HOW THE NEW MEMOREX 3680 MEASURES UP.

A DP MANAGER'S GUIDE TO DISC DRIVE PERFORMANCE EVALUATIONS.

TABLE OF CONTENTS

INTRODUCTION

Historically, plug compatible products developed by Memorex have been enhanced considerably with respect to existing systems. In many minds, however, the advent of thin film technology has created a new set of circumstances. Given the difficulties of developing thin film capabilities in the first place, is it logical to expect a Plug Compatible Manufacturer (PCM) to also advance the art with their initial offering?

Obviously it depends on the company. Memorex thin film systems, for example, have been under development since 1977. And as the second in a series of thin film head systems, the Memorex 3680 already reflects considerable product evolution.

In terms of storage capacity, data transfer rate and access speeds, it is identical to the 3380 system. However, once this parity position with the 3380 was reached, the designers of the 3680 turned their attention to enhancements that would improve throughput dramatically. The stakes for such an achievement are high, because even a modest improvement can have astonishing consequences for users.

TABLE OF CONTENTS

INTRODUCTION

Historically, plug compatible products developed by Memorex have been enhanced considerably with respect to existing systems. In many minds, however, the advent of thin film technology has created a new set of circumstances. Given the difficulties of developing thin film capabilities in the first place, is it logical to expect a Plug Compatible Manufacturer (PCM) to also advance the art with their initial offering?

Obviously it depends on the company. Memorex thin film systems, for example, have been under development since 1977. And as the second in a series of thin film head systems, the Memorex 3680 already reflects considerable product evolution.

In terms of storage capacity, data transfer rate and access speeds, it is identical to the 3380 system. However, once this parity position with the 3380 was reached, the designers of the 3680 turned their attention to enhancements that would improve throughput dramatically. The stakes for such an achievement are high, because even a modest improvement can have astonishing consequences for users.

A study of interactive user productivity by IBM's Arvind Thadhani[1] identified a significant increase in productivity as system response time decreased from five seconds to one second. Moreover, as system response time dropped into the subsecond range, the increase in user productivity skyrocketed (see Figure 1). In fact, in the range between 3.0 seconds and 0.3 seconds, user productivity increased almost three times as fast as the system response time decreased.

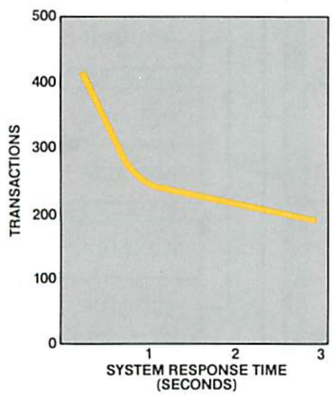

Figure 1: Plot of System Response Time vs. User Transaction.

The implications in terms of bottom line benefits are impossible to ignore. In one engineering study, for example, users working at on-line graphics terminals with a light pen improved productivity by as much as 550% when system response time decreased from 1.5 seconds to 0.4 seconds.

In still another study involving code generation, a test project was completed in just 61% of the time projected when system response time was decreased from 2.22 seconds to 0.84 seconds. What's more, there was a 57% reduction in the error rate.

Clearly, if user productivity increases 2 to 5 times, the cost savings quickly soar to significant proportions, but there are other ways to capitalize on throughput improvement. You could, for example, add more terminals without degrading your present system response time. And you could obviously improve the turnaround time on your batch processing. Scheduling also improves because even peak loads can be dispatched within predictable and consistent time frames. All told, the effect on DP operations is similar to adding a more powerful processor...an expensive eventuality most companies would prefer to defer as long as possible.

Historical data storage access architecture.

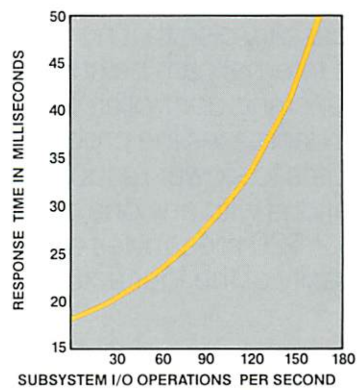

Figure 2: Plot of I/O's vs. Time.

As processing power dramatically increased the number of I/O's executed, accessing the burgeoning quantities of data continued to be possible only at the expense of processing speed. The fundamental problem was simply that the I/O's couldn't get out of their own way. (Figure 2.)

Various access architectures were developed in an effort to improve the throughput. The main objective was to prevent a single I/O from blocking access to the entire string as it obviously would do if all the actuators were simply connected in series. Thus, there were solutions that provided multiple paths, various path switching arrangements, parallel-series hookup arrangements and so forth. Some worked better than others but all of them improved data access.

In 3380 class subsystems, the sophistication of the data access architecture is a constant reminder of the importance of the problem.

[1] "FASTER IS BETTER – A Business Case For Subsecond Response Time," COMPUTERWORLD, April 18, 1983; Vol. 37, No. 16.

The IBM 3380 Subsystem utilizes a configuration called Dynamic Path Selection (DPS) whereby groups of four actuators share an internal path. (Figure 3.) The four internal paths are connected to the two string controllers through a "Two by Four Switch" which allows the paths to be connected to either string controller

In order to establish a connection between a Storage Director, a String Controller and an actuator, it is necessary not only that these be free, but necessary also that no other actuator on the same internal path be communicating with the other string controller. This additional requirement decreases the probability of access and contributes to slower response because communications with any one actuator effectively ties up 2,520 megabytes of data, the storage capacity of the four actuators on that path.

To minimize this problem, the 3380 is frequently configured as two half strings. This improves the overall response time compared with a full string because it reduces the contention for internal paths. This solution has a downside, however, in that each string, whether short or long, requires an expensive string controller. And for every two string controllers, an equally expensive storage controller is required.

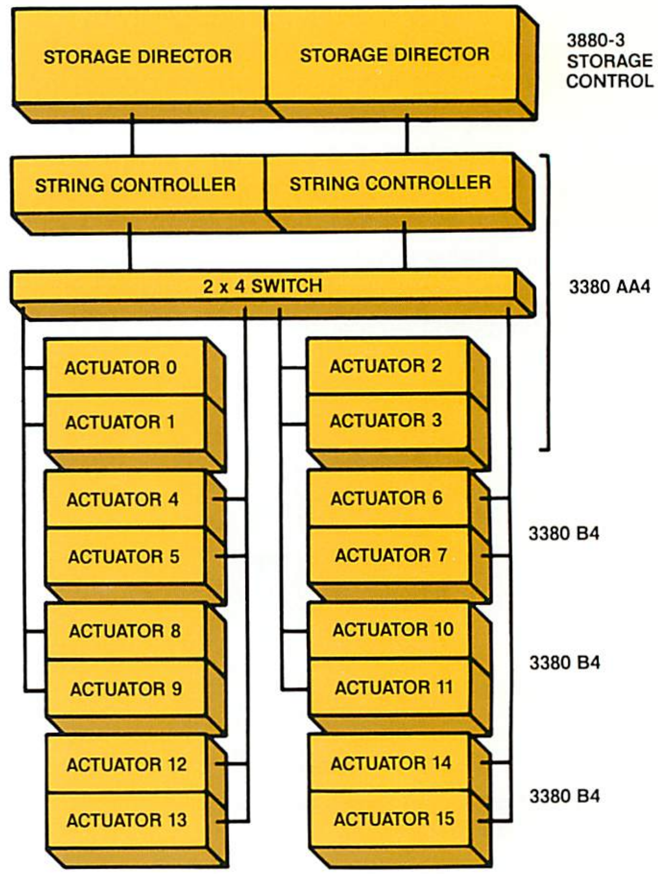

Figure 3: 3380 Pathing Method.

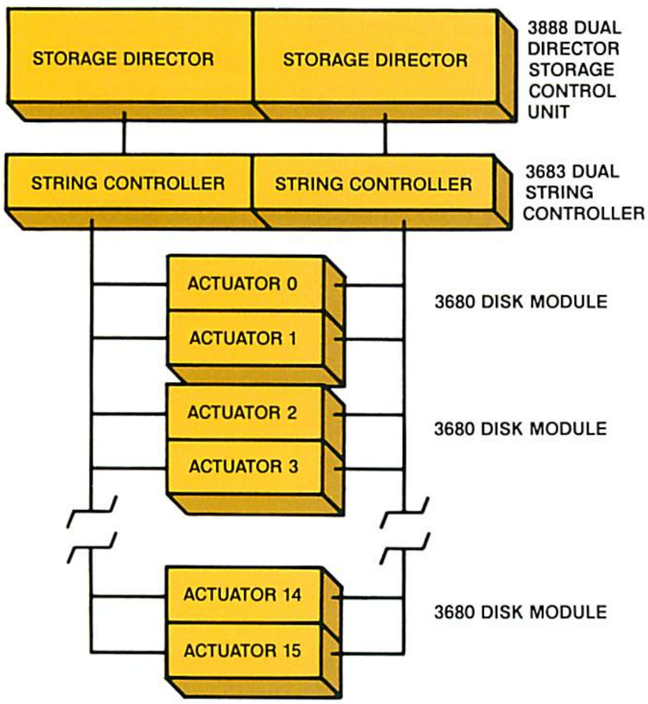

Figure 4: 3680 Pathing Method.

The Memorex 3680 Subsystem (Figure 4.) employs a pathing arrangement called "Maximum Availability Path Selection" or "MAPS" whereby each actuator is dual pathed and totally independent from each other. Together with separate actuator microprocessors, this enables simultaneous data transfer from any two actuators within a string, including to and from the same HDA.

Since the two data paths have the ability to access all of the actuators within the string, the data availability of the 3680 is far superior to that of the 3380, whether it is configured as a half string or a full string. In fact, our DASD Performance Analytical Model shows that a 3680 full string will outperform a 3380 half string simply because the maximum unit of contention in a 3680 string is only 630 megabytes—the amount of data stored under a single actuator.

This same model currently in use also predicts that MAPS will provide I/O transaction rates comparable to those of a 3380 configured with the Dynamic Reconnect feature of MVS/XA. In fact, when 3680's are optimized for use with MVS/XA, our models predict they will be able to use Dynamic Reconnect to even greater advantage so that the improvement already achieved will be even greater.

The MAPS feature also incorporates an improved protocol capability, which enables more data to be transferred per control operation. For example, MAPS transfers seek and set sector information in one control operation and does not need additional control operations to verify the transfer. With the 3380 system, additional control operations are required to verify seek and set sector and to verify the parameters transferred from these first two operations.

Further improvement in throughput results by staggering a track's index mark 180 degrees apart on adjacent cylinders. This speeds sequential processing because less time is required for latency. The result is that data can be retrieved faster because the system can start looking for it sooner.

Described in the following are the initial benchmark testing results which compare 3680 subsystem performance with that of the 3380 subsystem.

Figure 5 summarizes the results of a benchmark run for a large U.S. customer. The benchmark jobstream consisted of a program which executed job steps writing sequential records using multiple blocksizes. The program then did random reads using the same block-sizes and data sets created by the sequential writes.

Multiple jobstreams were run against multiple actuators on identical configuration 3680 and 3380 subsystems. These jobstreams were executed on a 3083, model E16 processor complex controlled by an MVS operating system.

The 4x4 test demonstrates the comparison found on identical logical volumes when four copies of the jobstream were run against two HDA's or four actuators on both the Memorex 3680 and IBM 3380 configurations. In the 3380 case, each HDA (and thus its two actuators) was on a different Dynamic Path Selection (DPS) internal path, so the unit of contention was kept to the minimum 1260 megabytes. The 2x2 tests demonstrate the comparison found on identical logical volumes when two copies of the jobstream were run against one HDA.

The 4x4 test constitutes the best case for the 3380, in that the HDA's were on different DPS internal paths. However, by using this configuration an improvement close to what may typically be expected is produced. The 2x2 test constitutes the worse case for the 3380 in that the two actuators were on the same DPS internal path, while the Memorex case could fully utilize the unrestricted dual paths provided by MAPS.

You can see a further example of the MAPS benefit in Figure 6. It summarizes results of a program which provides the ability to select I/O transaction rates, block sizes, read or write operations, and access modes. This was executed on the same processor with the same software as the preceding benchmark, but on Memorex and IBM strings with 4 spindles and 8 actuators.

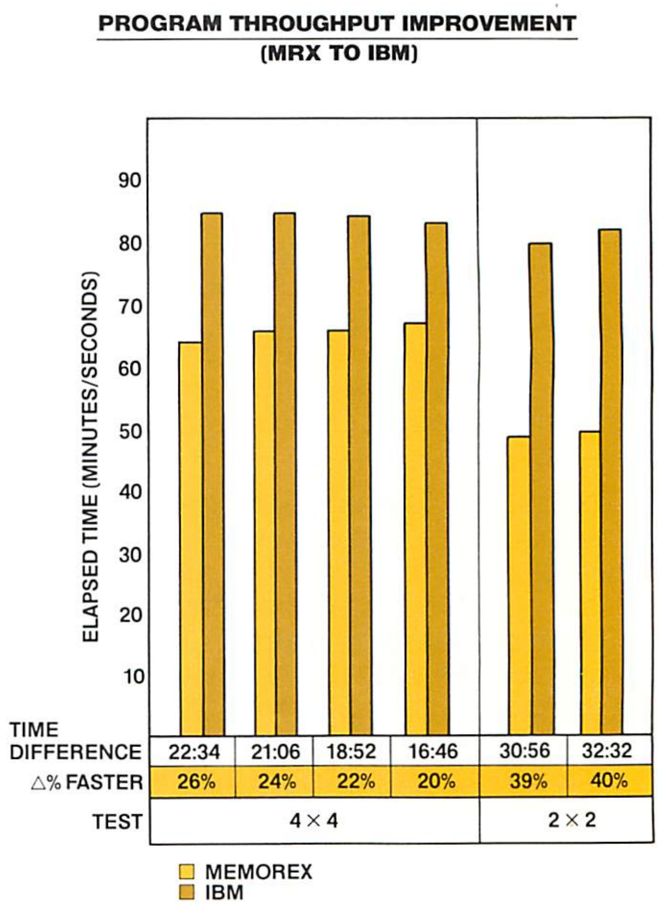

Figure 5: Benchmark Comparisons.

The results show that the throughput and performance improvements are sustained throughout common (4K and 6K) and less common (1K and 2K) blocksizes. They also show the improvements increase as I/O rates grow.

The most apparent observation that may be made with this benchmark is the Memorex subsystem's ability to handle significantly higher I/O transaction rates. This will provide benefits in peak transaction periods and when I/O's are skewed to a few actuators in the subsystem. It will also enable the Memorex subsystem to maintain a more stable transaction time at a terminal or turnaround for batch workload. For terminal response time this is particularly important, as we have earlier discussed.

As you can see from these results, the 3680 disc subsystem can significantly outperform the 3380 disc subsystem.

Performance comparisons are continuously being made by Memorex; if you would like to receive copies of future benchmark reports, just write "Benchmarks" on your business card and mail it to us at the address on the back of the brochure.

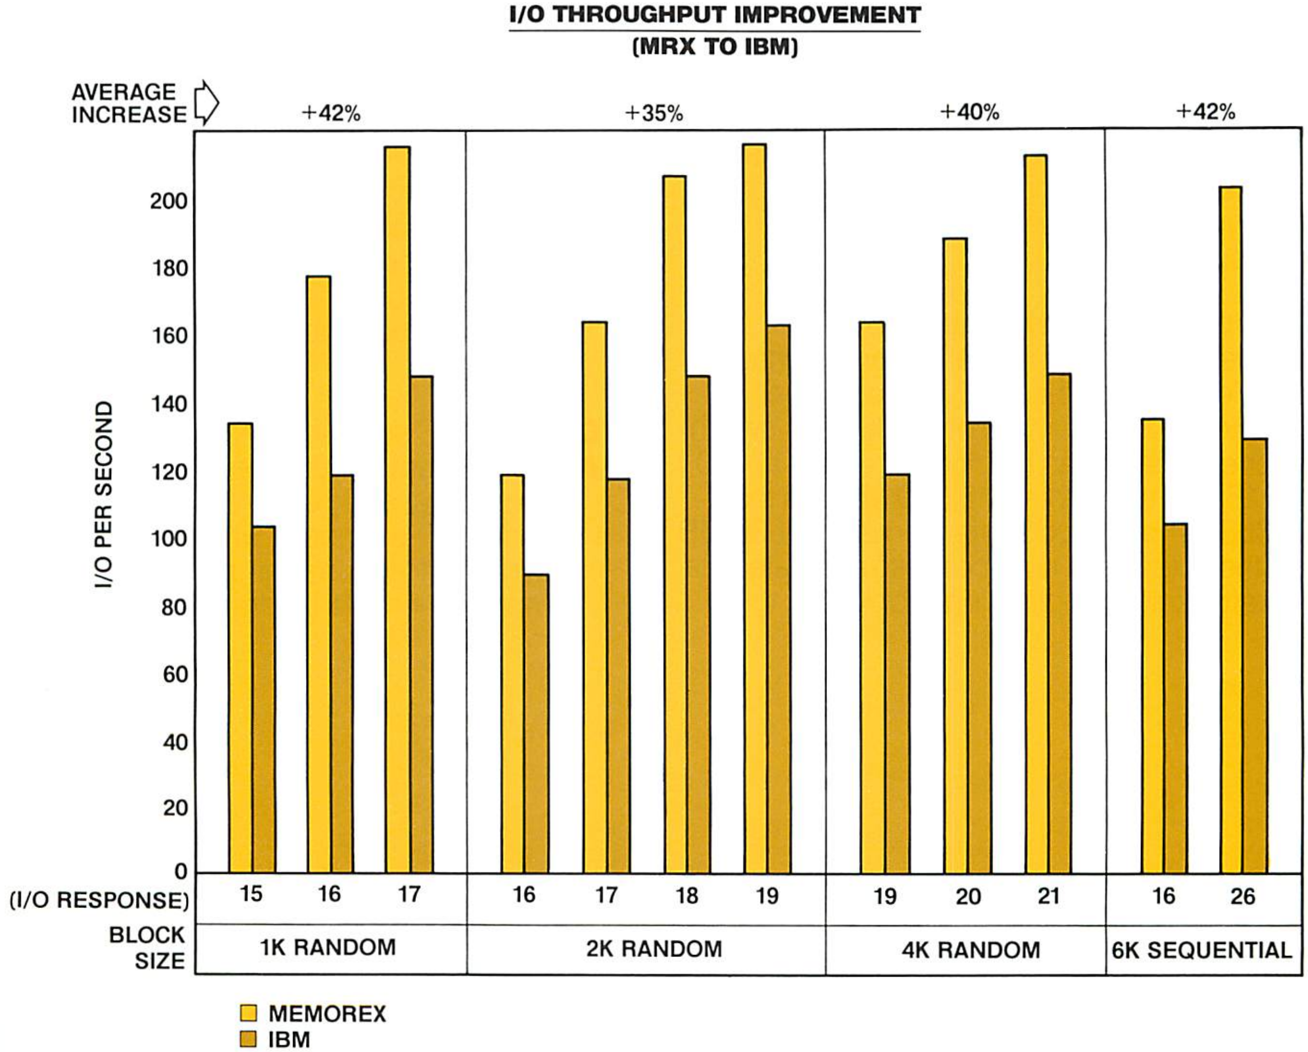

Figure 6: Benchmark Comparisons.

Q: Do I need additional or special software for MAPS?

A: No. MAPS is transparent to the operating system and will be supported by any operating system supporting 3380 subsystems.

Q: What is Dual Path?

A: Dual Path is a feature which provides a second path to each actuator within the string, permitting simultaneous read/write operations on any two different actuators within the string — even on the same Head Disc Assembly (HDA).

Q: How does it improve throughput performance?

A: Each of the two data paths can control any actuator in the string by means of the Maximum Availability Path Selection (MAPS). This feature improves data availability by reducing the Unit of Contention, False Device or Actuator Busy and missed RPS interrupts. Throughput is further facilitated by an improved protocol that speeds communications between components of the system.

Q: What is the Unit of Contention?

A: The amount of data that becomes unavailable due to blocking by a data transfer at any one actuator. With the 3380 2x4 switch, one actuator can block out 3 others at 630 MB each for a total of 2,520 MB. On the 3680, the Unit of Contention is only 630 MB (the actuator in use) because alternate data paths solve the blocking problem.

Q: What is False Device Busy?

A: False Device or Actuator Busy is a condition where a component constituting the path to the device or actuator is busy and cannot be recognized by the I/O control software. In this situation, it falsely appears that the device or actuator is busy. If, for example, an internal path of a 3380 is busy because a particular actuator is in use, attempts to utilize other actuators on the same path will result in a false busy.

Q: What is a Missed RPS Interrupt?

A: A Rotational Position Sensing Interrupt signals that a desired data location is about to pass beneath the read/write head. If the data location cannot be accessed due to the blocking problem discussed earlier, a Missed RPS Interrupt occurs. This introduces a delay of at least 16.6 ms which is the time required for an additional 360° rotation of the disc.

Q: What is Dynamic Reconnect?

A: Dynamic Reconnect is a feature of the MVS/XA operating system which uses hardware in the channel subsystem to provide a different reconnect path back to the processor.

Q: Will MAPS support Dynamic Reconnect?

A: Yes. Current 3680's offer comparable throughput and performance to 3380's with Dynamic Reconnect; when Memorex 3680 Dynamic Reconnect is implemented, the improvement will be dramatic.

Q: What is a 3380 internal path?
A: A path within a 3380 string to which up to four actuators may be attached serially.

Q: What is Extended Architecture?
A: Extended Architecture (XA) is a new hardware architecture developed by IBM that provides 31-bit addressing and moved I/O processing from the CPU to an External Data Controller (EXDC).

Q: Is there any internal path queueing?
A: No. I/O queues are not formed for the internal paths of a 3380 string. Should an I/O be issued for an actuator on an internal path already busy, a false busy will occur.

Q: Does MAPS alter the ratio of 3680's required to replace 3350's?
A: Yes. For the 3380, IBM at times has recommended a conversion ratio of 1.5 to 1; however, the 3680 ratio is 2 to 1. This is due to the throughput improvement of the 3680 which permits greater effective use of the storage capacity.

Q: Does block size affect the performance of MAPS?
A: Yes. The improved protocol of MAPS reduces the overhead associated with every I/O operation. Therefore should less effective smaller block sizes have to be used with 3380 technology, the 3680 benefit will increase.

Q: Is path utilization affected by MAPS?
A: Yes. Due to the enhanced protocol, utilization of components that constitute a path is reduced. Therefore, for the same amount of path utilization, the number of I/O's will be greater.

Q: What channel queueing algorithm should I use for MAPS?
A: To insure both paths are utilized effectively with a single processor, the Channel Rotate Balanced Scheduling algorithm should be used. When multiple processors are attached to the 3888, then the Last Channel Used Scheduling algorithm should be used.

Q: Do other PCM's have the equivalent of MAPS?
A: Not really. Other PCM's offer Dual Port capability which share similarities with MAPS, however the improved protocoling of MAPS together with staggered index marks improves the throughput potential considerably.

THE NEXT STEP

. . . is to prove it to yourself. In addition to the conclusions reached via benchmarks, it is possible to predict with considerable accuracy the improvement in throughput performance for any given DP operation. Using an advanced analytical model, Memorex Systems Engineers can perform on-site evaluations and provide you with documented projections, typically in one day.

An even more exacting evaluation can be performed in Santa Clara, California, using your software to develop comparisons on both 3380 and 3680 subsystems.

We invite you to be amazed.

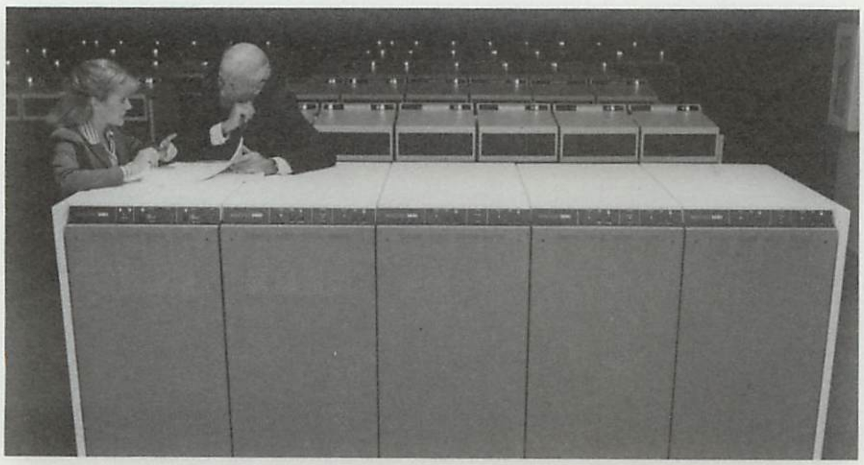

... is to prove it to yourself. In addition to the conclusions reached via benchmarks, it is possible to predict with considerable accuracy the improvement in throughput performance for any given DP operation. Using an advanced analytical model, Memorex Systems Engineers can perform on-site evaluations and provide you with documented projections, typically in one day.

An even more exacting evaluation can be performed in Santa Clara, California, using your software to develop comparisons on both 3380 and 3680 subsystems.

We invite you to be amazed.

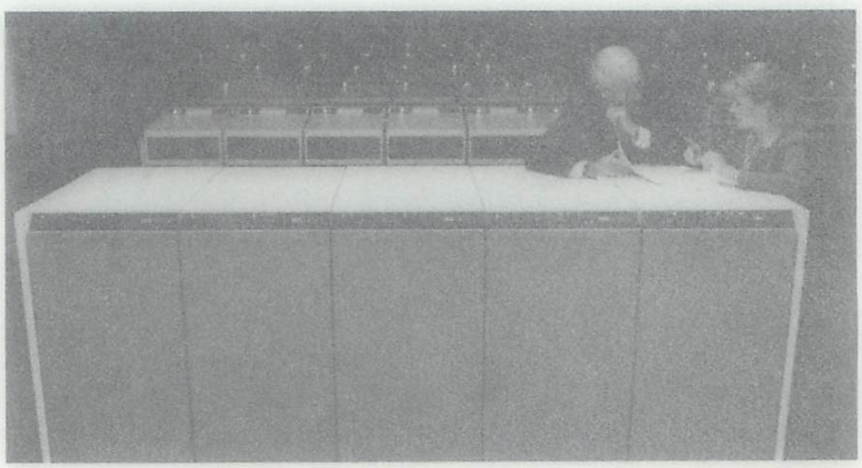

Memorex Corporation

San Tomas at Central Expressway
Santa Clara, California 95052
Telephone: (408) 987-2301

When it matters, make it Memorex T.M.

MEMOREX