


HP 3000 Computer Systems



Performance Guide



HP Computer Museum
www.hpmuseum.net

For research and education purposes only.

The performance of HP 3000 Computer Systems and related products is application dependent. The performance obtained in a user environment will depend on the particular application workload. The throughput rates and response times indicated in this guide relate only to the application workloads described herein.

HEWLETT-PACKARD DOES NOT WARRANT OR REPRESENT THAT THE PERFORMANCE DATA STATED HEREIN WILL BE ACHIEVED FOR ANY PARTICULAR CUSTOMER APPLICATION.



CONTENTS

Introduction	2
System Performance Test Description	4
Current Product Line Comparison	8
"Box Swap" Upgrade Comparisons	10
"Field" Upgrade Comparisons	13
Effects of Changes to System Variables	16
Effects of Additional Memory	16
Effects of Disc Caching and Versions of MPE V	18
Effects of Additional Discs	22
Questions and Answers	23
Data Communications Products Performance	27
Point-to-Point Spooled Printers on ATP and ADCC	27
MTS Printer Performance	29
Remote Job Entry (RJE)	31
Multileaving Remote Job Entry (MRJE)	33
SNA Network Remote Job Entry (NRJE)	35
Interactive Mainframe Facility (IMF)	37
X.25 Block Mode Terminal Communication	39
Distributed Systems (DS)	41

Introduction

Hewlett-Packard is committed to providing the right business system for your application needs.

The purpose of this guide is to show you the performance of the different HP 3000 systems to help you determine the best Hewlett-Packard solution for your needs.

This guide is intended to provide information for two groups of customers:

1. Customers who are considering purchasing HP business computers for the first time. The results provide an understanding of the performance capabilities of the HP 3000 product line and the relative position of each HP 3000 hardware product.
2. Existing customers who require this information in order to economically plan their future growth needs. These needs may be met through system expansion, system upgrades, or the purchase of additional systems.

The performance data shown here should be used only as a general comparison between different HP 3000 configurations and versions of MPE. This information is intended to serve as a guideline in sizing systems. Actual system performance will vary depending on application design, load, and configuration.

Products Tested

The products which have been performance tested for the purposes of this guide include HP 3000 Computer Systems and Data Communications products. The HP 3000 System Performance section includes data which positions the Series 30R/33R, the Series III, the Series 39, the Series 40, the Series 42, the Series 44, the Series 48, the Series 64, and the Series 68.

Hewlett-Packard Consulting Services

Hewlett-Packard offers on-site Systems Engineering consulting, both as standardized services and on a time-and-materials basis, in order to help maximize the productivity of HP 3000 systems. The standardized consulting products presently available for the HP 3000 include:

Installation Management

This service is designed to help the System Manager of a new installation become immediately productive by reducing start-up problems, customizing procedures to your objectives, and planning for long-term operational success. This includes discussion and recommendations on performance variables, methods, problem management, disaster recovery, job scheduling and load management.

System Performance Evaluation

This service can help to optimize system performance by identifying bottlenecks and their causes, and by recommending a strategy for corrective action. An HP 3000 Performance Specialist uses HP performance tools to analyze the system workload, identify system bottlenecks (both current and potential), and suggest improvements for better system performance. For more information on Hewlett-Packard Consulting Services, please contact your HP Sales Representative.

Hewlett-Packard Performance Products

In addition to Hewlett-Packard system performance consulting, there are several Hewlett-Packard software products available which enable you to measure and tune the performance of your HP 3000 system:

OPT/3000

On-line Performance Tool (OPT/3000) is an interactive software program which enables the user to identify system bottlenecks and observe how the system resources are being utilized. OPT/3000 uses graphical displays to show what percentage of the CPU is busy, how memory is being used, and how many I/O's are being performed on each disc. It also shows which users are consuming critical system resources.

APS/3000

Application Program Sampler (APS/3000) is an on-line software tool which helps programmers identify CPU bottlenecks in the source code of application programs. APS/3000 will graphically show which areas of a program are consuming the most CPU resources. The programmer can then tune his application to run more efficiently and increase overall system performance based on the results obtained from APS/3000.

For more information on Hewlett-Packard performance products, please contact your HP Sales Representative.

System Performance Test Description

Nature of the Tests

Test Environment

The standard application workload which was used for testing HP 3000 systems is the same workload which has been used in previous HP 3000 performance guides. By using the same mix, we are able to compare this set of results with our earlier test results. This mix may be representative of a general purpose EDP environment, but the main intent is to provide a consistent set of test results across the HP 3000 product line. The application mix consisted of on-line and batch jobs running concurrently. All databases and updated files were restored to their initial condition before each test. File placement was by default.

On-line mix:

The mix consisted of 60% application programming transactions using VPLUS/3000 and IMAGE/3000, 20% QUERY/3000 transactions, and 20% program development transactions and MPE commands.

The time between transactions varied from three to ten seconds. The tests on the older systems, the Series IIR and Series 30/33R, were run with the terminals at 2400 baud while all other tests were completed with the terminals running at 9600 baud.

- **VPLUS/3000 and IMAGE/3000:** This application consisted of a 2500 line COBOL program which performs 48 file inquiries, 14 files updates and 7 file add transactions in each loop of the script. This includes regular screen changes which are typical of a database inquiry/update application. Data set locking was used. The database consisted of 2500 entries organized into 10 master datasets and 8 detail datasets. VPLUS/3000 block mode data entry consisted of six screens and one main menu screen. Reviews were performed in groups of three or four. Adds and updates were done one at a time. The application is an on-line patient admission and charge collection system. The think times varied from four to ten seconds.
- **Data Inquiry:** QUERY/3000 was used for on-line database inquiry and reporting. The database used was the same as described above with the inquiries consisting of ten FINDs and nine REPORTs being performed on one of the detail datasets. The detail dataset consisted of 445 entries. The FINDs are performed on a search item which is accessed randomly through the master data set. A QUERY/3000 procedure file is used to define the report format. The think times varied from four to ten seconds.
- **Program Development/System Management:** A 1200-line COBOL program was compiled and linked but was not executed. EDIT/3000 was used to TEXT, MODIFY, and KEEP a 500-line program source file. System status commands were also executed in some of the scripts (SHOWTIME, SHOWJOB, SHOWME, REPORT, LISTF and HELP) as well as file control commands (BUILD, FILE, SAVE, and PURGE) and one utility program (LISTEQ). The think times varied from three to six seconds.

Batch mix:

Two or more jobs were run in batch mode, while the on-line sessions were active. The jobs were the same except that the second job logged on and waited for seconds before executing. Each job compiled and linked the 1200-line COBOL program but did not execute it. Then, each job executed a 1000-line COBOL program which consisted of a 500-line program and called a 500-line subroutine. The main program was an order, shipment, and inventory program which sorted 2500 120-byte records twice and created one report for each sort. The COBOL internal sort was used for the sorting in these jobs. The data file had a blocking factor of 16. Both sorts were ascending order sorts using two keys. The report which was generated was written to a null file. The number of batch jobs run simultaneously during one test varied with the number of terminals being tested. Two additional batch jobs were run for approximately every 50 terminals tested. For the 10 to 50 terminal tests, two batch jobs were run; for 60 to 100 terminals four batch jobs were executing, etc.

Test Limitations

As mentioned earlier, the mix used for this set of tests is intended to provide information on the relative performance of systems in the HP 3000 product line. System performance is *extremely* application dependent and caution should be exercised in making comparisons between this mix and other system loads.

This mix provides a consistent set of performance information for the HP 3000 from the older Series III and Series 30/33 models through the current Series 68. However, this mix may not be as typical for the current Series 68 customer as it was or is for the Series III customer. As HP introduces larger systems and additional software, it becomes increasingly difficult to define a typical HP 3000 workload.

The scenario used for these system performance tests is subject to the following limitations:

1. For larger configurations, the system mix was relatively light compared to the workloads used in many actual production environments.
2. The average transaction completed approximately ten disc I/Os, which is much lower than many large database applications.
3. The database used was smaller than most databases in the larger (Series 48 and Series 68) HP 3000 configurations. This smaller database provided exceptional results on some of the larger memory configuration tests with disc caching enabled, due to large portions of this database being cached.

Current Product Line Comparison

HP 3000 Performance Testing Results

The HP 3000 family consists of a broad range of commercial systems. The Series 30R/33R, Series IIIR, Series 40R, and Series 44R (the older systems), are now being re-marketed. The Series 30R performance is identical to that of the Series 33R with the same memory size, the primary difference being that the Series 33R offers more I/O expandability. The Series 30R, Series 33R, Series IIIR, and Series 39 represent the low-end of the HP 3000 product line. The Series 40, 42, 44, and 48 perform in the mid-range, and the Series 64 and Series 68 are the highest performing systems.

Figures 1 and 2 show the relative performance of the HP 3000 family by comparing both the response times and the transaction rates of the systems running the EDP mix. Typical customer configurations for each system are represented:

Series IIIR: 2 Mb Memory, four HP 7925 disc drives (1 master and 3 slaves), MPE IV, 10 - 50 terminals.

Series 39: 1 Mb Memory, one HP 7914 disc drive (132 Mb), MPE V/P (disc caching not included), 10 - 30 terminals.

Series 42: 2 Mb Memory, two HP 7933H disc drives (404 Mb), MPE V/E (disc caching enabled), 10 - 30 terminals, ADCC terminal connections.

Series 48: 4 Mb Memory, two HP 7933H disc drives, MPE V/E (disc caching enabled), 10 - 90 terminals, ATP terminal connections.

Series 68: 6 Mb Memory, four HP 7933H disc drives, MPE V/E (disc caching enabled), 50 - 180 terminals.

Conclusions

For the Series 39, with a 1 Mb configuration, acceptable performance was achieved with 10 to 20 terminals. With the addition of more terminals, the system became increasingly memory bound. For the Series 39, I/O was also a bottleneck at 30 terminals since in these tests there was only one disc drive and no disc caching.

The Series 42 performance was equal to or better than the Series III performance in all tests completed.

In these tests, the Series 48 was able to run 90 terminals simultaneously with just over two second response time. However, between 70 and 90 terminals, CPU saturation was reached so that adding additional terminals beyond 90 with this mix would result in a performance degradation.

The Series 68, as shown here, completed over 40,000 transactions an hour with 2 second response time at 180 terminals. Up to 400 terminals can be logged on simultaneously to the Series 68, but it is recommended that no more than 200 be actively processing at any one time for acceptable performance under most workloads.

Figure 1:





Figure 2:

Upgrade Performance

"Box Swap" Upgrades

Figures 3 - 6 show the performance that can be achieved by upgrading older HP 3000 models to larger, higher performance models. Performance of some of the more common upgrade paths are shown here.

Box Swap Upgrade Configurations:

Series 30/33: 1 Mb memory, two HP 7925 disc drives (one master and one slave), MPE IV.

Series III: 2 Mb memory, four HP 7925 disc drives (one master and three slaves), MPE IV.

Series 40: 2 Mb memory, two HP 7933H disc drives, MPE IV.

Series 44: 2 Mb memory, two HP 7933H disc drives, MPE IV, ADCC terminal connections.

Series 48: 4 Mb memory, two HP 7933H disc drives, MPE V/E, ATP terminal connections.

Series 68: 4 Mb memory, four HP 7933H disc drives, MPE V/E.

Figures 3 and 4 show the increases in performance that were achieved by upgrading a Series 30/33, Series III, or Series 40 to a Series 48. The Series 48, which was tested for configurations of 30 to 90 terminals, provided up to a 70% improvement in response time over the older systems at the 30 terminal level. The Series 48 was able to efficiently run almost twice as many terminals as the Series III with no degradation in response time.

Figures 5 and 6 show the relative performance improvements that may be observed when upgrading a Series 30/33, Series III, or Series 44 to a Series 68. Response time was reduced by up to 79% when a Series 44 was compared to a Series 68 running the EDP mix. The Series 68 was able to run three times as many terminals and complete three times as many transactions as the Series III with almost the same response time.

Figure 3:





Figure 4:

Figure 5:





Figure 6:

"Field" Upgrades

Field upgrades include a new, higher-performance version of MPE with disc caching software and 1 Mb of memory. Upgrading a Series 40, 44, or 64 to a 42, 48, or 68 respectively involves changing two system variables. Figures 7 - 10 show the performance effects of this type of upgrade.

Series 40 to Series 42 Upgrades

Figures 7 and 8 show the performance improvement that might be achieved by upgrading a Series 40 to a Series 42. In our testing, adding the additional 1 Mb of memory and the caching software provided as much as a 92% improvement in response time and a 46% improvement in throughput. Both the Series 40 and Series 42 were tested with two HP 7933H disc drives. The Series 42 can run this mix with very good performance at 30 terminals but the Series 40 was under severe memory pressure at the 30 terminal point.

Series 44 to Series 48 Upgrades and Series 64 to Series 68 Upgrades

Figures 9 and 10 show the performance improvements resulting from upgrades to the Series 48 and Series 68.

A 2 Mb Series 44 running MPE IV was compared to a 3 Mb Series 48 running MPE V/P. The Series 44 and Series 48 were both configured with two HP 7933H disc drives on the same GIC and ADCC's as the terminal connections. This data shows that response time was improved by up to 50% with this upgrade. The performance improvement was due to the addition of disc caching and additional memory which relieved the I/O bottleneck by taking advantage of available CPU resources.

The Series 64 to Series 68 upgrade data was taken from a 4 Mb Series 64 running MPE IV which was upgraded to 6 Mb of memory running MPE V/E. Both systems were configured with four HP 7933H disc drives split across two GIC's. The expanded tables version of MPE allowed double the number of terminals to run with two second response time. Under MPE IV, because of table limitations, only 90 sessions of this mix could be running concurrently.

Figure 7:





Figure 8:

Figure 9:





Figure 10:

Effects of Changes to System Variables

The data presented in the previous section shows the differences when HP 3000 systems are upgraded either through "box swaps" or field upgrades. These changes involve modifying more than one system variable. In this section, one system variable at a time is isolated to determine its effect on overall system level performance. The variables discussed are memory, disc caching, operating system versions, and number of disc drives.

Effects of Additional Memory

Figures 11 - 13 show how additional memory affected response time for the EDP mix. As can be seen from the graphs for both the Series 42 and Series 48, additional memory allowed considerably more terminals to be added before significant performance degradation occurred.

The Series 42 configuration consisted of two HP 7933H disc drives on one GIC, with ADCC's for terminal connections, and was operating under MPE V/P. For the Series 42, 2 Mb is a sufficient memory configuration until more than 30 terminals are running simultaneously.

The Series 48 was configured for two HP 7933H disc drives on one GIC, using the ATP for terminal connections, and was running MPE V/E. On the Series 48, there is negligible difference between the 3 and 4 Mb configurations until more than 50 terminals are running. The Series 68 was running MPE V/E and was configured for four HP 7933H disc drives split across two GIC's. On the Series 68, the system becomes memory bound with 4 Mb at 120 terminals, but there was virtually no difference in response time between the 6 and 8 Mb configurations.

These figures show that more memory is not always significantly better—even when disc caching is enabled. In the mix used here, the systems are functioning satisfactorily with less than maximum memory, except where heavily loaded (for example, the Series 48 tests with 70 and 90 terminals).

Figure 11:




Figure 12:





Figure 13:

Effects of Disc Caching and Versions of MPE V

Figures 14 to 19 show the differences in performance when disc caching is enabled for systems running both versions of MPE V: MPE V/P and MPE V/E.

Note: The following graphs indicate a difference in response time between MPE V/E and MPE V/P for some configurations and system loads. However, for those situations which indicated a significant difference, the response time was only one to two seconds and the difference would probably not be noticed by the users.

Figures 14 and 15 show the effects of disc caching on the Series 42. The Series 42 was configured with 3 Mb of main memory, two HP 7933H disc drives on one GIC, and ADCC's for the terminal connections. The tests run under MPE V/E operating system did not use any of the expanded table features. Disc caching provided up to a 34% improvement in response time using this mix. All response times were less than two seconds except for the MPE V/E, non-cached 50 terminal test. There was very little difference in throughput between MPE V/P and MPE V/E on the Series 42 as shown by Figure 14.

Figure 14:





Figure 15:



Caching improvements on the Series 48 are shown in Figures 16 and 17. The Series 48 configuration consisted of 4 Mb of main memory, two HP 7933H disc drives on one GIC, and ATP's for terminal connections. In all cases tested with the EDP mix, disc caching reduced response time and increased transaction throughput. Disc caching performs well with application mixes such as the EDP mix which do not require an excessive amount of CPU or main memory resources.

The effects of disc caching and the different versions of MPE V on the Series 68 are shown in Figures 18 and 19. The Series 68 was configured with 6 Mb of main memory and four HP 7933H disc drives split between two GIC's. As can be seen from these two figures, there was almost no performance difference between the two operating systems on the high-end 3000 system. With MPE V/E, the Series 68 ran more than twice as many terminals simultaneously than with MPE V/P with almost no difference in response time between 90 terminals and 180 terminals. Note that it appears as if only two lines are shown in Figure 19. The MPE V/E lines do not show up because there is virtually no difference in throughput between the two versions of the operating system.

Figure 16:





Figure 17:

Figure 18:






Figure 19:

Effect of Additional Discs

Figure 20 shows the effect on system performance as measured by response time when the number of discs was varied on a Series 68 running MPE V/E with 6 Mb of main memory. The tests with four discs were run with two HP 7933H's per GIC on two GIC's. The eight disc tests were run using HP 7933H disc drives, two per GIC, four per IMB. Without disc caching enabled, an I/O bottleneck existed with four discs. When four additional discs were added, the response time was reduced by as much as 74%. However, when disc caching was enabled, even with only four disc drives and 180 terminals active, no I/O bottleneck existed. With disc caching enabled, the additional four disc drives did not provide much of a performance improvement.

Figure 20:



Questions and Answers

As Hewlett-Packard introduces larger and more powerful HP 3000 systems and peripherals it becomes increasingly complex to size and configure these systems. This section answers some of the commonly asked questions regarding HP 3000 sizing and configuration under the MPE V operating system.

Systems



Will the Series 68 actually run 400 terminals?

The Series 68 *will* support 400 sessions *logged on* at one time. However, for most typical application mixes, only about 200 sessions can be *simultaneously accessing* the system with acceptable performance. The actual number of active sessions with acceptable performance will depend on the application mix and the desired response time.

What are the MIP ratings of the various HP 3000 systems?

Hewlett-Packard does not use MIP's (Millions of Instructions per Second) as a measure to position HP 3000 CPU's. A CPU's MIP rating will provide an absolute value for how many instructions can be performed by that CPU in one second, but it does not provide any indication of overall system performance. There is also danger in attempting to compare MIP ratings across different vendors' systems because the number of instructions executed depends upon the CPU's architecture and the complexity of its instruction set. It is entirely possible for one CPU to have a lower MIP rating than another and significantly outperform that system in terms of throughput. For this reason, we prefer to performance position our systems on the basis of *system* throughput rather than CPU instruction speeds. The chart below shows the relative systems performance for each system mentioned in this guide.

RELATIVE SYSTEMS PERFORMANCE

SERIES 30/33R	.30
SERIES III	.60
SERIES 39	1.00
SERIES 42	1.00
SERIES 48	1.00
SERIES 68	3.00

How can I optimize my CPU utilization?

The way to optimize CPU utilization is to reduce contention for the CPU. There are a variety of ways that CPU contention can be reduced. Load management techniques which would shift heavy loads to non-peak timeframes or between systems can help alleviate the contention. APS/3000 (Application Program Sampler) can be used to identify user code which inefficiently uses the CPU.

Memory

When is more memory needed?

If more than 10% of all disc I/O's performed are I/O's performed to manage memory, then the system is probably experiencing a memory shortage and additional memory would improve performance. One way of monitoring the time spent managing main memory is to use the software tool OPT/3000 for memory analysis.

In general, memory pressure is most likely to occur on systems with less than 2 Mb of main memory. On larger systems (4 Mb and above), other resources (such as CPU or I/O) will probably be saturated before memory.

How much main memory does MPE V/E require?

The smallest memory configuration supported under MPE V/E is 1/2 Mb (512 Kb). The increased capabilities of MPE V/E require larger code segments than MPE IV, some of which *must* be resident in main memory at all times. Also, with the expanded tables features of MPE V/E, some of the system tables may require more main memory than in the past.

It is recommended that your system contain at least 1 Mb of memory above what is needed to efficiently run your applications if disc caching is enabled. In many applications, especially those which do a large amount of sequential I/O, extra memory will provide significant performance improvements with disc caching.

I/O and Disc Caching

How is I/O performed on the HP 3000?

There are several steps to the I/O process on the HP 3000. If disc caching is not enabled on the system, the I/O is performed as follows:

1. The user application program requests an I/O through the MPE file system.
2. The file system handles any necessary buffering, blocking, etc., then passes control to the MPE I/O system to perform the actual I/O.
3. A SEEK command, which moves the heads of the disc to the proper cylinder where the required data resides, is executed. This takes about 25 ms on the average.
4. Once the SEEK is completed, the I/O system issues a command to read or write data to the disc drive. The rotational latency, (the time it takes for the disc platter to spin around to the position where the required data is under the disc heads), takes approximately 11 ms. The data is then transferred to the CPU. The transfer time is dependent upon how much data is being transferred, but the average time for typical MPE transfers is about 2 ms.
5. When the data is transferred from disc, the I/O system returns control to the MPE file system which handles any blocking or buffering and then returns to the user's application program.

The entire process of completing a physical I/O takes about 40 ms.

For many HP 3000 systems, I/O has been a bottleneck. Disc caching has greatly reduced this constraint. When disc caching is enabled, the MPE file system calls the caching system which checks the areas in main memory where cached data may reside. If the required data is found in main memory (a cache hit), the time to complete an I/O is reduced from 40 ms to approximately 5 ms. This reduction in I/O time greatly improves overall system performance. If the data is not already resident in main memory (a cache miss), a physical I/O to disc must then occur. This lengthens the time slightly for the individual physical I/O; however, each I/O transfers larger amounts of data when caching is enabled so that the overall number of I/O's is still greatly reduced in most cases.

In what environments does disc caching provide the greatest improvement in performance?

Disc caching will significantly improve the performance of HP 3000 systems when:

1. The CPU utilization is low before caching is enabled. For the Series 42 and Series 48, a CPU that is less than or equal to 60% busy is best. For the Series 68, the CPU should be not more than 70% utilized before caching for the most benefit.
2. Excess main memory is available. Memory management I/O should be not exceed 5% of the total I/O.
3. The main memory size is at least 2Mb on a Series 42, 3 Mb on a Series 48, and 4 Mb on a Series 68.
4. The data locality is high. Disc caching will work best when the same files are accessed frequently or in a sequential access environment.
5. The ratio of READs to WRITEs is high -- 3:1 or higher. The greatest benefit from caching is derived on READs.

Are there some environments where disc caching does not provide a performance improvement or causes a degradation?

Yes, it is possible that disc caching may not provide a performance improvement but cause a degradation. This can occur when:

1. CPU utilization is greater than 70% on a Series 4X or greater than 85% on a Series 6X before caching is enabled. Disc caching requires some CPU resources to execute effectively.
2. There is no excess main memory while running a pre-cached operating system. (Memory management I/O is greater than 10% of all system I/O). If there is no excess main memory, disc caching has no "room" to operate.
3. The system is configured for less than 2 Mb of memory on a Series 42, 3 Mb of memory on a Series 48, or 4 Mb of memory on a Series 68.
4. Data locality is poor. If the I/O is highly random, it is possible that the extra overhead associated with searching through main memory and incurring a cache miss and then having to do a physical I/O will cause system performance degradation.
5. Ratio of READs to WRITEs is low -- less than 2:1 When an application is write intensive, there may not be as significant of a performance improvement when caching is enabled because every cached WRITE must be posted to disc. However, once a WRITE has logically completed, the process does not have to wait for that WRITE to be posted to disc; this can be handled by disc caching in the background. The process can move to the next transaction without waiting for the physical update to occur.

Do I have to give up some of this performance gain if I want database integrity?

While most users of I/O intensive applications will experience significant performance gains with disc caching, it is most effective in applications in which multiple users share the same information. Applications which share databases usually benefit greatly. Disc caching not only reduces the contention for these databases, but it also reduces the lock service times of the database control blocks because I/O requests can be serviced more quickly. Because so much information may be stored in main memory at one time, new intrinsics and commands have been developed to preserve data base integrity when disc caching is enabled. These features may have an impact on system performance; however, their use is optional. For more information on the disc caching commands and IMAGE/3000 changes for disc caching, please refer to *Communicator 3000*, Volume 2, Issue 1.

If I'm using disc caching, should I still split heavily accessed files between discs?

In the past, before MPE V the spreading of related files (for example, master and detail data sets) across master disc drives could provide significant performance improvements. With disc caching we still recommend doing this to some extent, however it is not as much of an issue as it was in the past. One recommendation is to put those files which are randomly accessed or heavily updated on the same disc and not enable caching for that disc. Those files which are accessed sequentially or are not updated should be placed together on a disc and that disc cached.

Should I still follow the same rules in adding disc master disc drives, GIC's, and IMB's?

The limitations for physical I/O rates are the same as they have been with previous versions of MPE. The maximum number of physical I/O's a disc can sustain for extended periods of time is around 30 per second. As a guideline, the maximum number of physical I/O's sustained per channel is around 45 - 50 per second. If your physical I/O rate is at that level, adding a second GIC may improve performance. When the number of sustained balanced physical I/O's is above 85 per second per Intermodule Bus (IMB) an additional IMB should be added. For many applications, however, the number of physical I/O's has been greatly reduced with disc caching and therefore the need to add additional master disc drives, GIC's for discs and IMB's is also reduced. We used to operate under the theory that the way to maximize I/O throughput was to have as many smaller master disc drives as possible so that there are as many paths to the data as possible. With disc caching, because the need to perform physical I/O's is greatly reduced and because locality of data is important, the need for extra spindles for performance reasons no longer exists.

How much performance difference, if any, should I expect in moving from MPE V/P to MPE V/E?

For most environments, there will be at most only a slight degradation in overall performance as measured by transaction throughput and response time with MPE V/E over MPE V/P (5 to 12 %). Most users should not notice this difference. The degradation may be greater than that if:

1. MPE V/E is run on a system with less than 2 Mb of main memory or
2. MPE V/E is run on a system with CPU utilization greater than 85% with MPE IV or MPE V/P.

The upgrade from an MPE IV system to MPE V/E should provide performance improvements as well as offering additional functionality.

Data Communications Products Performance

The performance data presented here indicates some of the differences in performance between alternatives for data communications. This data by no means provides sufficient information to design complete applications or system configurations. Data communications performance evaluation requires consideration of many factors simultaneously, because they interact. Among the design variables affecting throughput and CPU utilization are: line speed, line quality, protocol choice, packet or line buffer size, X.25 level 2 and 3 window sizes, user buffer size, type of specific data communication product activity, directions of use, numbers of users, data compressability, basic CPU speed, MPE file blocking factors, and other system loads. This list is representative of the types of factors which must be considered in evaluating data communications performance. For more details we recommend that you consult your data communications specialist SE, who will have more extensive performance information available with which to design or evaluate your proposed application.



POINT-TO-POINT SPOOLED PRINTERS ON ATP AND ADCC

With the release of MPE V/E several new spooled printers can now be attached to the HP3000 in a point-to-point configuration using the ATP or the ADCC. This enhancement allows the placement of a shareable printing resource at a greater distance from the CPU than an HP-IB connection would allow. These spooled printers fall into three groups:

Character Printers: HP 2631B, HP 2932A, HP 2933A, HP 2934A
Page Printer : HP 2687A
Line Printer : HP 2563A

These connection of these printers was tested for throughput and CPU overhead at different baud rates. The differences between printers in a group was minimal; therefore the data from the testing of the printers is summarized in Tables 1-3.

Test Environment

The tests were run on a Series 48 (and Series 68 also for the HP 2932A), running MPE V/E. These systems were not running any other jobs during the tests. Files large enough to ensure 10 minutes of active printing were sent to the printers. The line length was varied; for the HP 2932A and HP 2563A, F132 was a file with 60 lines of 132 characters + CR + LF on each page. F60 was a file with 60 lines of 60 characters + CR + LF on each page. For the HP 2687A, F80 contained 60 lines of 78 characters + CR + LF on each page. F50 was 60 lines of 50 characters + CR + LF on each page.

Conclusions

The data indicates that connection of these printers to a serial port does not necessarily limit the total throughput of the printer. However, the higher speed page and line printers do consume a significant percentage of the CPU when running at full speed, particularly when connected to the ADCC. Care should be taken when configuring systems with these connections, and ATP ports should be used to connect the HP 268X and HP 2563A printers whenever possible.

Table 1. HP 2932 Character Printer Performance

speed	test file	ADCC Series 4x		ATP Series 4x		ATP Series 6x	
		% cpu util.	thruput char/sec	% cpu util.	thruput char/sec	% cpu util.	thruput char/sec
2400	F132	5.6	176	2.3	176	1.7	176
2400	F60	8.8	152	3.8	153	2.9	153
1200	F132	3.2	112	1.3	112	1.0	112
1200	F60	5.7	106	2.6	107	1.9	107

Note: Performance figures for HP 2932A and HP 2934A were similar;
HP 2631B results were slightly lower.

Table 2: HP 2687A Laser Printer Performance

speed	test file	ADCC Series 4x		ATP Series 4x	
		% cpu util.	thruput pages/min	% cpu util.	thruput pages/min
19200	F80	not appl.	not appl.	9.8	8
19200	F50	not appl.	not appl.	14.8	12
9600	F80	21.0	8	9.5	8
9600	F50	27.7	12	14.4	12
4800	F80	13.1	6	6.3	6
4800	F50	16.7	7	9.8	7

Table 3: HP 2563A Line Printer Performance

speed	test file	ADCC Series 4x		ATP Series 4x	
		% cpu util.	thruput lines/min	% cpu util.	thruput lines/min
19200	F132	not appl.	not appl.	8.02	300
19200	F60	not appl.	not appl.	7.65	300
9600	F132	19.05	300	7.84	300
9600	F60	16.65	300	7.68	300
4800	F132	12.32	194	5.28	198
4800	F60	16.30	300	7.36	300

MTS PRINTER PERFORMANCE

The performance and CPU utilization of line printers on a multipoint link were tested for both the HP 3000 Series 48 and 68. The tests measured the number of lines printed per minute and the CPU utilization at various line speeds. Figures 1 and 2 depict the results of these tests.

Test Environment

The tests were performed with a single printer on the multipoint line printing 132 character lines. Neither the type of the multipoint line (asynchronous or synchronous) nor how the printer was configured (as a termtyp 18 or 22 device) had much of an impact on the printer performance or CPU utilization. The effects of multiple printers on a multipoint line (using the HP 2333A cluster controller) are basically linear extrapolations of the single printer data. However, a large number of devices on a HP 2333A may use up the buffer space of the HP 2333A. If this situation occurs, additional polling and control commands will be transmitted between the HP3000 and the HP 2333A, resulting in higher CPU utilization. Refer to the MTS manual for the rules and guidelines of configuring printers on a multipoint line.

Conclusions

The speed (lpm) of line printers on a multipoint line is primarily a function of the multipoint line speed. The type of printer and the model of CPU have very little to do with the actual throughput. While printers will perform very close to their maximum speed (at appropriate line speeds), they will run a few lines per minute less than the rated speed due to interactions between the spooler, the printer's buffering mechanism and the bisync protocol.

As would be expected, the CPU utilization of printers (the HP 2563A in this example) is less on the Series 68 than on the Series 48, since the constraining factor is how much data can be transmitted across a multipoint line at a given speed. Generally, the CPU utilization on the Series 68 was about half that of the Series 48 in the tests performed.

Figure 1:





Figure 2:



REMOTE JOB ENTRY (RJE)

Figures 3 and 4 illustrate CPU utilization and effective CPU time per record processed for standalone operation of RJE on an HP 3000 with an HP 2680A Laser Printer. Tests were performed on 9.6 kbps and 56 kbps leased lines. Disc caching was not enabled for the MPE V/E tests.

Test Environment

- Series 48
- 2 MB memory
- Line speeds: 9.6 kbps and 56 kbps
- 400 byte buffer
- Dataset 1 (2501 characters/page)
- Light host load

Conclusions

- RJE utilized 1.0% - 2.5% more of the CPU under MPE V/E than under MPE IV (11.8% vs. 10.8% at 9.6 kbps and 43.1% vs. 40.6% at 56 kbps).
- Effective CPU time per record processed was approximately 10% greater under MPE V/E than under MPE IV.
- Under both MPE V/E and MPE IV, RJE operation at 56 kbps increased CPU utilization by a factor of approximately 4, while CPU time per record was affected only slightly. Throughput at 56 kbps was approximately 4 times throughput at 9.6 kbps.

Figure 3:





Figure 4:


MULTILEAVING REMOTE JOB ENTRY (MRJE)

Figures 5 and 6 illustrate CPU utilization and effective CPU time per record processed for standalone operation of MRJE on an HP 3000 with an HP 2680A Laser Printer. Tests were performed on 9.6 kbps and 56 kbps leased lines. Disc caching was not enabled for the MPE V/E tests.

Test Environment

- Series 48
- 2 MB memory
- Line speeds: 9.6 kbps and 56 kbps
- 400 byte buffer, 1 print stream at 9.6 kbps
- 2000 byte buffer, 2 print streams at 56 kbps
- Dataset 1 (2501 characters/page)
- Light host load


Figure 5:



Conclusions

- MRJE utilized 0.7% - 2.0% more of the CPU under MPE V/E than under MPE IV (22.7% vs. 20.7% at 9.6 kbps and 81.2% vs. 80.5% at 56 kbps).
- Effective CPU time per record processed was 5% - 10% greater under MPE V/E than under MPE IV.
- Relative to RJE, MRJE with one print stream and 400-byte buffers at 9.6 kbps had 92% greater CPU utilization, 59% greater CPU time per record, and 21% greater throughput.
- Increased throughput by a factor of 6 relative to one print stream and 400-byte buffers at 9.6 kbps. This was due to increasing CPU utilization and reducing CPU time per record to 12 msec (comparable to RJE with 400 byte buffers at 9.6 kbps).
- At 56 kbps, MRJE can effectively utilize the capacity of both the HP 3000 and the HP 2680A Laser Printer. Using 2000-byte buffers at 56 kbps, MRJE can support laser printing of two logical pages per physical page.

Figure 6:



SNA NETWORK REMOTE JOB ENTRY (NRJE)


Figures 7 and 8 illustrate CPU utilization and effective CPU time per record processed for standalone operation of SNA NRJE under MPE V/P on an HP 3000 with an HP 2680A Laser Printer. Tests were performed on a 56 kbps leased line. (Note: SNA NRJE will be supported on MPE V/P Delta 1, E Delta 1 and T-Mit.)

Test Environment

- Series 42
- 2 MB memory
- 56 kbps line speed
- 256 byte buffer using 1 print stream
- 512 byte buffer using 2 print streams
- Dataset 1 (2501 characters/page)
- Light host load



Figure 7:




Conclusions

- Using two print streams and 512 byte buffers increased CPU utilization to 85% compared to 78% with one print stream and 256-byte buffers.
- Using two print streams and 512 byte buffers decreased CPU time per record from 26.5 msec to 15.6 msec.
- After adjusting for buffer size differences, performance measured with SNA NRJE was comparable to performance with MRJE. SNA NRJE performance under MPE V/E may differ from performance under MPE V/P.
- At 56 kbps, SNA NRJE can effectively utilize the capacity of both the HP 3000 and the HP 2680A Laser Printer. SNA NRJE can support laser printing of one logical page per physical page. (MRJE, by using 2000-byte buffers at 56 kbps, can support laser printing of two logical-pages per physical page.)

NOTE

HP recommends additional memory with SNA NRJE for good performance. Generally, approximately 1 MB is required above what would be required without NRJE. (Each system should have at least 2 MB main memory.)


Figure 8:



INTERACTIVE MAINFRAME FACILITY (IMF)

Figure 9 illustrates CPU utilization under MPE IV and MPE V/E for standalone operation of IMF in Programmatic Mode. Pass-Thru operation is not affected by differences between MPE IV and MPE V/E.

Figure 9:



Test Environment

- Series 48
- 1 Mb main memory
- 9.6 kbps line speed
- 500 characters from host
- 250 characters from IMF
- 5 devices open
- Light host load

Conclusions

- CPU utilization reached saturation at 11.9% for MPE V/E and 10.4% for MPE IV.
- The CPU load when using either bisync or SDLC protocols is nearly identical. The user need not consider CPU performance in choosing the line protocol.
- No conclusions can be drawn for response times, which are very application-specific and dependent upon the environment in which IMF is operating. For a given application and environment, response times under MPE IV and MPE V should not differ significantly.
- In general, IMF performance under MPE V should not differ significantly from performance under MPE IV.

X.25 BLOCK MODE TERMINAL COMMUNICATION

With the release of MPE IV Q-Delta 2, HP 3000 users can now run their VPLUS applications over X.25 Public Data Networks (PDN's) and on the HP 2334A X.25 Terminal Cluster Controller. Users can further enhance the performance of their VPLUS applications by migrating to MPE V. The data below, which was collected by connecting the indicated number of terminals to an HP 3000 at 9600 baud through a HP 2334A Cluster Controller over an X.25 PAD, indicate that terminal response time and throughput improve with almost no impact on CPU utilization.

The most favorable results were achieved in configurations using four terminals. Terminal response time with this configuration running MPE V is 7% faster than the MPE IV configuration. Throughput, measured in characters per second, improves 7%. Despite these improvements, CPU utilization only increased 0.1%.

X.25 performance is a function of the PDN to be used, network configuration and network congestion. The data summarized below was collected using GTE/Telenet. The block size used was 128. To optimize performance when running block mode X.25 over a PDN, the block size should be an integral multiple of the packet size being used on the network. The tests were run on an HP 3000 Series 48 with HP 2624B terminals and no other applications on the system.

Figure 10:




Figure 11:





Figure 12:

DISTRIBUTED SYSTEMS

Figure 13 illustrates the increased throughput provided by DSCOPY -- the network file transfer facility of Distributed Systems (DS) under MPE V versus MPE IV. Figure 14 illustrates the improved CPU time per 1000 characters transferred.


Test Environment

These DS performance tests were run between a Series 68 and a Series 48 processor, each with 2 MB of main memory and one master disc drive. Communication was through directly connected INP's, with a line speed of 56 kbps. For MPE V tests, disc caching was enabled.

Additional test parameters are:

- Blocking factor = 16
- Compression enabled
- Line Buffer = 1024 words
- Lightly loaded system, single user

Figure 13:



Conclusions

- The DSCOPY throughput under MPE V is improved by 60% over the throughput under MPE IV.
- Communication over X.25 provides slightly better throughput than over the bisynchronous protocol link.
- CPU time per character transferred is decreased by 25% on the Series 64 and by 17% on the Series 44.
- Under MPE V, DSCOPY over X.25 requires up to 20% more CPU time per character than DSCOPY over the bisynchronous protocol link.

Figure 14:

