG

Full Metrics

2010 83 Partitions			
	10%	50%	90%
MIPS	403	2004	6247
S	1123	4221	11779
S/MIPS	1.201	2.349	3.707
DG/MIPS	2.236	6.593	52.539
PS/MIPS	6.766	14.055	37.306
D	1.534	1.630	1.894
DASD Resp	0.952	1.865	3.626
DASD Serv	0.681	1.564	2.980
Resp/Serv	1.227	1.827	1.232
Nacts	1305	4084	12105
NNTacts	73	264	882
DASD GB	3693	11558	34256
Used DG	585	2114	7065
AD	0.059	0.360	0.909

If the Model Works, What Should I See?

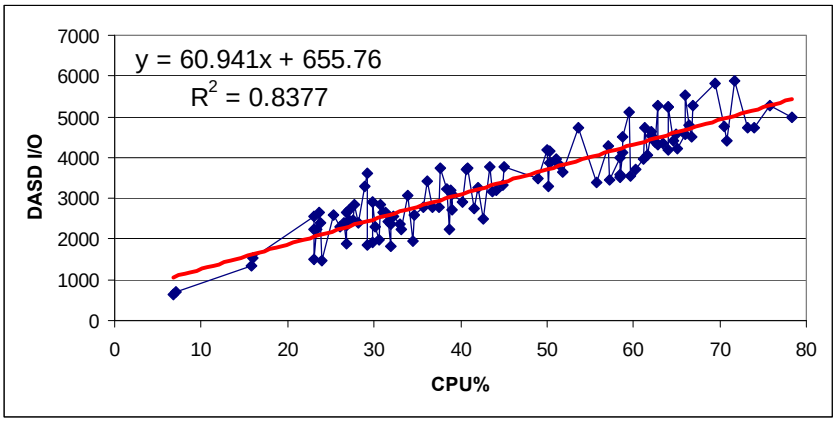

If the model is valid, the forced flow law would prescribe that the variation of CPU service would be proportional to the I/O service. In this graph that translates into a linear relationship (R^2>0.7?).

When It Doesn't Work

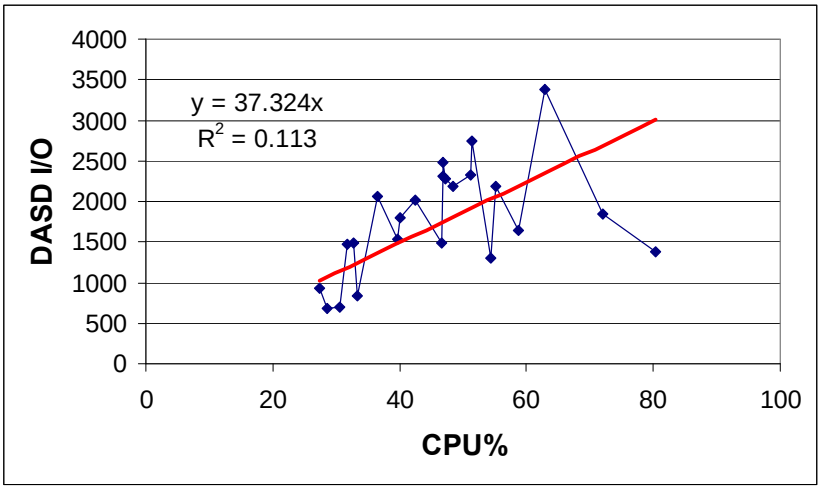

If the model Doesn't work, you would see a non-linear relationship.

RIOC=F(Workload Composition)

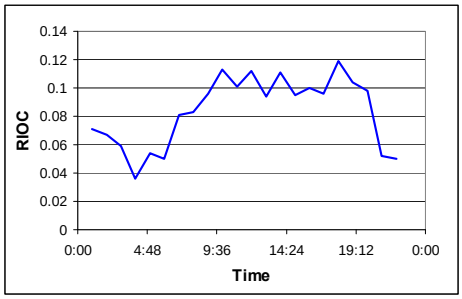

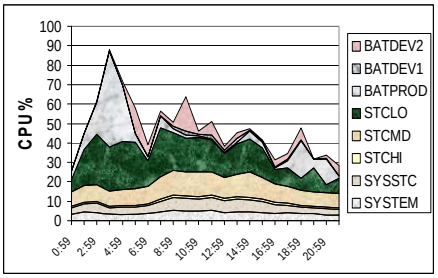

Compare those intervals in the RIOC plot where the RIOC is stable or unstable with the workload mixture for the same intervals in the CPU% plot. The workload mixture often explains the RIOC variation.

Use of Conceptual Framework like BS Model

- **Establishes Relationships**
  - **Balanced System, resource ratios**
- **Builds Expectations**
  - **Linear graph**
- **High lights Exceptions**
  - **Non Linearity**
  - **Outliers**
- **Generates Questions**
  - **Especially if not as expected**
- **But... it may cause you to see Framework (model) interactions that just aren't there!**

Another Example

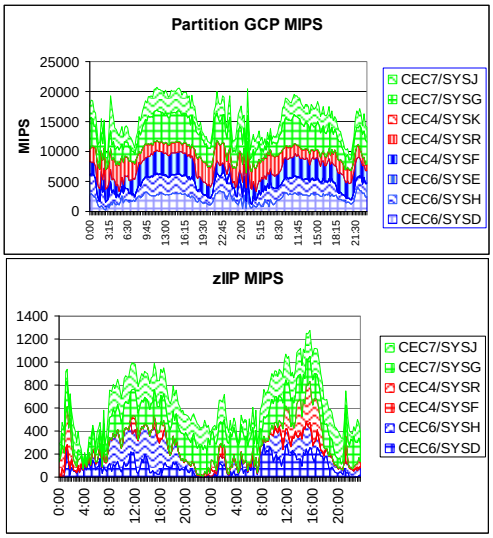

CEC 6 Util

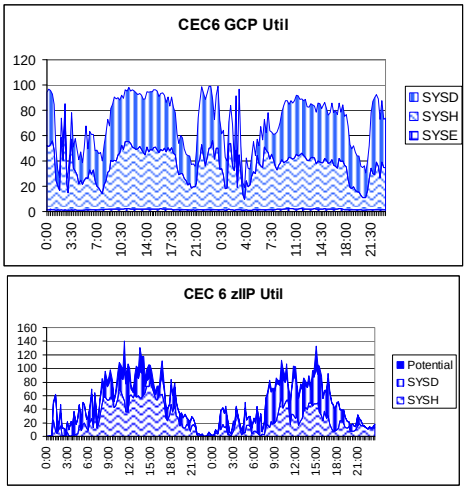

SYSD zIIP = F(GCP)?

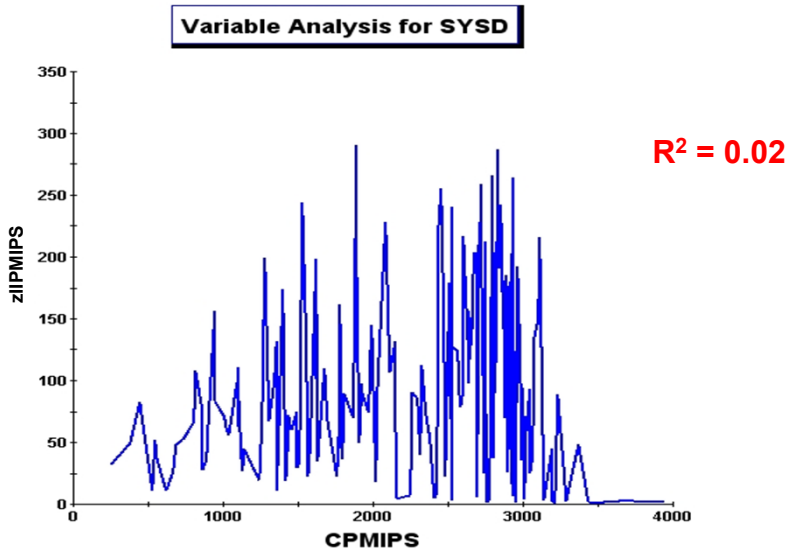

Workload zIIP = F(GCP)?

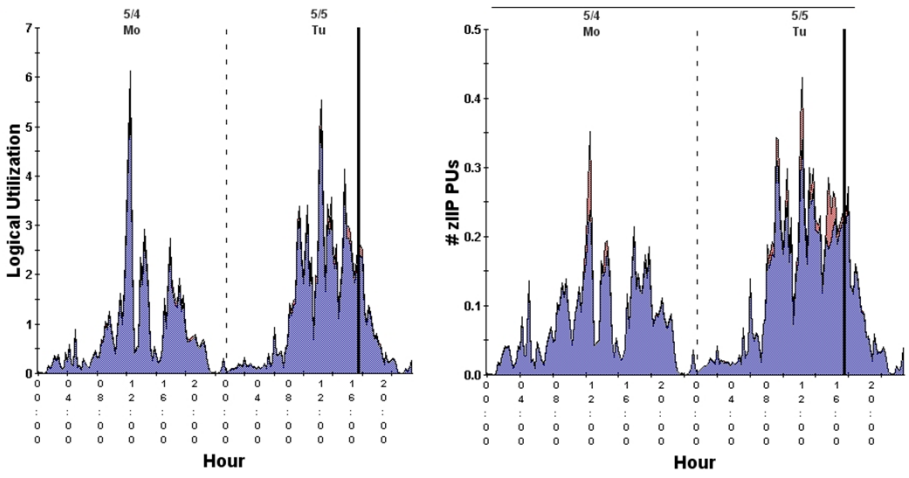

SYSD Workload zIIP = F(GCP)?

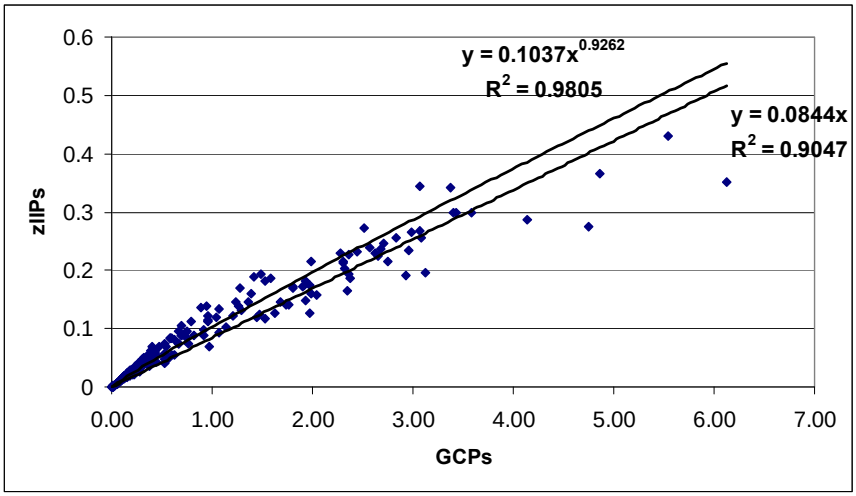

Processing Unit Logic

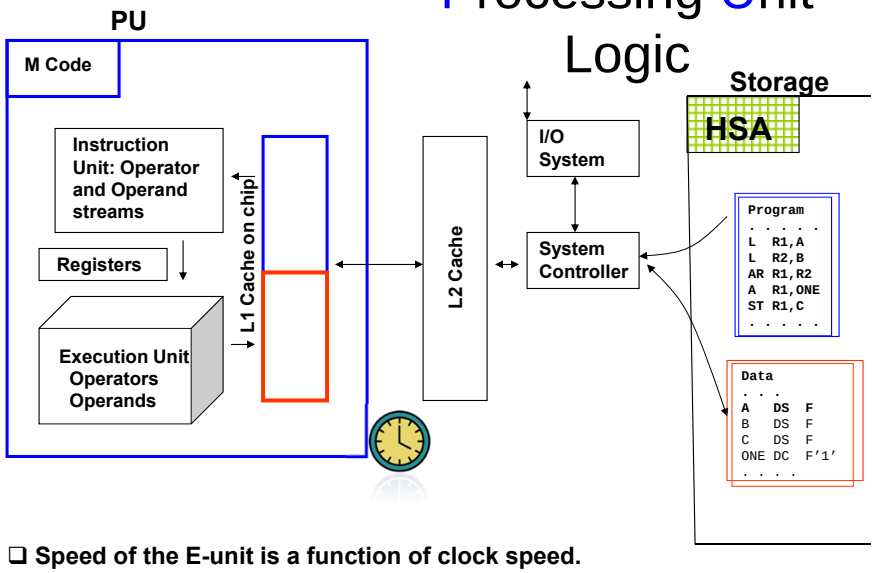

- Speed of the E-unit is a function of clock speed.
- Effective E-Unit speed is a function of the data location.

Processing Unit Logic

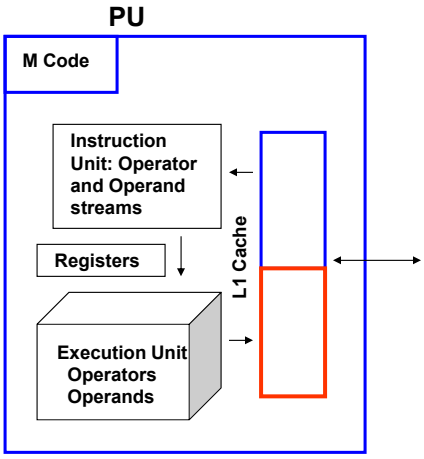

Roughly:

1. **Instruction fetch (from??) cycle**
2. **Instruction decode/Register fetch cycle**
3. **Data fetch cycle**
4. **Execution/Effective address cycle**
5. **Memory access cycle**
6. **Write-back cycle**

- **Speed of the I & E-unit is a function of clock speed.**
- **Effective E-Unit speed is a function of the architecture & data location.**

Z10 Memory – a simple view

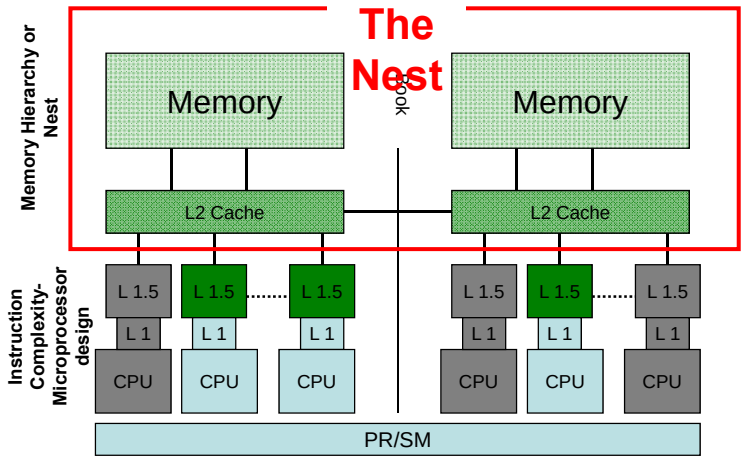

Reference: John Burg's presentation at SHARE 3/3/2011

http://www/ibm.com/support/techdocs/atsmastr.nsf/Webindex/TC000066

z196 versus z10 hardware comparison

- z10 EC
  - CPU
    - 4.4 GHz
  - Caches
    - L1 private 64k i, 128k d
    - L1.5 private 3 MB
    - L2 shared 48 MB / book
    - book interconnect: star

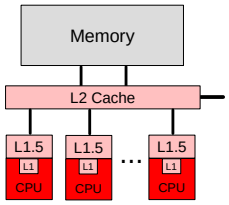

- z196
  - CPU
    - 5.2 GHz
    - Out-Of-Order execution
  - Caches
    - L1 private 64k i, 128k d
    - L2 private 1.5 MB
    - L3 shared 24 MB / chip
    - L4 shared 192 MB / book
    - book interconnect: star

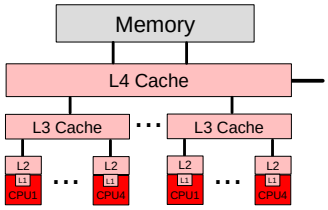

27

Introducing the new Relative Nest Intensity (RNI) metric

(SMF 113 Data)

Relative Nest Intensity

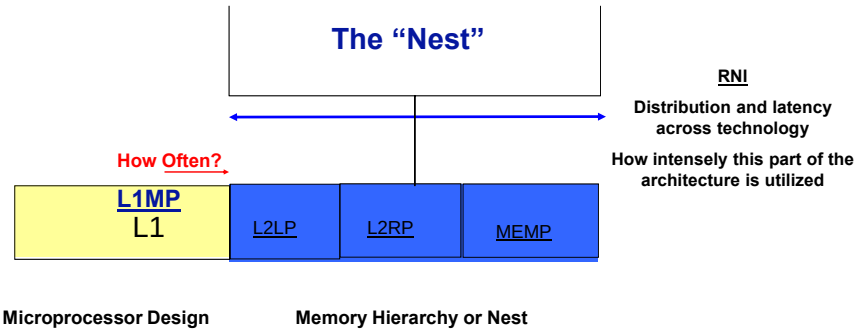

Note these Formulas may change in the future

Level 1 (L1) Miss Percent

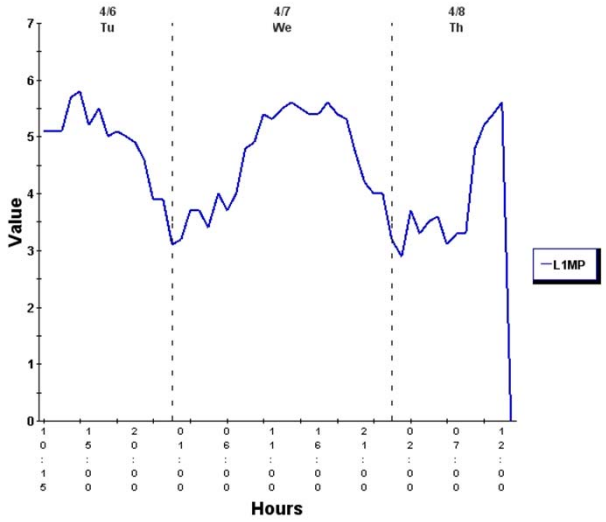

If not from L1, from Where?

(SMF 113 Data)

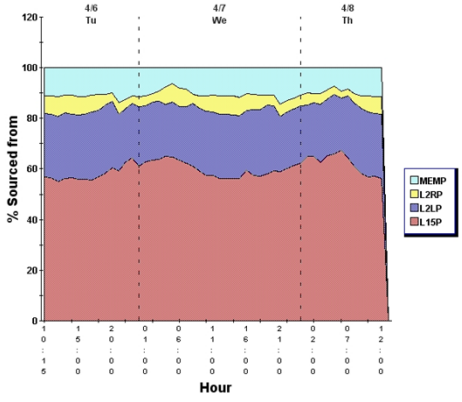

Here's the plot of percent sourcing from different levels of cache. As the sourcing moves from the highest level of cache (percent=L15P) to the slowest memory source (percent=MEMP), the performance degrades. Level 1 cache is the fastest and closest to the processing unit.

The sourcing shown in the graph is for data not found in level 1 cache. You can check the level 1 cache miss % by graphing variable L1MP.

- L1.5P=%sourced from level 1.5 cache
- L2LP=%sourced from level 2 cache same book
- L2RP=%sourced from level 2 cache different book
- MEMP=%sourced from memory

Remember than as more and more of the instructions and data has to be fetched from more distance caches, the machine effectively runs slower.

sas | Enterprise Guide
The Power to Know.

2097-E54 Sample

1 With Variables: CPI
5 Variables: L1MP L15P L2LP L2RP MEMP

Simple Statistics

Variable	N	Mean	Std Dev	Sum	Minimum	Maximum	Label
CPI	220	6.70909	0.82921	1476	3.90000	8.60000	CPI
L1MP	220	4.21091	0.61397	926.40000	1.90000	5.60000	L1MP
L15P	220	68.80273	2.73381	15137	62.50000	79.00000	L15P
L2LP	220	24.60364	1.92242	5413	17.70000	27.90000	L2LP
L2RP	220	0.09636	0.01876	21.20000	0	0.10000	L2RP
MEMP	220	6.51364	1.18811	1433	3.20000	10.40000	MEMP

Pearson Correlation Coefficients, N = 220
Prob > |r| under H0: Rho=0

	L1MP	L15P	L2LP	L2RP	MEMP
CPI	**0.79329**	-0.53761	0.41295	0.09019	**0.56472**
CPI	<.0001	<.0001	<.0001	0.1826	<.0001

Cycles per Instruction

- ❑ **Instruction complexity**
- ❑ **Memory activity**

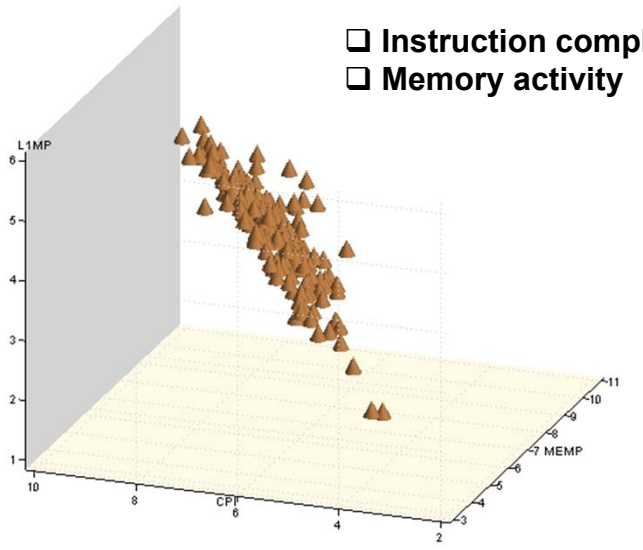

Stepwise Regression

Stepwise Analysis
Table of Results for General Stepwise

L1MP entered.

	df	*SS*	*MS*	*F*	*Significance F*	***Rsquare***
Regression	1	94.76340452	94.76340452	370.1004889	7.13123E-49	**0.629315051**
Residual	218	55.81841366	0.256047769			
Total	219	150.5818182				

	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*
Intercept	2.197521837	0.236981894	9.272952449	1.82816E-17	1.730452906	2.664590768
L1MP	1.071400255	0.055691885	19.23799597	7.13123E-49	0.96163681	1.181163699

MEMP entered.

	df	*SS*	*MS*	*F*	*Significance F*	***Rsquare***
Regression	2	144.0333909	72.01669544	2386.469636	1.8304E-148	**0.956512497**
Residual	217	6.548427298	0.030177084			
Total	219	150.5818182				

	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*
Intercept	-0.432760399	0.104193631	-4.153424666	4.70956E-05	-0.638121488	-0.22739931
L1MP	1.07847415	0.019120016	56.40550326	1.2808E-131	1.040789434	1.116158866
MEMP	0.399238545	0.009880522	40.40662553	6.1031E-103	0.379764469	0.418712621

Predict CPI in terms of Memory Activity

$CPIfit = a_0 + a_1 * L1MP + a_2 * MEMP$

	Coefficients
Intercept	-0.43276
L1MP	1.078474
MEMP	0.399239

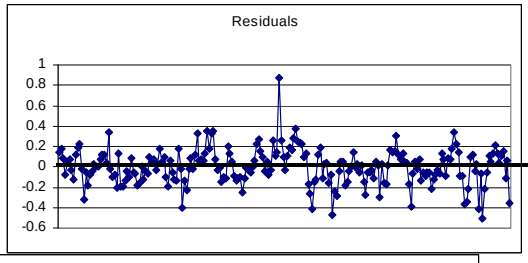

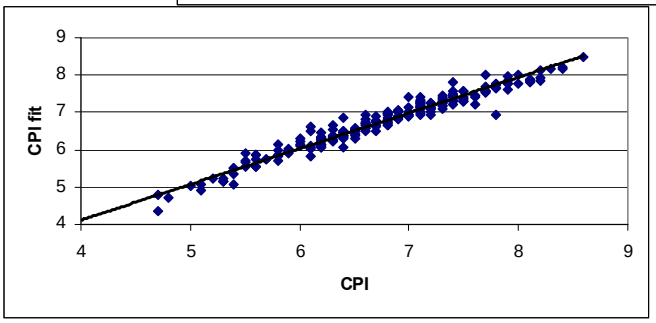

Relative Nest Intensity (RNI)

z10 RNI=(1.0*L2LP + 2.4*L2RP + 7.5*MEMP) / 100

Z196 RNI=1.6*(0.4*L3P + 1.0*L4LP + 2.4*L4RP + 7.5*MEMP) / 100

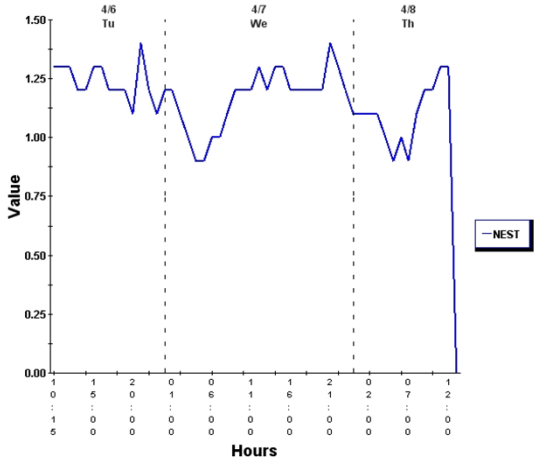

Memory and Workload Characteristics

L1MP	RNI	Workload Hint
<3%	>= 0.75 < 0.75	AVERAGE LOW
3% to 6%	>1.0 0.6 to 1.0 < 0.6	HIGH AVERAGE LOW
>6%	>=0.75 < 0.75	HIGH AVERAGE

Note that these are initial values and may change.

Storage Projection

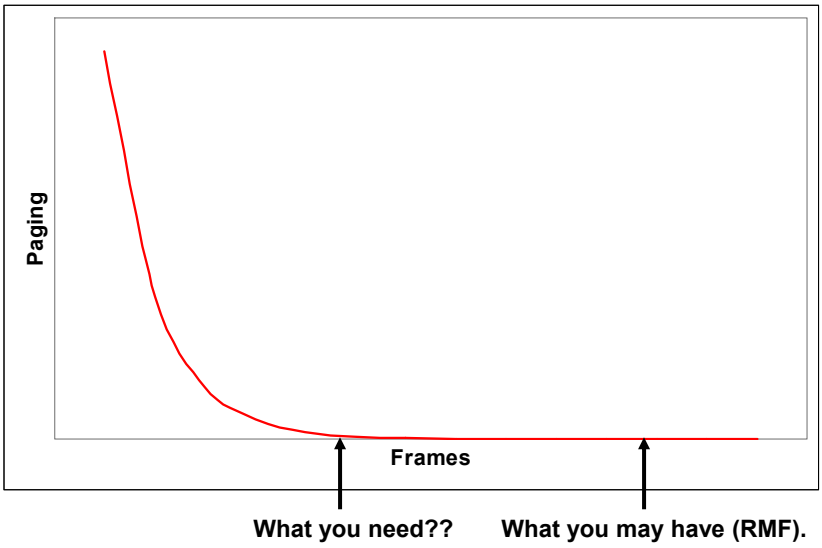

Storage = F(CPU%)?

Who is Using Storage?

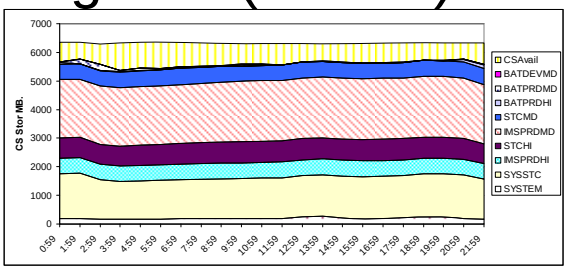

Who is Using CPU?

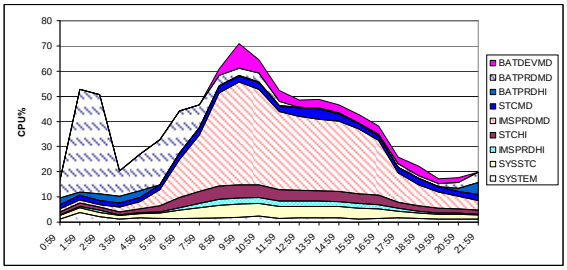

Expectations for Batch

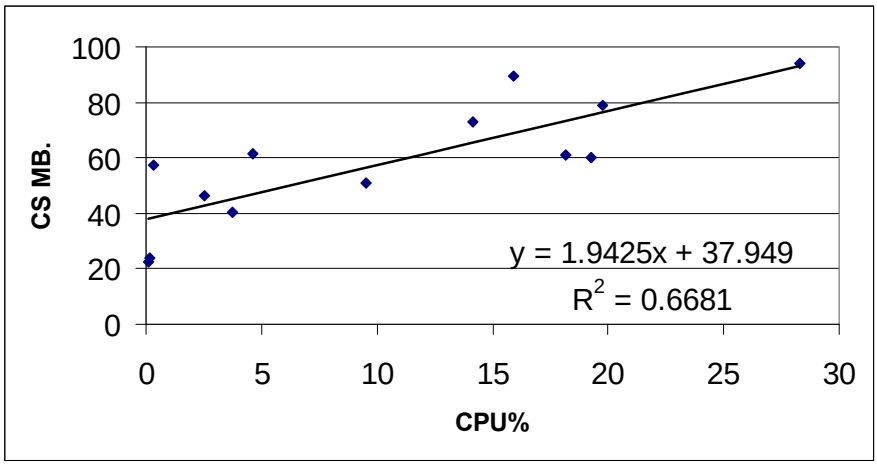

Expectations for DB/DC

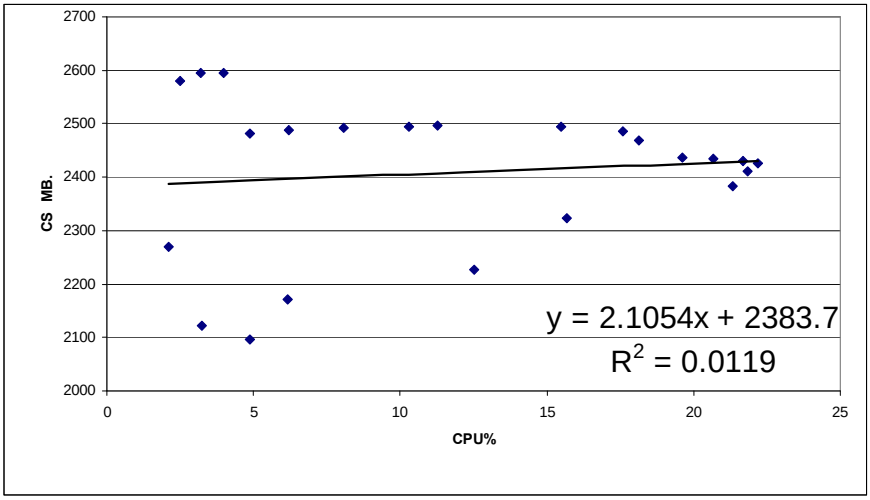

DB/DC Scaled to 0

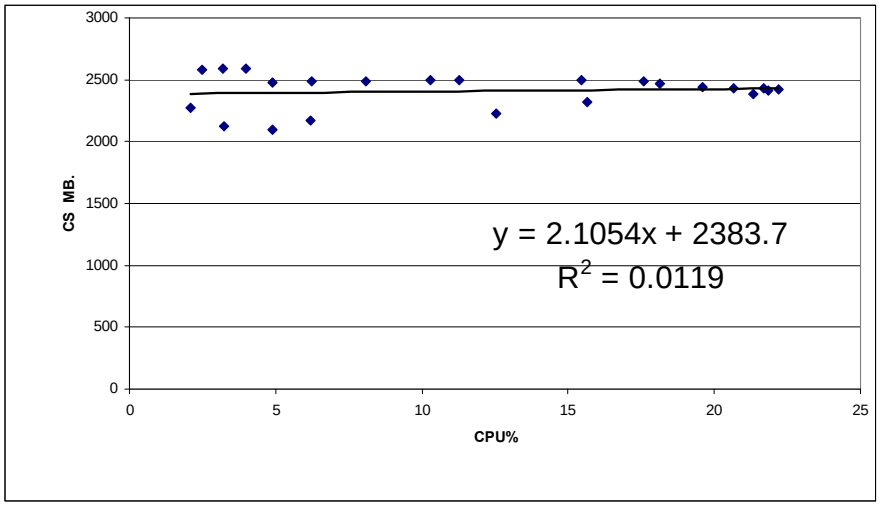

Note: R^2 doesn't change whereas perception does.

Storage Growth

- **Size depends upon the number of address spaces**
- **If growth ➡ More Address Spaces ➡ more storage**

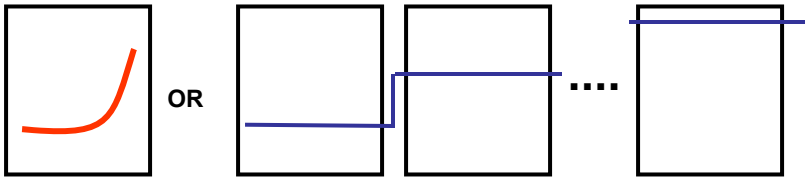

Some applications storage requirements grow as the load grows. Typically it's those applications where growth means more address spaces (TSO, batch). DB/DC often has workload growth without an increase in storage. Only after the number of address spaces increase does the storage change.

RMF Partition Report

CENTRAL STORAGE

	MIN	MAX	AVG
AVAILABLE	1450630	1917651	1638857
SQA	17,478	18,443	17,840
LPA	9,842	9,842	9,842
CSA	43,884	43,946	43,913
LSQA	178,761	180,850	179,946
REGIONS+SWA	4379708	4845902	4658493
TOTAL FRAMES	6553600	6553600	6553600

FIXED FRAMES

NUCLEUS	2,629	2,629	2,629
SQA	15,528	16,493	15,890
LPA	90	90	90
CSA	12,187	12,187	12,187
LSQA	65,380	67,270	66,451
REGIONS+SWA	60,380	73,810	61,791
BELOW 16 MEG	76	98	78
BETWEEN 16M-2G	38,193	39,699	39,097
TOTAL FRAMES	157,643	170,842	159,039

Sample Partitions

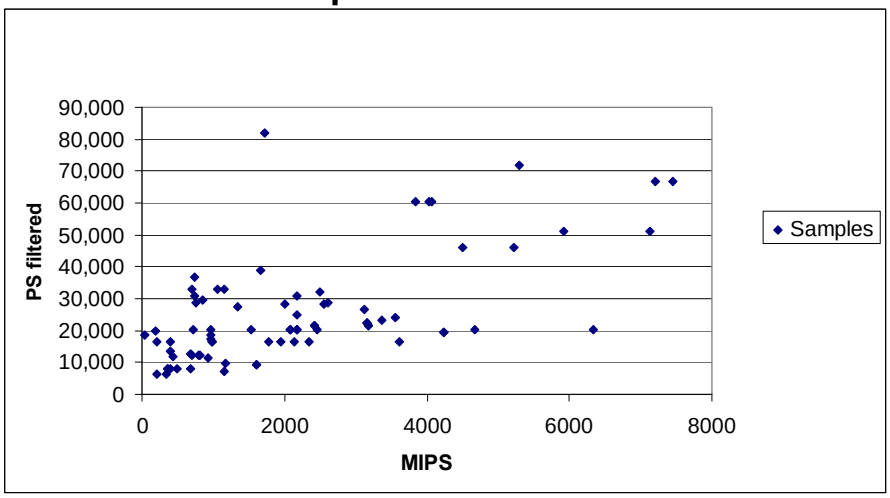

Projection Method 1

Use an exponential function $Y = a + b*X^p$

For any partition CS = 4000 + 0.04*(MIPS_Used)^D

Known Known

Therefore, in EXCEL D=LOG(((CS-4000)/0.04),10)/LOG(MIPS_Used,10)

Projection: Approach 1

Use and exponential function CS=4000 + 0.04*(MIPS_Used)^D

Known Known

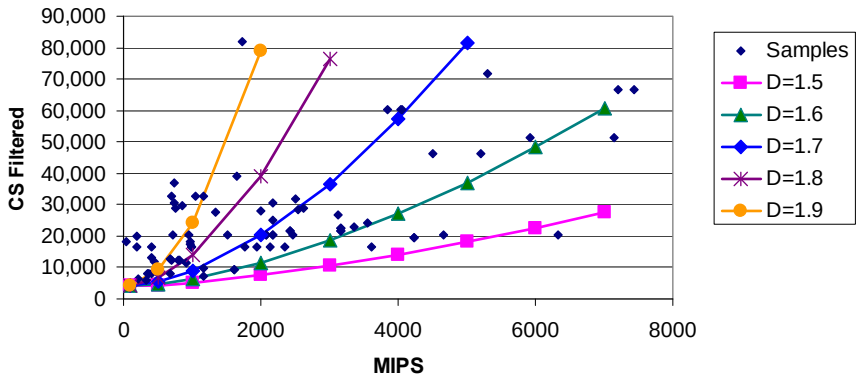

Projection Approach 1

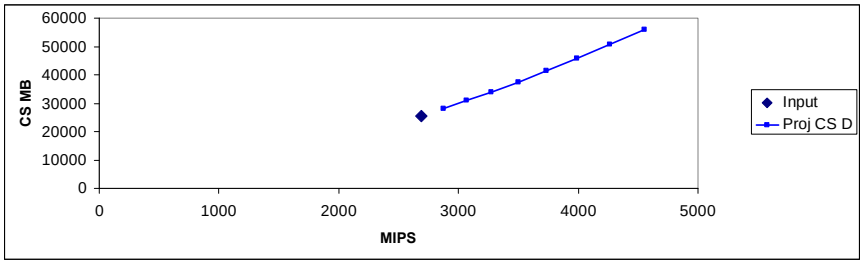

CS=25600, MIPS=2689
IF 25600 = 4000 + 0.04*(2689)^D
Then D=LOG(((25600-4000)/0.04),10)/LOG(2689,10)
D= 1.671451
D metric = [1.534, 1.630, 1.894]
Apply 30% per annum growth to MIPS
Use formula CS=4000+0.04*(MIPS)^1.671451 for future MIPS to project CS.

Projection Approach 2

Use ratio CS/MIPS
CS=25600, MIPS=2689; CS/MIPS = 9.52
CS/MIPS metric is [10%, 50%, 90%] = [6.77, 14.06, 37.31]
CS = 9.52 * MIPS

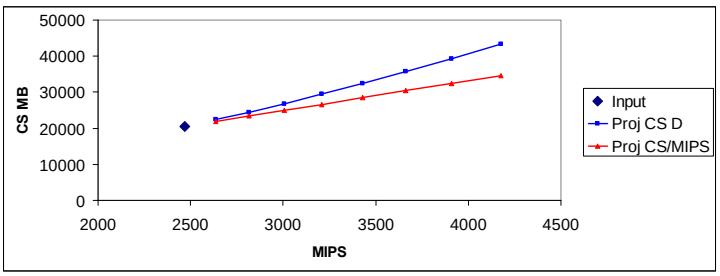

	Quarter	CS	MIPS	Ratio	Proj CS D	Ratio	Proj CS/MIPS	Ratio	D	CS/MIPS
Input	0	25600	2689		25600		25600		1.671451	9.520268
	1		2872	1.068	28111	1.098	27341	1.068		
	2		3067	1.068	30913	1.100	29200	1.068		
	3		3276	1.068	34041	1.101	31186	1.068		
	4		3498	1.068	37533	1.103	33306	1.068		
	5		3736	1.068	41431	1.104	35571	1.068		
	6		3990	1.068	45782	1.105	37990	1.068		
	7		4262	1.068	50638	1.106	40573	1.068		
	8		4552	1.068	56059	1.107	43332	1.068		

Note 6.8% quarterly growth is 30% per annum.

I'm a ~~Fuzzy~~ Thinker