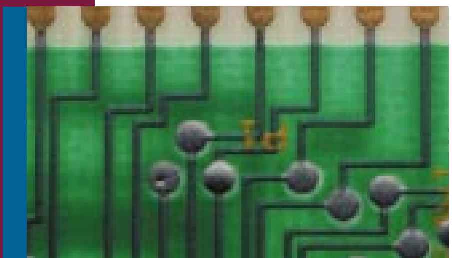

BY CHERYL WATSON

Vendor Performance Ratings: Why Your Results May Differ – Part II

It's important to understand that no measurement exists that can provide a single rating for a processor model that is indicative of its speed and capacity for a variety of workloads. However, these vendor claims can be used for sizing a machine.

PART I of this series (***Technical Support***, December 1997) presented my definition of various terms, examined the use of vendor claims and why you should use them, and described how vendors meet their claims. This article examines IBM's LSPR, Amdahl's and HDS's performance ratings, and why you should use these claims. Parts III and IV will examine why your installation's experience may differ and conclude with some tips on what you can do if this happens.

IBM'S LSPR

To confirm the speed and capacity of their processor models and to help customers understand what to expect from different models processing their workloads, IBM publishes its *Large Systems Performance Reference* [REF001] as a manual and a performance tool. You can also find the LSPR numbers on the Web [REF002]. The manual contains both the techniques and the results; I strongly recommend that you become familiar with IBM's methodology. Below, I have provided my summary and observations of IBM's 50-page discussion on the technique.

Workloads

IBM has designed and accumulated a series of traditional workloads it feels are representative of customers' workloads. The workloads consist of the following:

CB84 - Commercial Batch Workload: This set of 130 jobs with 610 unique steps provides a typical, traditional view of batch work. This workload consists of COBOL, Assembler H, and PL/I programs, along with compilers and utilities such as DFSORT. Access methods for BSAM, QSAM, BDAM, and VSAM are used. This is representative of the traditional batch applications running in most installations today.

CBW2 - Commercial Batch Workload 2: This set of jobs was created with the introduction of SP 4.2.2. It has 32 jobs with 157 steps and is representative of new applications that exploit such ESA functions as data in memory. It consists of programs written in C, COBOL, FORTRAN, and PL/I. The steps perform sorting, use DFSMS utilities, compilers, VSAM and DB2 utilities, SQL processing, SLR processing, GDDM graphics, and FORTRAN engineering routines. There is more JES processing, and about 50 percent of workload time is spent performing DB2 activities.

FPC1 - Engineering/Scientific Batch Workload: This workload is an engineering and manufacturing jobstream that includes static analysis, dynamic analysis, computational fluid dynamics, nuclear fuel calculations, and circuit analysis. This will be representative of much of the SAS work in commercial installations due to SAS's heavy use of floating point.

TSO: The TSO workload is representative of TSO program development using ISPF/PDF. This includes editing, browsing, foreground compiles, testing, graphics, and info/management. Twenty-five different scripts are used and driven by an internal driver to meet the activity required to drive the system to 70 percent or 90 percent utilization. Although IBM periodically uses TPNS to confirm the consistency of its own internal driver, TPNS is not used in this case.

CICS: In SP 4.2, the CICS workload consisted of 102 transactions, and in SP 5.1, the workload consists of 204 transactions. CICS is run in an MRO (Multiple Region Option) configuration with a TOR (Terminal Owning Region), an AOR (Application Owning Region), and a FOR (File Owning Region). In SP 4.2.2, an additional AOR/FOR region was added. As many of these "MROplexes" are run as needed to achieve 70 percent and 90 percent utilization, usually one MROplex per one or two CPUs. COBOL and Assembler are used for the programs, and VSAM is the primary access method. The work is designed to be representative of order entry, stock control, inventory tracking, production specification, hotel reservation, banking, and teller systems.

IMS: The IMS workload is similar to the CICS workload from DLI applications. There are 17 transaction types. Enough Message Processing Regions (MPRs) are run to bring the system to the desired utilization (70 percent and 90 percent) without causing contention between MPRs. DLI HDAM and HIDAM access methods are used with VSAM and OSAM databases. In SP 5.1, two IMS control regions are used, and data sharing occurs using the IMS Resource Lock Manager (IRLM). BMPRs (Batch Message Processing Regions) are not included.

DB2: The DB2 workload consists of seven transactions applied to two applications: inventory tracking and stock control. The DB2 requests are driven by IMS/DC. Enough regions are created to eliminate contention within the subsystems. There are two DB2 databases comprised of 11 tables for inventory and five tables for stock control, with one to five indexes for each table. Since the two workloads don't invoke DB2 sorts, the DB2 sort assist feature available on some models is not exercised.

Workload Testing

Only one type of workload is run during an LSPR test and the systems are run at fairly high CPU utilization (close to 100 percent for batch and FPC1, and at both 70 percent and 90 percent utilization for onlines and TSO. For the online work (e.g., CICS, IMS, and DB2), the IBM team waits until the system has stabilized before starting the measurement phase. Two very important items to note are that only one type of workload is run in each test and the tests are run in totally unconstrained environments. That is, CICS is not tested with TSO and IMS is not tested with batch. Also, in order to accurately determine the effect of the processor capacity, IBM must ensure that no other constraints exist on the system. That is, there is virtually no paging due to the abundance of all types of storage. There is no I/O constraint (almost 100 percent cache). There isn't a lack of VTAM buffers or JES initiators, and even the CPU is not run until it is constrained (it is never run at over 100 percent busy).

> It is fairly dangerous to try to compare service units between models from different vendors, and it's also dangerous to compare service units between models of widely different ages.

Measurements

From the measurements made while running these benchmarks, IBM calculates an Internal Throughput Rate (ITR) that is equal to the units of work (jobs or transactions) divided by the processor busy time. Models with higher capacities – thus, higher ITRs — will be able to process more work in the same amount of processor busy time as models with lower capacities.

Each workload will have its own ITR. To be able to compare two models, IBM uses an ITRR, Internal Throughput Rate Ratio that is calculated by taking the ITR for a base model and dividing it into the ITR for the new model. Thus, a model that can process 50 percent more work in the same amount of CPU time as compared to the 520 will have an ITRR of 1.5. Prior to June 1997, IBM published a list of the ITRRs using their 9021-520 as a base model with the ITR for each 520 workload being set to 1.0.

In June 1997, IBM published preliminary LSPR performance comparisons for their newest models using the CMOS 9672-R15 as a base. In August 1997, they republished their LSPR ratings for all models using the R15 as the new base. These new ratings were quite a bit different than the 520 ratings because the operating system and subsystem releases used in the LSPR runs were changed at the same time. This led to more than a little confusion. If we take IBM's statement that the R15 is equivalent to the 9021-711, and we also accept the 711 as a 62 MIPS machine, all other machines would see a corresponding 2 percent to 6 percent increase in MIPS ratings based on the new LSPR ratings!

Figure 1 shows an extract from IBM's LSPR charts for three models as compared to the 9672-R15. You can interpret the chart as saying that their TSO workload achieved 4.40 times as many transactions in the same amount of processor busy time on the 9672-R55 as compared to the 9672-R15. This is based on the total capacity of the model, not necessarily the speed of a CPU, as we'll see later.

In order to help people consider the capacity based on a mix of workloads, IBM derived an estimated ITRR called #MIX that consists of 20 percent of the ITRR of each of the five workloads: CICS, IMS, DB2, TSO, and CB84. This is a calculated value only, and is not confirmed by running 20 percent of each workload, which would be next to impossible to achieve consistently.

Figure 1: IBM LSPR ITRRs [from REF001 & REF002]

Machine	TSO	CICS	DB2	IMS	CB84	CBW2	FPC1	#MIX
9021-982	6.00	7.23	6.07	6.18	7.13	7.20	6.95	6.48
9672-R15	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
9672-R55	4.40	4.34	4.02	3.98	4.28	5.02	4.86	4.20
9672-RX5	6.48	6.61	5.92	6.11	6.76	9.09	9.61	6.36

How These Measurements Are Used

The #MIX, or an early expectation of #MIX, is used to derive the SRM service unit coefficient, the service units per second as published for each model. The SU/Sec is used by SRM to compensate for different speed CPUs when determining the frequency of invoking certain functions. The SU/Sec used to be a fairly good indicator of CPU speed because it is related to the speed of a single CPU. A CPU with an SU/Sec of 400 is roughly twice the speed of a CPU with an SU/Sec of 200. For several reasons, this number is becoming less effective as an indication of CPU speed.

First, the SU/Sec value is published and made available often before final LSPR tests have been completed. While the published ITRRs might change, the SU/Sec values are seldom changed.

Second, in older models, the difference in speed between workloads was fairly close. With modern processors, the difference in speed between workloads can be more than 30 percent. As an example, in Figure 1, the FPC1 workload on a 9672-RX5 has an ITRR (9.61). That's more than 60 percent higher than DB2 (5.92), and more than 50 percent higher than #MIX (6.36). It would be very difficult to use a single number to indicate the speed of the RX5 for these differing workloads. There's a 14 percent variation in just the five workloads used to derive the #MIX.

Changes After General Availability

If after General Availability (GA) of a model there are significant performance improvements through microcode or other means, IBM has indicated that they will rerun the test and republish the changed ITRRs. However, they do not expect to alter the SU/Sec values, the processor group ratings, or the MSU ratings.

AMDAHL PERFORMANCE CLAIMS

Amdahl has a set of internal benchmark jobs similar to IBM's, but they do not publish a description of their workloads or specific performance claims for each type of workload. They normally publish a range of performance that can be expected for a given model compared to their 5995-4570M. For example, their newly announced CMOS Millennium series contains a model GS745, which is listed as having a performance rating of 1.16 to 1.28 of the Amdahl 5995-4570M.

Since Amdahl does not publish their workloads, we can't be certain which workloads are at which end of the range, although we might expect them to be similar to IBM's workloads. Most analysts take the midpoint of the high and low to be the average and relate that to IBM's #MIX workload. Whether this is valid or not remains to be seen.

> There is no measurement in existence that can provide a single rating for a processor model that is indicative of its speed and capacity for a variety of workloads.

Amdahl has always derived their SU/Sec value a little differently than IBM, however. Their logic has been to provide consistent TSO response across a hardware change. In order to do this, the same percent of TSO transactions need to complete in first period. For this to be true, the durations must be adjusted to match the CPU speed. Amdahl assigns a value to the SU/Sec to ensure that the same percent of TSO transactions complete in first period. This has meant that the Amdahl SU/Sec values for bipolars have been 6 percent to 8 percent higher than corresponding IBM and HDS bipolars. The Amdahl models had SU/Sec values that resulted in calculations of about 52 SU/Sec per MIPS, while IBM and HDS had closer to 48 SU/Sec per MIPS.

With CMOS models, however, the vendors are getting closer to 51 SU/Sec while the Amdahl models vary from 48 to 52 SU/Sec (with a strange anomaly in the GS535 that results in almost 55 SU/Sec per MIPS). This means two things to you: it is fairly dangerous to try to compare service units between models from different vendors, and it's also dangerous to compare service units between models of widely different ages.

HDS PERFORMANCE CLAIMS

HDS uses two techniques for publishing performance ratings. Two series of HDS models, the GXxxxx series and their CMOS Pilot models, are designed to be directly competitive to corresponding IBM models and therefore use comparable IBM ratings. The Skyline models, which are based on the fastest CPU speed available today, are not comparable in speed to any IBM or Amdahl CPU, so HDS publishes separate ratings for the Skylines (as well as a few other models that don't have corresponding IBM matches).

The HDS models that are comparable to the IBM models are published as having "equivalency" to the IBM models, and HDS offers performance claims that are equivalent to IBM's. For the few models in these series that do not have a direct equivalent model within the IBM range, HDS publishes a performance range; e.g., one model might provide 1.2 to 1.4 times the performance of an HDS GX8110.

As mentioned, the Skyline models, which are really combinations of bipolar and CMOS technology, don't relate to an IBM model, but performance claims are published that indicate, for example, that a Skyline is 2.0 times faster than the HDS GX8114. HDS has derived these performance claims by running their own set of benchmark jobs. However, neither a description of the jobs or the resulting measurements are published.

I've noticed that Skyline SU/Sec values range from 48 to 52 SU/Sec per MIPS, so the SU/SEC values might appear higher or lower than service units from other vendors.

WHY YOU SHOULD USE THESE CLAIMS

The Bad News

It's important to understand that no measurement exists that can provide a single rating for a processor model that is indicative of its speed and capacity for a variety of workloads. This is similar to buying a car based on expected mileage. A car might be rated for 20 miles to the gallon, but that is seldom what you will achieve. You will drive the car much different than the testers who came up with the initial rating. For example, if you happen to have a lead foot (i.e., drive too fast!), you'll never achieve the mileage your car is rated for. If you drive it according to the testers' recommended speeds, and in their type of traffic, and on their types of roads, and with the same amount of weight in the car, and with all of the extra equipment turned off, you might be able to come close to achieving their estimate. The same is true of processor models. With that said, however, I strongly recommend that you use the vendor's claims for sizing a machine because it as close as you can get initially.

Performance Guarantees

I also believe that you should not obtain any hardware without some contractual commitment from the vendor about the performance that you expect to receive from the processor model. Since I know that it's possible to obtain a performance guarantee from a vendor (but also know that it won't be offered unless asked for), I'd recommend that with every installation, you plan to obtain such a guarantee. These guarantees can only be based on the vendor's claims; therefore, I think you should trust the vendor to provide the right capacity estimates, but do get it in writing!

The trick in any contract is to identify how you and the vendor will agree to the performance that you're getting. This often requires very knowledgeable people on both sides who can understand the difference in performance since your workloads may not match the vendor's.

The trick in any contract is to identify how you and the vendor will agree to the performance that you're getting.

Industry Charts Based on Vendor's Claims

Since most of the industry charts of MIPS are based on vendor's claims, almost every company is indirectly using what the vendor has provided. In Part III, I'll discuss why your results may differ from your vendor's.

REFERENCES

[REF001] *Large Systems Performance Reference*, IBM, SC28-1187

[REF002] Web site for LSPR Description and Numbers - **<http://www.s390.ibm.com/sections/lspr.html>** ts

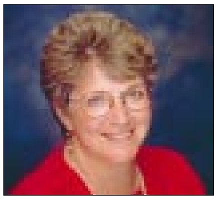

NaSPA member Cheryl Watson is an internationally-known consultant and author who has been working with IBM mainframes since 1965. She is the president of Watson & Walker, Inc., and is the principal author of Cheryl Watson's TUNING Letter, *a journal of OS/390 advice published since 1991. She will soon be releasing BoxScore, the first in her ADVISOR Series of low-cost performance analysis software. Visit her web page at www.watsonwalker.com.*