SIZING A FEDERATED COMPUTER SYSTEM

Joseph B. Major

IBM Canada Ltd.

SUMMARY

Discrete event simulation of a federated hardware-software system under a single resource manager is performed. Workload projections, system logic and performance information are merged into an integrated model. Data Base/Data Communication. Time-Shared and Batch work under OS/VS2 on a multisystem is modeled. System feasibility and equipment trade-offs are studied for stand-alone and concurrent subsystem operations. Statistical model validation and calibration thru benchmarks/measurements is discussed.

INTRODUCTION

In late 1972 IBM Canada began work on a comprehensive proposal. Preliminary analysis indicated that a combination of /370-158's would meet the performance requirements. Scarcity and incompleteness of OS/VS2 and /370-158 performance data at the time has significantly reduced the confidence in the accuracy of rough sizing estimates.

Quantitative analysis to present system resource utilization and turnaround time projections was required. Discrete event simulation of the likely systems seemed to offer multiple advantages:

A. accurate representation of stated probability distributions, reduced need to simplify system logic;
B. ease of implementation under tough delivery constraints;
C. ease of communicating results;
D. increased credibility of design.

THE SYSTEM

Capability to concurrently process batch (local, remote), interactive (time-shared) work and On-Line-Real-Time (OLRT) transactions and associated batch utilities that access OLRT databases was required. Asymmetric CPU configurations under a single resource manager expected to offer the solution. CPU.A (/370-158) would normally service regular batch and interactive jobs. CPU.B (/370-158, uni-or multiprocessor) would handle OLRT and related jobs. Software logic of OS/VS2-ASP, TSO and IMS had to be simulated.

The following types of data were given:

A. observed frequency distributions of job interarrival and turnaround times, of jobstep dataset, main storage, CPU time, physical I/O (sequential, indexed, direct) requirements, of CPU times between I/O's;
B. average print/punch/read loads;
C. OLRT transaction groupings, peak and off-peak loadings, CPU and logical I/O requirements, related database extents and dependencies, message lengths, application descriptions, program sizes;
D. interactive work definition (SHARE script, (1), p.102), workload, think time distribution.

Stated resource utilizations did not include operating system overhead (CPU, I/O) nor definitions of same. Thus, substantial performance data extrapolations were called for.

SIMULATION APPROACH

System-subsystem relationships lead to a natural division of work. Data reduction and model construction could proceed in parallel on each subsystem: batch (VS-ASP), interactive (TSO) and OLRT (IMS). Subsequent system integration necessitated uniformity of design, notations, naming conventions, subroutine usage and documentation methods.

GPSS-V (cf.(2)).was chosen as the simulation vehicle. Modular programming, subprogram integration thru the use of the parallel, asynchronous event flow were the natural components to build on. Even so, the overall effort required 68 man-weeks of work.

To investigate contention for various system resources and to accurately project workload component turnaround times simulation of individual I/O events and associated CPU processing was undertaken. Available measurement information allowed for representation of unique CPU servicing steps even in greater detail. Such was the choice for the on-line components (TSO, IMS) as well as for the scheduling logic of ASP. Dataset definitions were based more on judgement than on factual evidence. Major components of system overhead were estimated on the basis of limited measurements.

The objective was to identify system components that would result in significant queuing delays due to contention. Detailed representation of these was contrasted to cheaper, rougher models.

Transaction (job) processing was viewed as the sum of successive queuing steps. Aggregation of these would result in fewer and longer steps with total queue (turnaround) time differing from that of the non-aggregate case. The practicality of service steps aggregation into macro steps was investigated.

Hardware operation considerations, such as rotational position sensing on disk drives, contention by multiple CPU's for disk (drum) control units were covered. Various stipulations on service time distributions were investigated. The effect of dataset placement in terms of device/cylinder definitions was monitored in order to realistically assess seek time delays for moveable head devices.

Model details that resulted in slight perturbations only on the key projections were removed from the production runs. Model predictions were more sensitive to changes in data interpretation than to degree of detail chosen.

VALIDATION

Subsystem runs with sequences of antithetic random numbers were performed. The on-line components showed stable results for practical run lengths even at very high CPU utilization (92-95%). The skewed nature of batch characteristics, however, could not be reliably studied in less than 55-60 hours of simulated system operation. That was impractical in terms of available computer power and elapsed time. Thus, batch performance results were considered to be transient.

Parameter variations gave rise to performance predictions that could be 'rationalized' thru practical knowledge of data-processing systems.

Some component calibration was possible before finalizing the production runs. A significant bug - omission of CPU time usage by logical database calls - was found when model predictions were compared to live IMS (OS/MVT) measurements. Results were very consistent while in error. Antithetic tests showed that the system was stable. Without actual measurements the results would have been of no value. Subsequent measurements (/370-158, VS2 system) have confirmed the reasonability of model predictions.

CONCLUSIONS

Having worked with the model, we find in retrospect that:

A. performance projections are more sensitive to input data interpretations than to system logic simplifications - economies of representation must be, nevertheless, considered and justified;
B. implementation of discrete event models can be very time-consuming, particularly so if deadlines are to be met;
C. effective communication of results requires substantial documentation effort;
D. credibility of design is not improved by the 'accuracy', cleverness or neatness of the model, only by subsequent validation measurements.

ACKNOWLEDGMENT

F. Friedmann and J. Hebert of IBM Canada have significantly contributed to most phases of the simulation effort.

BIBLIOGRAPHY

1. Doherty, W.J. 'Scheduling TSS/360 for responsiveness', Fall Joint Computer Conference, 1970.
2. General Purpose Simulation System V, User's Manual, SH20-0851, IBM.