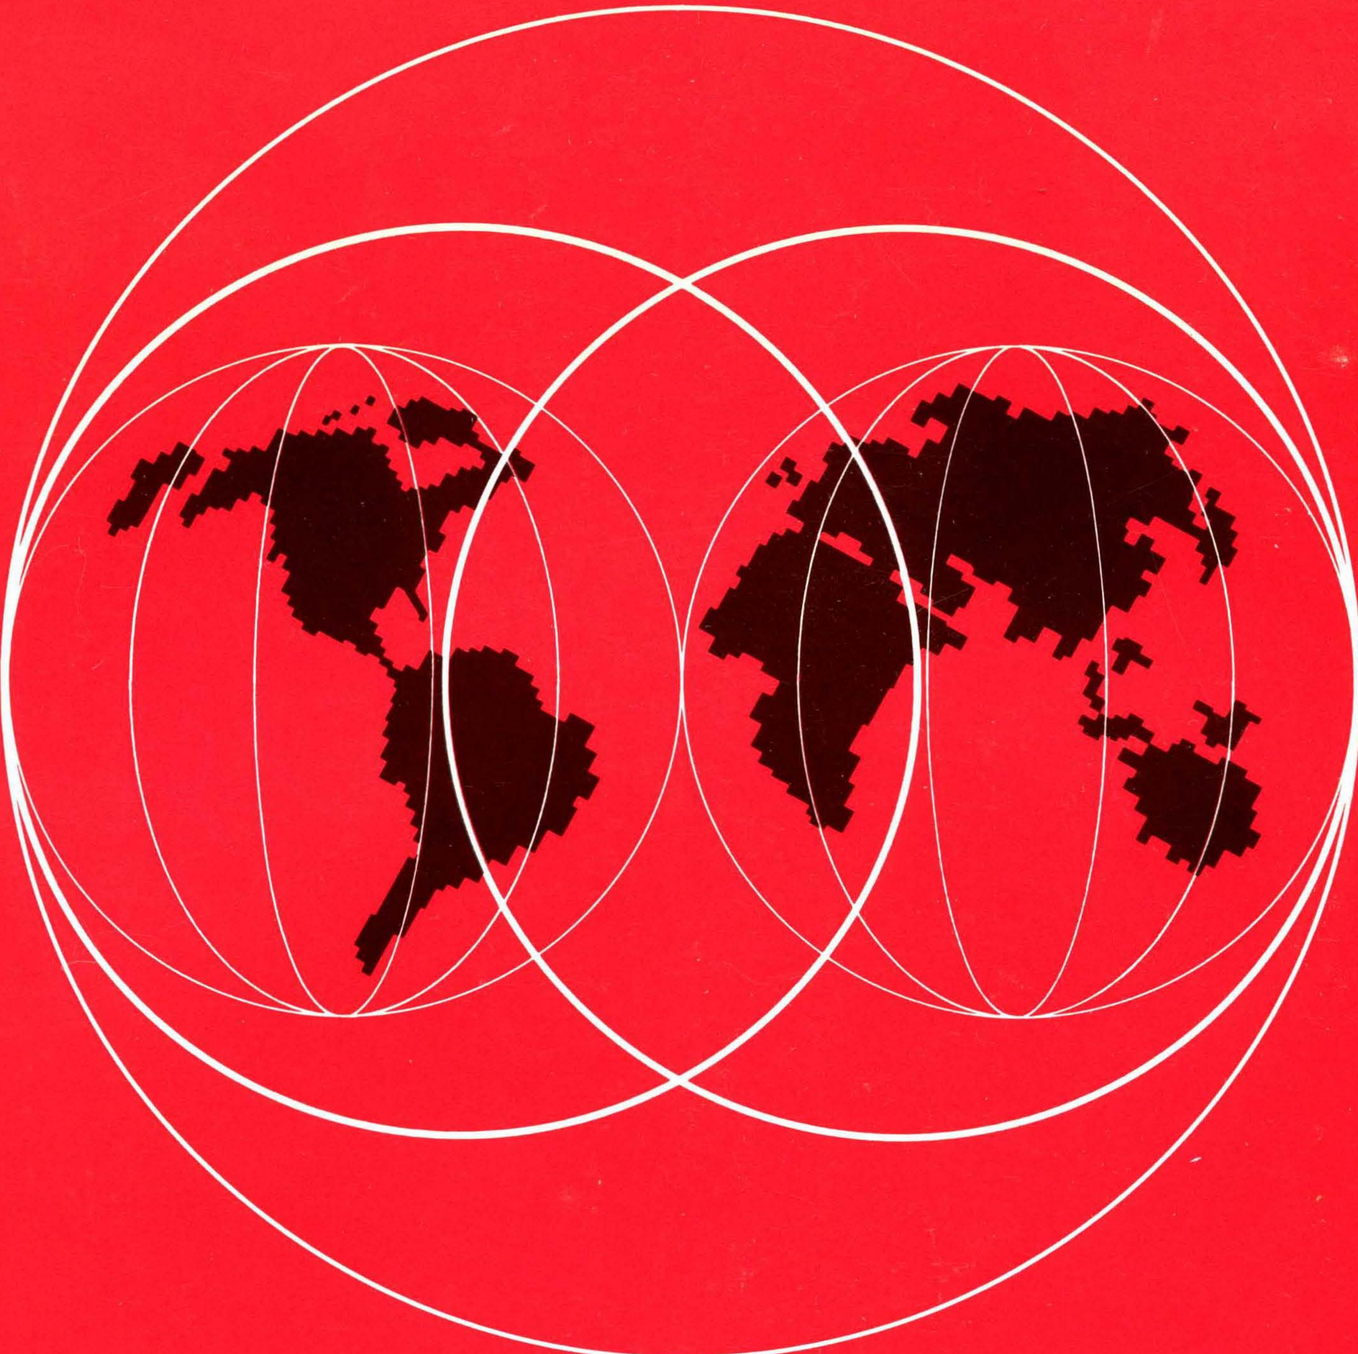

IBM

International Systems Centers

MVS/XA DASD Configuration Guidelines

IBM INTERNAL USE ONLY ZZ10-5049-00

MVS/XA DASD Configuration Guidelines

Document Number ZZ10-5049-00

June 1984

IBM World Trade Corporation
International Systems Center
Poughkeepsie, New York
USA

This document was produced using the following products:
IBM Document Composition Facility (program number 5748-XX9)
IBM Graphical Data Display Manager (program number 5748-XXH)
IBM 3800-3 Page Mode: Text and Image Merge System (program number 5798-DPK)
The master was printed on the IBM 3800 Printing Subsystem Model 3.

First Edition (June 1984)

This bulletin is an **IBM INTERNAL USE ONLY** document and, therefore, its use is confined to IBM employees. The intent is to provide insight into the capability of IBM products for IBM Marketing Representatives and System Engineers.

You may discuss pertinent information from this bulletin with a customer, and you can abstract pertinent information for presentation to your customer. However, you may **not** leave a copy of this information with them **unless the information has also been distributed in formal IBM documents.**

The opinions expressed, together with the suggestions and recommendations made in this document, are strictly those of the authors and have not been subjected to any formal IBM review. They must, therefore, be individually evaluated for each particular situation.

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "as is" basis **without any warranty expressed or implied**. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM program product in this publication is not intended to imply that only IBM's program product may be used. Any functionally equivalent program may be used instead.

A form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM World Trade Corporation, International Systems Centers, Department H52, Building 930, PO Box 390, Poughkeepsie, N.Y. 12602. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

PREFACE

This document presents guidelines for the configuration of DASD in an MVS/XA environment. The benefits that may be realized from the Dynamic Path Reconnection capability of the Dynamic Channel Subsystem and the IBM 3380 DASD are explored in detail. These benefits are particularly relevant to a migration from MVS/370 to MVS/XA in terms of 3380 DASD performance, configuration, and availability. Considerations are also presented for the configuring of IBM 3350 DASD in an MVS/XA environment.

The guidelines presented are based on extensive measurements performed in a benchmark environment. The measurements were carried out on a large number of possible 3380 DASD configurations using IMS/VS, TSO, and Batch workloads.

Organization and Content

The organization and content of each chapter are as follows:

- **Chapter 1, Introduction** describes briefly the objectives of the document and the scope of the measurement project from which the guidelines were developed.

- **Chapter 2, Configuration and Workload Descriptions** briefly describes the DASD configurations and workloads that were measured in order to develop the 3380 guidelines.

- **Chapter 3, Technical Concepts** discusses the technical concepts surrounding the Dynamic Path Reconnection capability of the 3380 DASD and the 370-XA Dynamic Channel Subsystem. The groundwork for performance expectations is established as well as defining the terminology used throughout the document.

- **Chapter 4, Configuring 3380 for MVS/XA** presents through three scenarios the expected benefits that arise from the implementation of Dynamic Path Reconnection. These benefits are expressed primarily in terms of decreased Device Response Time and increased Actuator I/O Rate. Guidelines are presented relative to the range of expected benefits and possible configuration alternatives.

- **Chapter 5, Configuring 3350 for MVS/XA** presents considerations for the use of Preferred Path specification and the configuration of mixed 3350 and 3380 DASD.

Contributors

This guide has been produced by the following staff members and residents of the International Systems Center - Poughkeepsie:

- Craig Aitken - IBM Australia
- Dean Beavis - IBM New Zealand
- Dries de Jong - IBM Netherlands
- Heinrich Ilg - IBM Switzerland
- Terry Jones - IBM UK
- Yoichi Kitazume - IBM Japan
- Stewart Lunn - IBM UK
- Piero Mancini - IBM Italy
- Ramon Ramirez - IBM Spain
- Fritz Staheli - IBM Canada
- Ron Wanless - International Systems Center

Comments by readers of this guide will be greatly appreciated. Please feel free to return the comment form attached to the back of this document.

Acknowledgements

The authors are indebted to the many individuals who took time from their busy schedules to review this document and provide technical corrections and suggested improvements, in particular the following:

- J.E. Andrada - GPD Los Gatos
- E.I. Cohen - DSD Poughkeepsie
- S. Goldstein - GPD Santa Teresa
- D.W. Hunter - DSD Poughkeepsie
- A.C. Munce - GPD San Jose
- E.J. Nauss - GPD San Jose
- G. Neaga - ISC Poughkeepsie
- U. Pimiskern - IBM Austria
- L. van Bommel - IBM Netherlands
- R.J. Wicks - WSC Gaithersburg

TABLE OF CONTENTS

Introduction 1

Configuration and Workload Descriptions 3
Configuration Descriptions 3
Workload Descriptions 10
  IMS/VS Workload Description 10
  Batch (CB80) Workload Description 12
  TSO/SPF Workload Description 13

Technical Concepts 15
Introduction 15
Terminology 15
Dynamic Path Selection 18
What Can be Expected from DPR? 21
Configurable Units 24
Workload Characteristics 27
Measurement Interpretation 29

Configuring 3380 for MVS/XA 31
Introduction 31
Scenario One - Improving the Device Response Time 33
  Objectives 33
  IMS/VS Workload 34
    Summary of Results 35
    Observations 36
    Guidelines 37
  Batch Workload 38
    Summary of Results 39
    Observations 39
    Guidelines 40
  TSO Workload 41
    Summary of Results 42
    Observations 42
    Guidelines 44
Scenario Two - Increasing the Actuator I/O Rate 45
  Objectives 45
  IMS Workload 46
    Summary of Results 47
    Observations 48
    Guidelines 51
  Batch Workload 52
    Summary of Results 53
    Observations 53
    Guidelines 54
  TSO Workload 55
    Summary of Results 56
    Observations 56
    Guidelines 57
Scenario Three - Expanding the Configuration 58

Objectives 58
Observations and Guidelines 61
A-type versus B-type Configurations 61
A-type versus C-type Configurations 62
B-type versus C-type Configurations 63
A-type versus E-type Configurations 64
Channel Path Utilization 64
Channel Capacity in Terms of I/O Requests per Second 66
Actuator I/O Rate 67
Size of a Configurable Unit 68
Summary 69

Configuring 3350 for MVS/XA 73
Introduction 73
Preferred Path Specification 73
Preferred Path and 3350 75
3380 and 3350 Mixed Configurations 77

Appendix A. Bibliography 81

LIST OF ILLUSTRATIONS

Figure 1. A-type Configurations 5
Figure 2. B-type Configurations 6
Figure 3. C-type Configurations 7
Figure 4. D-type Configuration 8
Figure 5. E-type Configurations 9
Figure 6. 3380 Configurable Unit 19
Figure 7. 21
Figure 8. Single Configurable Unit - MVS/370 and MVS/XA 24
Figure 9. Multiple Configurable Units - MVS/XA 25
Figure 10. Comparing Configurable Units 26
Figure 11. Connect Time and Access Density vs Number of Actuators 27
Figure 12. Device Connect Time and Disconnect Time 28
Figure 13. Improvement in Device Response Time - IMS/VS Databases 35
Figure 14. 38
Figure 15. Improvement in Device Response Time - Batch Datasets 39
Figure 16. Improvement in Device Response Times - Batch Datasets 40
Figure 17. Improvement in Device Response Time - TSO User Datasets 42
Figure 18. TSO Terminal Response Time Improvements 43
Figure 19. Improvement in Actuator I/O Rate - IMS/VS Databases 47
Figure 20. Correlation between Actuator I/O Rate and IMS/VS Transaction Rate 48
Figure 21. Device Response Time vs IMS/VS Transaction Rate 49
Figure 22. Actuator I/O Rates in DPR mode - IMS/VS Database Accesses . . . 50
Figure 23. Improvement in Actuator I/O Rate - Batch Datasets 53
Figure 24. Device Response Time vs Ended Transactions- Batch Workload . . 54
Figure 25. Improvement in Actuator I/O Rate - TSO User Datasets 56
Figure 26. Comparison of the Configurations 59
Figure 27. Comparison of an equivalent A-type and C-type configuration. . 62
Figure 28. Spread in channel path utilization for 16 actuator configurations. 66
Figure 29. The possible effect of the size of a Configurable Unit 68
Figure 30. Overview of 3380 Configurable Units 70
Figure 31. 3350 String Switched Configuration 74
Figure 32. 3350 String Switched Configuration - IOCP Example 76
Figure 33. 3380 and 3350 DASD Sharing a Channel Path 77
Figure 34. 3380 and 3350 DASD Sharing Two Channel Paths 79

INTRODUCTION

The introduction of the System/370 Extended Architecture (370-XA) provides a number of significant new facilities beyond System/370. The most significant change from System/370 is in the new I/O facilities provided by the Dynamic Channel Subsystem. These facilities include among others:

- **Path-independent Addressing** of I/O devices which permits the initiation of an I/O operation to any device from any central processor in the complex.

- **Path Management**, whereby the channel subsystem determines path availability, selects a path, and manages any busy conditions encountered while attempting to initiate the I/O operation with the device.

- **Dynamic Path Reconnection (DPR)**, which permits any I/O device using this capability to reconnect to any available channel path to which it has access in order to continue execution of a channel program.

These capabilities complement one other, permitting the channel subsystem and the I/O device to choose the first available path to initiate or continue execution of the chain of operations involved in an I/O request.

MVS/XA is the operating system that takes advantage of the new facilities provided by the System/370 Extended Architecture.

The purpose of this document is to provide insight into the behavior of DASD subsystems in this MVS/XA environment in order that installations may make the most effective use of the new capabilities. In particular, the major objectives are to:

- Provide information on the effectiveness of Dynamic Path Reconnection in several application environments

- Provide guidelines for effective DASD configurations

- Provide a means of assessing alternative approaches for increasing the capacity of a DASD configuration

The development of expectations and guidelines can be accomplished in a number of ways:

- Through analytic modeling and simulation

- Through the use of synthetic I/O 'driver' programs

- Through measurement of 'typical' workloads

With the availability of both the hardware and software, we feel the measurement approach is the most appropriate and would provide the most realistic representation of I/O in a customer environment. This document is based on measurements performed in a 'benchmark' environment using three different workloads and a number of different IBM 3380 DASD configurations. The workloads include: two interactive workloads, TSO and IMS/VS, and a Batch oriented workload. All measurements were performed in an MVS/SP 2.1.1 environment.

The objective of these benchmark runs was to compare results of a number of different 3380 DASD configurations, with DPR and without DPR capability, running the three different workloads. The Resource Measurement Facility (RMF) Program Product and IMSPARS were used to provide the data and make the comparisons.

The development and use of the guidelines is not a strictly scientific approach, but rather one that requires:

- judgement,
- experience, and
- discussion.

The authors believe that the discussion that follows coupled with the reader's good judgement and continuing experience will contribute significantly to the effective design, planning, and management of DASD subsystems in an MVS/XA environment.

CONFIGURATION DESCRIPTIONS

The guidelines presented in this document were developed after measuring what we consider to be a representative sample of the 3380 DASD configurations found in common use today. The purpose of this chapter is to define the measured configurations and the naming convention we have used throughout the document.

The 3880/3380 DASD configurations that were measured in this study have been divided into five 'types' of configuration. Within each 'type' of configuration, a range of actuators was measured. For the purpose of referencing the configurations throughout this document, each configuration has been assigned an arbitrary name. The naming convention used is as follows: a single letter (A - E) denoting the general type of configuration followed by a two digit number (e.g. A16) indicating the total number of 3380 actuators in the configuration. A brief description of each configuration follows:

A-type Configurations

- Two Channel Paths
- Two 3880 Storage Directors
- A single 3380 string varying in length from eight to sixteen actuators

B-type Configurations

- Two Channel Paths
- Two 3880 Storage Directors
- Two 3380 strings, each varying in length from eight to sixteen actuators

C-type Configurations

- Two Channel Paths
- Four 3880 Storage Directors
- Two 3380 strings, each varying in length from eight to sixteen actuators
- Each 3380 string is connected to a pair of 3880 Storage Directors

D-type Configurations

- Four Channel Paths
- Four 3880 Storage Directors
- Two 3380 strings of eight actuators
- Each 3380 string is connected to a pair of 3880 Storage Directors
- This configuration is equivalent to a pair of A08 configurations

E-type Configurations

- Four Channel Paths
- Four 3880 Storage Directors
- Two 3380 strings, each varying in length from eight to sixteen actuators
- Each 3380 string is connected to a pair of 3880 Storage Directors
- Each 3880 Storage Director is connected to a pair of Channel Paths

These configurations are illustrated in the following figures.

Configuration A08

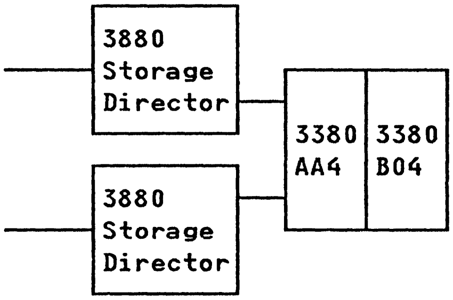

Configuration A12

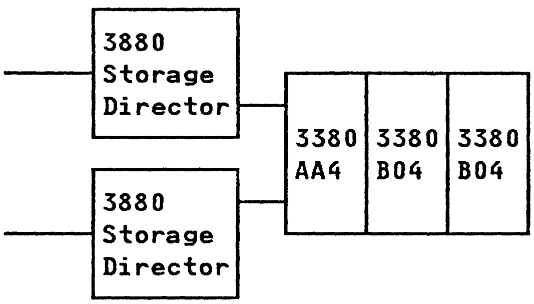

Configuration A16

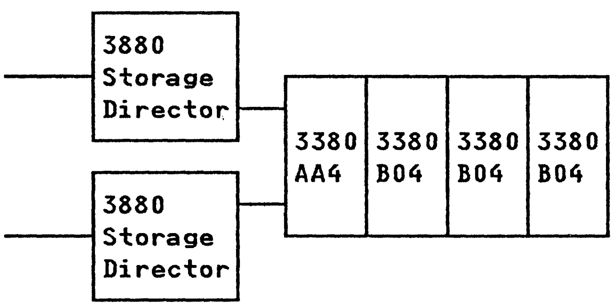

Figure 1. A-type Configurations

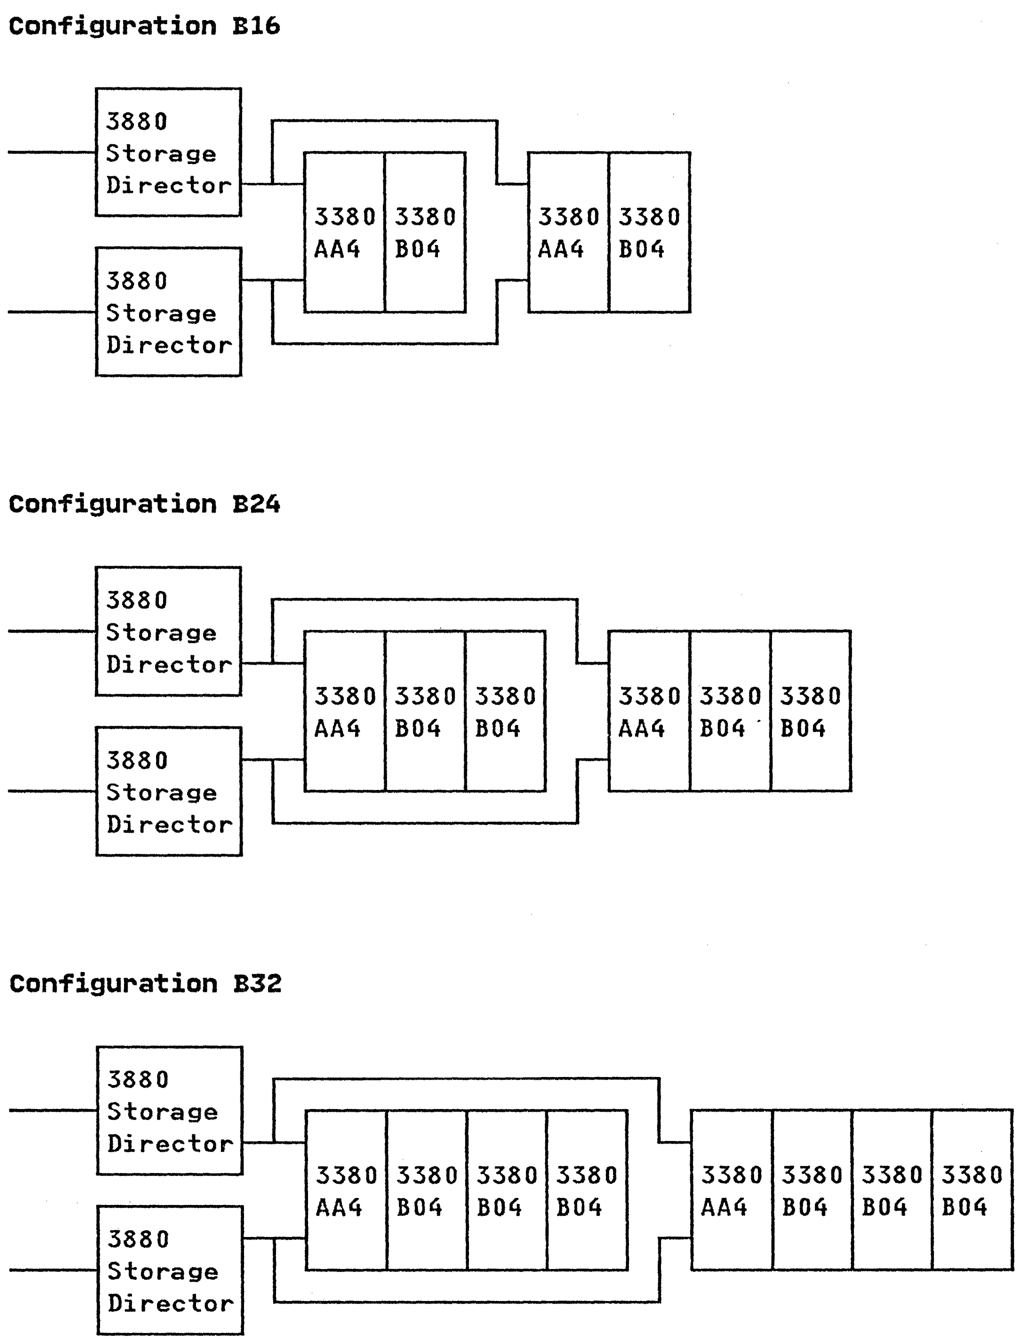

Figure 2. B-type Configurations

Configuration C16

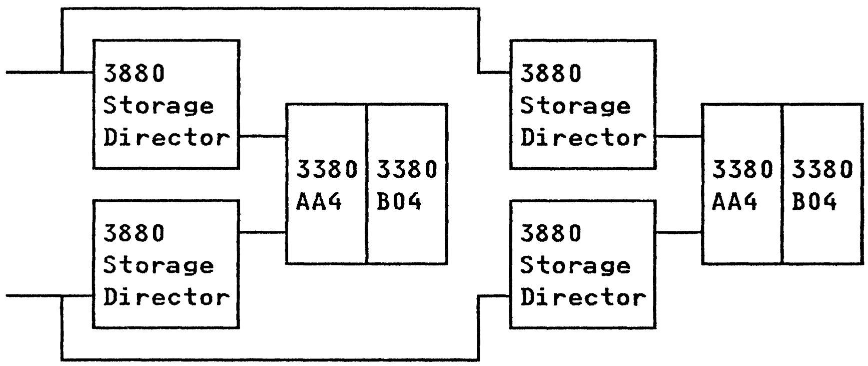

Configuration C24

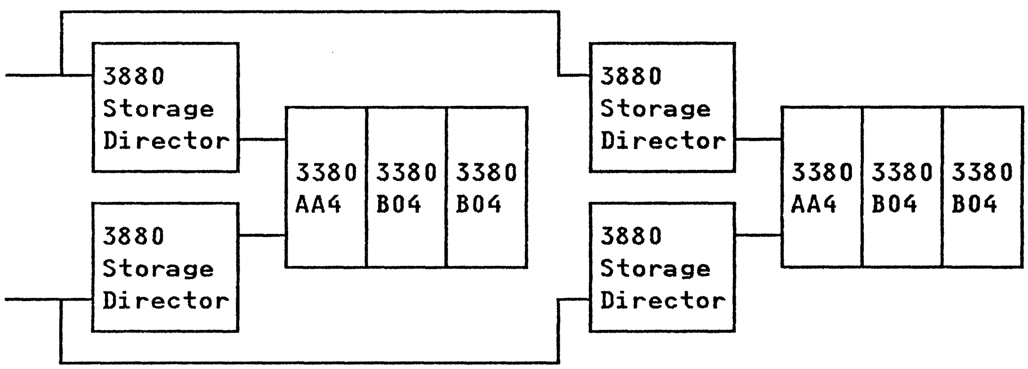

Configuration C32

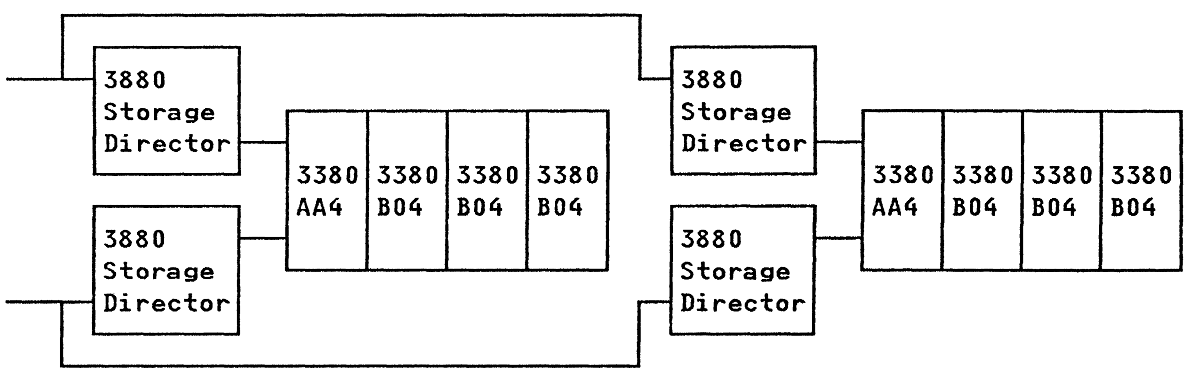

Figure 3. C-type Configurations

Configuration D16

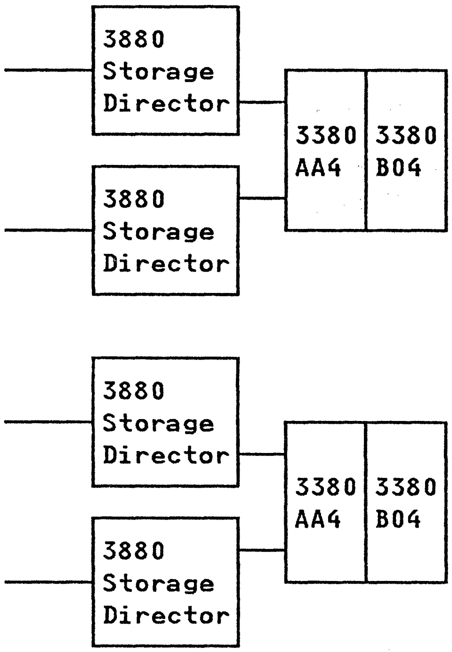

Figure 4. D-type Configuration

Configuration E16

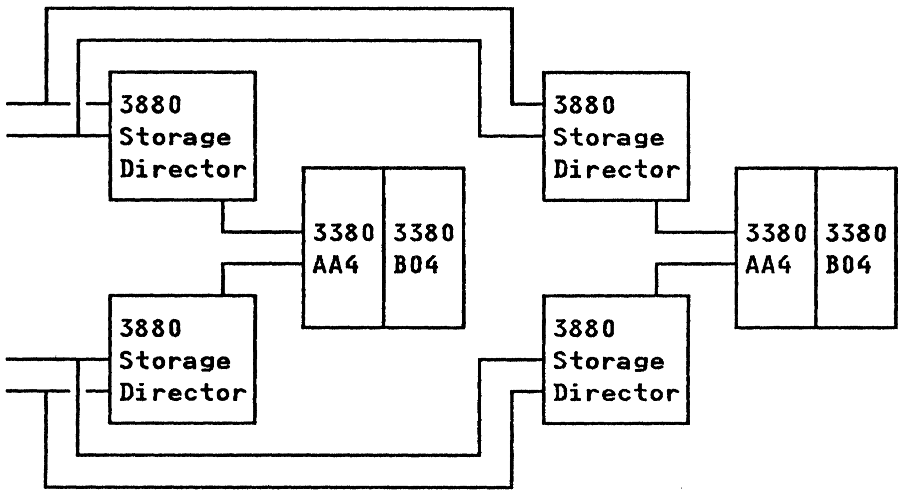

Configuration E32

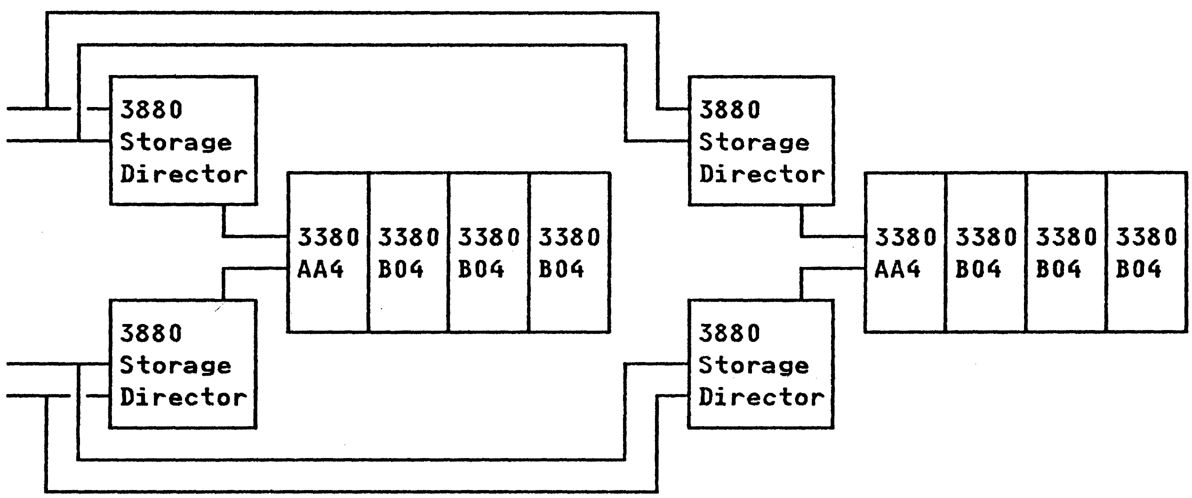

Figure 5. E-type Configurations

WORKLOAD DESCRIPTIONS

IMS/VS Workload Description

The IMS workload used in all of the measurements is called the Data Systems Workload (DSW). It was developed within IBM in 1979 to represent a cross-industry orientation of IMS customer work with regard to intent, complexity, and external characteristics and has been used in many performance measurements including the Large Systems Performance Reference (LSPR) bulletins. Particular attention was paid to the representation of large screen displays of the terminals and the Message Format Services of the workload. Thirty-four transactions were developed to represent the total scope of IMS application activity. From these 34 transactions, 17 specific transactions were chosen and an input distribution was appropriately weighted to match the GUIDE/SHARE IMS Users Survey. This workload issues an average of 11.3 data base calls per transaction and 4 data communication calls per transaction.

The resulting workload has the following general characteristics:

- Conversational applications
- Use of logical relationships
- Use of secondary indices
- Representative data bases
- High-level language programs
- Representative TP characteristics
- Two wait for input transactions
- VSAM/OSAM data base

The following applications are included:

- Order entry
- Receiving / Stock control
- Inventory tracking
- Production specification
- Banking
- Reservations

IMS Program Module Fetching

There are several ways in which IMS can obtain the program modules it needs to process transactions.

- Program Fetch
- Virtual Fetch (available only with IMS 1.2 and MVS/SP1.3 and MVS/XA)
- Preload

In this series of measurements Virtual Fetch was used throughout for all load modules.

Virtual fetch is started via a procedure, and it runs in its own address space. All program modules within data sets referred to in the procedure are assigned to virtual fetch, so that if a module is needed, a virtual fetch is done instead of a program fetch. When virtual fetch is initialized, the programs are copied into a read-only VIO page data set. The program library modules are then moved into the MPRs, once per schedule as needed. The fetch process by which an MPR obtains a program module is a parallel block page-in of the entire program. Virtual fetch can use real storage greater than 16 MB when available. It has a shorter path length than program fetch, so it gives a performance benefit when compared with program fetch. (The IMS Control Region does not use virtual fetch.)

Virtual fetch was chosen over program fetch because of the performance benefit that virtual fetch provides.

Multiple Transactions per Schedule

In IMS, the Control Region selects a Message Processing Region (MPR) to handle each transaction. The MPR then invokes a copy of the appropriate program to process the transaction. This program could be preloaded, loaded when invoked, or already loaded because the last transaction handled by the MPR was the one being scheduled.

If the program is already loaded due to a previous identical transaction, multiple transactions can be processed by each "schedule" (loading of the processing program); hence I/O and processing time can be saved. The higher the transaction rate and the higher the incidence of identical transactions, the greater the probability that more than one transaction will be scheduled for each program load. This phenomenon is also called "coat-tailing". Coat-tailing is possible with either program fetch or virtual fetch if the process limit count (PLMCT) is not set to 1.

The measurements performed for this study did not use coat-tailing. (i.e., PLMCT was set to 1). This meant that a module already in storage could not be reused, but had to be refetched every time it was needed by a new transaction. Coat-tailing can cause problems in data interpretation when comparing performance between different runs since a consistent amount of I/O per transaction will not be done. To avoid this distortion, coat-tailing was not used.

Batch (CB80) Workload Description

CB80 (Commercial Batch 1980) is a jobstream intended to represent a commercial batch workload. It consists of 30 jobs, including sorts, COBOL programs, COBOL compiles, an assembly, link-edits, and synthetic jobs. The synthetic jobs use several access methods (BSAM, QSAM, BDAM, and VSAM). They also do fixed point arithmetic, GETMAIN, and FREEMAIN. Their private area reference patterns are designed to represent those observed in customer workloads. There are no tight loops within any of the synthetic programs. the components have appeared in previous measurement jobstreams.

The CB80 jobstream consists of a set of synthetically generated jobs of a predetermined mix. The basic jobstream consists of 30 jobs, some of which are repeated within the basic set. Included are several sorts, an assembly and two COBOL/VS compilations along with an assortment of synthetic programs. The synthetic programs are derived from a single set of subroutines. The characteristics of the various programs depend on the way in which they have been assembled, and the various parameters that have been coded within them. Their operation cannot be changed without re-assembly.

CB80 uses both temporary and cataloged data sets. The latter contain sort data, master file data, read-only data, and data for updates. User catalogs and all permanent data sets are created and loaded by utility programs which are run prior to a series of measurement runs.

Each copy of CB80 consists of 30 jobs and 72 job steps:

- 50 Parameterized steps
- 12 COBOL go steps (including 4 COBOL sorts)
- 4 ICEMAN sort steps
- 3 Linkedit steps
- 2 COBOL compile steps
- 1 Assemble step

The jobstream uses a range of dataset types, including VSAM, BSAM, QSAM and BDAM, but not ISAM. A total of some 372 datasets are used, including SYSOUT, but excluding JOBCAT and JOBLIB statements.

In the 66 GO steps, there are 146 permanent datasets, with the following distribution:

- 99 QSAM
- 36 BSAM
- 33 BDAM
- 20 VSAM with a CI size of 2048 (includes ESDS, KSDS and RRDS)
- 13 VSAM with a CI size of 4096 (includes ESDS, KSDS and RRDS)

TSO/SPF Workload Description

The TSO/SPF workload used in the measurements is the same workload used for the Large System Performance Reference (LSPR) measurements. The functions of the thirteen scripts used in the measurements are shown below:

Script	Description
ALCDEL	Allocates, edits, and deletes a sequential data set.
ASMEDT	Edits ASM source code and then cancels the edit.
ASMFOR	Assembles using ASM(XF) under PDF and then does a LOADGO via a CLIST.
BROWSE	Browses listings from PLIFOR.
COBEDT	Does line by line edit insert of a COBOL program and then cancels the edit.
COBFOR	Compiles and Linkedits using the COBOL Prompter and Linkage Editor under PDF, then does a TESTCOB using a CLIST. Three PDS datasets are also allocated and deleted.
COPYPRT	Copies and prints a data set using PDF menus.
LOGISPF	Exits PDF and logs the user off the system; then it logs the user back on and issues the PDF command.
PLIEDT	Edits a PL/I source then cancels the edit.
PLIFOR	Invokes the interactive PL/I Checkout compiler under PDF and re-uses or allocates a list data set.
SUBMIT	Edits a job, submits it, and then cancels the edit. The submitted job is scanned and the output put in a 'throw away' class.
TXTEDT	Edits 6 half screens of SCRIPT/VS text, and then cancels the edit.
TXTFOR	Invokes SCRIPT/VS under PDF, and scrolls the output. The listing is then deleted.

The type and frequency of the TSO transactions and in which user work scripts they occur may be found in Appendix D.4 of the Large Systems Performance Reference (LSPR - MVS/XA), form number ZZ05-0409-00.

INTRODUCTION

Before discussing the results of the measurements, it is appropriate to review the hardware features, configuration techniques, and considerations, which lie behind the need for performing the measurements. It is then possible to develop expectations about the nature of any possible improvements in performance, and how to configure taking advantage of the new facilities of MVS/XA and DPR on the 3380 device.

After introducing some of the terminology used we will discuss the following topics:

- Dynamic Path Selection and Dynamic Path Reconnect
- What can be expected from DPR
- The concept of Configurable Units
- Workload characteristics
- Measurement interpretation

TERMINOLOGY

In this section we define some of the terms that are used throughout this document. While we will be using these terms in the normally accepted manner, some qualification is required relative to our discussion. For example, when we are talking about the 'Total I/O Rate', we will not be referring to the total 'system I/O rate' but rather the total I/O rate of the DASD subsystem under discussion.

Total I/O Rate in the context of this document, is the total number of I/O requests per second for the measured 3380 DASD configuration.

This value is obtained from the 'DEVICE ACTIVITY RATE' field of the RMF Direct Access Device Activity report.

Actuator I/O Rate is the average number of I/O requests per actuator for a measured DASD configuration. This value is equal to the 'Total I/O Rate' for a measured 3380 DASD configuration divided by the number of actuators in the configuration.

Device Response Time is the time required for a device to complete an I/O request. This time includes the total hardware service and software queueing time for an average I/O request to the device.

This time is obtained from the 'AVG RESP TIME' field of the RMF Direct Access Device Activity report and is reported in milliseconds.

IOS Queue Time is the time an I/O request must wait on an IOS queue before a SSCH instruction can be issued. This delay occurs when a previous request to the same device or subchannel is in progress.

This time is obtained from the 'AVG IOSQ TIME' field of the RMF Direct Access Device Activity report and is reported in milliseconds.

Pending Time is the time an I/O request must wait in the channel subsystem after acceptance of the SSCH and before acceptance of the first CCW by the DASD subsystem.

This time is obtained from the 'AVG PEND TIME' field of the RMF Direct Access Device Activity report and is reported in milliseconds.

Disconnect Time is the time a device is disconnected from the channel path during an I/O operation. This reflects the time when the device is in use but not transferring data. It includes the time the device might disconnect to perform functions such as Seek/Set Sector, as well as any reconnection delay.

This time is obtained from the 'AVG DISC TIME' field of the RMF Direct Access Device Activity report and is reported in milliseconds.

Device Connect Time is the total time the device was connected to a channel path for the purpose of searching and transferring commands or data between the device and main storage. In addition to the actual data transfer time, this includes the search time needed to maintain channel path, control unit, and device connection.

This value is obtained from the 'AVG CONN TIME' field of the RMF Direct Access Device Activity report and is reported in milliseconds.

Device Utilization is the percentage of time when the device was in use. It includes both Connect and Disconnect Time.

Channel Path Utilization is the percentage of time that the channel path was busy. This value is calculated by RMF from sampled data collected by the SRM.

This value is reported in the 'PERCENT CH PATH BUSY' field of the RMF Channel Path Activity Report.

In our analysis, we have averaged the utilizations of the Channel Paths connected to the measured 3380 DASD configuration. This averaging was possible since the channel paths were well balanced due to the symmetric nature of the 3380 DASD configurations.

Terminal Response Time

For the TSO workloads, we have used the 'AVG TRANS TIME' as reported in the RMF Workload Activity Report.

For the IMS/VS Workloads, we have obtained the value from the 'Average Transit Time (Total)' field in the IMSPARS Transit Time Analysis Report.

Transaction Rate

For the TSO workloads, we have used the 'RATE PER SECOND' as reported for 'INPUT TERMINAL WAIT' from the RMF Paging Activity Report.

For the IMS/VS Workloads, we have obtained the value from the 'Transactions Processed by the Above Programs' field in the IMSPARS IMS Internal Resource Usage Report.

For the Batch workload, we have used 'Service Units' as reported in the RMF Workload Activity Report.

DYNAMIC PATH SELECTION

Dynamic Path Selection (DPS) is a feature that is unique to the IBM 3380.[1] It provides four functions:

- String Switch
- Alternate Controller
- System Related Device Reserve
- Dynamic Path Reconnection

The first three functions are supported by MVS/SP Version 1 and MVS/SP Version 2 (MVS/XA). The fourth function, Dynamic Path Reconnection, is supported by MVS/SP Version 2 only.

String Switch

The 3380 DPS is implemented with two head of string controllers, each of which is attached to a 3880 Storage Director to provide two paths to a string of 3380 devices or actuators. This provides simultaneous data transfer from two 3380 spindles on the same string. The reader is referred to Figure 6 on page 19.

The actuators in a 3380 string are attached to four internal data paths, which in turn are attached to the two string controllers by means of a 2x4 switch array. A full string has sixteen actuators, and is organized in the following manner:

Path 0 - Actuators 0,1,8,9	Path 1 - Actuators 2,3,A,B
Path 2 - Actuators 4,5,C,D	Path 3 - Actuators 6,7,E,F

From this it can be understood that a string consisting of only four actuators (two spindles) will have only two internal data paths, one with actuators 0 and 1, and the other with actuators 2 and 3.

At any instant in time, only one head of string controller can have access to a particular internal path. However, two of the four internal paths to the access mechanisms can be actively transferring data at the same time.

Alternate Controller

There are two head of string controllers with 3380 DPS, which are always active. If one fails, its path can be taken offline, and the other controller will manage the four internal paths to the string on its own. Access to all actuators in the 3380 string will still be possible, however only one actuator may transfer data at a time.

---

[1] For more detailed information on the 3380, as well as an aid for installing the devices, refer to Washington Systems Center Technical Bulletin 3380 DASD Features, Installation and Conversion, GG22-9308-00.

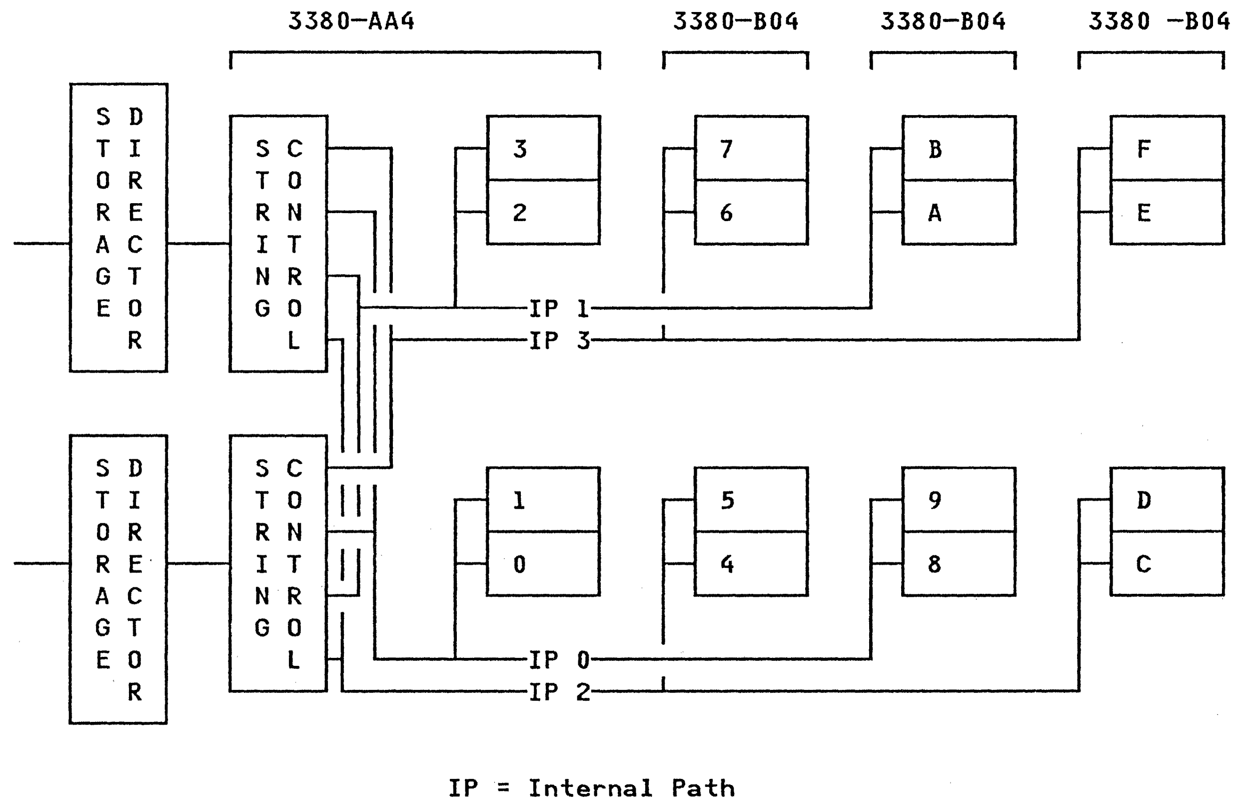

Figure 6. 3380 Configurable Unit

System Related Device Reserve

3380 DPS provides a system wide reserve facility, whereby a device reserved to a particular system can be accessed by that system over any of the paths provided to that system. If a path fails, then alternate paths may be used to access the device. Any other systems are still denied access until the reserve is released.

Dynamic Path Reconnection

The Dynamic Path Reconnection (DPR) capability allows disconnected DASD operations, such as Seek/Set Sector, to reconnect over any available channel path from the device to the MVS/XA system initiating the I/O request. This capability is currently implemented in the 3880/3380 DASD and the 370-XA Dynamic Channel Subsystem.

The effect this capability has on 3380 I/O operations in an MVS/XA environment is that the **'apparent'** Channel Path utilization (as seen by the device) will be lower than the **'physical'** Channel Path utilization (as seen by RMF). This has a direct impact on the Set Sector Reconnect Delay encountered when a device attempts to reconnect to a channel path. The average delay will be shorter since the device has a higher probability of reconnecting to any available channel path than to a specific path.

System Related Device Reserve and Dynamic Path Reconnection are implemented through Path Grouping. Path Grouping allows the 3880 Storage Director to know which MVS systems are attached to its channel interfaces. At MVS initialization time or at VARY ONLINE, MVS/XA or MVS/370 creates a Path Group consisting of the available paths to the device. When running under MVS/XA, a 3380 device may reconnect to the first available path in the group of paths belonging to the system, instead of only to the path that originated the operation, as is the case when running under MVS/370. This increase in opportunities for reconnection potentially represents an improvement in performance.

Example

Such performance improvement is readily quantified for the Set Sector reconnect situation, which usually occurs once per access. An actuator that is at its target sector raises a very brief signal. If no reconnection path is free, then 16.6 milliseconds (one revolution) is lost until the target sector is again in sight and reconnection is again attempted. If **p** = probability of no free path, it turns out that the average number of lost revolutions per Set Sector is equal to **p/(1-p)**. Here then is how access times gets extended, depending on **p**:

p	Avg Reconnection Delay
0.1	1.8 msecs
0.2	4.2 msecs
0.3	7.1 msecs
0.4	11.1 msecs
0.5	16.6 msecs

Thus, if for a particular workload, DPR were to pull **p** down from 0.4, say, to somewhere between 0.2 and 0.3 (a realistic expectation), that would have a dramatic effect on average DASD access time.

One of the primary objectives of this document is to explore the performance consequences of this fourth item: Dynamic Path Reconnection.

WHAT CAN BE EXPECTED FROM DPR?

In technical terms the effect of DPR can be stated very briefly: it reduces the Set Sector Reconnect Delay time. A secondary effect is the reduction of the front-end queueing time in IOS. There are several ways an installation can realize benefits from these effects:

1. Improve or decrease the Device Response Time for an I/O request
2. Increase the Actuator I/O Rate to an existing configuration
3. Expand the capacity of a 3380 configuration as a consequence of the second point.

One of these three approaches, or a combination, would provide significant benefits to an installation when migrating from MVS/370 to MVS/XA, i.e. migrating from a 'Non-DPR' environment to a 'DPR' environment.

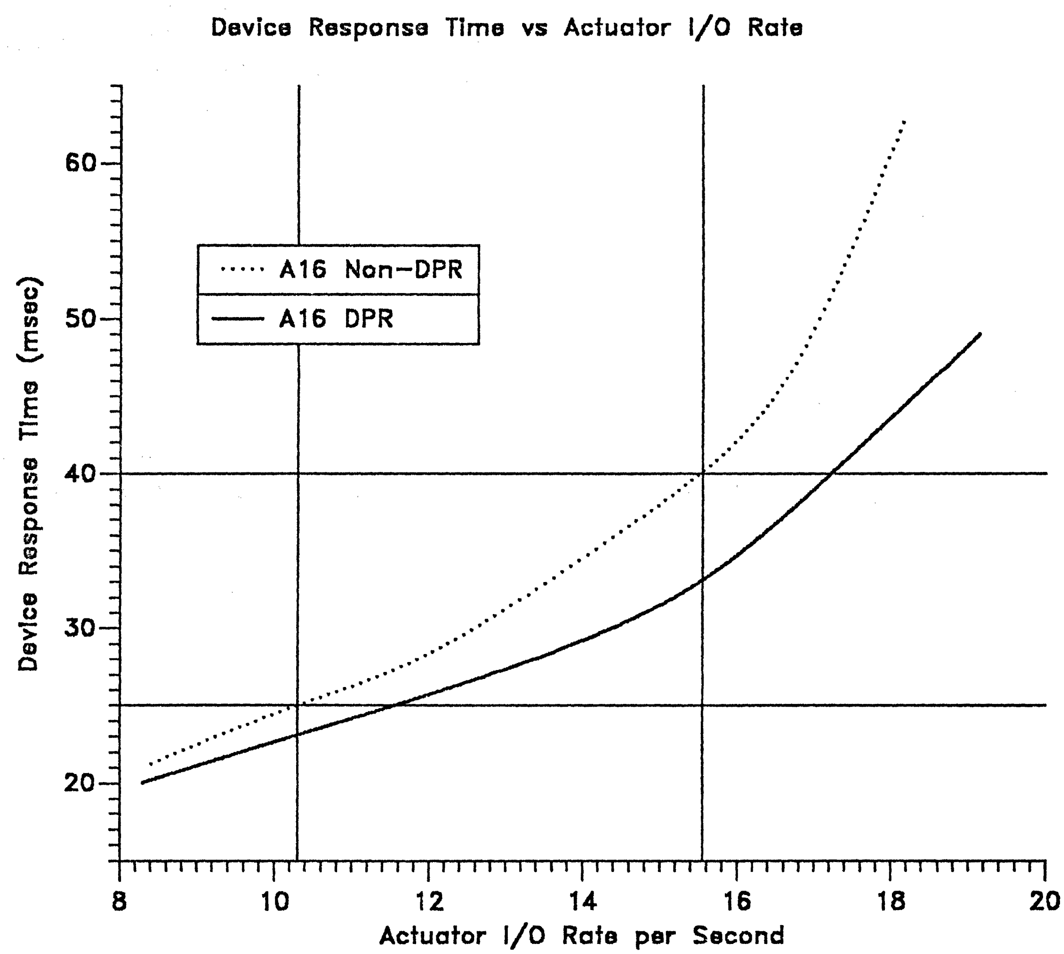

Figure 7.

Figure 7 on page 21 is included here to understand better the the controlling parameters. It illustrates two curves reflecting the results of measurement runs performed using the same configuration, one running with DPR and the other running without DPR.

Improving the Device Response Time

An installation may wish to decrease the Device Response Time for the average I/O request while:

- maintaining the same Actuator I/O Rate, and
- without changing their current configuration.

Clearly this type of improvement would have a positive influence on interactive applications such as TSO, IMS/VS, or CICS. It would help to achieve sub-second response times for TSO and/or the ability to provide consistent application response times.

As well, the better Device Response Time should lead to better DB/DC response time. This would mean less time spent using buffers and message regions leading to lower virtual storage requirements or more transactions from the same available virtual storage.

In the terms of Figure 7 on page 21 it would mean that you try to realize an improvement in the vertical direction.

Increasing the Actuator I/O Rate

Alternatively, an installation may wish to increase the number of I/O requests to a given set of actuators while:

- maintaining the same average Device Response Time per I/O request, and
- without changing their current configuration.

This would be seen as an application growth benefit. The introduction of DPR may provide a means of increasing the **access density** of the 3380 devices. Increasing the access density means that the I/O request rate to a given number of megabytes of stored data can be raised to a higher level while maintaining the same Device Response Times. As an example, more IMS/VS users can access the same databases at the same time. As well, it could be used to store more active data on each device.

In the terms of Figure 7 on page 21 it would mean that you try to realize an improvement in the horizontal direction.

Expanding the Current Configuration

Another option would be to expand the current 3380 configuration by either increasing the string length with additional 3380-B04 units or adding additional strings to the same channel paths while:

- maintaining the same Actuator I/O Rate, and
- maintaining the same Device Response Time.

This would mean that the number of actuators per channel path would be increased as well as the number of I/O requests. Realization of this option would be seen as both a financial benefit and a system benefit, lowering the overall cost of an I/O request and making better use of existing channel paths and storage directors. At the same time, the applications would receive an equivalent level of service.

In the graphical representation in Figure 7 on page 21 this approach implies that the current level of service is satisfactory and you wish to remain on the upper curve and increase the overall capacity of the system.

It is through these three 'scenarios' that we will present the measurement results and configuration guidelines in the next chapter.

CONFIGURABLE UNITS

A Configurable Unit is a convenient way of organizing the overall DASD configuration of a system into manageable units. It is a building block used for setting up and planning a system configuration, and, afterwards for controlling and reporting on the I/O activity and performance. It is frequently used to separate dissimilar types of I/O workload, such as IMS/VS databases, TSO user datasets, paging, and system datasets from one another.

The concept of Configurable Units is not new. It is well understood and has been implemented by many installations for a number of years.

Two of the principal areas to consider in implementation of Configurable Units are availability considerations and the sizing or string length of the configuration.

Availability Considerations

From an availability point of view, the Configurable Unit must be constructed in such a way as to satisfy the requirement of data availability. Normally the Configurable Unit will consist of two or more of the same components at each level, e.g. two channel paths, two storage directors, and one or two string controllers through which the data can be accessed. In this respect nothing is new in MVS/XA. Figure 8 illustrates a 3380 Configurable Unit for either MVS/370 or MVS/XA.

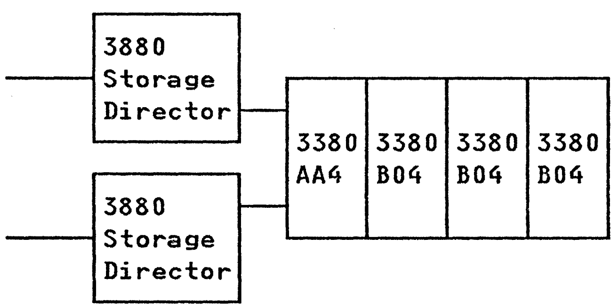

Figure 8. Single Configurable Unit — MVS/370 and MVS/XA

Size of the Configurable Unit

The sizing of the Configurable Unit must be such that performance problems are avoided. It is in this area that the introduction of Dynamic Path Reconnection may change the traditional approach to system configuration. From a performance point of view, it is now possible and valid to combine two Configurable Units as illustrated in Figure 9. This configuration is not recommended in an MVS/370 environment for performance reasons. However, it performs very well when connected to a Dynamic Channel Subsystem running under MVS/XA.

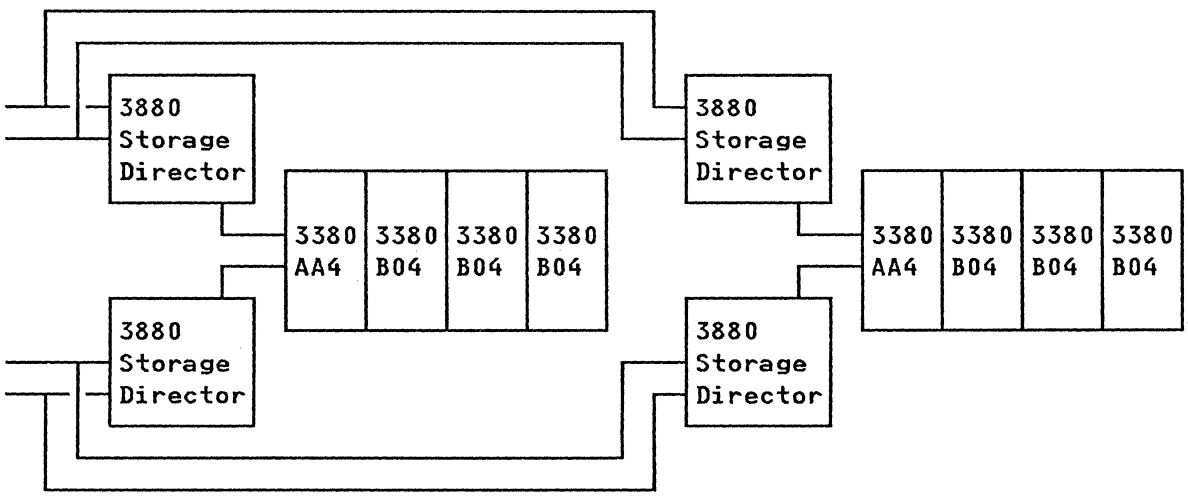

Figure 9. Multiple Configurable Units - MVS/XA

In an MVS/370 environment, this configuration would only be valid when attached to a multiprocessor or dyadic system with two Channel Sets configured. For normal operation many installations would VARY OFF two of the four channel paths because of the negative effect the additional level of DASD switching has on performance. In the event of a component failure, operator recovery action would be required to enable the alternate 'availability' paths.

The negative performance impact is illustrated in Figure 10 on page 26. It illustrates measurement results for the Configurable Units depicted in Figure 4 on page 8 and Figure 5 on page 9. The curves identified as 'E16' represent the situation where all four channel interfaces are enabled, whereas the 'D16' curves are the result of measurement runs with two different channel interfaces per string enabled.

It is clear that the performance results of the D16 and E16 configurations are for all practical purposes identical when running in DPR mode. It does not matter whether or not two paths are varied offline for each string. However, in Non-DPR mode, there is a very definite difference in performance.

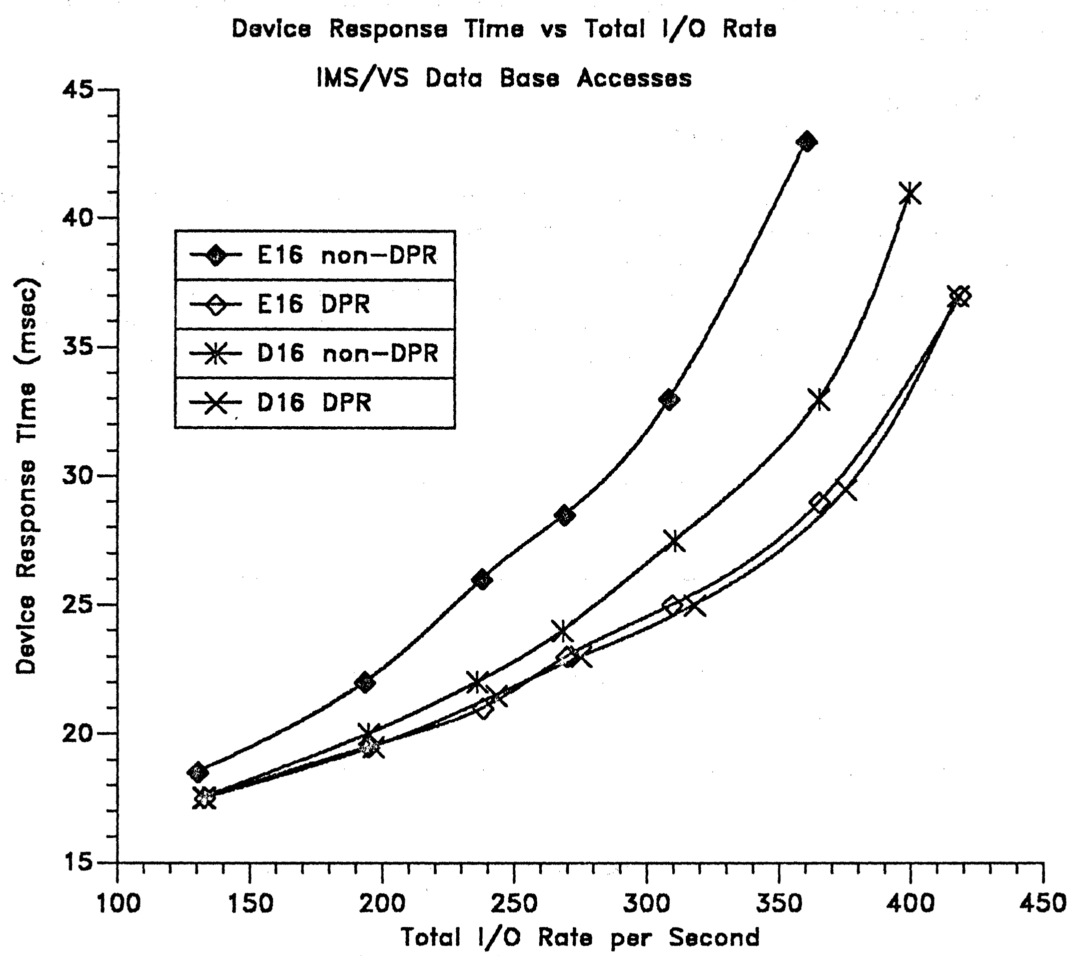

Figure 10. Comparing Configurable Units

The actual sizing of a Configurable Unit is related to the type of workload for which it will be used. For each environment a balance has to be found between the number of channel paths, control units, and devices, while retaining all of the availability requirements.

WORKLOAD CHARACTERISTICS

There are two factors that should be considered when trying to qualify a type of workload in technical terms:

- Access Density
- Duration of an I/O Operation

A short description of these factors is included here, because understanding them helps both with interpretation of the measurement results and most of all in applying these results to a specific environment. Figure 11 illustrates the relationship between these two factors. It attempts to define an expectation in very general terms and the same approach will be used again in the chapter "Scenario Three - Expanding the Configuration" on page 58 to summarize the results of the measurements.

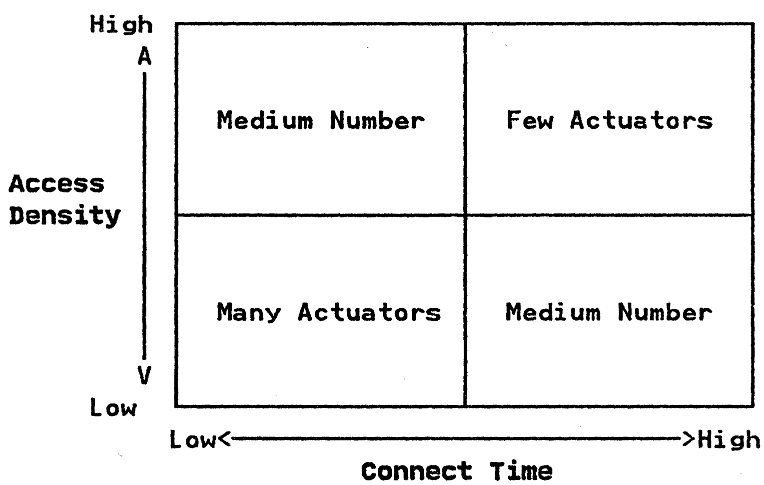

Figure 11. Connect Time and Access Density vs Number of Actuators

Access Density

A workload will generate a high access density when the actuator is asked to perform many I/O operations per second. The result is that relatively few actuators will fill the capacity of the channel. When the access density is low, more actuators can be configured, because each of them is generating less work for the channel.

Duration of an I/O Operation

The duration of an I/O operation can be subdivided into two major portions:

- **Connect Time** - The time during which the attention of the channel is needed.

- **Disconnect Time** - The time during which only the device is busy handling the I/O operation.

Figure 12 represents the overall duration of an I/O from the channel subsystem point of view. SEEK represents the expected Seek Time for the operation. Rotational Delay for the 3380 normally averages approximately 8.3 milliseconds. SSRD stands for Set Sector Reconnect Delay, the duration of which is closely related to the path utilization. Finally, Data Transfer refers to the Search and Data Transfer time, either read or write.

IS	SEEK	Rotational Delay	SSRD	Search	Data Transfer
CT	Disconnect Time			Connect Time	

IS — Initial Selection
CT — Connect Time
SSRD — Set Sector Reconnect Delay
Search — For some channel programs

Figure 12. Device Connect Time and Disconnect Time

Both Connect and Disconnect time are reported in the MVS/XA RMF Direct Access Device Activity Report.

In the measurements performed, Connect Time ranged from 2.4 msec to 5.4 msec, while Disconnect Time ranged from 13 msec to 40 msec. Hence, the relative length of the two portions in the figure is indicative for a real situation. As contrasted with devices that have lower data transfer rates, the time spent transferring data with a 3380 is a smaller portion of the total busy time. Thus, an increase in the amount of data transferred has much less effect on Actuator I/O Rate.

Therefore, using 3380 DASD you may plan a very high Access Density even when the Connect Time is relatively long. However, this can only be achieved by keeping the Channel Path utilization under control. At longer Connect Times, fewer actuators can be configured on a set of channel paths.

This is the second dimension in Figure 11 on page 27. These two elements clearly show up in the measurement results, and will be discussed in more detail later.

MEASUREMENT INTERPRETATION

The measurement runs have produced a large number of RMF reports. Interpretation, trend analysis, and arriving at conclusive statements on the relative improvement of one environment over another is not easily done.

To be able to come to conclusions, a common reference point is needed. For example, Scenario Two investigates the improvement in Actuator I/O Rate at the same Device Response Time in both DPR and Non-DPR mode. It will only be a coincidence when RMF provides such a common reference point.

To address this problem, extensive use has been made of the Graphical Data Display Manager (GDDM), and in particular the Interactive Chart Utility (ICU) of the Presentation Graphics Feature (PGF).[2]

The Interactive Chart Utility was used to generate charts from the appropriate RMF values and provide the common reference base that is required. As an example, Figure 7 on page 21 illustrates Actuator I/O Rate versus Device Response Time for the Non-DPR and the DPR measurements of the same configuration. It is this type of graph that was used to provide the basic information for arriving at the guidelines for Scenarios One and Two.

For Scenario One, Improving the Device Response Time, the difference between the two curves in the vertical direction at the same Actuator I/O Rate represents the improvement of DPR over Non-DPR.

For Scenario Two, Increasing the Actuator I/O Rate, the difference between the two curves in the horizontal direction at the same Device Response Time represents the improvement of DPR over Non-DPR.

---

[2] Graphical Data Display Manager (GDDM), program number 5748-XXH.

INTRODUCTION

We intend to present guidelines for configuring 3380 DASD through the presentation of three scenarios. They are:

- Scenario One - Improving the Device Response Time

  In this scenario we will address the improvements in Device Response Time that might be expected when migrating from a Non-DPR environment to a DPR environment.

- Scenario Two - Increasing the Actuator I/O Rate

  In this scenario we will address the increase in Actuator I/O Rate that might be expected when migrating from a Non-DPR environment to a DPR environment.

- Scenario Three - Expanding the Configuration

  In this scenario we will compare the measured configurations against each other and explore the possibilities for expanding a 3380 DASD configuration in a DPR environment.

It should be pointed out that we are separating the improvements in Device Response Time and Actuator I/O Rate solely for purposes of discussion. In reality, both improvements will generally occur together.

For Scenarios One and Two, the results are summarized in three tables, one for each type of workload. The improvement of DPR over Non-DPR is expressed as a range of percentage improvements, covering a range of typical rates and utilizations. The rates and utilizations can apply to several items:

- Channel Path Utilization
- Actuator I/O Rate
- Device Response Time
- Device Service Time

The utilizations given in the tables apply to the Non-DPR mode of operation. They are intended to be reference values, so that each installation will be able to determine the particular improvement to be expected in its own case. All percentages have been calculated relative to the Non-DPR mode values.

Device Response Time is not a directly reported item in the MVS/370 version of RMF as it is in the MVS/XA version of RMF. You can, however, arrive at the Device Response Time with the following formula:

$$\text{DEVICE RESPONSE TIME} = \frac{\text{AVG Q LNGTH}}{\text{DEV ACTIVITY RATE}} + \text{AVG SERV TIME}$$

The assumption for Scenario One and Two is that the configuration does not change. Therefore, these scenarios would be analogous to an initial migration from MVS/370 to MVS/XA. The topic of Configurable Units will not be discussed there, nor will the configurations be evaluated against each other.

Only in Scenario Three, Expanding the Configuration, will the question be raised as to the practical limits of a configuration, and what type of configuration is to be preferred over the other.

When looking at the tables, two important issues should be kept in mind:

- In the measurement runs, each actuator in the measured configuration contributed to the total activity. This means that, measured over a period of 30-60 minutes, all actuators have performed more or less the same number of I/O operations. This is generally referred to as a non-skewed type of workload.

  Thus, in all configurations the number of **configured** actuators was equal to the number of **active** actuators. This may be realistic to some extent for IMS/VS databases, but certainly not for TSO or Batch datasets. This should be kept in mind when trying to understand the information in the following sections.

- Although three different workloads (IMS/VS, TSO, and Batch) were measured, it is nevertheless true that each represents a **specific** workload. It is highly unlikely that any of these workloads will match your workload exactly. This implies that all percentages and numbers presented should be considered as examples of the behavior of 3380 DASD with and without the Dynamic Path Reconnection capability.

SCENARIO ONE - IMPROVING THE DEVICE RESPONSE TIME

Objectives

In Scenario One, our objective will be to improve the **Device Response Time** of an I/O request for an existing configuration. Opting for this goal implies that the configuration will not be changed. It also means that it is not our intention to complete more I/O requests within the same amount of time. So it should be expected that the Channel Path Utilization will be what it was before.

An MVS/370 installation with 3380 DASD and running a type of workload with similar Set Sector Reconnect Delays that will not have the tendency to issue its I/O operations at a higher rate, can expect to realize this particular benefit merely by migrating their System Control Program (SCP) from MVS/370 to MVS/XA.

The improvement that can be expected will be presented over a range of Device Response Times. The percentage improvement has been calculated relative to the Non-DPR Device Response Time. The actual choice of the range of values differs for each workload. The ranges selected are from a technical point of view arbitrary, but are considered realistic for the particular types of workload.

- IMS/VS Databases - 25 msec to 40 msec
- Batch Datasets - 25 msec to 60 msec
- TSO User Datasets - 25 msec to 50 msec

The importance of presenting the results over a range of Device Response Times is that the low and high values of the range may greatly influence the degree of improvement that can be expected. In general, the higher the Device Response Time is in Non-DPR mode, the greater the improvement that can be expected in DPR mode. Figure 7 on page 21 illustrates this point very well.

In the following sections we will discuss the results of the three workloads: IMS/VS, Batch, and TSO. For each workload, we will present a summary of the improvement in Device Response Times for the measured configurations, observations from the measurement runs, and finally guidelines for the expected benefits.

IMS/VS Workload

The following chart illustrates the range of Device Response Time improvements that were observed for the measured configurations in an IMS/VS environment. The percentage improvements indicated were calculated relative to the Non-DPR Device Response Time values.

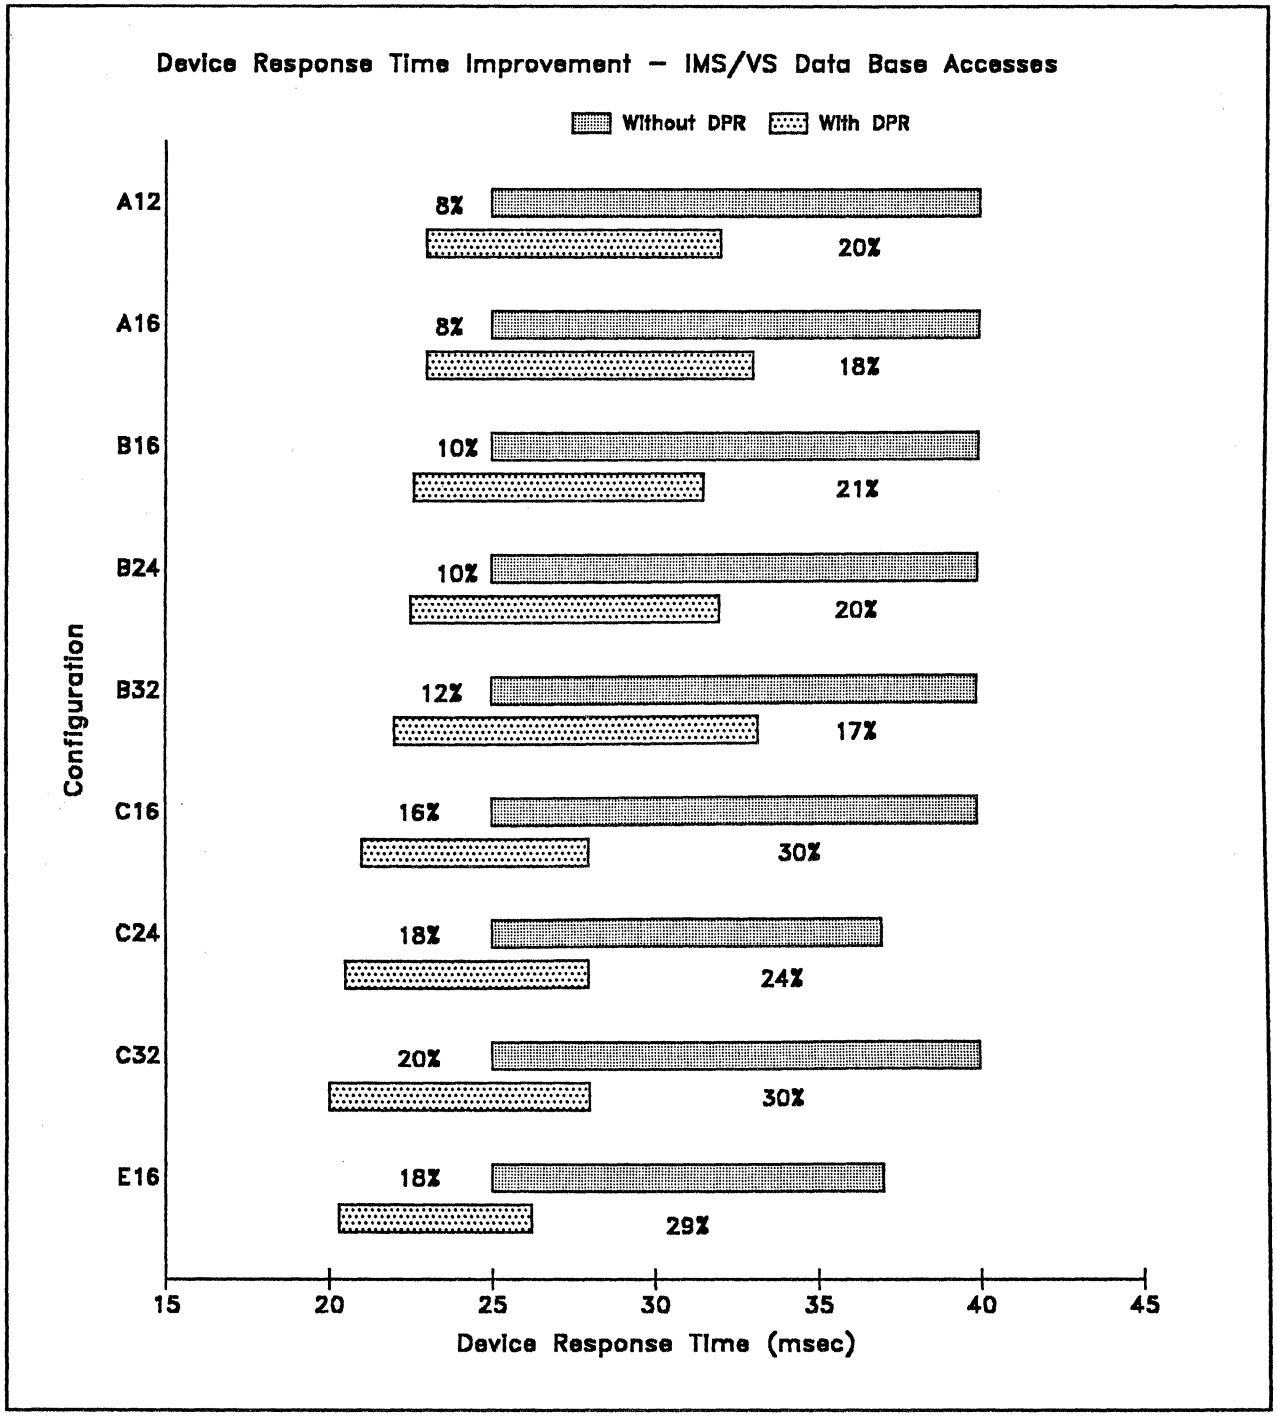

Chart Interpretation

The preceding chart and the similar Batch and TSO charts should be interpreted in the following manner. For a given configuration, the upper bar represents a range of Device Response Times for the workload run without DPR. The lower bar represents the same workload run with DPR at the same Actuator I/O Rates. As an example, consider the C16 Configuration. In the Non-DPR run, an Actuator I/O Rate of 11.9 was observed at a Device Response Time of 25 msec. When the same workload was run with DPR, a Device Response Time of 21 msec was observed at the same Actuator I/O Rate of 11.9 I/Os per second. This represents an improvement of 16% in Device Response Time between the two runs as indicated on the chart.

Summary of Results

The measurement results are summarized numerically in Figure 13.

Config-uration	Actuator I/O Rate per Second [1]	Channel Path Utilization (%) [1]	Device Response Time (msec)		% Improvement in Device Response Time
			Non-DPR	DPR	
A12	12.1 - 17.9	16.2 - 24.6	25.0 - 40.0	23.0 - 32.0	8 - 20
A16	10.3 - 15.6	18.6 - 29.2	25.0 - 40.0	23.0 - 33.0	8 - 18
B16	11.3 - 16.4	20.2 - 29.7	25.0 - 40.0	22.6 - 31.5	10 - 21
B24	8.5 - 12.0	23.6 - 33.1	25.0 - 40.0	22.5 - 32.0	10 - 20
B32	6.8 - 9.8	24.6 - 36.2	25.0 - 40.0	22.0 - 33.2	12 - 17
C16	11.9 - 17.6	21.4 - 33.3	25.0 - 40.0	21.0 - 28.0	16 - 30
C24	9.0 - 12.8	25.0 - 34.6	25.0 - 37.0	20.5 - 28.0	18 - 24
C32	7.3 - 10.6	27.4 - 38.2	25.0 - 40.0	20.0 - 28.0	20 - 30
E16	14.2 - 20.7	12.5 - 19.1	25.0 - 37.0	20.5 - 26.2	18 - 29

[1] Non-DPR reference values

Figure 13. Improvement in Device Response Time - IMS/VS Databases

Observations

1. The lower value in the range of improvement was **always** observed at the lower utilization or device response time; the higher value was **always** observed at the higher utilization or device response time.

   Thus, it can be concluded that the improvement in Device Response Time is more significant when the activity in Non-DPR mode is at a higher level. This conclusion matches our expectation: the Set Sector Reconnect Delay Time tends to be more significant at higher channel path utilizations. This is exactly the place where the benefits of DPR are the greatest.[3]

2. Figure 13 on page 35 does not show which configuration delivers the **best** response times. It only shows the range of improvements for a given configuration.

   The configuration that gives the **best** response times in DPR mode is the E16 configuration.

   For an evaluation of configurations relative to each other in DPR mode, see "Scenario Three - Expanding the Configuration" on page 58.

3. The C-type configurations compare very well with the B-type configurations, although the same number of channel paths is used in both. The C-type configurations differ in that each 3380-AA4 unit is attached to its own pair of storage directors.

   The technical background for the better performance of the C-type configurations is related to the Storage Director-String Controller protocol definition. When in DPR mode, the components in the data path have to exchange more status information than in Non-DPR mode. Since this overhead appears per I/O operation, the effect of it will eventually be noticed at high I/O rates.

   The old adage is that during an I/O operation, all path components are busy at the same time: Channel Path, Storage Director, String Controller, and Device. To a large extent this still is true, except for this 'post disconnection' information exchange.

   This is the reason that the C-type configurations have a better performance potential in MVS/XA than the B-type configurations.

4. In the measured environment, there was almost a 100% correlation between Device Response Time and IMS/VS Terminal Response Time, up to a Device Response Time of 35-45 msec. Above that Device Response Time, the IMS/VS Terminal Response Time appeared to be very sensitive to a slight increase in Device Response Time. Thus, the better Device Response Time provided by DPR trans-

---

[3] For a more detailed discussion on this effect see the section "Dynamic Path Selection" on page 18.

lates directly into a better Terminal Response Time, and also helps to avoid erratic Terminal Response Times.

5. The shortest Device Response Time in DPR mode was measured for the C32 and E16 configurations. An average Device Response Time of 17 msec. at an Actuator I/O Rate of 5 per second was observed for the C32 Configuration. An average Device Response Time of 17.5 msec. at an Actuator I/O Rate of 8 per second was observed for the E16 Configuration.

Guidelines

The A-type and B-type Configurations are probably most representative of the configurations currently installed. Therefore, for an IMS/VS type of workload, the **expected improvement** in Device Response Time for these configurations, when following Scenario One, is between **8% and 21%.**

The C-type Configurations with two additional Storage Directors and the E-type Configuration with two additional Storage Directors and two additional Channel Paths have the potential to provide for significantly greater improvements in Device Response Time. Therefore, for an IMS/VS type of workload, the **expected improvement** in Device Response Time for these configurations, when following Scenario One, is between **16% and 30%.**

Clearly, substantially greater benefits are obtainable at higher levels of utilization when required by the workload.

Batch Workload

In considering the improvement in Device Response Times for the Batch environment, the results need to be viewed somewhat differently from the interactive environments. An interactive system is driven primarily by the user's think time and accordingly the user will see the improvement in Device Response Time reflected in an improved Terminal Response Time.

A Batch system is self driven and throughput oriented and does not fit perfectly into Scenario One. Assuming unused CPU resources are available, it immediately takes advantage of the increased channel capacity and will try to get more I/O operations completed in the same time, thereby increasing the Actuator I/O rate. The conclusion is that Scenario One is not very likely to occur by itself for Batch type of workloads. However, we will present the results in a similar manner to better understand the improvements.

The following chart illustrates the range of Device Response Time improvements that were observed for the measured configurations in a Batch only environment. The percentage improvements indicated were calculated relative to the Non-DPR Device Response Time values.

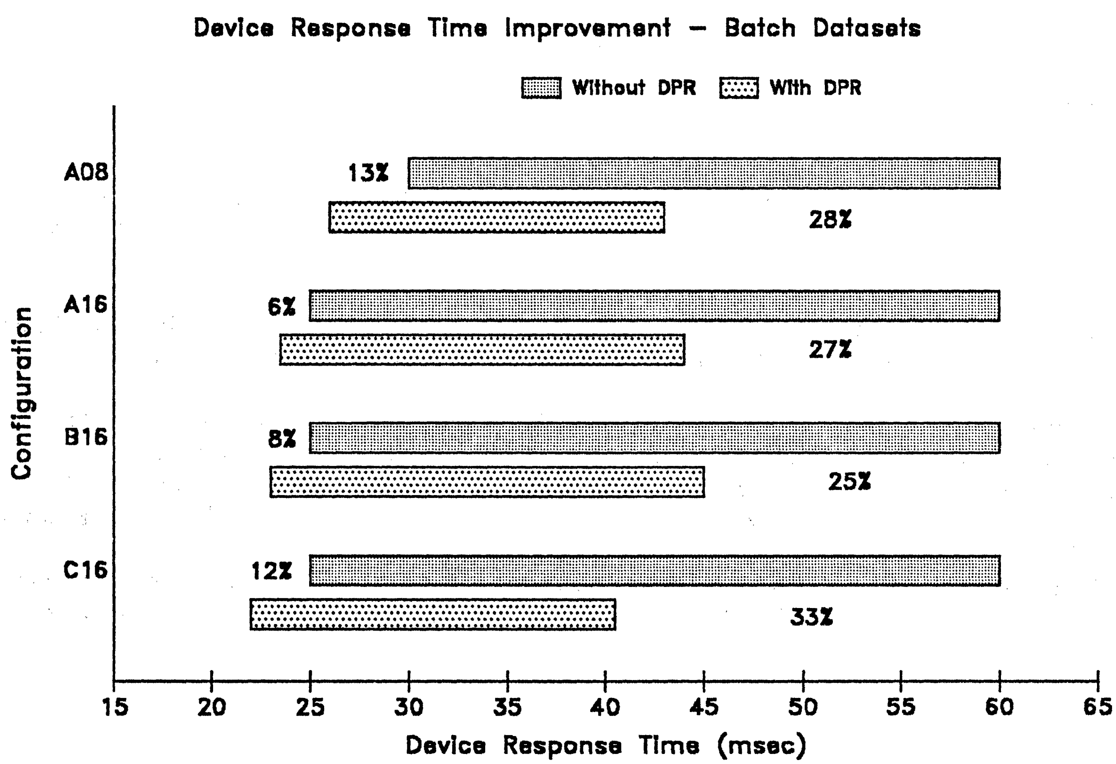

Figure 14.

Summary of Results

The measurement results are summarized numerically in Figure 15 on page 39.

Configuration	Actuator I/O Rate per Second [1]	Channel Path Utilization (%) [1]	Device Response Time (msec)		% Improvement in Device Response Time
			Non-DPR	DPR	
A08	13.9 - 19.9	18.2 - 26.3	30.0 - 60.0	26.0 - 43.0	13 - 28
A16	7.4 - 13.5	18.2 - 34.6	25.0 - 60.0	23.5 - 44.0	6 - 27
B16	8.1 - 14.0	19.9 - 35.6	25.0 - 60.0	23.0 - 45.0	8 - 25
C16	8.5 - 14.8	20.9 - 35.5	25.0 - 60.0	22.0 - 40.5	12 - 33

[1] Non-DPR reference values

Figure 15. Improvement in Device Response Time - Batch Datasets

Observations

1. The first observation is the same as for the IMS/VS workload: the improvement in Device Response Time is more significant when the activity in Non-DPR mode is at a high level.

2. For an evaluation of configurations relative to each other when working in DPR mode, see "Scenario Three - Expanding the Configuration" on page 58.

3. The impression is that the level of improvement is better for the Batch workload than for the IMS workload. However, this impression may be misleading; when the percentages are calculated for the same Device Response Time, the results are very similar. Refer to the chart in Figure 16 on page 40. This chart presents the improvement in Device Response Times at various Non-DPR values. The value at a Device Response Time of 40 msec. allows you to compare the Batch and IMS/VS workloads. See also Figure 13 on page 35.

Configuration	Device Response Time (msec)		Improvement (%)
	Non-DPR	DPR	
A08	40.0	33.0	18%
	60.0	43.0	28%
	80.0	53.0	34%
A16	40.0	33.0	18%
	60.0	44.0	27%
	80.0	52.0	35%
B16	40.0	30.5	24%
	60.0	45.0	25%
	80.0	50.5	37%
C16	40.0	30.0	25%
	60.0	40.5	33%
	75.0	45.0	40%

Figure 16. Improvement in Device Response Times - Batch Datasets

Guidelines

Due to the nature of Batch workloads, the actuators will tend to be stressed very heavily. Therefore, for Batch type of workload, the **expected improvement** in Device Response Time, when following Scenario One, is between **25% and 35%.**

TSO Workload

The following chart illustrates the range of Device Response Time improvements that were observed for the measured configurations in a TSO environment. The percentage improvements indicated were calculated relative to the Non-DPR Device Response Time values.

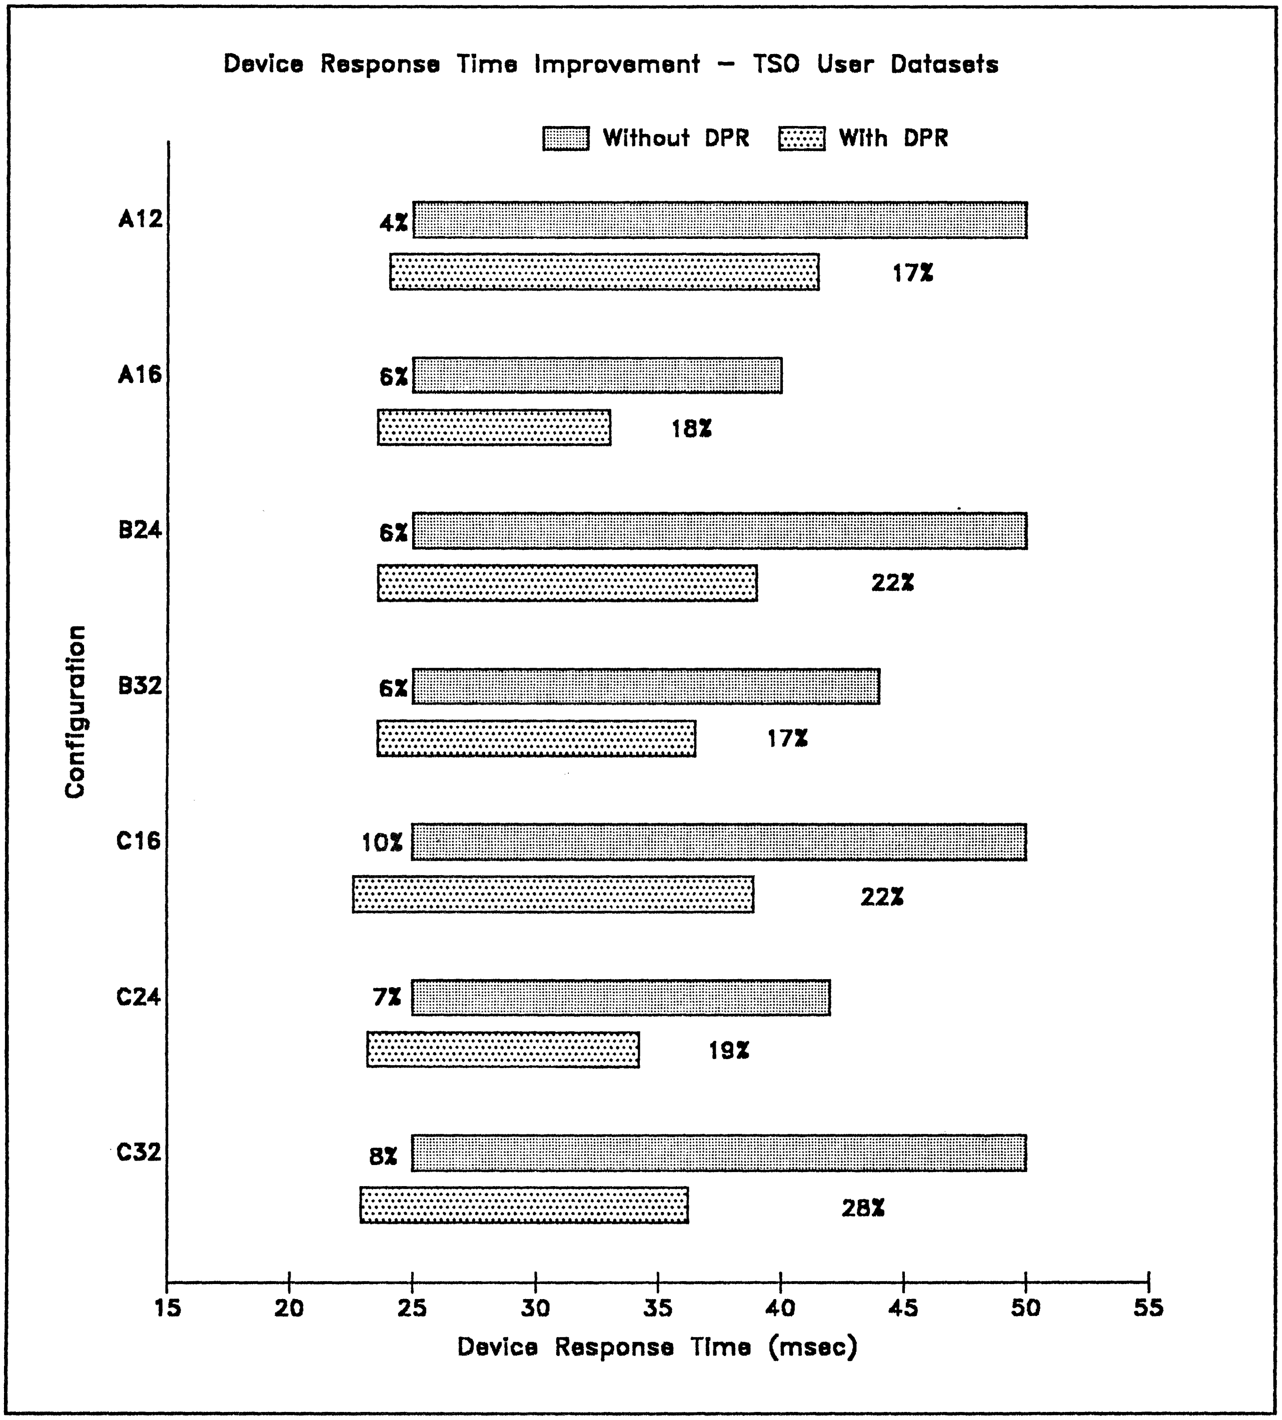

Summary of Results

The measurement results are summarized numerically in Figure 17.

Configuration	Actuator I/O Rate per Second[1]	Channel Path Utilization (%) [1]	Device Response Time (msec)		% Improvement in Device Response Time
			Non-DPR	DPR	
A12	5.1 - 10.7	16.9 - 37.0	25.0 - 50.0	24.1 - 41.5	4 - 17
A16	4.1 - 7.5	17.4 - 34.4	25.0 - 40.0	23.6 - 33.0	6 - 18
B24	3.7 - 6.2	24.6 - 40.9	25.0 - 50.0	23.6 - 39.0	6 - 22
B32	2.7 - 4.6	24.1 - 41.0	25.0 - 44.0	23.6 - 36.5	6 - 17
C16	4.6 - 9.6	20.9 - 44.2	25.0 - 50.0	22.6 - 38.9	10 - 22
C24	3.5 - 6.2	22.9 - 40.2	25.0 - 42.0	23.2 - 34.2	7 - 19
C32	2.6 - 5.4	23.9 - 48.2	25.0 - 50.0	22.9 - 36.2	8 - 28

[1] Non-DPR reference values

Figure 17. Improvement in Device Response Time - TSO User Datasets

Observations

1. The improvement in Device Response Time is more significant when the activity in Non-DPR mode is at a higher level. The lower value in the range of improvement was **always** observed at the lower utilization or device response time; the higher value was **always** observed at the higher utilization or device response time. This effect was also observed in the case of the IMS/VS workload.

2. For an evaluation of configurations relative to each other when working in DPR mode, see "Scenario Three - Expanding the Configuration" on page 58.

3. The performance advantage of C-type configurations over B-type configurations is not as distinct as for the IMS/VS workload. Of course we would expect this result, since the total I/O rate for a configuration was considerably lower in the measured TSO environment than it was for IMS/VS. You may compare this by multiplying the number of actuators in a configuration by the Actuator I/O Rate. See also observation 3 in the discussion on IMS/VS for Scenario One.

4. The configurations can also be compared by evaluating their maximum throughput. When doing so, it is obvious that for the very large configurations (B24, B32, C24, C32) the C-type configurations offer a more significant performance improvement with respect to the B-type configurations. For the smaller configurations (A16, B16, C16) the advantage of the C-type configuration is less visible.

5. The improvement that the DPR configuration offers can be used in two ways:

   - Better Terminal Response Time may be achieved for the same transaction and I/O rate. This is illustrated by Figure 18. This graph shows that at the same Actuator I/O Rate (Scenario One assumes a consistent Actuator I/O Rate.) the DPR configuration provides better Terminal Response Times. The improvement expressed as a percentage is very much the same for Terminal Response Time as it is for Device Response time.

   - A greater number of transactions may be completed at the same Terminal Response Time. This aspect will be considered in Scenario Two.

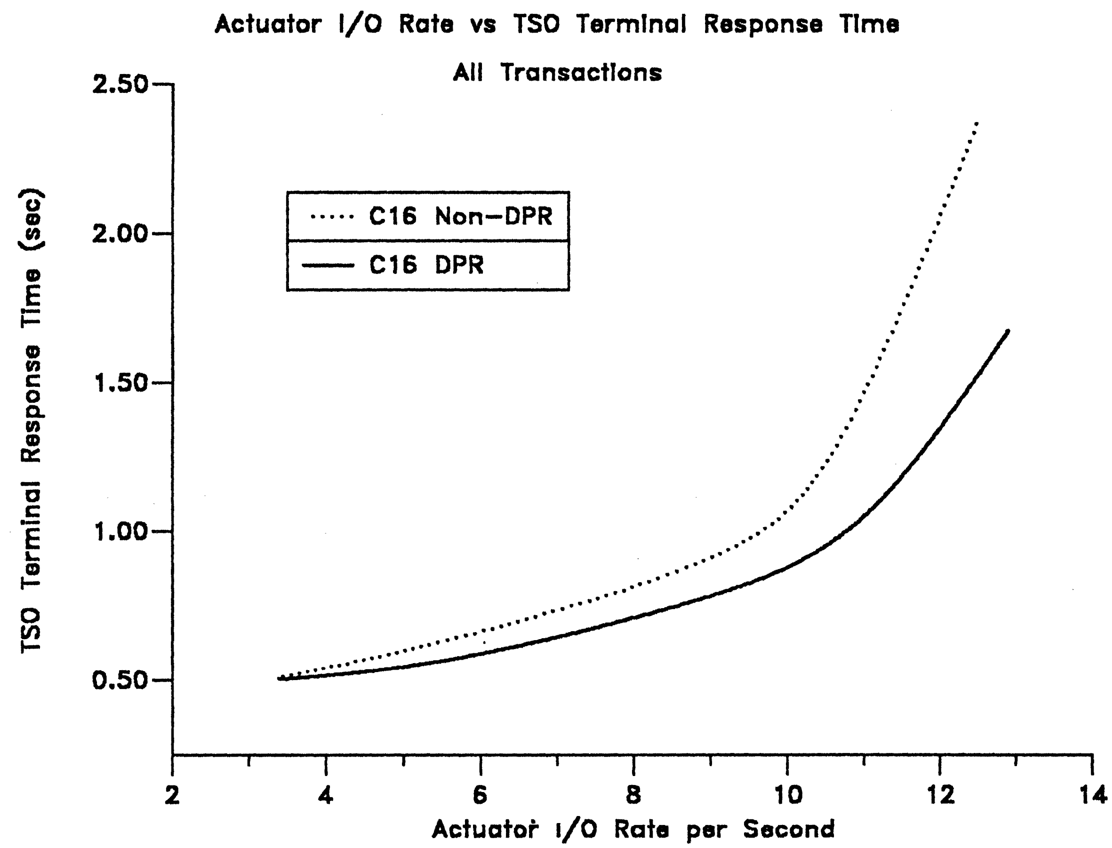

Figure 18. TSO Terminal Response Time Improvements

Guidelines

The A-type and B-type Configurations are probably most representative of the configurations currently installed. Therefore, for TSO type of workload, the **expected improvement** in Device Response Time, when following Scenario One, is between **6% and 22%.**

The C-type Configurations with two additional Storage Directors have the potential to provide for significantly greater improvements in Device Response Times. Therefore, for a TSO type of workload, the **expected improvement** in Device Response Time for these configurations, when following Scenario One, is between **7% and 28%.**

Clearly, substantially greater benefits are obtainable at higher levels of utilization when required by the workload.

SCENARIO TWO - INCREASING THE ACTUATOR I/O RATE

Objectives

In Scenario Two, our objective will be to increase the number of I/O requests to an actuator while maintaining an equivalent or constant Device Response Time. As in Scenario One, the 3380 DASD configuration will remain unchanged. It should be expected that this approach will cause a higher channel path utilization due to the higher I/O rate, although the number of devices does not change.

The effect of DPR is that the device running in DPR mode will not need as much time to complete an I/O operation as it did when running in Non-DPR mode. This will increase the capacity of the DASD configuration and allow us to drive the devices at higher Actuator I/O Rates. This in itself does not imply that the Device Response Time will be longer. The Seek Time, Rotational Delay and Search/Data Transfer will not take more time as the device gets busier. However, the increased number of I/O requests will cause an increase in channel path utilization. Because the RPS reconnect delay is dominated by channel path utilization, the average reconnect delay will also increase.

The net result will be that at the same Device Response Time, **the actuator will be doing more work**, which is our primary objective in this scenario.

It should be pointed out that the system designer will often set service times for DASD operations to meet service levels of end user response times. Thus the capacity of a subsystem is how many DASD operations can be sustained at a given Device Response Time.

For IMS/VS, Batch, and TSO workloads, an improvement in Actuator I/O Rate should translate into more work being completed. This will be verified in two ways:

- First, by investigating the correlation between the Actuator I/O Rate and the Transaction Completion Rate. Figure 20 on page 48 is an example of this. The straight line is an indication of a high degree of correlation.

  Note: The measurement runs were conducted in an environment where resources were generally unconstrained, except for the DASD configuration under consideration. In your installation the correlation between DASD performance and work completed may not be as clear as indicated in Figure 20.

- Secondly, by investigating the Transaction Completion Rate for a given Device Response Time. Remember that in Scenario Two, we assume a constant Device Response Time. The curves in Figure 21 on page 49 show that when the configuration is working in DPR mode, more transactions are completed.

IMS Workload

The following chart illustrates the range of Actuator I/O Rate increases that were observed for the measured configurations in an IMS/VS environment. The percentage increases indicated were calculated relative to the Non-DPR Actuator I/O Rate values.

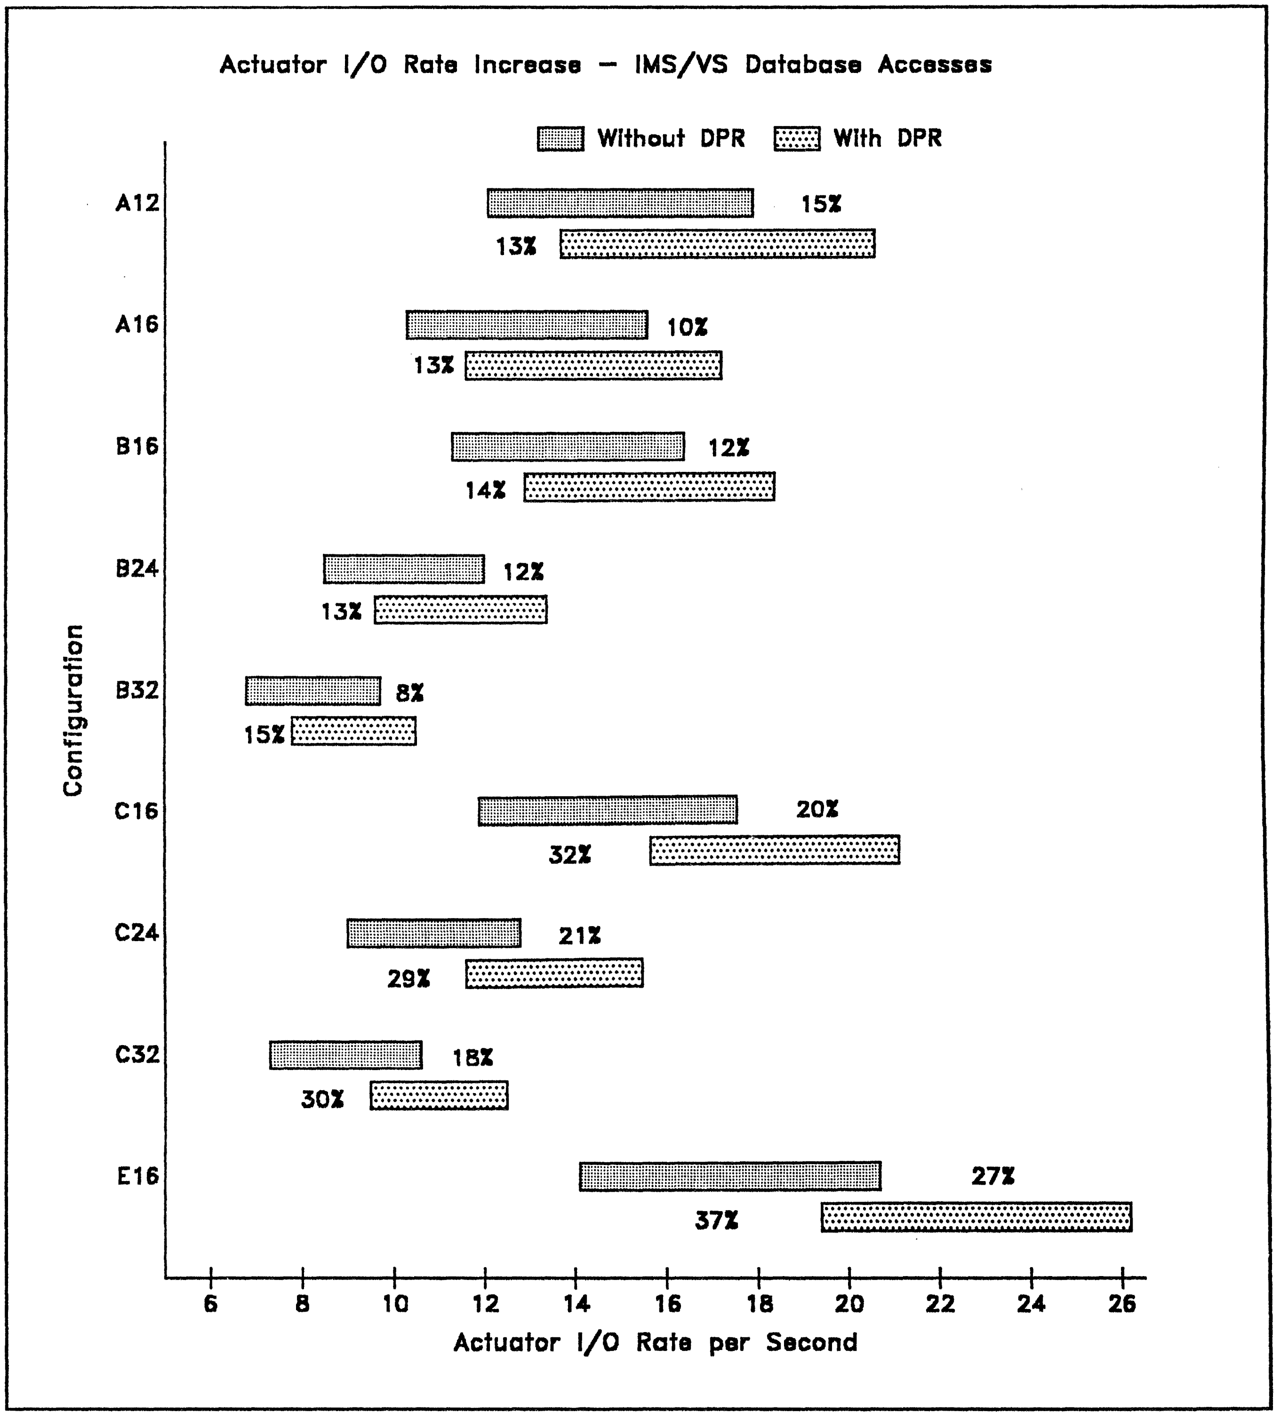

Chart Interpretation

The preceding chart and the similar Batch and TSO charts should be interpreted in the following manner. For a given configuration, the upper bar represents a range of Actuator I/O Rates per Second for the workload run without DPR. The lower bar represents the same workload run with DPR at the same Device Response Times. As an example, consider the C16 Configuration. In the Non-DPR run, an Actuator I/O Rate of 11.9 was observed at a Device Response Time of 25 msec. When the same workload was run with DPR, an Actuator I/O Rate per Second of 15.7 was observed at the same Device Response Time of 25 msec. This represents an improvement of 32% in Actuator I/O Rate between the two runs as indicated on the chart.

Summary of Results

The measurement results are summarized numerically in Figure 19.

Configuration	Device Response Time (msec)[1]	Channel Path Utilization (%) [1]	Actuator I/O Rate per Sec.		% Improvement in Actuator I/O Rate
			Non-DPR	DPR	
A12	25.0 - 40.0	16.2 - 24.6	12.1 - 17.9	13.7 - 20.6	13 - 15
A16	25.0 - 40.0	18.6 - 29.2	10.3 - 15.6	11.6 - 17.2	13 - 10
B16	25.0 - 40.0	20.2 - 29.7	11.3 - 16.4	12.9 - 18.4	14 - 12
B24	25.0 - 40.0	23.6 - 33.1	8.5 - 12.0	9.6 - 13.4	13 - 12
B32	25.0 - 40.0	24.6 - 36.2	6.8 - 9.7	7.8 - 10.5	15 - 8
C16	25.0 - 40.0	21.4 - 33.3	11.9 - 17.6	15.7 - 21.2	32 - 20
C24	25.0 - 37.0	25.0 - 34.6	9.0 - 12.8	11.6 - 15.5	29 - 21
C32	25.0 - 40.0	27.4 - 38.2	7.3 - 10.6	9.5 - 12.5	30 - 18
E16	25.0 - 37.0	12.5 - 19.1	14.2 - 20.7	19.4 - 26.2	37 - 27

[1] Non-DPR reference values

Figure 19. Improvement in Actuator I/O Rate - IMS/VS Databases

Observations

1. As in the tables for Scenario One, the values in the reported ranges relate left to left and right to right.

2. In general the percentage improvement appears to be more significant at the lower utilizations. However, the actual increase in the number of I/O requests is greater at the high end of the measured range.

   This comment is also valid for the TSO and Batch workloads.

3. A high degree of correlation was observed between the Actuator I/O Rate and the IMS/VS Transaction Rate. The straight line in Figure 20 is an indication of this high degree of correlation.

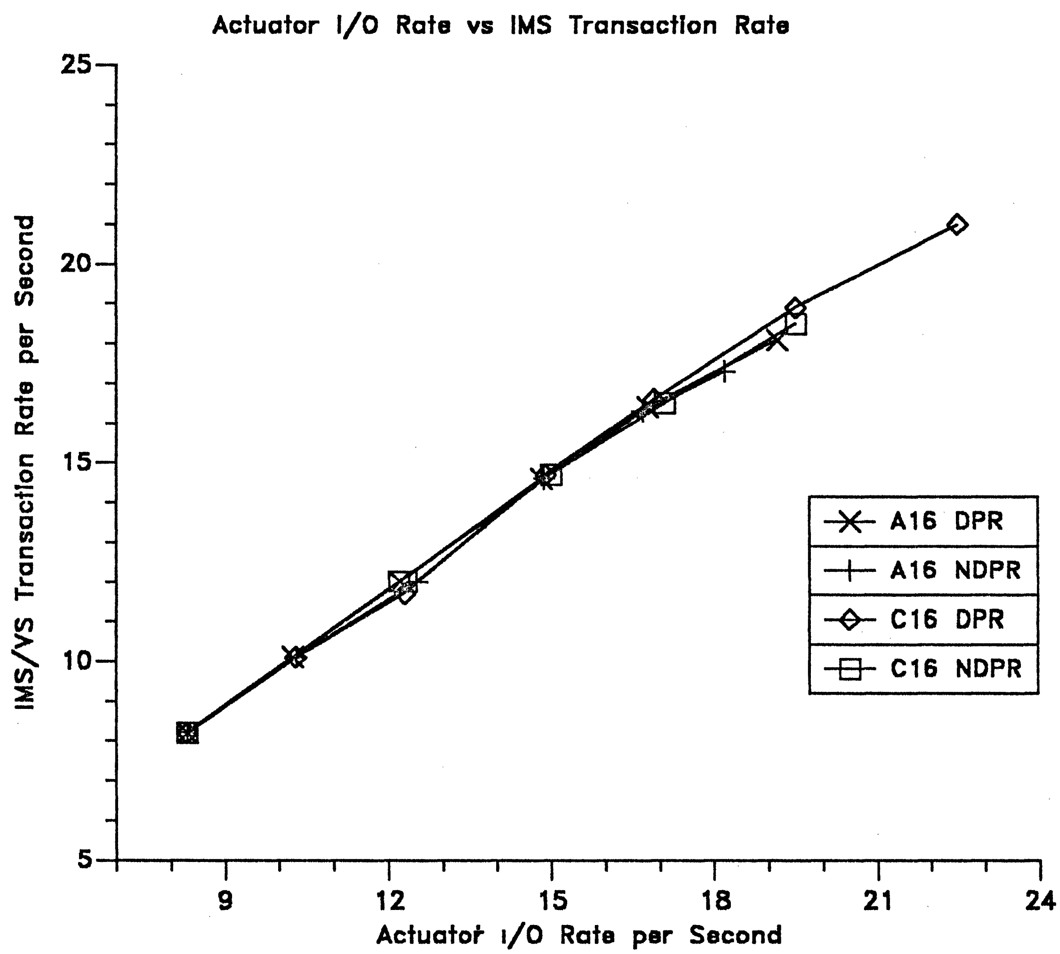

Figure 20. Correlation between Actuator I/O Rate and IMS/VS Transaction Rate

4. As well, the Transaction Rate in DPR mode was considerably better for equivalent Device Response Times as indicated in Figure 21.

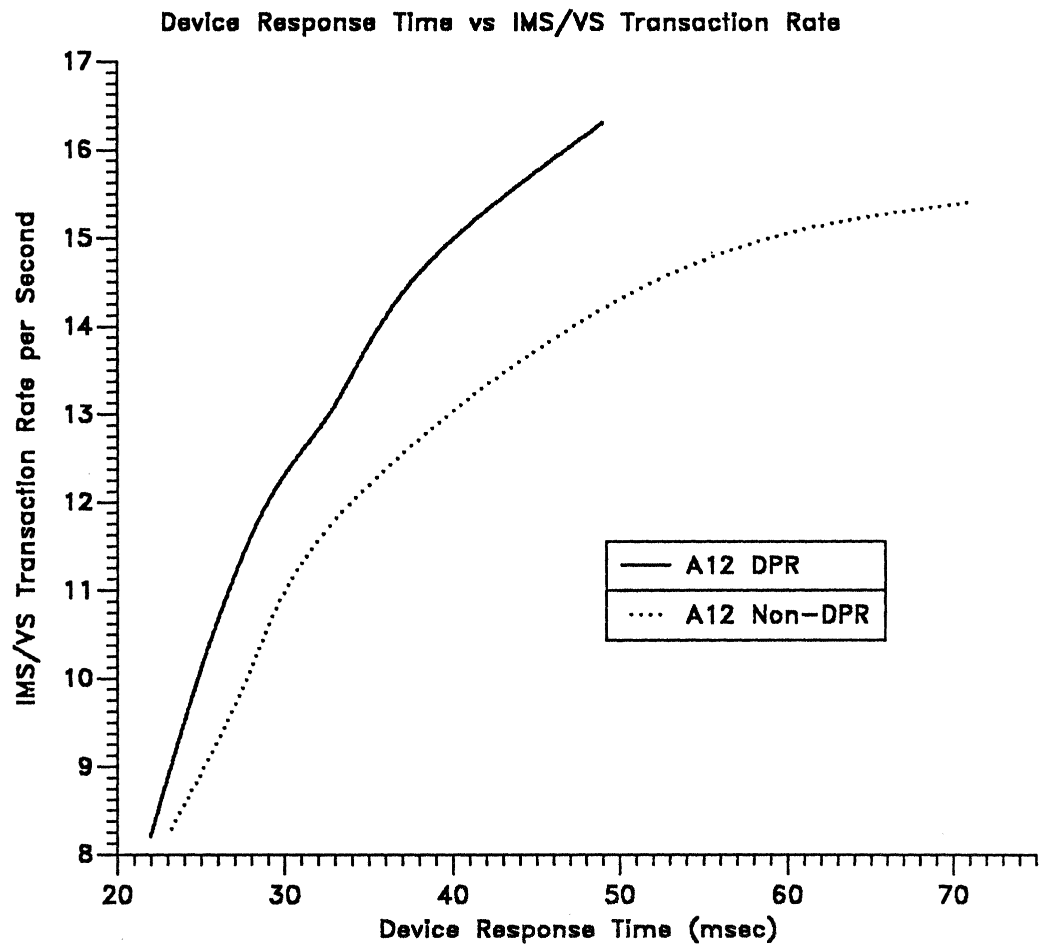

Figure 21. Device Response Time vs IMS/VS Transaction Rate

5. The Actuator I/O Rates in DPR mode are very high, especially for the smaller configurations. The Actuator I/O Rates for a range of Device Response Times are summarized in Figure 22 on page 50. These configurations are particularly well suited for meeting the demands of a workload with a very high access density.

Device Response Time	Actuator I/O Rate per Second		
	25 msec	30 msec	40 msec
Configuration: E16	19	23	27
C16	16	18	21
A12	14	17	20

Figure 22. Actuator I/O Rates in DPR mode — IMS/VS Database Accesses

It is believed that three factors are critical in achieving very high Actuator I/O Rates while maintaining acceptable Device Response Times:

- Minimize Path Contention

  A small Configurable Unit with a limited number of actuators attached to a pair of channel paths.

- Relatively short Seek times.

  Analysis of the measurement data suggests that the seek times experienced in these workload measurements was in the order of 2 to 3 milliseconds.

- Relatively short Connect Times.

  The influence of Connect Time can be appreciated by comparing IMS/VS, Batch, and TSO results.

As a consequence, the Actuator I/O Rates achieved during the IMS/VS measurements have to be considered as attainable and realistic, but also have to be regarded as perhaps too optimistic for the average real life IMS/VS workloads.

This is not an environment that just happens to exist; it must be planned for carefully in the design of the application I/O subsystem.

6. The effect of improved actuator performance or capability means:

   - that the work completes faster, or
   - that more active data can be placed on the same set of actuators.

   Both options provide benefits in an IMS/VS environment.

7. Again the C-type configurations compare very well with the others. For a further discussion on this effect, refer back to the third observation for the IMS/VS workload in Scenario One.

8. An evaluation of the configurations relative to each other will be presented in Scenario Three.

Guidelines

The A-type and B-type Configurations are probably most representative of the configurations currently installed. Therefore, for an IMS/VS type of workload, the **expected increase** in Actuator I/O Rate for these configurations is between **8% and 15%.**

The C-type Configurations with two additional Storage Directors have the potential to provide for significantly greater increases in Actuator I/O Rate in DPR mode. Therefore, for an IMS/VS type of workload, the **expected increase** in Actuator I/O Rate for these configurations is between **18% and 32%.**

The E-type Configuration with two additional Storage Directors and two additional Channel Paths has the potential to provide for significantly greater increases in Actuator I/O Rate in DPR mode. Therefore, for an IMS/VS type of workload, the **expected increase** in Actuator I/O Rate for this configuration is between **27% and 37%.**

Batch Workload

The following chart illustrates the range of Actuator I/O Rate increases that were observed for the measured configurations in a Batch environment. The percentage increases indicated were calculated relative to the Non-DPR Actuator I/O rate values.

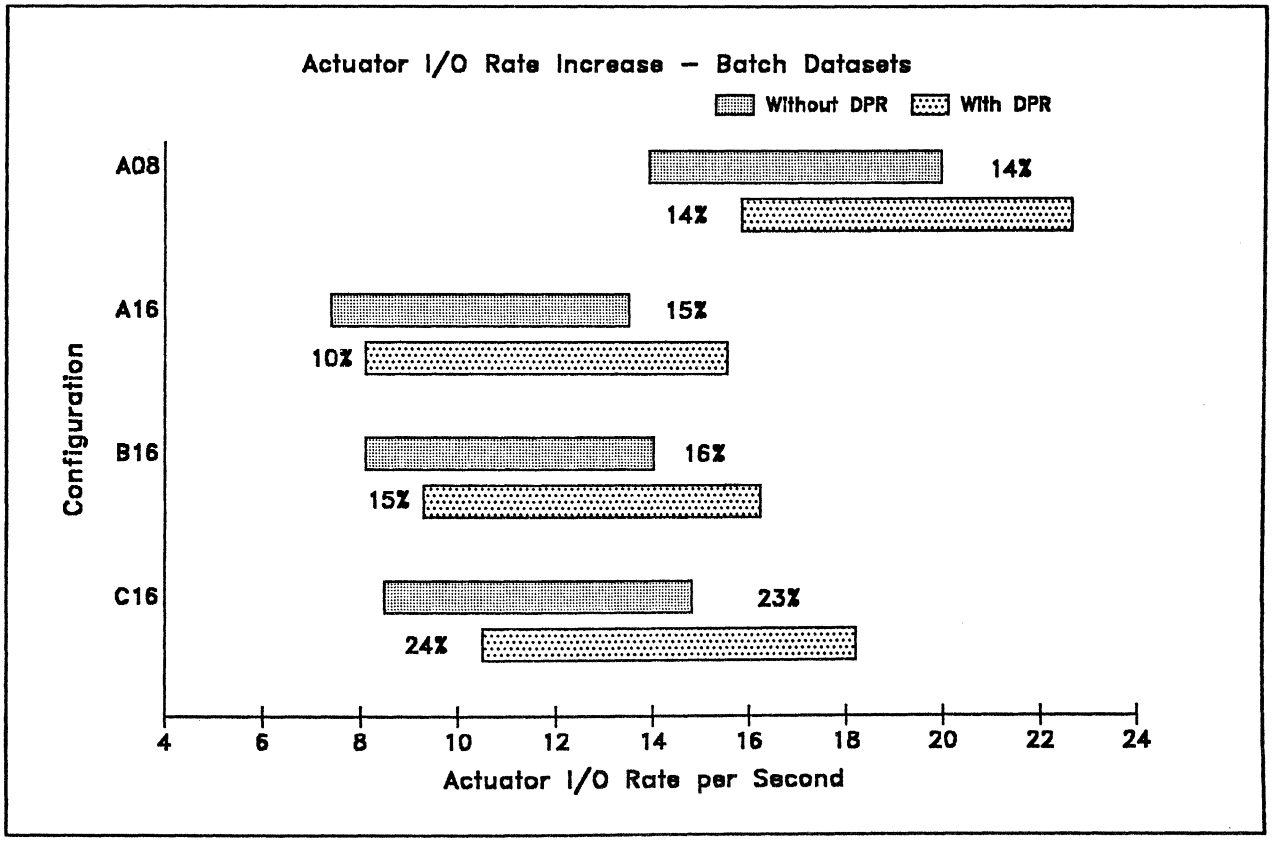

Summary of Results

The measurement results are summarized numerically in Figure 23.

Configuration	Device Response Time (msec)[1]	Channel Path Utilization (%) [1]	Actuator I/O Rate per Sec.		% Improvement in Actuator I/O Rate
			Non-DPR	DPR	
A08	30.0 - 60.0	18.2 - 26.3	13.9 - 19.9	15.8 - 22.6	14 - 14
A16	25.0 - 60.0	18.2 - 34.6	7.4 - 13.5	8.1 - 15.5	10 - 15
B16	25.0 - 60.0	19.9 - 35.6	8.1 - 14.0	9.3 - 16.2	15 - 16
C16	25.0 - 60.0	20.9 - 35.5	8.5 - 14.8	10.5 - 18.2	24 - 23

[1] Non-DPR reference values

Figure 23. Improvement in Actuator I/O Rate - Batch Datasets

Observations

1. As in the tables of Scenario One, the values in the reported ranges relate left to left and right to right.

2. The result of the A08 configuration is remarkable. Although it is certainly actuator, rather than channel path constrained, an improvement of 14% was observed.

   This means that a 3380 DPR configuration is able to respond to a very dynamic buildup in I/O demand, as is usually the case for Batch workloads.

3. The percentage increase in the Actuator I/O Rate remained relatively constant over a range of Device Response Time and workload levels for each individual configuration.

4. As with the IMS/VS workload, there was a high degree of correlation between the Actuator I/O Rate and the Service Unit Rate.

5. The relationship between Device Response Time and Service Rate shows that more work was completed at the same Device Response Time. See Figure 24 on page 54. These results were obtained by increasing the number of initiators to provide an increased workload level.

6. An evaluation of the configurations relative to each other will be presented in Scenario Three.

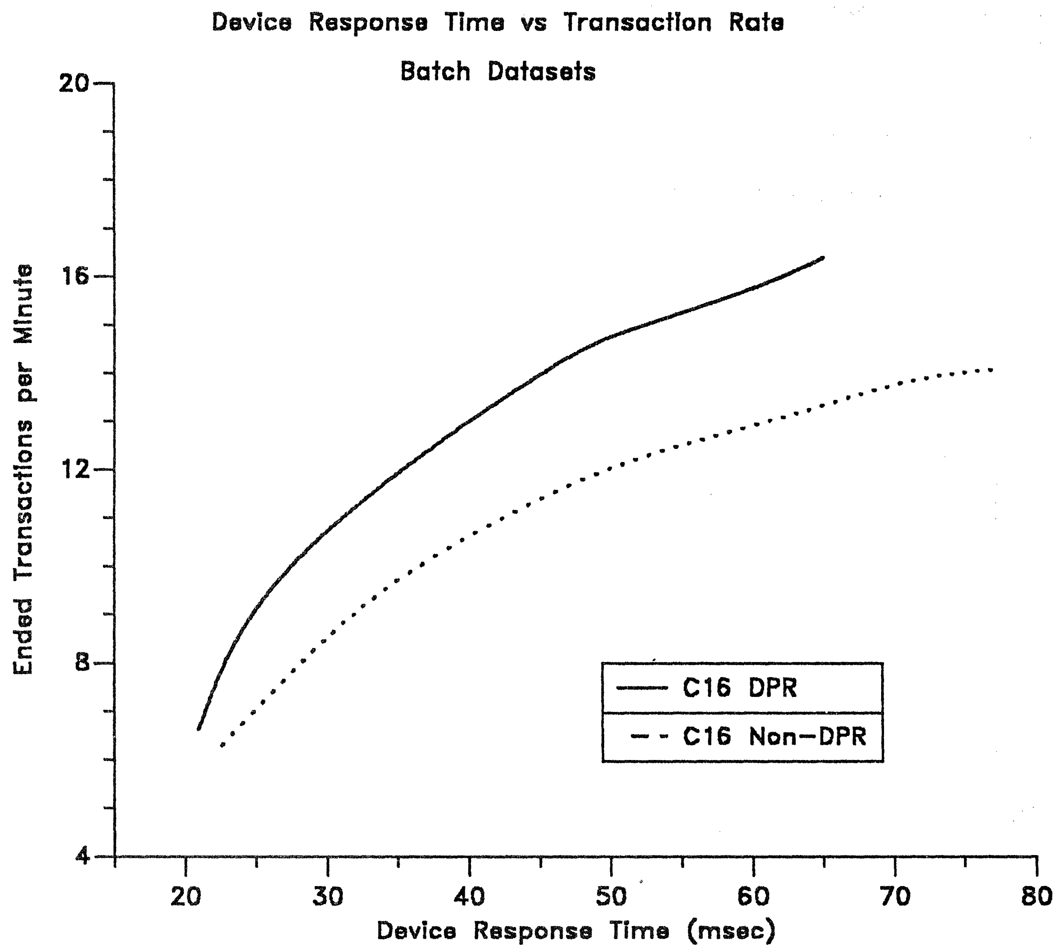

Figure 24. Device Response Time vs Ended Transactions— Batch Workload

Guidelines

Due to the nature of Batch workloads, the actuators will tend to be stressed very heavily. Therefore, for Batch type of workload, the **expected increase** in Actuator I/O Rate, when following Scenario Two, is between **10% and 24%.**

TSO Workload

The following chart illustrates the range of Actuator I/O Rate increases that were observed for the measured configurations in a TSO environment. The percentage increases indicated were calculated relative to the Non-DPR Actuator I/O Rate values.

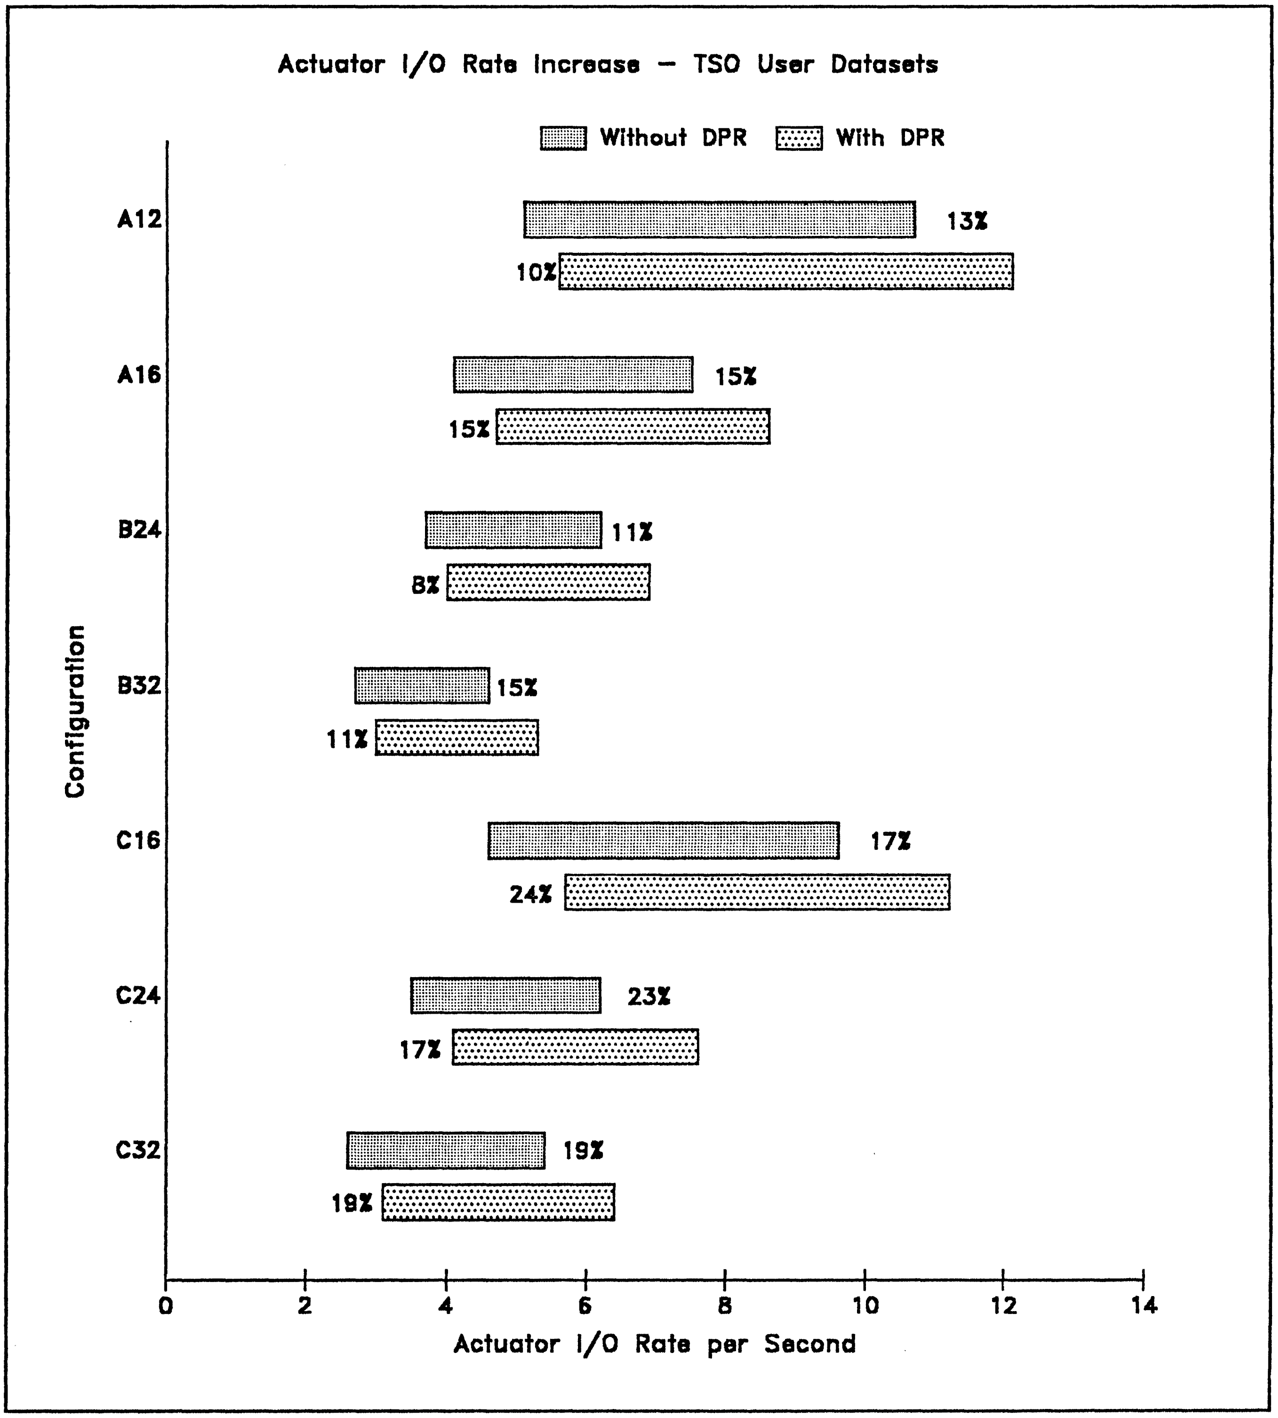

Summary of Results

The measurement results are summarized numerically in Figure 25.

Configuration	Device Response Time (msec)[1]	Channel Path Utilization (%) [1]	Actuator I/O Rate per Sec.		% Improvement in Actuator I/O Rate
			Non-DPR	DPR	
A12	25.0 - 50.0	16.9 - 37.0	5.1 - 10.7	5.6 - 12.1	10 - 13
A16	25.0 - 40.0	17.4 - 34.4	4.1 - 7.5	4.7 - 8.6	15 - 15
B24	25.0 - 50.0	24.6 - 40.9	3.7 - 6.2	4.0 - 6.9	8 - 11
B32	25.0 - 44.0	24.1 - 41.0	2.7 - 4.6	3.0 - 5.3	11 - 15
C16	25.0 - 50.0	20.9 - 44.2	4.6 - 9.6	5.7 - 11.2	24 - 17
C24	25.0 - 42.0	22.9 - 40.2	3.5 - 6.2	4.1 - 7.6	17 - 23
C32	25.0 - 50.0	23.9 - 48.2	2.6 - 5.4	3.1 - 6.4	19 - 19

[1] Non-DPR reference values

Figure 25. Improvement in Actuator I/O Rate - TSO User Datasets

Observations

1. As in the tables for Scenario One, the values in the reported ranges relate left to left and right to right.

2. For a TSO installation, Scenario Two could be translated into installing more terminals while maintaining the same I/O configuration, assuming the configuration has sufficient data storage capacity.

   For example, if you have an A16 configuration in MVS/370 and then go to MVS/XA, you should be able to satisfy the needs of 15 percent more users with the same configuration and at the same response times as in MVS/370.

3. An evaluation of the configurations relative to each other will be presented in Scenario Three.

Guidelines

The A-type and B-type Configurations are probably most representative of the configurations currently installed. Therefore, for a TSO type of workload, the **expected increase** in Actuator I/O Rate, when following Scenario Two, is between **8% and 15%.**

The C-type Configurations with two additional Storage Directors have the potential to provide for significantly greater increases in Actuator I/O Rate in DPR mode. Therefore, for a TSO type of workload, the **expected increase** in Actuator I/O Rate for these configurations is between **17% and 24%.**

SCENARIO THREE - EXPANDING THE CONFIGURATION

Objectives

Scenario Three is essentially a growth scenario. If the current Device Response Time and Actuator I/O Rate in Non-DPR mode are satisfactory, then more actuators can be accommodated on the same set of channel paths when running in DPR mode. One similarity between this scenario and Scenario Two is that in both cases the channel path utilization will be higher.

Channel Path Utilization is critical for Set Sector Reconnect Delay, but due to the effect of DPR, the device does not see the physical, but rather an apparently lower Channel Path utilization. As long as the apparent Channel Path utilization in a DPR configuration does not exceed the physical Channel Path utilization in a Non-DPR configuration, the Set Sector Reconnect Delay, and thus the Device Service Time, should be better or at least equivalent.

Of course there are limits to increasing the number of actuators on a set of channel paths. A balance has to be found between the capacity of a channel path and the load generated by the devices that are attached to it. As a result of this it will be necessary to review the existing guidelines.

The basic information to make this possible is summarized in Figure 26 on page 59. This table has the same structure as the one which was used to develop the expectations, Figure 11 on page 27.

In the vertical scale a range of Actuator I/O Rates is given, while the horizontal scale gives a range of Connect Times. These Connect Times relate to the three types of workload that have been measured. As in Scenarios One and Two, a range of realistic Device Response Times was chosen. Then each measured configuration was entered into the table at the appropriate Actuator I/O Rate.

The resulting table may be a disappointment for those who love simplicity. If so, it is comforting to know that the contributors to this document also love simplicity, but nevertheless have opted for providing enough information for you to be able to understand the results of the measurement runs, and to apply them to your own situation.

Actuator I/O Rate per Sec.	Device Response Times								
	25ms	30ms	40ms	40ms	50ms	60ms	30ms	40ms	50ms
27	—	—	E16	—	—	—	—	—	—
.									
24	—	E16	—	—	—	—	—	—	—
.						A08			
22	—	—	—	—	—	—	—	—	—
21	—	—	C16	—	A08	—	—	—	—
			A12						
20	—	—	—	A08	—	—	—	—	—
	E16								
19	—	—	—	—	—	—	—	—	—
18	—	C16	B16	—	—	C16	—	—	—
17	—	A12	A16	—	C16	—	—	—	—
16	C16	B16	C24	—	—	B16	—	—	—
						A16			
15	—	—	—	C16	B16	—	—	—	—
		A16			A16				
14	A12	—	—	—	—	—	—	—	E16
		C24	B24						
13	B16	—	—	B16	—	—	—	E16	—
			C32	A16					
12	A16	—	—	—	—	—	—	—	A12
	C24	B24							
11	—	C32	—	—	—	—	—	—	C16
			B32						
10	B24	—	—	—	—	—	—	A12	B16
	C32							C16	A16
9	—	B32	—	—	—	—	E16	B16	—
								A16	
8	B32	—	—	—	—	—	—	—	C24
							A12	C24	
7	—	—	—	—	—	—	C/A16	—	B24
							B16		C32
6	—	—	—	—	—	—	—	B24	—
							C24	C32	B32
5	—	—	—	—	—	—	—	B32	—
							B24		
4	—	—	—	—	—	—	C32	—	—
							B32		
Connect Times	IMS Databases 2.4 msec			Batch Datasets 3.2 msec			TSO User Datasets 5.4 msec		

Figure 26. Comparison of the Configurations

Using the Table

First of all, when the Actuator I/O Rate increases, fewer actuators can be configured on a channel path. This can be verified by analysing any column in the table, the number of actuators will be equal or lower when the Actuator I/O Rate increases.

Secondly, as the Connect Time increases, the number of actuators in a Configurable Unit will be smaller. To verify this, the information in the table can be interpreted in two ways:

1. For a given Actuator I/O Rate, compare the results at different Connect Times, but at the same Device Response Time. For example, compare IMS/VS and Batch at 40 msec Device Response Time and an Actuator I/O Rate of approximately 12.5 per second. The table shows in this case 32 actuators for IMS/VS, and 16 for Batch.

13	B16	-	C32	A16	-	-	-	E16	-
12	A16	-			-	-	-	-	A12

2. Determine what the Actuator I/O Rate is for a particular configuration in the three different environments, again at the same Device Response Time. For example, at a Device Response Time of 40 msec, the C16 configuration has an Actuator I/O Rate of 21 for IMS/VS, 15 for Batch, and 10 for TSO.

When using Figure 26 on page 59 for planning purposes, it is of the utmost importance to realize that only a subset of all possible and valid configurations has been measured. Furthermore, the table presents the results of configurations measured at their performance limits.

We caution you to not simply use this table as a configuration selection guide for your installation. More in particular, it must be emphasized that you should **NOT** design your configuration with the objective of attaining the very high Actuator I/O Rates unless you have a detailed understanding of both your workload and the hardware performance characteristics.

In the measured environments, these configurations had no 'spare capacity' which may not be compatible with your configuration design objectives. The Actuator I/O Rates in the table are the highest obtainable at a given Device Response Time objective for a particular workload.

Observations and Guidelines

The guidelines that will be formulated in this section are based on three common types of workload, but, for each workload only one cross section was measured at varying levels of intensity. So do not consider the guidelines to be prescriptions. You will get the best benefit from the results when you are able to position your type of workload relative to the measured workloads.

In this respect it is felt that the type of workload is qualified to a high degree by one factor: Connect Time, or in MVS/370 RMF terminology: Path Service Time. If you currently do not use 3380, you may calculate the Connect Time on 3380. The technique for doing so is described in "DASD Path and Device Contention Considerations".[4]

A-type versus B-type Configurations

A performance evaluation of the A-type and B-type configurations leads to the conclusion that the B-type are only marginally better. From a performance point of view, it is difficult to recommend the B-type over the A-type. As well, the B-type configuration requires more floor space than the equivalent A-type.

However, when you consider the availability aspect, there is another incentive to consider the B-type configuration. For example, when you plan a 16 actuator configuration, the B-type reduces the impact of a hardware malfunction, in particular the the loss of a 3380-AA4 unit.

When the 3380-AA4 unit of an A16 string is lost, you lose access to up to 10 gigabytes of data, recovery of which may be very difficult.

When one 3380-AA4 unit of a B16 configuration is lost, you only lose access to up to 5 gigabytes of data which may have to be recovered. Finding space in your configuration to recover 5 gigabytes is easier than it is for 10 gigabytes!

Of course this availability consideration will also apply to the evaluation of A-type vs C-type configurations. The B-type and C-type configurations give you the opportunity to better plan and control your availability requirements.

---

[4] Refer to Washington Systems Center Technical Bulletin DASD Path and Device Contention Considerations, GG22-9217.

A-type versus C-type Configurations

From a performance point of view, it must be emphasized that:

- A pair of A08 configurations will perform better than a single C16 configuration.
- The same is true for a pair of A12 or A16 configurations relative to a single C24 or C32 configuration respectively.

To illustrate this point, consider Figure 27. This chart compares a C16 with a D16 configuration, where the D16 configuration consists of two A08 configurations.

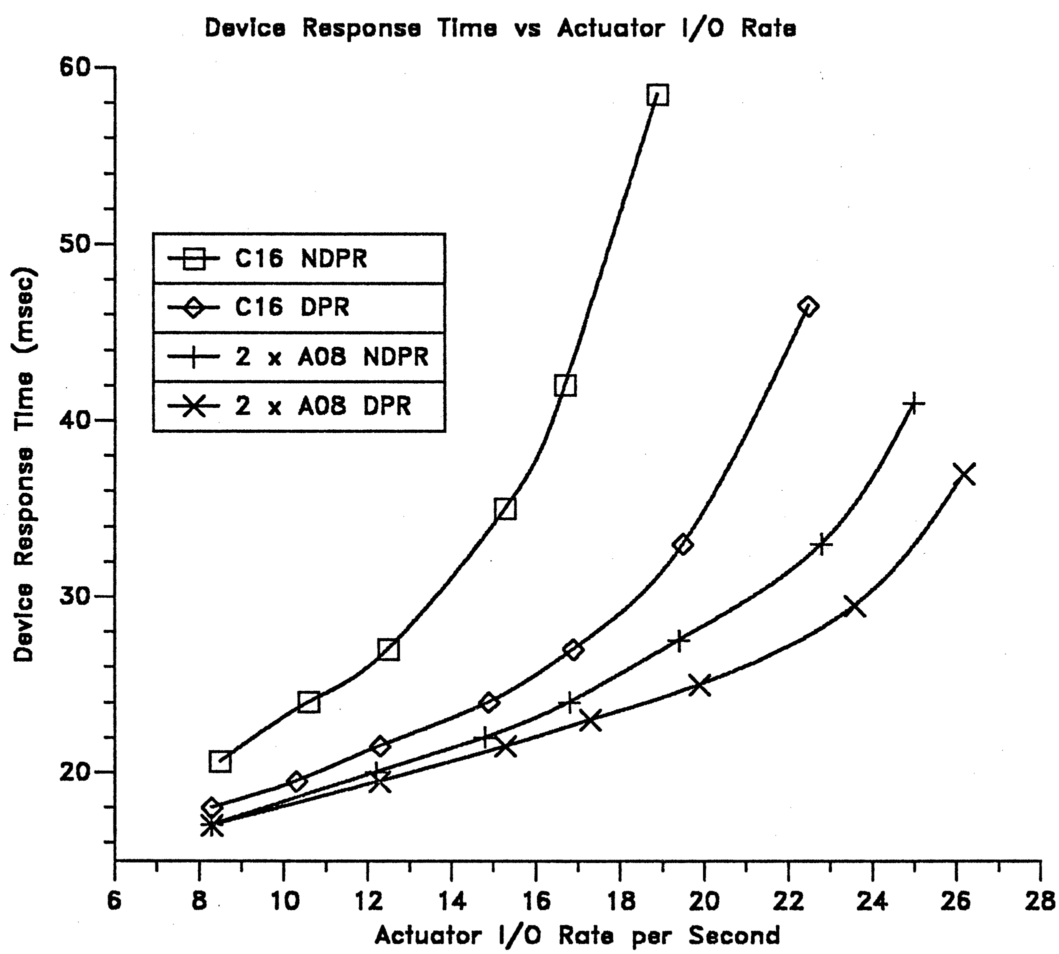

Figure 27. Comparison of an equivalent A-type and C-type configuration.

It is clear that the D16 configuration is superior to the C16 configuration, both in a DPR and Non-DPR environment. This is true for all points on the graph, in spite of the fact that the C16 configuration benefits more from DPR than does the D16 configuration.

From a hardware point of view, the only difference between a pair of A-type configurations and a C-type configuration is that the 2xA-type uses two additional channel paths. The C-type configurations should be considered as a valid alternative when there is a shortage of channel paths. In general, 308X processors can be configured with a sufficient number of high performance channel paths.

For this reason, it recommended that as a **first approach**, an I/O configuration should be constructed from **A-type Configurable Units**.

B-type versus C-type Configurations

If a decision is to be made between B-type and C-type configurations, then the most important factor to look at is the I/O activity rate.

- Decide in favor of the C-type configurations if the activity rate is high. Because there is a small Control Unit overhead for each I/O, the C-type is more attractive. This overhead will eventually be noticed, but only at very high activity rates, as can be seen from studying Figure 26 on page 59. Compare for this purpose the 'distance' between the B-type and C-type configurations for the IMS/VS and TSO workloads.

- Decide in favor of the B-type configurations if the Connect Time for each I/O operation is high. In this case relatively few I/O operations will fill the capacity of the two channel paths, and thus the overhead for each I/O will not be noticed. This is again confirmed in Figure 26 on page 59. The results for the corresponding B-type and C-type configurations in the TSO runs are to be considered as equivalent. This is certainly not true for the IMS/VS measurements.

In addition, the C-type configurations have a more attractive availability characteristic than the B-type. Should a Storage Director fail in a B-type configuration, all actuators can be accessed by only one path. Should a Storage Director fail in a C-type configuration, half of the actuators can be accessed via two paths, and the other half via one path. This might be a strong argument to decide in favor of C-type configurations, also when the performance improvement with respect to the B-type is not conclusive.

A-type versus E-type Configurations

The E-type configurations have the same performance level as a corresponding pair of A-type configurations, e.g. an E16 configuration is equivalent in performance to a pair of A08 configurations.

The difference between the two configuration types is that the E-type configurations consist of two A-type configurations or Configurable Units tied together by means of Two Channel Switches. The extra cost involved for the additional hardware features can be offset by the higher **availability** characteristic for this configuration. Should one out of the four channel paths in a pair of A-type configurations be lost, one of the strings can be accessed via only one path. If the same problem arises in a E-type configuration, all actuators in the two strings can be accessed via three channel paths.

The E-type configuration is the most obvious choice for 3084 complexes that have to run both in single-image mode and physically partitioned.

It is important to understand that there is **NO** performance penalty in bringing together two different types of workload into **one** E-type configuration, as long as the different workloads are not intermixed in the same string. All of us have learned in the past that every additional level of switching in a configuration has a performance impact. Not so in this case, since there are as many Channel Paths as Storage Directors, and DPR provides for reconnection on any available path, the activity in one string will never interfere with the activity in the other string.

The fact that you are able to have all channel interfaces enabled at the Storage Director level will simplify the recovery procedures, making them less prone to error, for 3380 DASD attached to an MVS/XA system.

In summary, the E-type configurations can be qualified as the most sympathetic to high and dynamic workloads.

Channel Path Utilization

The issue of channel path utilization is normally not as important in an MVS/XA environment as it has been in an MVS/370 environment. Some of the reasons for this are:

- The DPR capability increases the capacity of a set of channel paths.
- The channel paths of the Dynamic Channel Subsystem in combination with 3380 DASD can sustain a data rate of 3 Mb/sec. Many System/370 channels are working with a data rate of 1.2 Mb/sec. (e.g. 3350) or 1.5 Mb/sec. (e.g. 3380 with Speed Matching Buffer).

- Channel capacity, in terms of the number of configurable paths, is usually not as scarce a resource on a 308X processor running MVS/XA as on previous processors (e.g. 303x, 370/158, and 370/168).

Set Sector Reconnect Delay is a very important part of the average Device Response Time for an I/O request. In order to maintain control over this factor, channel path utilization is not the only thing to consider. When a device attempts to reconnect to the channel path it sees other things than just channel path utilization.

- In some cases a high internal path utilization is responsible for an undesirable Device Response Time. From an analysis of the measurements performed, internal path utilization did not turn out to be a major problem. For example, the A-type and B-type configurations provided almost equivalent performance.

  In a dynamic customer environment, the internal path contention may present a more significant problem than was observed in the stable benchmark environment. In this case, the B-type configuration with twice the number of internal paths may provide significantly better performance.

- As explained earlier, high access density types of workload are sensitive to Control Unit utilization.

- Only after the device has obtained the internal path and the Storage Director will it be affected by the actual Channel Path Utilization.

It depends on factors such as the type of workload, the amount skewness, and the size of the Configurable Unit if the device will meet any of these problems, and if it does, which one it will be.

This is one of the reasons for choosing Device Response Time as the final evaluation criterion rather than Channel Path Utilization.

You can get a more precise feeling for what channel path utilization is all about in a Dynamic Channel Subsystem by looking more closely at the Scenario Two tables. Since channel path utilization can be defined as: Total I/O Rate x Connect Time, it is valid to apply the improvement percentages of the tables to the channel path utilization which is also presented in the same table. You will then arrive at the new channel path utilization in MVS/XA. For example, consider the following table:

Configuration /Workload		Non-DPR Channel Path Utilization	Actuator I/O Rate Increase	DPR Channel Path Utilization
A16	IMS/VS	29.2%	10%	32.1%
	Batch	34.6%	15%	39.8%
	TSO	34.4%	15%	39.6%
C16	IMS/VS	33.3%	20%	40.0%
	Batch	35.5%	23%	46.2%
	TSO	44.2%	16%	50.7%

Figure 28. Spread in channel path utilization for 16 actuator configurations.

This table illustrates that the acceptable channel path utilization may range from 32% - 51% for a sixteen actuator configuration. Our conclusion is that while channel path utilization is an important performance issue in some environments, it is rather difficult to develop a **universal** guideline.

The most channel constrained configuration was the C32 configuration in the TSO environment. Its low Actuator I/O Rate avoids problems at the internal path and Control Unit level. The maximum channel path utilization measured for the C32 configuration was 56% while maintaining a Device Response Time of 50 msec.

Channel Capacity in Terms of I/O Requests per Second

As a result of DPR it is possible to sustain I/O Rates on a channel path which are unprecedented. The number of actuators of each configuration in Figure 26 on page 59 multiplied by the corresponding Actuator I/O Rate will give the I/O Rate for two channel paths, except for the E-type configurations. In the last case you will have the combined I/O Rate of four channel paths.

The resulting figures (when recalculated to the I/O Rate per channel path) range from as low as 36 I/O requests/second per channel path to as high as 192 I/O requests/second per channel path. This range is so large that it is very difficult, if not impossible to develop any coherent guideline or Rule of Thumb. The reason is of course the same as outlined in the previous observation about Channel Utilization: the channel path is no longer the most critical resource. Therefore, do not focus your attention in planning solely on that area.

If you do make the calculations with the objective of understanding more about channel capacity, then you must also understand that the resulting numbers represent the **maximum** that a configuration can deliver at a certain Device Response Time objective. This is probably not what you would like to use for DASD configuration planning, be it short, medium, or long term.

Actuator I/O Rate

The 3380 provides the possibility of achieving very high Actuator I/O Rates, thus allowing for very high **access density**. The maximum measured was 27 I/O requests per second for the IMS/VS workload. Exploit this capability if you need it!

However, you will certainly not achieve high actuator I/O Rates with a consistent performance behavior if you do not carefully plan for it. Benchmark environments such as used for this study are very stable and produce predictable results. Real life environments are generally not as stable and you will need a high level of control over the parameters that may have an influence. Consider the following:

- Limit the size of the Configurable Unit.
- Consider A-type, C-type, or E-type configurations.
- Explicitly plan dataset placement in order to avoid long Seek times.
- Avoid long Search Times. Examples of these are non-indexed VTOC searches and PDS Directory searches. Use Indexed VTOCs and the Linklist Look-Aside feature available with MVS/SP 2.1.1 and later.
- Investigate the nature of the application; will it generate equally distributed or skewed arrivals. If the arrival pattern of the I/O operations is very skewed, the Device Service Time will not suffer if you configured properly. But you could end up with long IOSQ Times.
- It has been explained earlier that long data transfers are not necessarily incompatible with high Actuator I/O Rates. (Refer back to the chapter on Configurable Units.) But in this case the channel path utilization could become the most critical factor.
- When very high I/O Rates are required, you should examine the applicability of a buffered Control Unit such as the 3880-13.
- For IMS/VS, acceptable I/O rates are very much dependent on the configuration.
- For Batch, try to get as much out of it as you can manage. It is throughput oriented!
- For TSO, any Actuator I/O Rate above 8-10 requests per second needs special attention.

Size of a Configurable Unit

It is possible to obtain a good Device Response Time with many actuators in a Configurable Unit, but only at relatively low Actuator I/O Rates. Figure 29 illustrates this very well. In the case of the B32 Configuration, the TSO Terminal Response Time increases rapidly while the actuators are still at a low activity level. It must be clear that such a configuration is not suitable for handling a dynamically changing workload. Hence, the B32 Configuration is probably not the first one to think of when planning a DASD configuration for TSO.

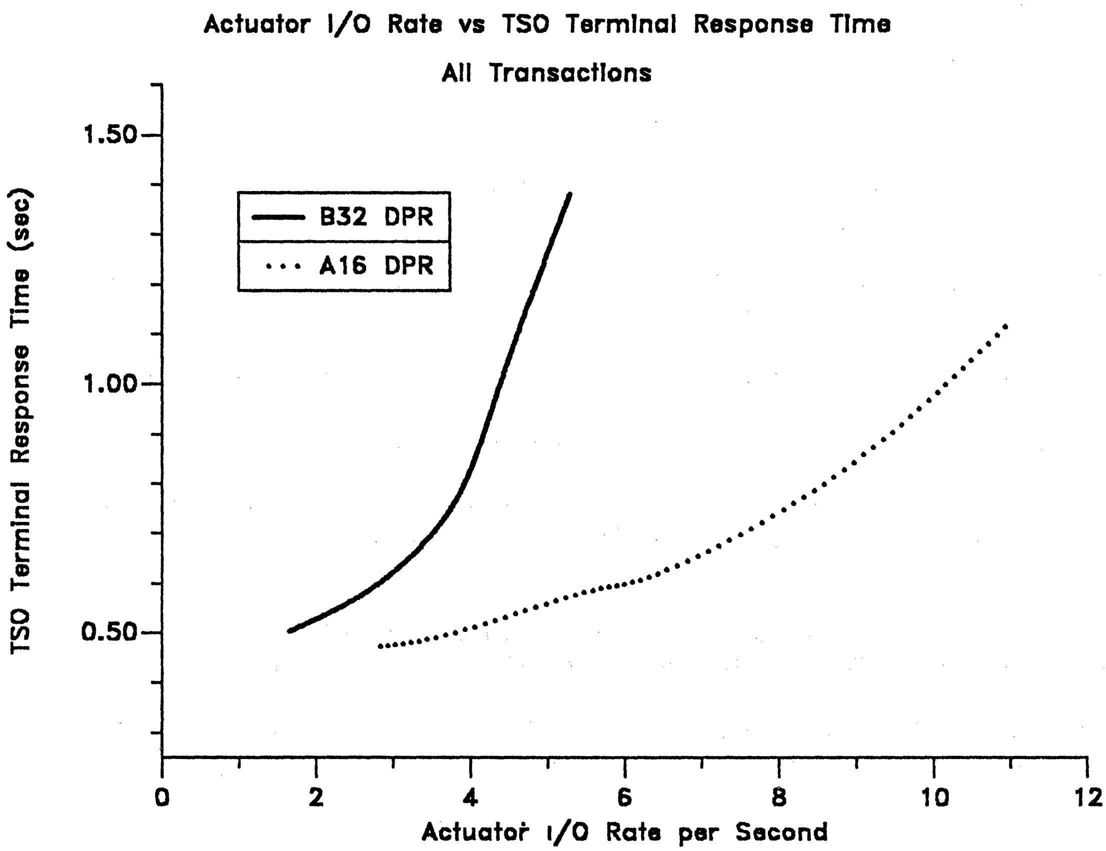

Figure 29. The possible effect of the size of a Configurable Unit on Terminal Response Time.

Especially for Batch workloads and to a lesser degree the TSO workload, there will be a difference between **active** actuators and **configured** actuators. If for example the outcome of the measurement indicates that an E16 configuration would be suitable to meet the performance requirements of a particular Batch workload, then quite possibly an E24 configuration would also fulfill the performance expectations.

Summary

As a summary, we propose the following approach to configuration planning:

1. Determine the applications, system and user, that will require dedicated Configurable Units. Think of Swapping, Paging, IMS/VS Databases, TSO, VSPC, and Batch datasets. Inevitably you will have to compromise as it will not always be possible to dedicate a Configurable Unit for a particular workload. If you combine workloads in a Configurable Unit, then try to avoid intermixing:

   - Different workloads on the same actuator.
   - Throughput oriented and response time oriented workloads.
   - Long Connect Times and short Connect Times.
   - Actuators with a critical I/O requirement with any other type of work.

2. Estimate the I/O rate that will be generated in the system for each Configurable Unit identified in point one above.

3. Formulate your Device Response Time objective for each Configurable Unit.

4. Measure or compute the Connect Time per Configurable Unit.

5. Use your Device Response Time objectives and Connect Times to define the expected Actuator I/O Rates. Use Figure 26 on page 59, but keep a safety margin.

6. Divide the expected Actuator I/O Rate into the Total I/O Rate expected for the Configurable Unit. The result is the required number of actuators.

7. Having arrived at the total number of actuators, a choice must be made for either A-, B-, C- or E-type Configurable Units.

   Use the table and criteria illustrated in Figure 30 on page 70.

Number of Channels Paths	2	2	2	4
Number of Storage Directors	2	2	4	4
Number of 3380-AA4 Units	1	2	2	2

Configurable Units	A04	–	–	–
A-, B-, and C-type Configurations are considered as a SINGLE Configurable Unit	A08	B08	C08	E08
	A12	B12	C12	E12
	A16	B16	C16	E16
E-type configurations are considered as a combination of TWO Configurable Units		B20	C20	E20
		B24	C24	E24
Note: This table is NOT exhaustive		B28	C28	E28
		B32	C32	E32

Figure 30. Overview of 3380 Configurable Units

A04 Configurable Unit

There is no other option for this Configurable Unit. It should be considered only for a very specific performance environment, such as Paging/Swapping or high use program libraries.

A08 - B08 - C08 Configurable Units

- These Configurable Units should only be considered for high performance applications.
- The first choice is the A08 Configurable Unit.[5]
- The B08 Configurable Unit is a valid option if a growth path is required.
- The C08 Configurable Unit is not a realistic option as the number of actuators is not enough to justify two additional Storage Directors. These could be added to a B-type Configurable Unit when the type and amount of work demand it.

---

[5] Combining two A-type Configurable Units into an E-type configuration is recommended.

A12 - B12 - C12, A16 - B16 - C16 Configurable Units

- The A-type is the most obvious choice.[6]

- The B-type or C-type Configurable Units are hard to justify for performance reasons. The double A-type, or the E12/E16 perform better than the B12/B16 or C12/C16.

- Decide in favor of B- or C-type if maximum control over availability is required.

B20 - C20, B24 - C24, B32 - C32 Configurable Units

- These Configurable Units should only be considered when a solution with A-type or E-type is not possible.

- The attractive availability characteristic of the B-type and C-type configurations is partially lost when the string length of the Configurable Unit is increased.

- Channel Utilization might become a critical factor on the larger C-type Configurations.

- Decide in favor of the C-type if the number of I/O requests per channel exceeds 100-120. If the number is lower, then the B-type will be adequate.

- These Configurable Units are most exposed to inconsistency in performance results. They are the least suitable for dynamically changing workloads.

[6] Combining two A-type Configurable Units into an E-type configuration is recommended.

INTRODUCTION

When migrating from MVS/370 to MVS/XA, 3350 and mixed 3350/3380 configurations require special attention. When you have both 3380 and 3350 devices sharing the same channel paths, you may end up with worse device response times, particularly when the increased capability of the 3380 devices is taken into consideration.

If you have string switches on the 3350 strings, you may have specified a certain channel selection algorithm in the IECIOSxx member of SYS1.PARMLIB to reduce the interference caused by a string switch. In MVS/XA, all channel selection algorithms have been removed, since path selection and management is done by the Dynamic Channel Subsystem. A Preferred Path specification has been introduced to address the string switch situation.

In this chapter, we will discuss the following topics:

- Preferred path specification
- Use of Preferred Path with 3350 DASD
- Considerations for 3380 and 3350 DASD sharing channel paths.

PREFERRED PATH SPECIFICATION

We include a brief discussion of the **Preferred Path** specification for an I/O device as the concept is important in understanding the configuration guidelines presented later.

In 370-XA mode, the channel subsystem attempts to initiate an I/O request over the **Preferred Path,** if one is specified. If the preferred path is busy or if no preferred path is specified, the channel subsystem uses the **rotation algorithm** to initiate the I/O requests.

Preferred Path

The objective of preferred path is to provide a method of controlling the path selection logic in the channel subsystem to reduce the contention for the head of string controller. This would have a direct effect on the average Device Response Time for an I/O request. Preferred path is usually used in a situation where there are two strings of actuators attached to a pair of Control Units or Storage Directors by means of String Switches. The negative influence of String Switches on device response time has been already discussed in many documents.

Let us review briefly why the string switch has a negative impact on device response time. We assume that I/O activity to the two strings is evenly distributed.

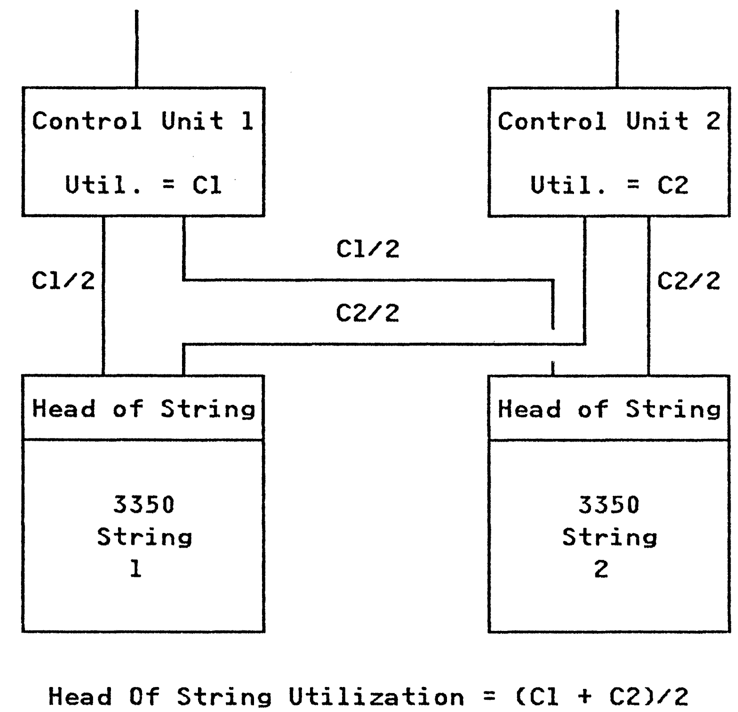

Figure 31. 3350 String Switched Configuration

Consider the switched configuration in Figure 31. There a String 1 device ready to reconnect with Control Unit 1 faces two points of blockage: head of string and control unit. Even if String 1 is free, Control Unit 1 might by tied up serving String 2. In that case, reconnection could not occur; a full revolution (16.6 msecs) would be lost.

If the left and right sides of this configuration were not connected through a string switch, there would be only one point of blockage; namely, head of string. Thus an non-switched configuration entails a lower probability of missed reconnections and generally better performance than does a switch configuration.

PREFERRED PATH AND 3350

In MVS/370, there are several algorithms for channel path selection. When IOS attempts to initiate an I/O request to a device, a channel path to the device is selected according the the specification in the IECIOSxx member in SYS1.PARMLIB. The possible algorithms are: SEQUENTIAL, REVERSED SEQUENTIAL, LAST CHANNEL USED, ROTATE and BALANCE. These algorithms are specified at the logical channel level. If no OPTCHAN is specified for the device, then there is only one channel path available and the channel algorithms are not used.

In MVS/XA, these algorithms have been removed, since the Dynamic Channel Subsystem performs the path selection function instead of IOS. The channel subsystem performs the channel path selection as follows:

- **Preferred path** - The channel subsystem attempts to use the Preferred Path, if one is specified, for the initiation of an I/O request to the device. If the preferred path is busy or if no preferred path is specified, the channel subsystem uses the rotation algorithm. Preferred path is specified via the PATH keyword on the IODEVICE macro in the IOCP generation process.

- **Rotation Algorithm** - The channel subsystem will first try the next channel path in the rotation order following the last channel path used for an I/O request to a device in that logical control unit.

When a preferred path is not specified, the utilizations of the channel paths in a logical control unit tend to be balanced. (There is a exception in the case of 3380 which we will discuss in "3380 and 3350 Mixed Configurations" on page 77.) From the performance point of view, there are some cases where you should specify a preferred path for devices. For example, in Figure 32 on page 76, you should specify a preferred path for all devices in String 1 connected to CHPID 05 and a different preferred path for all devices in String 2 connected to CHPID 25. This will reduce the interference between String 1 and String 2, and provide better Device Response Times.

Of course an alternative approach would be to switch off one of the channel paths so that each string has a dedicated channel path and control unit. But, if a channel path or a control unit goes to down, operator intervention will be required to restore access to the string of devices.

If both channel paths and both control units are available for access to a device, MVS can continue to access the device without any operator intervention. If you want this automatic recovery and still do not want to lose performance, you should specify different preferred paths for each string of devices.

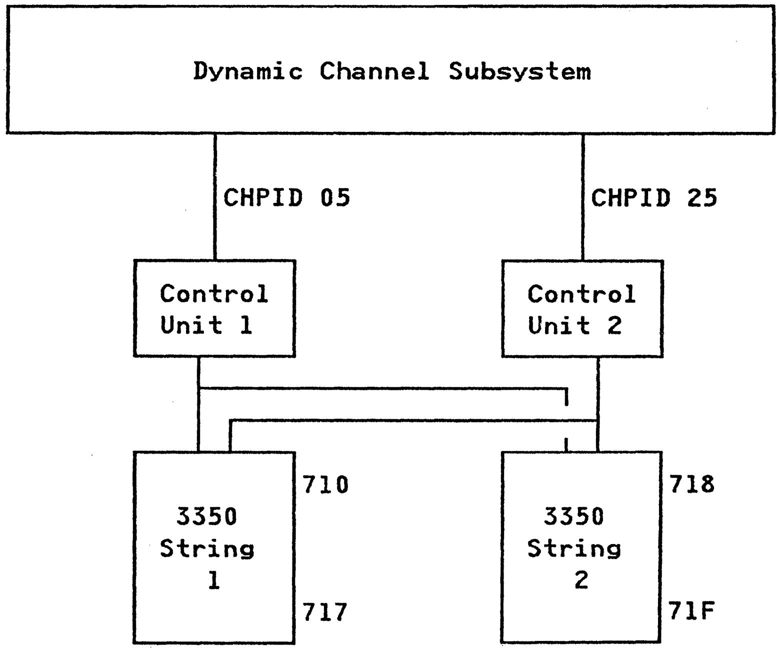

IOCP Example

```
    ID    MSG1= ......
    CHPID PATH = ........
    CNTLUNIT CUNUMBR=000,PATH=05,SHARED=N,UNIT=3880,        *
          UNITADD=((00,16))
    CNTLUNIT CUNUMBR=001,PATH=05,SHARED=N,UNIT=3880,        *
          UNITADD=((00,25))
    IODEVICE CUNUMBR=(000,001),UNIT=3350,ADDRESS=(710,8),   *
          PATH=05      <---- PREFERRED PATH
    IODEVICE CUNUMBR=(000,001),UNIT=3350,ADDRESS=(718,8),   *
          PATH=25      <---- PREFERRED PATH
```

Figure 32. 3350 String Switched Configuration - IOCP Example

3380 AND 3350 MIXED CONFIGURATIONS

In this section we discuss considerations related to the mixing of 3350 and 3380 DASD on the same channel paths. While we do not recommend the mixing of 3350 and 3380 devices in an MVS/XA environment, there are some situations where it may be unavoidable, such as a shortage of channel paths, during a 3350 to 3380 migration, or during a backup situation.

Example 1 - Single Shared Path

Figure 33 is an example of a possible mixed configuration that may require special consideration. We are assuming that the shared path is required either due to a shortage of channel paths and/or as an 'availability' path.

In an MVS/XA environment, you have the possibility of increasing the Actuator I/O Rate of the 3380 devices due to the DPR capability and therefore increasing the channel path utilization. This may have a negative affect on the 3350 devices which do not have the DPR capability.

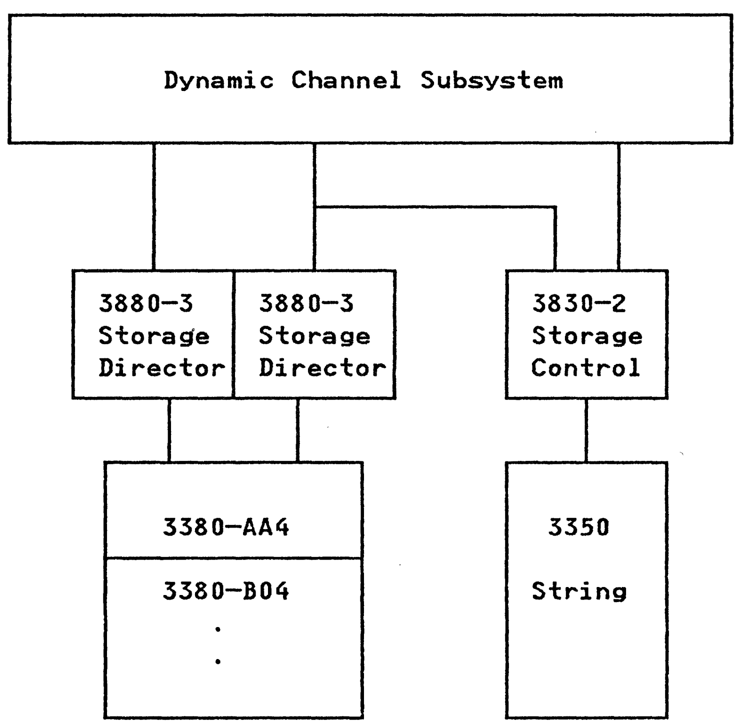

Figure 33. 3380 and 3350 DASD Sharing a Channel Path

In such a case, any 3350 I/O requests initiated across the shared channel path may experience worse device response times than under MVS/370, because the the probability of a reconnection miss will become higher. Under MVS/370, the solution would be to disable the interface switch to the 3350 string on the shared path. However, in the event of a path malfunction, this would require operator intervention to enable the 'availability' path.

In an MVS/XA environment, the recommended approach is to specify the unshared path to the 3350 string as the 'preferred path'. This will cause the channel subsystem to use this path almost exclusively for the 3350 I/O requests and avoid the extended device response times. This approach also permits the shared path interface to remain enabled during normal operation, thus providing an 'availability' path without operator intervention.

Example 2 - Multiple Shared Paths

For this example, refer to Figure 34 on page 79. If we consider the 3380 portion of the configuration first, two channel switch features have been installed on both storage directors providing two additional paths to the channel subsystem. These additional paths will not increase the performance of the 3380 devices but do provide 'availability' paths. In an MVS/XA environment, you will find that the channel paths connected to the 'A' interfaces of the storage directors will have a higher utilization than the 'B' interfaces. In fact, the utilization across the 'A' interfaces will be quite well balanced. The utilization across the 'B' interfaces will also be quite well balanced but lower.

The reason is that while it is possible for the I/O request to reconnect on any of the four channel paths, the 3880 Storage Director will favor the lower (A interface) channel path if available.

If we now consider including the 3350 devices in this configuration, they should be connected to the channel paths connected to the 'B' interfaces of the 3880 Storage Directors. This configuration would provide good device response time for all devices configured as well as providing 'availability' paths to all devices.

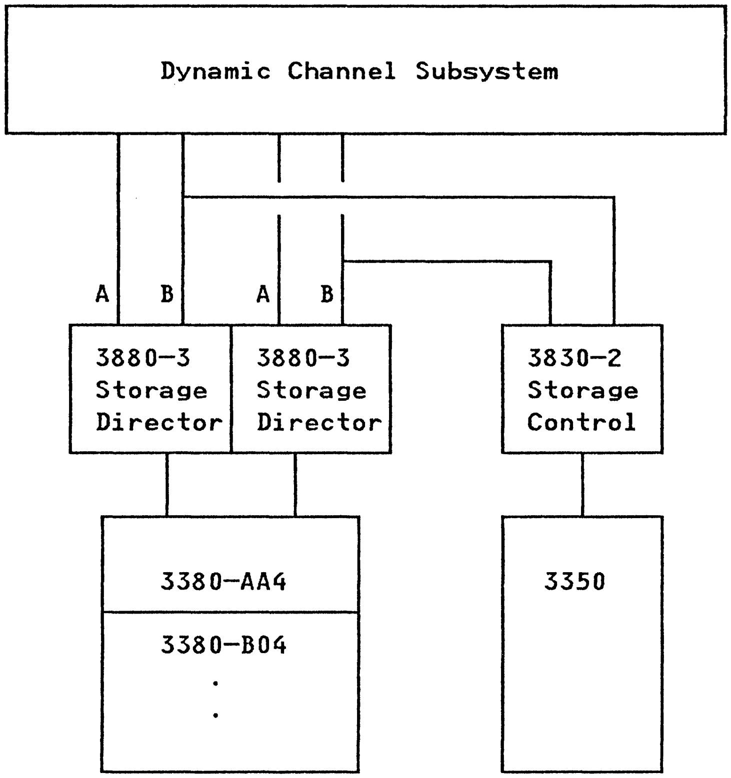

Figure 34. 3380 and 3350 DASD Sharing Two Channel Paths

APPENDIX A. BIBLIOGRAPHY

1. R.L. Cormier, R.J. Dugan and R.R. Guyette, **System/370 Extended Architecture: The Channel Subsystem**, IBM Journal of Research and Development, Vol. 27 No. 3 (May 1983).

2. P. Dorn, S.E. Friesenborg, and P. Hallsworth, **Performance and Capacity Considerations**, GG22-9232, IBM Corporation; available through IBM branch offices.

3. S.E. Friesenborg, **DASD Path and Device Contention Considerations**, GG22-9217, IBM Corporation; available through IBM branch offices.

4. S.E. Friesenborg, **DASD Access Method Performance Considerations**, GG22-9241, IBM Corporation; available through IBM branch offices.

5. **IOCP User's Guide and Reference**, GC28-1027, IBM Corporation; available through IBM branch offices.

6. **MVS/XA RMF Reference and User's Guide**, SC28-1138, IBM Corporation; available through IBM branch offices.

7. U. Pimiskern, **MVS/XA I/O Performance Considerations**, GG22-9346, IBM Corporation; available through IBM branch offices.

8. R. Wicks, **Balanced Systems and Capacity Planning**, GG22-9299, IBM Corporation; available through IBM branch offices.

9. R. Wicks, **DASD Expectations: The 3380, 3880-13, and MVS/XA**, GG22-9363, IBM Corporation; available through IBM branch offices.

READER'S COMMENTS

Title: MVS/XA DASD Configuration Guidelines

International Systems Center Technical Bulletin ZZ10-5049-00

You may use this form to communicate your comments about this publication, its organization or subject matter with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Comments:

Reply requested   Name : ______________________________

Yes / No   Job Title : ______________________________

Address : ______________________________

____

Cut or Fold Along Line

Reader's Comment Form

Fold and tape

Please Do Not Staple

Fold and tape

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE:

IBM International Systems Center
Department H52, Building 930
P.O. Box 390
Poughkeepsie, New York 12602
U.S.A.

Fold and tape

Please Do Not Staple

Fold and tape

IBM®

ZZ10-5049-00
IBM Internal Use Only

MVS/XA DASD Configuration Guidelines ZZ10-5049-00 IBM Internal Use Only