A major purpose of the Technical Information Center is to provide the broadest dissemination possible of information contained in DOE's Research and Development Reports to business, industry, the academic community, and federal, state and local governments.

Although a small portion of this report is not reproducible, it is being made available to expedite the availability of information on the research discussed herein.

1

LA-UR-83-3396

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

LA-UR--83-3396

DE84 003[illegible]

TITLE AN INTERACTIVE BENCHMARK COMPARISON BETWEEN A VAX 11/780 AND A VAX 11/782

NOTICE

PORTIONS OF THIS REPORT ARE ILLEGIBLE.

It has been reproduced from the best available copy to permit the broadest possible availability.

AUTHOR(S) Olaf Lubeck
Ronald Martinez

SUBMITTED TO Paper to be Presented in the Proceedings of the
1983 Fall DECUS U.S. Symposium
Las Vegas, NV
October 24-28, 1983

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes.

The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

Los Alamos Los Alamos National Laboratory
Los Alamos, New Mexico 87545

AN INTERACTIVE BENCHMARK COMPARISON BETWEEN A VAX 11/780 AND A VAX 11/782

Olaf Lubeck and Ronald Martinez*
Los Alamos National Laboratory
Los Alamos, New Mexico

ABSTRACT

In this paper, we present the results of a benchmark comparison between a single processor VAX 11/780 and a dual processor VAX 11/782. Using the VMS Monitor utility and available accounting data, we first parameterized the workload on the VAX 11/782. We developed synthetic scripts that matched both the current workload parameters and also generated four forecasted workloads. Using a remote terminal emulator (RTE), we submitted interactive user commands from the scripts to the single and dual processor VAXs. Response times were then measured as a function of script category and VAX model.

INTRODUCTION

The VAX 11/782 is a dual processor VAX with two identical 780 processors used in an asymmetric fashion with a common memory. The asymmetric nature of the processors lies in the way that the operating system schedules the CPUs. The primary processor is the only one allowed to run in kernel mode. User processes are scheduled on the attached processor before the primary. However, if a process running on the attached processor makes an I/O request or system service request (for example, page faults), then the primary processor is interrupted and must service the request since it alone operates in kernel mode.

The two processors do not cooperate to make a single job execute faster, rather the dual processor VAX should exhibit a higher throughput of varying degree (as compared with a VAX 11/780) dependent on the type of workload. It is expected that a workload that is more CPU intensive would experience a higher throughput than an I/O intensive workload.

WORKLOAD CHARACTERIZATION

The first and most important step of any benchmark must be to properly characterize the current workload. Using the VMS monitor utility and available accounting information, we gathered data during prime shift for a typical 5-day work week. The number of active users during an 8 to 5 day was found to vary between 5 and 12; the average number of users was 8. Table 1, gleaned from accounting data, shows the most frequently used DCL commands, the processes with highest CPU usage, and the processes with highest page fault rates. The information is indicative of a general Fortran development environment but also shows that the INGRES database package is often used. Also, VAXIMA (a LISP routine that manipulates mathematical equations symbolically) is a large CPU intensive process and page faults heavily.

*This work performed under the auspices of the US Department of Energy contract W-7405-ENG-36

TABLE 1 Accounting Data Summary

Frequently Used DCL Commands	Major CPU Processes	Major Page Faulting Processes
DIRECTORY	USER1	VAXIMA
DELETE	USER2	USER1
SEARCH	MASS	FORTRAN
SET	SHOW	MAPPER
SHOW	VAXIMA	
COPY	MAPPER	
TYPE		
EDT		
TEDI		
LINK		
FORTRAN		
INGRES		

Table 2 shows averages of selected monitor parameters over the 5-day period that was characterized (column 1). CPU usage is given in units of a single 780 processor. It is clear that the overall workload is quite compute intensive with a 93% CPU utilization. The main reason for this high utilization is that there were CPU bound jobs in low-priority batch queues that were always present and would soak up any additional CPU time left by the higher priority interactive processes.

TABLE 2 Monitor Data Comparison

Parameters	Actual Workload	Scenario
Direct I/O	7.9	7.3
Buffered I/O	10.6	10.3
Page Faults	26.0	34.8
Page Read Rate	7.4	10.8
Page Write Rate	2.3	2.6
System CPU	0.24	0.36
User CPU	1.62	1.64
Idle CPU	0.14	0.0

SYNTHETIC SCRIPTS

Based on the information obtained in the workload characterization, five separate scripts were generated. Table 3 summarizes the utilities and functions that were performed within each script.

TABLE 3. Synthetic Script Features	
Script 1:	TEDI, FORTRAN, LINK, SUBMIT TEDI, FORTRAN,LINK, RUN (CPU bound)
Script 2:	FORTRAN, LINK, RUN (I/O bound)
Script 3:	INGRES database queries
Script 4:	EDT, FORTRAN, LINK TYPE, DIR, SHOW, COPY, SEARCH, ETC.
Script 5	VAXIMA

Script 1 edited, compiled, linked, and submitted two batch jobs and then executed a short CPU bound program. Script 2 compiled, linked, and executed an I/O bound job. Script 3 executed a Fortran program that queried an INGRES database. Script 4 contained miscellaneous DCL commands to approximate the frequency of usage listed in Table 1. Script 5 used VAXIMA to perform various symbolic manipulations.

The scripts were combined into a 'scenario' and the scenario was subsequently tuned to match the average monitor parameters found in the workload characterization (second column of Table 2). Using this scenario as a basis, we then produced four 'forecasted' scenarios such that each represented an increased workload. Table 4 shows the make-up of each scenario.

TABLE 4 Scenarios	
Scenario I	8 users
Mix	2 script 1, 1 script 2, 2 script 3, 3 script 4
Scenario II	16 users
Mix	4 script 1, 2 script 2, 4 script 3, 6 script 4
Scenario III	8 users (1 VAXIMA)
Mix	2 script 1, 1 script 2, 1 script 3, 3 script 4, 1 script 5
Scenario IV	16 users (1 VAXIMA)
Mix	4 script 1, 2 script 2, 4 script 3, 5 script 4, 1 script 5
Scenario V	16 users (2 VAXIMA)
Mix	4 script 1, 2 script 2, 4 script 3, 4 script 4, 2 script 5

EXPERIMENTAL RESULTS

Table 5 lists the physical configuration of the two VAXs used in this experiment.

The major system generation parameters were identical during the benchmark runs. Additionally, since the configurations were similar but not identical, we also ran Scenarios IV and V on the VAX 11/782 system with one processor disabled. Figures 1-5 show the results from execution of scenarios I through IV on both computer systems. For our compute intensive workload, the dual

TABLE 5. Physical Configuration	
VAX 11/782:	two processors with floating-point accelerators 4 Mbytes of memory three RP07 disks DZ-11 Communications Boards
VAX 11/782:	one processor with floating-point accelerator 5.5 Mbytes of memory one RP07 disk, one RP06 disk DZ-11 Communications Boards

processor VAX had roughly 80% faster response times. The difference between single and dual processors decreased slightly as the workload increased. The exception was Scenario V where 16 users (two of which were executing the VAXIMA script) caused such a memory-I/O bottleneck that there was virtually no advantage to having the second processor. Another anomalous result was that the VAXIMA script performed no better on the dual processor as compared with the single processor (Figure 3) in one case, and response times were degraded in two other cases (Figures 4 and 5). To ascertain that this result was not caused by the difference in machine configurations, we ran these last two scenarios on a single processor VAX 11/782 where one processor was disabled. Figures 6 and 7 depict these measurements and show that the same anomalous result was obtained.

Our explanation for the degradation seen by the VAXIMA scripts lies in the fact that it page faults heavily. If the VAXIMA script has been scheduled on the attached processor, each page fault causes an interrupt and system service request to be issued to the primary processor. Thus, we reason that system overhead required to satisfy the page fault request is higher on the dual processor than on the single processor.

CONCLUSIONS

We have characterized our current workload on a VAX 11/782 as highly compute intensive and found that it experiences 80% faster response times than if it were run on a VAX 11/780. However, a very reasonably forecasted workload representing twice as many users (16) and including two active VAXIMA scripts would not run on the currently configured VAX 11/782 any faster than the single processor VAX 11/780. We were not able to examine the effect of larger memory size. Also, regardless of the number of users, we have found a routine VAXIMA (whose use is increasing among our VAX 11/782 user community) that did not experience faster throughput on the VAX 11/782 and in fact could experience slower throughput. We attribute this behavior to the added latency on the VAX 11/782 in servicing a process that page faults heavily.

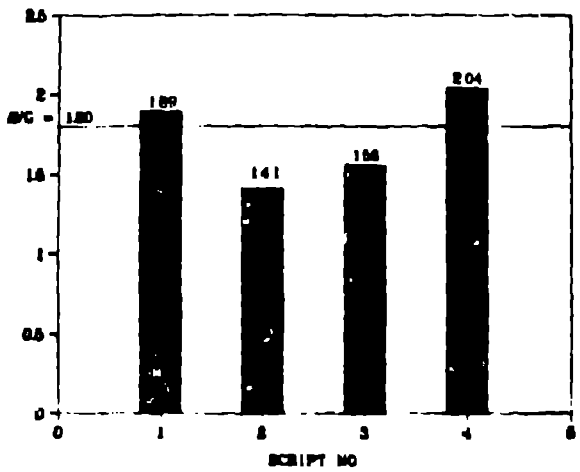

Figure 1. The ratio of 780 to 782 response times for each script category in Scenario I.

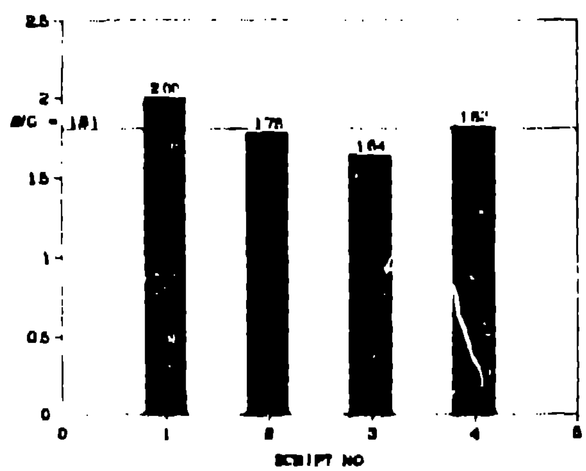

Figure 2. The ratio of 780 to 782 response times for each script category in Scenario II.

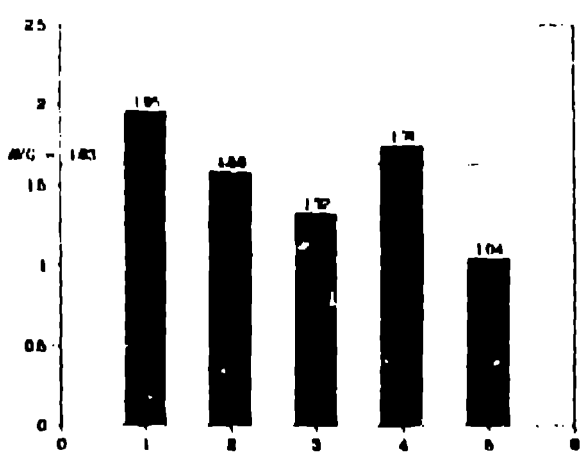

Figure 3. The ratio of 780 to 782 response times for each script category in Scenario III.

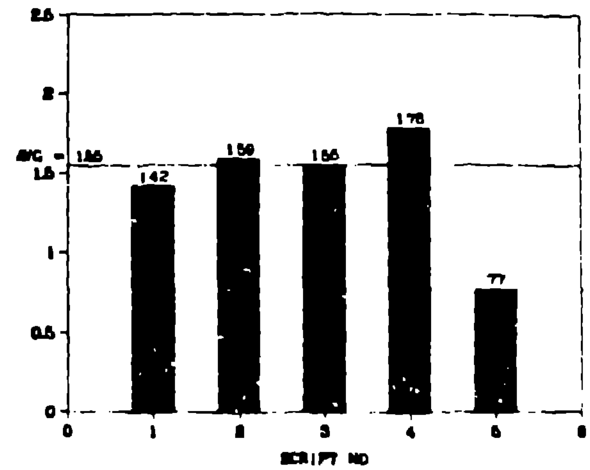

Figure 4. The ratio of 780 to 782 response times for each script category in Scenario IV.

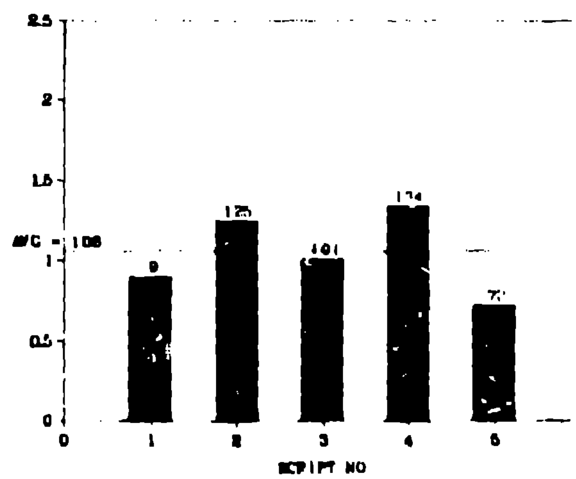

Figure 5. The ratio of 780 to 782 response times for each script category in Scenario V.

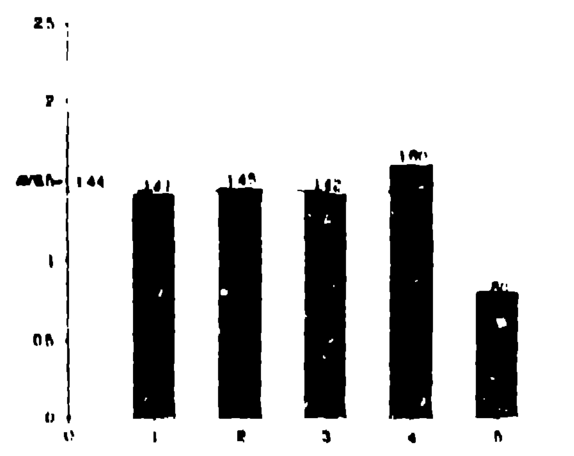

Figure 6. The response time ratio of a one processor 782 to a dual processor 782 for Scenario IV.

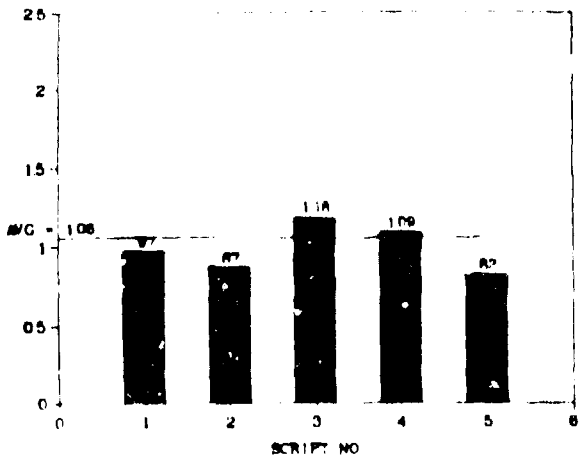

Figure 7. The response time ratio of a one processor 782 to a dual processor 782 for Scenario V.