AD-A059 938 COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON D C F/G 9/2
WORLDWIDE MILITARY COMMAND AND CONTROL SYSTEM (WWMCCS) H-6000 T--ETC(U)
FEB 78 B M WALLACK, G H GERO
UNCLASSIFIED CCTC-TM-180-78-VOL-2 SBIE-AD-E100 095 NL

1 OF 2
AD
A059938

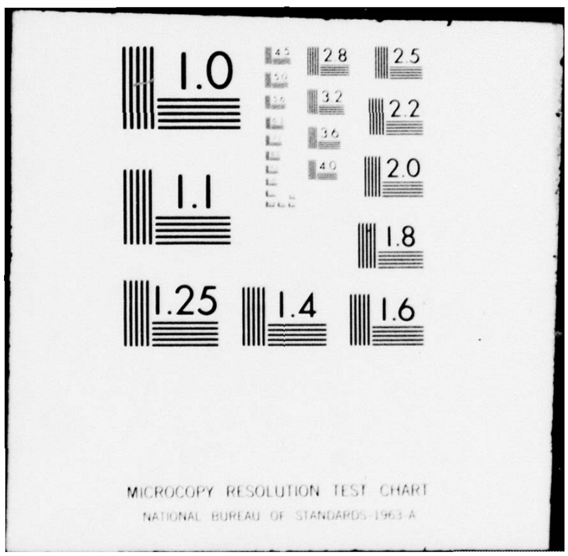
1.0
2.8
2.5
3.2
2.2
3.6
4.0
2.0
1.1
1.8
1.25
1.4
1.6
MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD A059938

CCTC

TECHNICAL MEMORANDUM
TM-180-78
VOLUME II
3 FEBRUARY 1978

LEVEL III

COMMAND & CONTROL TECHNICAL CENTER

DDC RECEIVED OCT 17 1978 B

DDC FILE COPY

WORLDWIDE MILITARY COMMAND AND CONTROL SYSTEM (WWMCCS)

H-6000 TUNING GUIDE

VOLUME II

BATCH TURNAROUND TIME ANALYSIS PROCEDURES

DEFENSE COMMUNICATIONS AGENCY

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

78 10 10 041

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

AD-E100 095

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER CCTC-TM-180-78-VOL-2	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Worldwide Military Command and Control System (WWMCCS). H-6000 Tuning Guide. Volume II - Batch turnaround time Analyses Procedures.		5. TYPE OF REPORT & PERIOD COVERED Technical memorandum 6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Barry M. Wallack George H. Gero		8. CONTRACT OR GRANT NUMBER(s)
9. PERFORMING ORGANIZATION NAME AND ADDRESS Command and Control Technical Center Attn: CCTC/CPE/C702 Washington, DC 20301		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Same as above		12. REPORT DATE 3 February 1978 13. NUMBER OF PAGES 165
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Computers, Test and evlauation, Performance tests, Time sharing

Honeywell 6000 computers, Computer systems, Federal Computer Performance and Evaluation and Simulation Center

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The Federal Computer Performance Evaluation and Simulation Center (FEDSIM) has developed a document for WWMCCS installations that can be used by site personnel to analyze the performance characteristics of their Honeywell 6000 (H-6000) computer systems. This document, called an H-6000 Tuning Guide, incorporates detailed analyses procedures that guide the analyst in applying specific techniques to improve system performance.

This volume presents a set of procedures for analysis of batch job turnaround time. It first presents a model of the processis and queue points associated

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered) with batch job turnaround time and then describes nine tests that use the model to direct the analysis of turnaround time.

COMMAND AND CONTROL TECHNICAL CENTER

Technical Memorandum TM 180-78

3 February 1978

H-6000 TUNING GUIDE

BATCH TURNAROUND TIME ANALYSIS PROCEDURES

VOLUME II

PROJECT PERSONNEL:

Barry M. Wallack
George H. Gero JR.
CCTC/CPE/C702
RM BE685, The Pentagon
Washington, D.C. 20301
AV - 692-2725
COMM - 697-2725

APPROVED BY:

TIPTON P. MOTT-SMITH
Colonel, USAF
Deputy Director for
Computer Services

PREFACE

This report is based on detailed analysis of a large amount of technical information. The results address procedures for the analysis of batch turnaround time and GCOS Time Sharing System response time in World Wide Military Command and Control Systems (WWMCCS). Because of the complexity of the anlaysis procedures and their dependence on the WWMCCS workloads and operational environments, generalizing the procedures beyond the environment described or extracting conclusions without their respective qualifying conditions is not possible. Questions related to this report or the possibility of extending the stated conclusions or recommendations should be addressed to the Computer Performance Evaluation Office, C702, The Pentagon.

To gain a general understanding of the approach of the H-6000 Tuning Guide, Volume I Section II, Volume II Section II, and Volume III Section II should be read. One or more of the hypothesis tests (search procedures) in Volume II Sections IV-XII and Volume III Sections III-X should also be read. Not all these tests have to be read at the start of a tuning effort. Each should be read as it needs to be applied. To start a tuning effort, Volume I should be read and applied. The procedure for anlaysis of batch turnaround time begins in Volume II Section III. The procedure for analysis of Time Sharing response time begins in Volume III Section II.

The H-6000 Tuning Guide has never been tested by a novice in performance evaluation, although field tests have been conducted by FEDSIM personnel. For this reason, it remains a preliminary version.

TABLE OF CONTENTS

TABLE OF CONTENTS (Cont'd)

TABLE OF CONTENTS (Cont'd)

TABLE OF CONTENTS (Cont'd)

TABLES

TABLES (Cont'd)

TABLES (Cont'd)

FIGURES

FIGURES (Cont'd)

TABLES (Cont'd)

FIGURES

ABSTRACT

The Federal Computer Performance Evaluation and Simulation Center (FEDSIM) has developed a document for WWMCCS installations that can be used by site personnel to analyze the performance characteristics of their Honeywell 6000 (H-6000) computer systems. This document, called an H-6000 Tuning Guide, incorporates detailed analysis procedures that guide the analyst in applying specific techniques to improve system performance.

The four volumes of the Tuning Guide present a precisely structured system of procedures for the analysis of the performance of WWMCCS computer services and systems:

Volume I — WWMCCS System Tuning Process. The first volume describes the overall structure and application of the Tuning Guide. It explains the approach, procedures, and processes taken by the Tuning Guide to provide analyses of batch job turnaround time and GCOS Time Sharing System (TSS) response time.

Volume II — Batch Turnaround Time Analysis Procedures. The second volume presents a set of procedures for analysis of batch job turnaround time. It first presents a model of the processes and queue points associated with batch job turnaround time and then describes nine tests that use the model to direct the analysis of turnaround time.

Volume III — TSS Response Time Analysis Procedures. The third volume serves the same general purpose and has the same general structure as Volume II. Volume III presents a complete set of procedures for investigating the response time of GCOS Time Sharing System (TSS) interactions. The volume first presents a model of the processes and queue points associated with TSS response time and then describes eight tests to direct an analysis of TSS response time.

Volume IV — H-6000 Tuning Guide Appendices. The fourth volume provides the appendices referenced by the other volumes of the Tuning Guide. The volume contains detailed descriptions of report formats and other references data.

I. INTRODUCTION

A. BACKGROUND

The Office of the Joint Chiefs of Staff (JCS) has directed that the Command and Control Technical Center (CCTC) develop a computer performance analysis capability to support the World Wide Military Command and Control System (WWMCCS).

CCTC, acting at the direction of the JCS, has specified that WWMCCS ADP managers are to apply various computer performance evaluation (CPE) tools and techniques to the systems now running at their sites. CCTC has also defined the need to instruct WWMCCS technical personnel in the selection and application of the CPE tools and techniques appropriate to individual WWMCCS ADP sites.

CCTC asked FEDSIM to plan and implement a document that could be employed by WWMCCS ADP personnel to diagnose problems and propose changes that would improve the performance of WWMCCS ADP systems.

B. PROJECT OBJECTIVE

The objective of the resulting FEDSIM project was to provide all WWMCCS installations with a document that could be used by staff personnel to analyze the performance characteristics of their ADP systems. This document, called an H-6000 Tuning Guide, was to contain sets of analysis procedures to improve system performance.

The product of the completed FEDSIM project is a four-volume H-6000 Tuning Guide (referred to hereafter as the Guide). The Guide volumes present a precisely structured system of procedures for the analysis of WWMCCS computer services and systems. The titles of the four volumes are: (1) WWMCCS System Tuning Process, (2) Batch Turnaround Time Analysis Procedures, (3) TSS Response Time Analysis Procedures, and (4) H-6000 Tuning Guide Appendices.

C. PURPOSE AND CONTENTS OF THE GUIDE

Computer jobs may be submitted by WWMCCS users to run either as batch jobs or as on-line jobs. Batch jobs, as processed by the WWMCCS systems, may be submitted by users at a site or may be initiated via a process called "job spawning" through the WWMCCS Time Sharing System. On-line jobs addressed by the Guide include the subsystems that run under control of the WWMCCS Time Sharing System. The performance of both batch and on-line jobs can be measured and analyzed with reference to the amount of elapsed time that the system takes to process them. Batch job elapsed processing time is called batch turnaround time. On-line job (or, more precisely, "terminal interaction") elapsed processing time is called response time.

Volume II of the Guide presents a complete set of procedures for analyzing the turnaround time of batch jobs at a WWMCCS site. The volume first presents a model of the processes and queue points associated with batch job turnaround time and then describes nine tests that use the model to direct an analysis. Both the model and the tests are designed to be applied only in the WWMCCS system environment.

Table I-1 lists the sections of Volume II and briefly identifies how each is to be used.

SECTION NO.	TITLE	APPLICATION OF SECTION
I	Introduction	
II	Turnaround Time Model	1. Define Potential Delay Points 2. Hypothesize Causes For Delay
III	Turnaround Time Model Scan	1. Direct Analysis of Delay Points 2. Initiate Other Tests
IV	Seek Elongation Test	1. Confirm Disk Seek Delay 2. Identify Files to Move
V	Memory Constraint Test	1. Confirm High Memory Wait Time 2. Identify Source of Wait
VI	Device Errors Test	1. Determine Tape Handlers In Error 2. Determine Disk Units In Error
VII	Pathway Utilization Test	1. Validate I/O Service Queuing Delay 2. Isolate Libraries And Files To Move
VIII	CPU Utilization Test	1. Determine CPU Utilization Level 2. Determine Dominant CPU User 3. Identify High CPU Jobs 4. Develop Execution Activity Map
IX	Insufficient Device Test	1. Confirm Tape-Caused Delay 2. Confirm Disk Space Delay
X	"Few Activities In System" Test	1. Substantiate Light Job Scheduling 2. Identify Scheduling Alternatives
XI	IOS Delays Test	1. Confirm GCOS IOS Code As Delay 2. Propose Alternative Solutions
XII	Urgency Codes Test	1. Confirm Delay Caused by Urgency Codes 2. Determine If Initiated by Operator

CONTENTS OF VOLUME II

TABLE I-1

II. BATCH TURNAROUND TIME MODEL

The Batch Turnaround Time Analysis Procedures described in the remaining sections of this volume use a model of WWMCCS system processes to guide analysis. The model, documented in this section, assists a CPE analyst in formulating and testing hypotheses concerning the cause of batch turnaround time elongation.

A. DEFINITION OF MODEL STRUCTURE

The Batch Turnaround Time Model is an organized description of the components of batch turnaround time in a WWMCCS system. The Model identifies potential points where batch jobs and activities could be delayed (i.e. elongated) as they are processed. The Model is used to: (1) guide analysis and (2) provide specific batch turnaround time elongation hypotheses. The Turnaround Time Model Scan (see Section III) provides guidance; nine batch turnaround time analysis tests (see Section IV through XII) are used to confirm or deny hypothesized causes of job elongation.

In a general sense, batch turnaround time is a measure of the total elapsed time required by a data processing installation to process an individual batch job. This includes both computer processes and the physical handling of media that are used by a job. Batch turnaround time can also be measured for all batch jobs processed during a given operating period (e.g., shift, day, week) to give the aggregate processing characteristics of a computer site. In the former case, the analysis is conducted to examine a site's handling of an individual job; in the latter case, the analysis is conducted to examine the processes through which jobs are directed as they pass through the site.

The Batch Turnaround Time Model views batch jobs as entering the computer system from either a local or a remote source. Jobs submitted for processing at a site job receipt counter are considered as Local Batch jobs. Jobs that have been submitted from remote batch terminals or spawned by GCOS Time Sharing System users are considered as remote batch jobs. If the output of these jobs arrives directly at the user's terminal with no computer center action, the jobs are considered Remote Batch "A" jobs. If their output must be removed from the system (e.g., printer) and physically delivered to the user, they are considered Remote Batch "B" jobs. These distinctions are useful when analyzing delays (to turnaround time) experienced outside the computer system.

Figure II-1 shows the three primary levels of the Batch Turnaround Time Model: (1) Model Selection Level, (2) Phase Selection Level, and (3) Process Selection Level. The Model Selection Level describes the three types of batch jobs dealt with by the Model. The Phase Selection Level describes the three phases of batch processing in the Model: (1) Pre-Processing, (2) System Processing, and (3) Post Processing. The Process Selection Level describes the main processes of System Processing. Also included, but defined by the analyst, are processes composing Pre-Processing and Post Processing. Sample processes for these phases are discussed below. The Model considers hardware, software and queue residency delays beneath the third primary level. Each process has a list of Elongation Hypotheses - one for each specific cause of delay to that process. Each Elongation Hypothesis can be confirmed or denied by one of the nine tests which comprise most of this volume. In this way, the tests identify the cause(s) of the delays to batch jobs.

B. USE OF MODEL STRUCTURE

A WWMCCS CPE analyst can use the Batch Turnaround Time Model as a guide to identify causes of elongated batch turnaround time. The Model provides a means of conducting a structured analysis. An analyst uses the Model to: (1) determine what types of batch jobs are being elongated, (2) identify WWMCCS system processes where elongation is occurring, and (3) hypothesize causes for that elongation. Once a set of possible causes has been identified, the rest of the analysis is structured by the set.

By selecting a Model (see Figure II-1), an analyst constrains the scope of an investigation. For example, if remote batch jobs are being elongated, Pre-Processing delay points do not have to be examined during analysis, because remote batch jobs enter the system directly from a terminal without experiencing Pre-Processing delays.

The Guide directs an investigation to include only processes and phases that are exhibiting measurable elongation. For example, tests need not be conducted for GCOS System Scheduling delays if measured job elongation in that process is within reasonable limits.

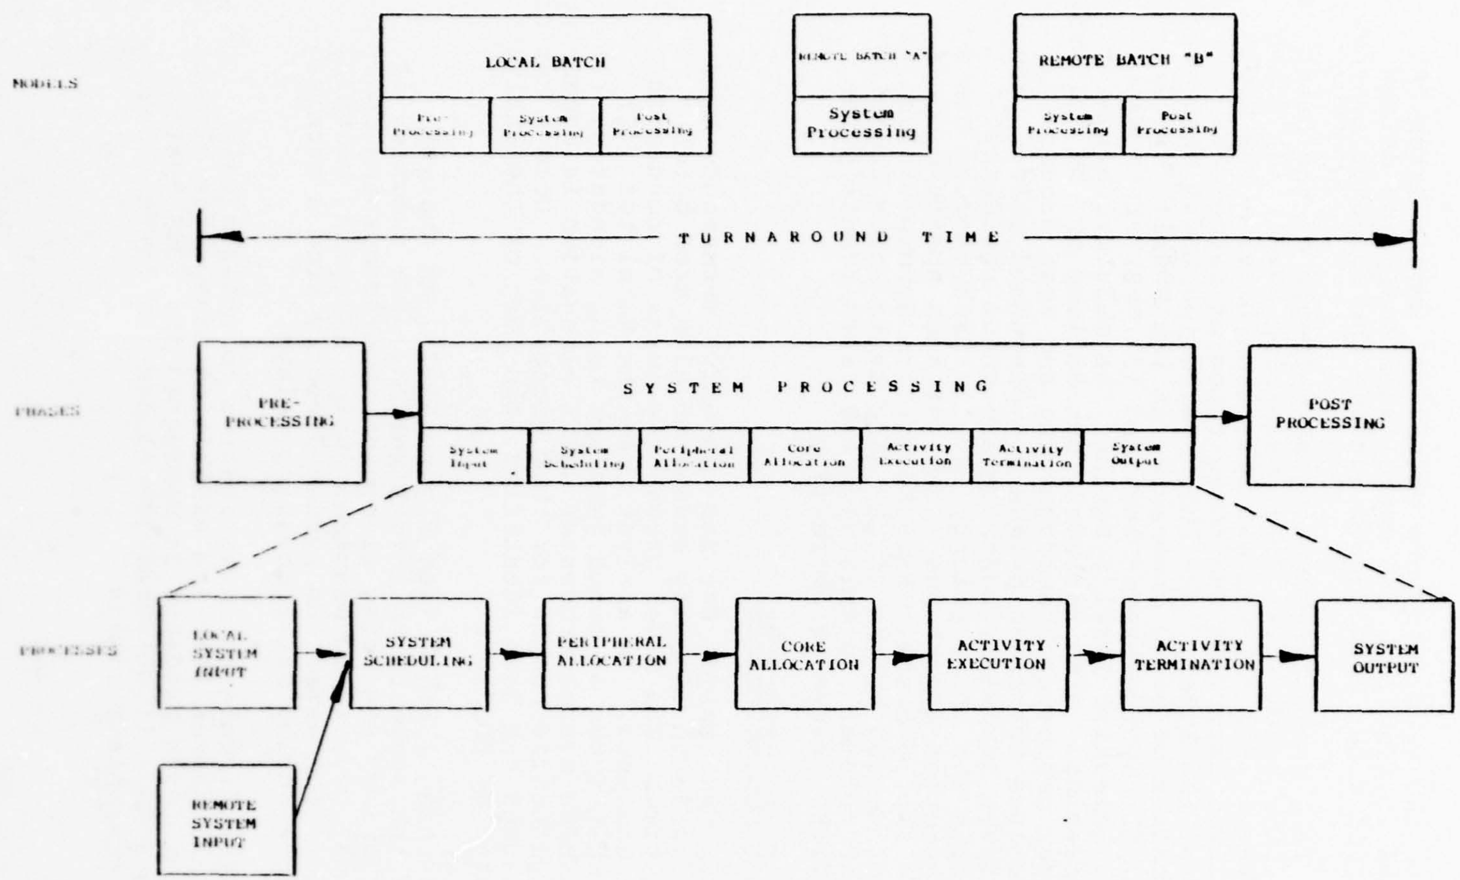

BATCH TURNAROUND TIME MODEL STRUCTURE

FIGURE II-1

In selecting particular hypotheses, an analyst chooses particular causes of process elongation that can be tested to confirm or deny that they are the source(s) of the measured job delay. Once confirmed, the source(s) of the delay can be removed or their effect diminished by any of several techniques.

C. BATCH TURNAROUND TIME MODELS

Figure II-2 portrays the three models by which an analyst begins use of the Model structure: (1) Local Batch Model, (2) Remote Batch "A" Model, and (3) Remote Batch "B" Model.

1. Local Batch Model

The Local Batch Model defines the components of batch turnaround time for batch jobs that are locally submitted and processed.

Three phases of batch turnaround time are described by the Local Batch Model: (1) Pre-Processing, (2) System Processing, and (3) Post Processing. The processes of each of these phases are defined in Section II.E. Lower levels of the Turnaround Time Model elaborate on these three phases. Procedures direct an analyst to examine the processes of all three phases to determine which takes the largest amount of batch turnaround time.

2. Remote Batch "A" Model

The Remote Batch "A" Model defines the components of batch turnaround time for batch jobs that (1) are spawned through the GCOS Time Sharing System or remote batch terminals and (2) produce output that does not need Post Processing (e.g., removing from printer and sorting into output bins) before the user can access it. The GCOS system processes are the only ones investigated when the analysis is guided by this model.

3. Remote Batch "B" Model

The Remote Batch "B" Model defines the components of batch turnaround time for batch jobs that are (1) spawned through the GCOS Time Sharing System or from remote batch terminals and (2) produce output that needs Post Processing before the user can access it. GCOS system processes and Post Processing processes are investigated when the analysis is guided by this model.

MODEL	LOCAL BATCH MODEL		
PHASE	PRE-PROCESSING	SYSTEM PROCESSING	POST PROCESSING

MODEL	REMOTE BATCH "A" MODEL
PHASE	SYSTEM PROCESSING

MODEL	REMOTE BATCH "B" MODEL	
PHASE	SYSTEM PROCESSING	POST PROCESSING

BATCH TURNAROUND TIME MODELS

FIGURE II-2

D. BATCH TURNAROUND TIME PHASES

The three batch turnaround time phases have already been defined: Pre-Processing, System Processing, and Post Processing. These terms will be further clarified in Section II.E. Using the Batch Turnaround Time Analysis System and manual logs (see Section III.B) the analyst determines what part of turnaround time is spent in each phase and selects for further investigation the phase with the most turnaround time.

E. BATCH TURNAROUND TIME PROCESSES

1. Process Definitions

Each phase is divided into processes, as shown in Figure II-3. The Pre-Processing and Post Processing processes vary from site to site and must be defined by the analyst. The System Processing processes are relatively constant, and the Guide recognizes seven; (1) System Input, (2) System Scheduling, (3) Peripheral Allocation, (4) Core Allocation, (5) Activity Execution, (6) Job Termination, and (7) System Output (SYSOUT).

The Pre-Processing and Post Processing processes shown in Figure II-3 are only samples. They are discussed for the benefit of the analyst who must define his own site's processes. At the hypothetical site whose Pre-Processing and Post Processing processes are represented in Figure II-3, Local Batch jobs are handled as cards or tape over a counter. A receipt is issued and the jobs wait on a table for a period of time. This receipting and wait is termed the "Job Receipt Table" process in the figure. The jobs then are transferred to a table in the machine room where they wait again. This transfer and wait is termed the "Machine Room Table" process in the figure. The "File Library" process starts when the tape (or file) librarian looks at the job and begins to assemble any tapes on other input needed by the job. When a job is passed with its associated tapes to a final table near the console, the "File Library" process ends and the "Console Queue" process begins. The "Console Queue" process ends when the job's card deck or IMCV tape is read into the system.

Post Processing processes are defined in a similar manner. A Local Batch or Remote Batch "B" job enters the "Printer Output Bin" process when its output has been printed. It waits in this process until it has been removed from the printer, at which time it enters the "Decollation"

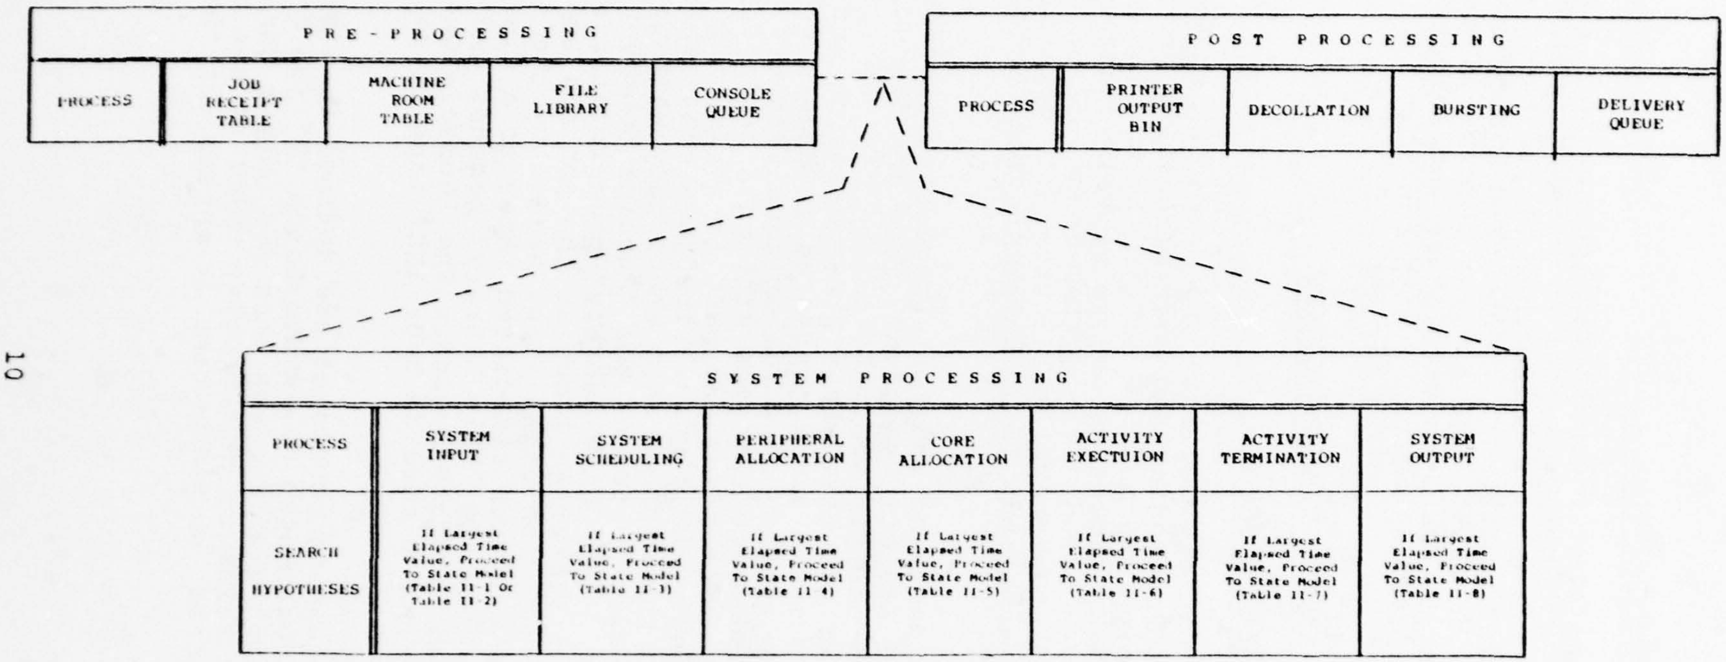

PROCESS SELECTION LEVEL

FIGURE II-3

process. Here jobs wait for a period of time and then the output is decollated. At this point they are transferred to a different room and enter the "Bursting" process. After the output is bursted, the jobs are transferred to another table, from which they are sorted into bins. This last process is called the "Delivery Queue" process. As with Pre-Processing, these processes are samples; the analyst must define the specific processes jobs experience at his site.

2. Process Selection

The analyst investigates the elapsed time of all the processes of the phase selected above in Section II.D. The Batch Turnaround Time Analysis System (see Section III.B) is used to report these times. The analyst selects the process associated with the most elapsed time for further investigation below.

F. BATCH TURNAROUND TIME ELONGATION HYPOTHESES

Each process has one or more "elongation hypotheses." These hypotheses are statements indicating a possible reason why a process might take longer than site management would wish. The object of the selection of a model, phase, and process is to narrow the possible causes of long turnaround time to a small list of elongation hypotheses. The hypotheses are then tested using one or more of the tests contained in the remainder of this volume. If one of the hypotheses proves true, instructions for reducing its effect on turnaround time are given at the end of the associated test. No sample elongation hypotheses for the Pre-Processing and Post Processing phases are given. Since the processes of these phases and possible causes of their elongation vary from site to site, and the causes for delays are easily observable, the elongation hypotheses in these areas are left to the analyst to formulate and test.

The System Processing processes are complicated enough that, for the purposes of listing and clarifying elongation hypotheses, each process has been divided into "states." The states and the elongation hypotheses for each state are given in Table II-1 through Table II-8. System Input has two tables: Table II-1 would be used if the Local Batch Model had been selected, and Table II-2 would be used if one of the remote batch models had been selected. Each table breaks its process into several states. Objectives, queue

STATE	INITIALIZE	BEGIN	READ	START	NXTBF	.MGNAT MODULE	.MSCAN MODULE
OBJECTIVE/ CHECK POINT	1. .MGEIN Setup 2. SSA Cell Setup 3. Input Media Determination	1. Initialize I/O for Input Device 2. Initialize I/O for *J and J* Files 3. Call Read Module	1. Initiate Device Read 2. (Disk): Get Next Link for Mass Store Link	1. Process First Buffer 2. Check for Presence of $SNUMB Card	Process Buffers After First One: 1. Obtain IAS Space (*J & J* Files) 2. Buffer to Device 3. Check Dumper (.CRDMP) 4. Make .CRRQQ Entry	Check Validity of $IDENT Control Card	Classify Jobs by: 1. Resource Requirements 2. Activity Types 3. Class/ Priority 4. $LIMITS (CPU Time & Core Size)
ELONGATION HYPOTHESES	Insufficient Memory For GEIN	Input Device Errors	1. Long Seeks on *J File 2. Long Seeks on J* File 3. Long Queues to Devices and/or Channels Containing J* and/or *J Files 4. Device Errors on Disk Devices containing J* and/or *J Files 5. IOS* Code Delays 6. Insufficient Disk Space: J* and/or *J Files 7. Input Device Errors	Insufficient Disk Space: J* and/or *J Files	1. Long Seeks on *J File 2. Long Seeks on J* File 3. Long Queues to Devices and/or Channels Containing J* and/or *J Files 4. Device Errors on Disk Devices containing J* and/or *J Files 5. IOS* Code Delays 6. Insufficient Disk Space: J* and/or *J Files 7. .CRDMP Set Often	Inefficient Code	Inefficient Code

*I/O Supervisor

SYSTEM INPUT PROCESS ELONGATION HYPOTHESES

TABLE II-1

STATE	NEW JOB REQUEST INFORMATION SETUP	RQ TABLE RESIDENCE	SEARCH FOR NEW JOB REQUESTS	T1 QUEUE RESIDENCE	JOB DATA WRITE TO DISK	BUILD .CRRQQ ENTRY	.MONAT MODULE	.MSCAN MODULE
OBJECTIVE/ QUEUE POINT	1. Determine if Space Available to Hold RQ Entry 2. Initialize Remote Batch Input Processing	Job Request Waiting in RQ Table	1. Check For T1 Entry 2. Request IAS Space	Job Held in T1 Queue for More Data Blocks	1. Initiate Write to Disk 2. Normal Disk Write Return 3. Maintain Disk Space 4. Error Report Processing	1. Build Entry 2. Place in .CRRQQ	Check Validity of $IDENT Control Card	Classify Jobs by: 1. Resource Requirements 2. Activity Types 3. Class/ Priority 4. $LIMITS (CPU Time and Core Size)
ELONGATION HYPOTHESES	1. .CRGSP Gate Shut Often 2. Insufficient Memory for RQEIN	Long RQ Table Entry Residence Time	1. Long Seeks to *J and/or J* Files 2. Long Queues to Devices and/or Channels Containing J* and/or *J Files 3. Device Errors on Disk Devices Containing J* and/or *J Files 4. Disk IOS* Code Delay 5. Insufficient Disk Space: J* and/or *J Files	Long T1 Queue Entry Residence Time	1. Long Seeks to *J and/or J* Files 2. Long Queues to Devices and/or Channels Containing J* and/or *J Files 3. Device Errors on Disk Devices Containing J* and/or *J Files 4. Disk IOS* Code Delay 5. Insufficient Disk Space: J* and/or *J Files 6. .CRRQQ Gate Shut Often	.CRRQQ Gate Shut Often	Inefficient Code	Inefficient Code

*I/O Supervisor

REMOTE SYSTEM INPUT PROCESS ELONGATION HYPOTHESES

TABLE II-2

STATE	.CRRQ QUEUE RESIDENCE	CHECK .CRRQ	NEW JOB ENTERING	CATALOG JOB	CATALOG RESIDENCE	UPDATE CATALOG	ATTEMPT JOB START
OBJECTIVES/ QUEUE POINT	Job Entry Waiting in System Scheduler Input Queue	SCAN .CRRQ for Job Entries	1. Check Input Traffic Rate 2. If Number of Jobs < MIN, Start Job (Bypass Scheduling)	1. Link Job to TCB Chain 2. Initiate Catalog I/O	Job Entry Waiting on System Scheduler Catalog	Determine if Task Entry is Available in the Catalog Block	1. Get TCB Entry 2. Assign Buffer 3. Assign Program Number 4. Get SYSOUT Space 5. Copy J* to SYSOUT 6. Make .CRJOB Entry
ELONGATION HYPOTHESES	.CRRQ Entry Residence Time Long	Insufficient Core Available for Task Control Block Chain	1. MAX Parameter Setting Too Low 2. CLASS MAX Parameter Setting Too Low 3. MIN Parameter Setting Too High	1. Long Seeks to System Scheduler Catalog(s) 2. Long Queues to Devices and/or Channels Containing System Scheduler Catalog(s) 3. Device Errors on Devices Containing System Scheduler	Job Catalog Entry Residence Time	1. Long Seeks to System Scheduler Catalog(s) 2. Long Queues to Devices and/or Channels Containing System Scheduler Catalogs 3. Device Errors on Devices Containing System Scheduler Catalogs 4. Disk IOS* Code Delays	1. Long Seeks to Devices Containing SYSOUT Files 2. Long Queues to Devices and/or Channels Containing SYSOUT Files 3. Device Errors on Devices Containing SYSOUT Files 4. Disk IOS* Code Delays 5. Insufficient Disk Space for SYSOUT 6. Insufficient Buffer Space 7. Insufficient Program Numbers

*I/O Supervisor

SYSTEM SCHEDULING PROCESS ELONGATION HYPOTHESES

TABLE II-3

STATE	.CRJOB QUEUE RESIDENCE	PROCESS .CRJOB ENTRIES	NEW JOB	NEW ACTIVITY	PERIPHERAL ALLOCATION	PALC/CALC DAMPER PROCESSING	WRITE SSA TO *J FILE
OBJECTIVE/ QUEUE POINT	Job Entry Waiting in Peripheral Allocator Input Queue	SCAN .CRJOB for Job Entries	Scan JCL Cards	Set up Control Stack for Needed Resources	1. Attempt Peripheral Allocation 2. If Allocation not Complete, Release Peripherals, and *J & Control Stack	Scan to Determine if Activity is to be Denial Access to Core Allocator	1. Write SSA & Device Data to *J. 2. Make Entry in .CRPOQ
ELONGATION HYPOTHESES	.CRJOB Entry Residence Time Long			Urgency Code for Jobs Too Low	1. Tapes Requested not Available 2. Disk Space Requested not Available 3. Long Seeks to PALC Scratch File 4. Device Errors on Device Containing PALC Scratch File 5. Long Queues to Devices and/or Channels Containing PALC Scratch File 6. Disk IOS* Code Delay	1. Tapes Requested not Available 2. Disk Space Requested not Available 3. PALC/CALC Damper Set Too Often	1. Long Seeks to *J File 2. Device Errors on Device Containing *J File 3. Long Queues to Devices and/or Channels Containing *J File 4. Disk IOS* Code Delay

*I/O Supervisor

PERIPHERAL ALLOCATION PROCESS ELONGATION HYPOTHESES

TABLE II-4

STATE	.CRPOQ QUEUE RESIDENCE	PROCESS .CRPOQ ENTRIES	NEW JOB	PRIORITY ADJUSTMENT	CORE ALLOCATION	START JOB
OBJECTIVE/ QUEUE POINT	Activity Entry Waiting in Core Allocator Input Queue	1. Scan .CRPOQ for Activity Entries 2. Call Appropriate Module (i.e., Start New Job/Activity or Unswap a Job)	1. Update CALC Tables and Slave Area Definitions 2. Update and Relink Core Urgency Queue for all Jobs	Modify Selected Activity's Urgency Code	1. Try to Fit Program Without Adjustment 2. Initiate Compaction Based on Urgency 3. Initiate Swap Based on Urgency	1. Build Bootstrap to Bring Slave to Core 2. Zero Slave Core 3. Enable Bootstrap by Dispatcher Entry
ELONGATION HYPOTHESES	.CRPOQ Entry Residence Time Too Long			Operator Changes to Urgency Codes Too Frequent or Too Great	1. Insufficient Memory 2. Urgency Codes of Other Jobs in System Too High 3. "MXCNT" Parameter Too Low	Time to Zero Slave Core Too Long

CORE ALLOCATION PROCESS ELONGATION HYPOTHESES

TABLE II-5

STATE	CPU PROCESS	GECALL PROCESS	SSA PROCESS	MME PROCESS	FAULT PROCESS	NON-DISPATCH STATE
OBJECTIVE/ QUEUE POINT	Activity is Dispatched and Executing	Activity has Issued a GECALL for Code	Activity has Requested SSA Code	MME Process, Other Than Fault, is Performed for Activity	Fault Process is Performed for Activity	Activity in Dispatcher Queue (or in non-Dispatch Candidate State), but Not in Control of the Processor
ELONGATION HYPOTHESES	Inefficient Program Code	1. Long Seeks to GECALL Library 2. Device Errors on Device that Contains GECALL Library 3. Long Queues to Devices and/or Channels Containing GECALL Library 4. Disk IOS* Code Delays 5. Inefficient Called Code	1. Long Seeks to SSA Library 2. Device Errors on Device that Contains SSA Library 3. Long Queues to Devices and/or Channels Containing SSA Library 4. Disk IOS* Code Delays 5. Inefficient SSA Module Code 6. Other Modules Had to be Pushed Down	1. Long Seeks to SSA Library 2. Device Errors on Device that Contains SSA Library 3. Long Queues to Devices and/or Channels Containing SSA Library 4. Disk IOS* Code Delays 5. Inefficient SSA Module Code 6. Other Modules Had to be Pushed Down	1. Long Seeks to SSA Library 2. Device Errors on Device that Contains SSA Library 3. Long Queues to Devices and/or Channels Containing SSA Library 4. Disk IOS* Code Delays 5. Inefficient SSA Module Code 6. Other Modules Had to be Pushed Down	1. Waiting for I/O Because of Long Seeks to Files 2. Waiting for I/O Because of Queues to Devices and/or Channels 3. Delays due to Tape and/or Disk IOS* Code 4. Too Little Memory 5. Deadlock or Single-Thread Situations In Unique Site Systems and Patches 6. Urgency Codes of Other Activities are Causing Elongation of Selected Jobs 7. Few (or no) Activities is Cause for High Value

*I/O Supervisor

ACTIVITY EXECUTION PROCESS ELONGATION HYPOTHESES

TABLE II-6

STATE	SWAP/COMPACTION STATE	ACTIVITY I/O PROCESS: SERVICE	ACTIVITY I/O PROCESS: TAPE	ACTIVITY I/O PROCESS: DISK	ACTIVITY I/O PROCESS UNIT RECORD
OBJECTIVE/ QUEUE POINT	Activity has been Swapped or is Undergoing a Swap/ Move Operation	I/O Processing Being Provided by GCOS; Logical Part of I/O Operation for Activity	Tape I/O Operation is Being Executed for Activity	IAS I/O Operation is Being Executed for Activity	Unit Record I/O Operation Being Executed for Activity
ELONGATION HYPOTHESES	1. Long Seeks to Swap File(s) 2. Activities Wasting Memory 3. Device Errors on Devices That Contain Swap File(s) 4. Queues to Devices and/ or Channels Containing Swap File(s) 5. Disk IOS* Code Delays 6. Urgency Codes of Some Jobs Too High	Delays on Tape, Disk, and/or Unit Record Done by Activities Due to Inefficient IOS* Code	1. Excessive Tape I/O Errors 2. Queues to Activity Tape Devices and/or Channels	1. Long Seeks to Activity Files (Including SYSOUT) 2. Excessive Disk I/O Errors 3. Queues to Activity Disk Devices and/or Channels 4. Insufficient SYSOUT Buffers	1. Excessive Unit Record Device Errors 2. Queues to Activity Unit Record Devices and/or Channels

*I/O Supervisor

ACTIVITY EXECUTION PROCESS ELONGATION HYPOTHESES

TABLE II-6 (Cont'd)

STATE	TERM INITIATION PROCESSING	TERM CONTINUATION PROCESSING	REMOTE LINE TERM PROCESSING	ACTIVITY COMPRESSION PROCESSING	ACTIVITY RELEASE & ACCOUNTING PROCESS	USER ACCOUNTING PROCESSING
OBJECTIVE/ QUEUE POINT	1. Set General Termination Type in .STATE Word 2. Determine Type of Activity Terminating 3. Determine Type of Termination	1. Determine if Compression to be Executed 2. Normal Wrapup Processing (e.g., ABORT, Util., Outstand I/O, etc.) 3. Call Next Module	1. If TSS Line, Reconnect to TSS 2. If not TSS Line, Disconnect it	1. Read *J to Get SSA Copy 2. Read CTL Cards for Next Activity; OK to Compress? 3. If so: Mark Program "In Execution" 4. Go to Location 32_8 of Slave	1. Clear Fault Vectors 2. Read *J; Get Activity Report 3. Update Accounting Data; User Modifications 4. Release Unwanted Peripherals 5. Inform CALC of Termination	Determined by User Code
ELONGATION HYPOTHESES		1. Long Seeks to System Module Files 2. Excessive Errors on Devices with System Module Files 3. Queues to Devices and/or Channels Containing Module Files		1. Long Seeks to *J File 2. Queues to Devices and/or Channels Containing *J File 3. Device Errors on Devices Containing *J File 4. Disk IOS* Code Delay	1. Long Seeks to *J File 2. Queues to Devices and/or Channels Containing *J File 3. Device Errors on Devices Containing *J File 4. Disk IOS* Code Delay 5. I/O Errors on Accounting File 6. Long Queue to Accounting File	1. Inefficient Code 2. If User I/O: a. Queues b. Long Seeks c. I/O Errors d. IOS* Code Delay

*I/O Supervisor

TERMINATION PROCESS ELONGATION HYPOTHESES

TABLE II-7

STATE	.CRSYQ QUEUE RESIDENCE	SYSOUT DISPERSAL INITIALIZATION	START SYSOUT PROCESSING	BUFFER INPUT PROCESSING	BUFFER OUTPUT PROCESSING
OBJECTIVE/ QUEUE POINT	Activity Entry Waiting in SYSOUT Input Queue	1. Scan .CRSYQ for Job Entries 2. Call SYSOUT Executive and Overlay Extensions	1. Initiate SYSOUT Device Subroutines 2. Load First Input Buffer 3. Initialize Banner Point	Process SYSOUT Input Buffer	Process SYSOUT Output Buffer
ELONGATION HYPOTHESES	.CRSYQ Entry Residence Time Long	1. Long Seeks to SSA Library 2. Queues to Devices and/or Channels Containing SSA Modules 3. Device Errors on Devices that Contain SSA Modules 4. Disk IOS* Code Delay	1. Long Seeks to SYSOUT Files 2. Queues to Devices and/ or Channels Containing SYSOUT Files 3. Device Errors on Devices that Contain SYSOUT Files 4. Disk IOS* Code Delay	1. Long Seeks to SYSOUT Files 2. Queues to Devices and/or Channels Containing SYSOUT Files 3. Device Errors on Devices that Contain SYSOUT Files 4. Disk IOS* Code Delay 5. Lack of Sufficient SSA Buffers	1. Pathway Utilization: Devices and/or Channels Containing SYSOUT Output Devices 2. Delays on IOS Service To SYSOUT Output Devices 3. .MPOUT Serving Other Devices

*I/O Supervisor

SYSTEM OUTPUT PROCESS ELONGATION HYPOTHESES

TABLE II-8

points, and elongation hypotheses are given for each state. Current instrumentation does not allow the breakdown of turnaround time by state, so that the elongation hypotheses for a particular process must be considered as a whole. Section III.G gives instructions for choosing elongation hypotheses to be tested.

1. Local System Input

Table II-1 defines seven states through which a local batch job must pass as it enters the system. The Model includes jobs entered via the card reader or IMCV tape. The Local System Input Process begins with the initialization of .MGEIN and terminates with the execution of the .MSCAN module optionally implemented at a particular site.

2. Remote System Input

Table II-2 defines eight states through which a remote batch (either type "A" or type "B") job must pass as it enters the system. Remote System Input begins with initialization and terminates with the execution of the optional .MSCAN module.

3. System Scheduling

Table II-3 defines seven states of the WWMCCS system scheduling process. The process begins with the entry of a scheduling request in the GCOS System Scheduler Input Queue (.CRRGQ) and terminates with the attempt to start a particular job by placing an entry in the GCOS Peripheral Allocator Input Queue (.CRJOB).

4. Peripheral Allocation

Table II-4 defines the seven states of the GCOS peripheral allocation process. The process begins with the entry of a job request in the GCOS Peripheral Allocator Input Queue (.CRJOB) and ends with the entry of a job request in the GCOS Core Allocator Input Queue (.CRPOQ).

5. Core Allocation

Table II-5 defines the six states of the GCOS core allocation process. The process begins with the entry of a job request in the GCOS Core Allocator Input Queue (.CRPOQ) and ends when the GCOS Core Allocator attempts to start a job by making an entry for it in the GCOS Dispatcher Queue. Core allocation is provided from the same Model process for new jobs and for executing activities.

6. Activity Execution

Table II-6 defines the eleven states that an activity can assume during execution. Unlike the mainly sequential states defined for the other system processes, the activity execution states are entered in a random fashion by the executing batch programs. Note from Table II-6 that some states involve code execution. Others involve: (1) processes executed by GCOS for the job, (2) I/O operations performed for the job, or (3) non-dispatch conditions such as waiting in queues or swapping.

7. Activity Termination Processing

Table II-7 defines the six states through which an activity or job passes as it is terminating execution. The process begins when an activity requests termination and ends when an entry is placed in the GCOS System Output Queue (.CRSYQ) by the GCOS Activity Terminator.

8. System Output Processing

Table II-8 defines the five states through which local batch and remote batch jobs pass as they receive system output service. The process begins with an entry placed in the GCOS System Output Queue (.CRSYQ) and ends with the completion of the printing or punching operation.

G. GENERAL HYPOTHESIS TEST TECHNIQUES

The truth or falsehood of the elongation hypotheses given in Tables II-1 through II-8 is established by using the test procedures in Sections IV through XII of the volume. Instructions for choosing which tests to apply first are given in Section III. Not all hypotheses defined by the Batch Turnaround Time Model are subject to confirmation because of limited instrumentation. In many cases, instrumentation to test these hypotheses has not been sought due to the overhead involved and the small likelihood that the problem could be removed with significant shortening of turnaround time.

1. Format of the Tests

Each section of the rest of this volume contains an individual test description. Tests are comprised of procedures and steps. A procedure contains a series of analytical steps that are directed toward an analytic objective. For example, one of the procedures under the Seek Elongation Test documented in Section IV has the following objective: "Identify High-Use Extents" on the disk units in the system being monitored. The steps of each procedure involve: (1) examining specific reports to obtain performance data, (2) entering metric values on a form, (3) calculating a ratio or percentage from the entered values, and (4) making certain decisions (and subsequent recommendations) from the calculated values.

One test (i.e., CPU Execution Characteristics) uses a hardware monitor to gather data because no other technique can be applied. Some tests use the outputs from more than one CPE tool in their procedures. When this is the case, it is the CPE analyst's responsibility to run the data collection software for all monitors at the same time.

2. Representative Values of Frequency Distribution

Some test procedures require an analyst to pick a "Representative Value" to describe a frequency distribution. Two Memory Utilization Monitor Reports (MUM) are involved: the Total Elapsed Time an Activity Was in Memory, and the Elapsed Wait Time For Memory Requests in 1/10 Second. The following paragraphs give guidelines for choosing Representative Values.

a. Introduction. A Representative Value is a single number which attempts to describe either (1) the amount of memory wait time or (2) the amount of memory residence time experienced by the "typical" or "most important" job. Each time the two MUM reports named above are produced, one Representative Value is derived from each report. The memory wait time Representative Value is divided by the memory residence time Representative Value, deriving a ratio of memory wait time to memory residence time. If the ratio is large (i.e., jobs wait for memory for long periods compared to the length of time they use the memory requested), memory is probably an important constraint to performance. If the ratio is small, memory is probably not a constraint.

Frequently the data will be such that every possible Representative Value will yield the same type of ratio--always large or always small. This is the case when memory waits are very small and memory residence times very large. In these cases, the analyst does not need to investigate further before choosing Representative Values. He may simply note that the ratio will always be large (small). The analyst should insure, however, that it is not a case of a large number of jobs with short wait times obscuring the fact that the jobs with bad turnaround have long memory wait times.

Any single value is an over-simplification of a full frequency distribution; however, the fact that the Representative Values are used only as a ratio helps make their use valid. This is especially true when the two distributions have similar shapes (e.g., both skewed to the right--see below for descriptions of typical shapes). In this case, systematic errors in picking Representative Values will tend to cancel when the ratio is formed.

The type of turnaround problem being analyzed governs whether the Representative Value should be for the whole distribution (i.e., "typical") or for a subset of the distribution (i.e., "most important"). If all jobs are experiencing poor turnaround, a "typical" value should be used. If only a few jobs are experiencing poor turnaround, only the values experienced by these jobs should be used to pick the Representative Value. [This can be accomplished by noting which SNUMB's had poor turnaround and reducing the data again using only those SNUMB's (see the Batch Turnaround Time Analysis System User's Guide supplied with this volume).]

b. Types of Distributions. Figure II-4 shows a hypothetical "Total Elapsed Time an Activity Was in Memory" report. Variations of this report will be used to illustrate types of distributions. The report contains a row for each amount of time waited, from 1.0 seconds to 2.7 seconds. The number of activities which spent a certain amount of time in memory are reported in the row whose time range (in the "Tenths Second" column) includes that amount of time. For example, 1,670 activities spent 1.7 seconds in memory. The percentage of activities

```
DISTRIBUTION COLLECTED ON SYSTEM         AT 14141124 ON 77-02-13

              THE TOTAL ELAPSED TIME AN ACTIVITY WAS IN MEMORY

INDIV.   CUMUL.   CUMUL.   INDIV.  TENTHS              PERCENT PROBABILITY OF OCCURRENCE                 REPORT
NUMBER   NUMBER   PROB.    PROB.   SECOND   00   05   10   15   20   25   30   35   40   45   50           71
                                            I....I....I....I....I....I....I....I....I....I....I
     1        1   0.000    0.000   10-  10  I
    13       14   0.002    0.001   11-  11  I
    47       61   0.007    0.005   12-  12  IX
   164      225   0.024    0.018   13-  13  IXX
   411      636   0.069    0.044   14-  14  IXXXX
   883     1519   0.164    0.095   15-  15  IXXXXXXXXXX
  1353     2872   0.310    0.146   16-  16  IXXXXXXXXXXXXXXX
  1670     4542   0.490    0.180   17-  17  IXXXXXXXXXXXXXXXXXXX
  1663     6205   0.669    0.179   18-  18  IXXXXXXXXXXXXXXXXXX
  1353     7558   0.815    0.146   19-  19  IXXXXXXXXXXXXXXX
   826     8384   0.904    0.089   20-  20  IXXXXXXXXX
   483     8867   0.956    0.052   21-  21  IXXXXX
   227     9094   0.980    0.024   22-  22  IXX
   117     9211   0.993    0.013   23-  23  IX
    45     9256   0.998    0.005   24-  24  I
    15     9271   0.999    0.002   25-  25  I
     4     9275   1.000    0.000   26-  26  I
     1     9276   1.000    0.000   27-  27  I

9276 ENTRIES TOTAL      AVERAGE =    17.62193     VARIANCE =      4.856     STANDARD DEVIATION =      2.203
```

SAMPLE SYMMETRIC DISTRIBUTION

FIGURE II-4

falling in a certain row is reported in the "Individual Probability" column. The sum of these percentages for all rows above (and including) a certain row is reported in the "Cumulative Probability" column for that row. For example, the 1,670 activities in the 1.7 second row comprised 18.0% of all activities. These activities, combined with all activities with elapsed time values less than 1.7 seconds, comprised 49.0% of all activities. The "Elapsed Wait Time for Memory Requests in 1/10 Second" report has the same format as Figure II-4. See the MUM User's Guide (CCTC.613.UG) for more detail.

(1) Symmetric Distribution Closely Clustered Around a Single Point. The Representative Value for the distribution shown in Figure II-4 is easy to pick. The values are closely clustered around the "Average": 17.6 tenths of a second, or 1.76 seconds. The shape of the distribution is symmetric--about the same number of activities had values over 1.76 seconds as under 1.76 seconds. The absence of a second line under the "Entries Total" line indicates that no activities stayed in memory longer than 2.7 seconds. When a distribution resembles this example, use the "Average" printed at the bottom as the Representative Value.

(2) Skewed Distribution. Figure II-5 shows another distribution. This distribution is "skewed" (i.e., not symmetric), because most of the activities spent around 0.1 to 0.7 seconds in memory, while some spent as much as four or five. Care should be taken when picking a Representative Value from this distribution. If the analyst wants to emphasize the "typical" activity, which stayed in memory 0.3 seconds of less, he could pick the "median"[1] of the

---

[1]The median is the value which evenly divides the activities in the distribution--half spent less time in memory, and half spent greater time in memory. In Figure II-5, the median is about 0.29 seconds. The median can be estimated from these reports by descending down the "Cumulative Probability" column until the value first exceeds 0.50. The median falls within the time range of this row. In Figure II-5, the second row has a Cumulative Probability

DISTRIBUTION COLLECTED ON SYSTEM AT 14:41:26 ON 77-02-13

THE TOTAL ELAPSED TIME AN ACTIVITY WAS IN MEMORY

REPORT 21

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PROB.	INDIV. PROB.	TENTHS SECOND	PERCENT PROBABILITY OF OCCURRENCE (00 05 10 15 20 25 30 35 40 45 50)
					I....I....I....I....I....I....I....I....I....I....I
36628	36628	0.307	0.307	0- 1	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
31525	68153	0.575	0.268	2- 3	IXXXXXXXXXXXXXXXXXXXXXXXXXXX
18677	86830	0.732	0.158	4- 5	IXXXXXXXXXXXXXXXX
11540	98370	0.830	0.097	6- 7	IXXXXXXXXXX
7393	105763	0.892	0.062	8- 9	IXXXXXX
4724	110487	0.932	0.040	10- 11	IXXXX
2981	113468	0.957	0.025	12- 13	IXXX
1861	115329	0.973	0.016	14- 15	IXX
1198	116507	0.983	0.010	16- 17	IX
756	117261	0.989	0.006	18- 19	IX
479	117740	0.993	0.004	20- 21	I
297	118037	0.996	0.003	22- 23	I
180	118217	0.997	0.002	24- 25	I
122	118339	0.998	0.001	26- 27	I
76	118415	0.999	0.001	28- 29	I
56	118471	0.999	0.000	30- 31	I
28	118499	0.999	0.000	32- 33	I
20	118519	1.000	0.000	34- 35	I
11	118530	1.000	0.000	36- 37	I
4	118534	1.000	0.000	38- 39	I
4	118539	1.000	0.000	40- 41	I
5	118543	1.000	0.000	42- 43	I
7	118550	1.000	0.000	44- 45	I
1	118551	1.000	0.000	46- 47	I
1	118552	1.000	0.000	48- 49	I
2	118554	1.000	0.000	50- 51	I
3	118557	1.000	0.000	52- 53	I
1	118558	1.000	0.000	54- 55	I
1	118559	1.000	0.000	56- 57	I

118559 ENTRIES TOTAL AVERAGE = 4.24719 VARIANCE = 18.519 STANDARD DEVIATION = 4.303

SAMPLE SKEWED DISTRIBUTION

FIGURE II-5

distribution. Otherwise, he may pick the average printed at the bottom of the report. In Figure II-5, the average is 0.42 seconds. When the distribution is skewed, the average is always closer than the median to the few "outlying" values.

The analyst may believe the larger time values (e.g., four to five seconds) to be more representative of the jobs that are experiencing long turnaround times. In this case, it is best to obtain the reports again with the jobs experiencing acceptable total turnaround times eliminated from the data.

If both MUM reports have the same type of skewed distribution, the ratio of the two Representative Values will probably be the same whether the medians or averages are used; however, the analyst should not use the median of one distribution and the average of the other if the distributions are similar.

---

value of 0.575, so the median is between 0.2 and 0.3 seconds (more specifically, between 0.15 and 0.35 seconds because 0.15 is the dividing line between 0.1 and 0.2 and 0.35 is the dividing line between 0.3 and 0.4). The actual value can be estimated by interpolation: Let A be the row in which the median occurs (the second row in Figure II-5). Let B be the "Cumulative Probability" value for the row immediately above A. Let C be the "Cumulative Probability" value for row A. Let D be the start of the range reported in row A, and E the end of the range. (The activities in row A spent more than D seconds, but less than E seconds in memory.) Then an estimate of the median would be

$$D + \frac{0.500 - B}{C - B} (E - D)$$

In the Figure II-5, the estimate would be

$$0.15 + \frac{0.500 - 0.307}{0.575 - 0.307} * (0.35 - 0.15)$$

$$= 0.15 + \frac{0.193}{0.268} * 0.20$$

$$\approx 0.29 \text{ seconds.}$$

(3) Skewed Distribution with Outliers. Figure II-6 shows a distribution which is more skewed than the one in Figure II-5. Note the bottom line, which states that 33 activities were in memory longer than the longest time range on the plot. These activities (the "outliers") are over 15% (33/207) of the total number of activities, and their average time in memory is 1,508 seconds--over four times the last plotted value of 369.9 seconds. The wide variation of time in memory may indicate that two or more kinds of jobs or types of time periods are being reported. The analyst should investigate the individual activities to determine if either the long-running activities or the short-running activities can be eliminated from the distribution (either type of activity may not experience poor turnaround). The analyst may find that the two (or more) types of activities are multiple activities of the same job, such as a short compilation followed by a long execution. In this case, the average should be used as the representative value, since it is a better index of the total time the job (i.e., all activities) spent in memory.

(4) Poor Reduction Parameters. Distributions such as Figure II-7 and Figure II-8 need to be reduced again with an expanded plot (a larger time interval on each row). Figure II-7 shows only one or two observations at the most in each row, with most of the observations off the end of the plot. This gives the distribution the appearance of a multi-spiked distribution (a distribution with widely separated clusters of activities) when it may be only that there were few activities run during the data collection period and/or the time interval per row is too small. Reducing the data with longer time intervals (i.e., 0-10 seconds, 11-20 seconds, etc.) per row will merge the artificial "clusters" of only one or two activities, and will bring most of the "out of range" data onto the plot. Figure II-8 shows only 20% of the activities "out of range", but the average for these activities is so low that a doubling of the time interval per row might almost eliminate the "out of range" category. Once the data in Figure II-7 and Figure II-8 are reduced again, the rules for Figure II-6 will probably apply.

THIS PAGE IS BEST QUALITY PRACTICABLE
FROM COPY FURNISHED TO DDC

DISTRIBUTION COLLECTED ON SYSTEM AT 14141124 ON 77-02-13

THE TOTAL ELAPSED TIME AN ACTIVITY WAS IN MEMORY

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PROB.	INDIV. PROB.	TENTHS SECOND	PERCENT PROBABILITY OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50 I....I....I....I....I....I....I....I....I....I....I	RECORD 21
50	50	0.242	0.242	25- 99	IXXXXXXXXXXXXXXXXXXXXXXXX	
19	69	0.333	0.092	100- 199	IXXXXXXXXX	
16	85	0.411	0.077	200- 299	IXXXXXXXX	
7	92	0.444	0.034	300- 399	IXXX	
15	107	0.517	0.072	400- 499	IXXXXXXX	
9	116	0.560	0.043	500- 599	IXXXX	
9	125	0.604	0.043	600- 699	IXXXX	
9	134	0.647	0.043	700- 799	IXXXX	
4	138	0.667	0.019	800- 899	IXX	
1	139	0.671	0.005	900- 999	I	
3	142	0.686	0.014	1000-1099	IX	
0	142	0.686	0.	1100-1199	I	
4	146	0.705	0.019	1200-1299	IXX	
1	147	0.710	0.005	1300-1399	I	
1	148	0.715	0.005	1400-1499	I	
3	151	0.729	0.014	1500-1599	IX	
3	154	0.744	0.014	1600-1699	IX	
1	155	0.749	0.005	1700-1799	I	
2	157	0.758	0.010	1800-1899	IX	
3	160	0.773	0.014	1900-1999	IX	
1	161	0.778	0.005	2000-2099	I	
2	163	0.787	0.010	2100-2199	IX	
1	164	0.792	0.005	2200-2299	I	
1	165	0.797	0.005	2300-2399	I	
0	165	0.797	0.	2400-2499	I	
1	166	0.802	0.005	2500-2599	I	
0	166	0.802	0.	2600-2699	I	
1	167	0.807	0.005	2700-2799	I	
2	169	0.816	0.010	2800-2899	IX	
1	170	0.821	0.005	2900-2999	I	
0	170	0.821	0.	3000-3099	I	
0	170	0.821	0.	3100-3199	I	
0	170	0.821	0.	3200-3299	I	
0	170	0.821	0.	3300-3399	I	
1	171	0.826	0.005	3400-3499	I	
2	173	0.836	0.010	3500-3599	IX	
1	174	0.841	0.005	3600-3699	I	

207 ENTRIES TOTAL AVERAGE = 2960.04971 VARIANCE =43484971.560 STANDARD DEVIATION = 6594.706
33 OUT OF RANGE AVERAGE FOR THESE = 15380.81974

SAMPLE SKEWED DISTRIBUTION
WITH OUTLIERS

FIGURE II-6

DISTRIBUTION COLLECTED ON SYSTEM AT 14141124 ON 77-02-11

THE TOTAL ELAPSED TIME AN ACTIVITY WAS IN MEMORY

REPORT 21

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PROB.	INDIV. PROB.	CENTISECOND	PERCENT PROBABILITY OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50
					I....I....I....I....I....I....I....I....I....I....I
2	2	0.051	0.051	6- 7	IXXXXX
1	3	0.077	0.026	8- 9	IXXX
0	3	0.077	0.	10- 11	I
0	3	0.077	0.	12- 13	I
0	3	0.077	0.	14- 15	I
0	3	0.077	0.	16- 17	I
1	4	0.103	0.026	18- 19	IXXX
0	4	0.103	0.	20- 21	I
1	5	0.128	0.026	22- 23	IXXX
1	6	0.154	0.026	24- 25	IXXX
0	6	0.154	0.	26- 27	I
1	7	0.179	0.026	28- 29	IXXX
2	9	0.231	0.051	30- 31	IXXXXX
0	9	0.231	0.	32- 33	I
1	10	0.256	0.026	34- 35	IXXX
0	10	0.256	0.	36- 37	I
0	10	0.256	0.	38- 39	I
0	10	0.256	0.	40- 41	I
0	10	0.256	0.	42- 43	I
0	10	0.256	0.	44- 45	I
1	11	0.282	0.026	46- 47	IXXX
0	11	0.282	0.	48- 49	I
0	11	0.282	0.	50- 51	I
0	11	0.282	0.	52- 53	I
0	11	0.282	0.	54- 55	I
0	11	0.282	0.	56- 57	I
0	11	0.282	0.	58- 59	I
0	11	0.282	0.	60- 61	I
0	11	0.282	0.	62- 63	I
1	12	0.308	0.026	64- 65	IXXX

39 ENTRIES TOTAL AVERAGE = 1317.28206 VARIANCE = 3274432.000 STANDARD DEVIATION = 1810.644

27 OUT OF RANGE AVERAGE FOR THESE = 1890.70370

SAMPLE DISTRIBUTION WITH
POOR REDUCTION PARAMETERS

FIGURE II-7

DISTRIBUTION COLLECTED ON SYSTEM [illegible]

THE TOTAL ELAPSED TIME AN ACTIVITY WAS IN MEMORY

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PROB.	INDIV. PROB.	CENTIS SECOND	PERCENT PROBABILITY OF OCCURRENCE (0 05 10 15 20 25 30 35 40 45 50)	REPORT
120	120	0.047	0.047	0- 4	IXXXXX	
214	334	0.131	0.084	5- 9	IXXXXXXXXX	
113	447	0.176	0.044	10- 14	IXXXX	
91	538	0.212	0.036	15- 19	IXXXX	
69	607	0.239	0.027	20- 24	IXXX	
69	676	0.266	0.027	25- 29	IXXX	
63	739	0.291	0.025	30- 34	IXX	
72	811	0.319	0.028	35- 39	IXXX	
78	889	0.350	0.031	40- 44	IXXX	
71	960	0.378	0.028	45- 49	IXXX	
85	1045	0.411	0.033	50- 54	IXXX	
70	1115	0.439	0.028	55- 59	IXXX	
63	1178	0.464	0.025	60- 64	IXX	
51	1229	0.484	0.020	65- 69	IXX	
60	1289	0.507	0.024	70- 74	IXX	
61	1350	0.531	0.024	75- 79	IXX	
42	1392	0.548	0.017	80- 84	IXX	
53	1445	0.569	0.021	85- 89	IXX	
43	1488	0.586	0.017	90- 94	IXX	
40	1528	0.602	0.016	95- 99	IXX	
50	1578	0.621	0.020	100- 104	IXX	
30	1608	0.633	0.012	105- 109	IX	
27	1635	0.644	0.011	110- 114	IX	
27	1662	0.654	0.011	115- 119	IX	
33	1695	0.667	0.013	120- 124	IX	
24	1719	0.677	0.009	125- 129	IX	
29	1748	0.688	0.011	130- 134	IX	
26	1774	0.698	0.010	135- 139	IX	
28	1802	0.709	0.011	140- 144	IX	
34	1836	0.723	0.013	145- 149	IX	
26	1862	0.733	0.010	150- 154	IX	
17	1879	0.740	0.007	155- 159	IX	
16	1895	0.746	0.006	160- 164	IX	
28	1923	0.757	0.011	165- 169	IX	
25	1948	0.767	0.010	170- 174	IX	
23	1971	0.776	0.009	175- 179	IX	
14	1985	0.781	0.006	180- 184	IX	
17	2002	0.788	0.007	185- 189	IX	
19	2021	0.796	0.007	190- 194	IX	
16	2037	0.802	0.006	195- 199	IX	

2540 ENTRIES TOTAL AVERAGE = 119.52954 VARIANCE = 25647.663 STANDARD DEVIATION = 160.149
503 OUT OF RANGE AVERAGE FOR THESE = 340.12127

SAMPLE DISTRIBUTION WITH
POOR REDUCTION PARAMETERS

FIGURE II-8

(5) Other Distributions. Distributions which do not fit the examples contained in this section need special treatment. If some of the data cannot be eliminated to make the distribution more like the examples, the analyst should consider requesting outside help in choosing a Representative Value or in determining if memory is a constraint by other means. The Representative Value concept may not be valid in this case.

3. Decision Values

The decision steps of the test procedures call for comparisons of calculated ratios to specific sample decision values. These comparisons are used to trigger changes in procedure execution sequence or to justify recommendations that the analyst is to make as a result of conducting the test procedure.

Decision values are bracketed ("[value]") in the test procedure documentation. Note that sample decision values are given for illustrative purposes. They will not always be correct for a particular site. Each site whould carefully refine the sample decision values to meet site criteria. These sample decision values have been selected with the following assumptions: (1) a turnaround time elongation problem exists, (2) the procedures are used to search for causes of the problem, and (3) the WWMCCS site will confirm the sample decision values or will modify them through continued evaluation.

H. CONSTRAINTS TO PROCEDURE DEVELOPMENT

No tuning guide can investigate all possible causes of turnaround time elongation. There remain many potentially unaddressed and uninstrumented performance variables in the WWMCCS system processes that may still constrain system performance. The analyst may not uncover all causes of elongation simply by executing the tests in this Guide. The test procedures in the Guide, however, should assist an analyst with the proper background in identifying unaddressed causes of batch turnaround time elongation.

III. TURNAROUND TIME MODEL SCAN

This section describes procedures for initiating the analysis of batch job turnaround time.

A. ANALYSIS SUMMARY

The procedures of the Turnaround Time Model Scan are used to: (1) direct the use of the Batch Turnaround Time Analysis System, (2) identify the delaying processes or queue points in a system being analyzed, and (3) direct the conduct of other tests to confirm the cause of the delay(s). These activities correspond to parts of the Problem Analysis Phase discussed in Volume I, Section IV. Running the Batch Turnaround Time Analysis System corresponds to Volume I, Section IV.B: Run the Appropriate Analysis System. Identifying the delaying processes corresponds to Volume I, Section IV.C: Evaluate Analyzer Output. Directing the conduct of other tests corresponds to Volume I, Sections IV.D through IV.F: Follow the Guide Test Procedures, Implement Guide Tuning Recommendations, and Evaluate Need to Continue. The detailed instructions for each test are given in Sections IV through XII of this volume.

Procedures executed under the Turnaround Time Model Scan include: (1) Run the Batch Turnaround Time Analysis System, (2) Determine Model Of Interest, (3) Select Phase, (4) Select Processes Of Phase To Be Investigated. (5) (Optionally) Determine Execution State Values, and (6) Select and Conduct Tests. Figure III-1 shows the procedures and steps to be executed. The Turnaround Time Model Scan Form (see Figure III-2) is used to record performance data. The procedures and use of the form are described in following subsections.

Reports of the Batch Turnaround Time Analysis System (see Guide Appendix B) used in the procedures of the scan include: (1) System Report, (2) Model Report, and (3) Activity Execution Model Report.

B. RUN THE BATCH TURNAROUND TIME ANALYSIS SYSTEM

The Batch Turnaround Time Analysis System consists of a data collector and a data reduction program. The data collector collects the data to start the investigation of long turnaround times. This tool uses the GCOS traces to track jobs thru the system. The data collector has been incorporated into the Generalized Monitor Facility and the reader should refer to the General Monitoring Facility Users Manual for directions on the use of this tool.

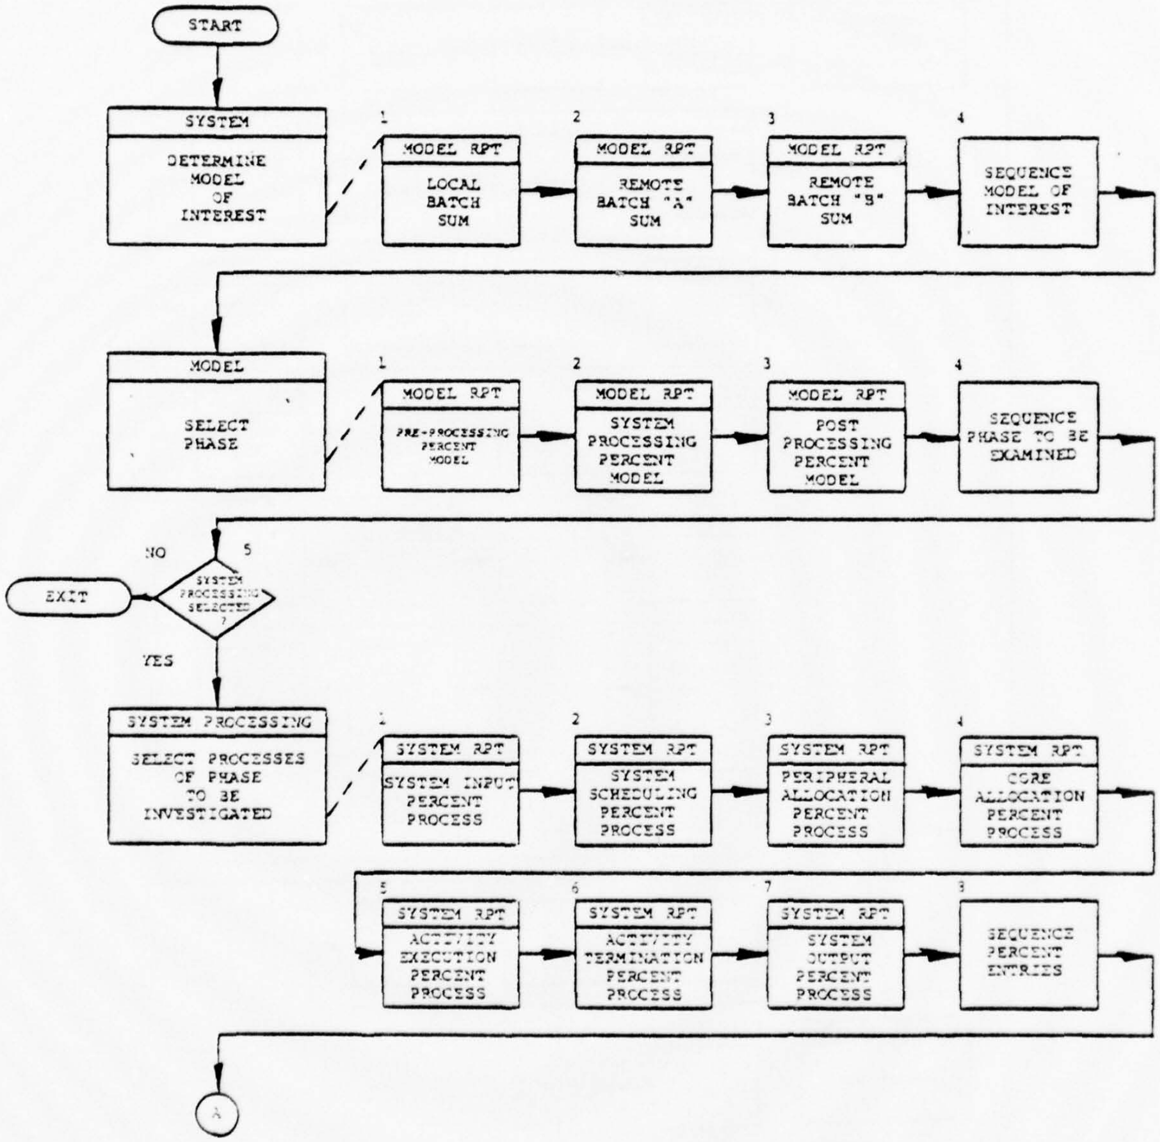

TURNAROUND TIME MODEL SCAN PROCEDURES

FIGURE III-1

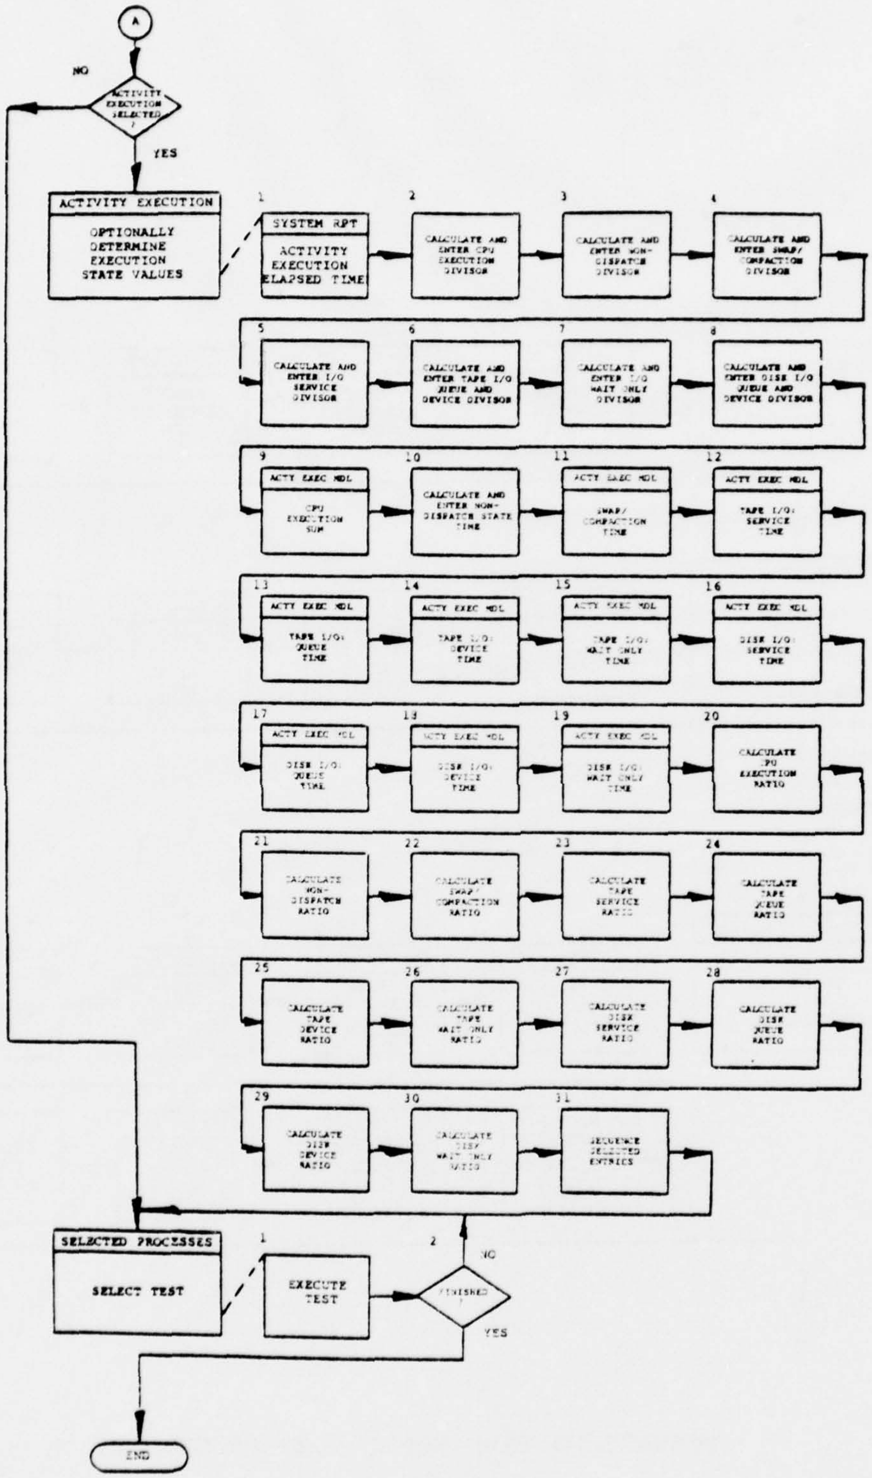

TURNAROUND TIME MODEL SCAN PROCEDURES

FIGURE III-1 (Cont'd)

TURNAROUND SCAN PROCEDURE	RUN		PAGE	

SYSTEM	ELAPSED TIME SUM	VALUE	SEQ
	LOCAL BATCH MODEL		
	REMOTE BATCH "A" MODEL		
	REMOTE BATCH "B" MODEL		

PHASE	PERCENT MODEL	VALUE	SEQ
	PRE-PROCESSING		
	SYSTEM PROCESSING		
	POST PROCESSING		

SYSTEM PROCESSING			HYPOTHESIS TESTS									
PROCESS	PERCENT OF SYSTEM PROCESSING	SEQ	SEEK ELONGATION	MEMORY CONSTRAINT	DEVICE ERRORS	PATHWAY UTILIZATION	CPU EXECUTION	INSUFFICIENT DEVICES	"FEW" ACTIVITIES	IOS DELAYS	URGENCY CODES	NOTE
SYSTEM INPUT			1		3	2		5		4		
SYSTEM SCHEDULING			1		3	2		5		4		
PERIPHERAL ALLOCATION			4	2	5	6		1		7	3	
CORE ALLOCATION				1							2	
ACTIVITY EXECUTION			2	5	7	1	3		4	8	6	
ACTIVITY TERMINATION			1		3	2				4		
SYSTEM OUTPUT			1		3	2				4		

ACTIVITY EXECUTION						HYPOTHESIS TESTS										
EXECUTION STATE			VALUE	CALCULATION	SEQ	SEEK ELONGATION	MEMORY CONSTRAINT	DEVICE ERRORS	PATHWAY UTILIZATION	CPU EXECUTION	INSUFFICIENT DEVICES	"FEW" ACTIVITIES	IOS DELAYS	URGENCY CODES	EXAMINE PROGRAMS	NOTE
DIVISORS	EXECUTION			—	—	—	—	—	—	—	—	—	—	—	—	
	CPU			—	—	—	—	—	—	—	—	—	—	—	—	
	NON-DISPATCH			—	—	—	—	—	—	—	—	—	—	—	—	
	SWAP/COMPACTION			—	—	—	—	—	—	—	—	—	—	—	—	
	I/O SERVICE			—	—	—	—	—	—	—	—	—	—	—	—	
	TAPE I/O QUEUE AND DEVICE			—	—	—	—	—	—	—	—	—	—	—	—	
	WAIT-ONLY			—	—	—	—	—	—	—	—	—	—	—	—	
	DISK I/O QUEUE AND DEVICE			—	—	—	—	—	—	—	—	—	—	—	—	
CPU EXECUTION										1						
NON-DISPATCH						3	4	6	2			1		5	7	
SWAP/COMPACTION							1							2		
I/O PROCESS TIME	TAPE	SERVICE											1			
		QUEUE						2	1							
		DEVICE						2	1							
		WAIT-ONLY						3	2			1				
	DISK	SERVICE											1			
		QUEUE				1		3	2							
		DEVICE				1		3	2							
		WAIT-ONLY				1		4	3			2				

TURNAROUND TIME MODEL SCAN FORM

FIGURE III-2

After the data has been collected, the data tape and manual logs are processed thru a data reduction program to track jobs from submission to users retrieval. The amounts of elapsed time spent in the intermediate phases and processes are summarized in two reports. Appendix B - Volume IV contains the User's Guide for the data reduction program. The GMF Users Manual should be consulted so that only data for the first two reports are captured. Note that manual logs must be kept while the Batch Turnaround Time Analysis System is running. These logs record the time of submission (start of Pre-Processing) and the time of return to the user (end of Post Processing) for each job to be included in the data. The times are then recorded on punched cards for input to the data reduction program (see the Batch Turnaround Time Analysis System User's Guide Appendix B). The analyst may wish to define processes subdividing Pre-processing and/or Post Processing, but this is not required until Pre-Processing or Post Processing has been chosen for further investigation as a result of Section III.D. If processes are defined, manual logs must record the start and end of each process. If Pre & Post processing do not need to be analyzed the use of manual logs can be deleted. (See Appendix B)

The analyst should make sure the Batch Turnaround Time Analysis System is collecting data during and only during the periods experienceing long turnaround times. At most sites, this will mean collecting data from 8 AM to 5 PM, since, at other times of day, short turnaround time is considered less critical. Jobs which do not run to completion during this time will contribute little or not data; therefore, the analyst should run the Batch Turnaround Time Analysis System until any backlog of jobs is gone. Jobs which entered the system while the backlog was being worked off, after the normal end of processing, do not have to run to completion before data collection is stopped. They should be eliminated from the data by not including their SNUMB's in input to the data reduction program (see the User's Guide Appendix B).

The analyst should also make sure that the data is representative (i.e., typical) of the data that would be collected on a "normal" day. This probably requires running the Batch Turnaround Time Analysis System for several time periods and comparing the outputs to see if they vary widely from one day or cycle to the next. Known workload cycles (e.g., end of month, end of year) should be taken into account. They should be eliminated from data collection periods if it is all right for turnaround times to be longer then. They should be the only data collection periods if they are the only periods experiencing poor turnaround times.

C. DETERMINE MODEL OF INTEREST

This procedure determines which one of the three batch system models is to be investigated during a particular analysis of turnaround time. The procedure uses the System Section of the Turnaround Scan Procedure Form.

1. Step No. 1: Local Batch Sum

The System Section of the Turnaround Scan Procedure Form is prepared for the sequencing decision in Step Number 4 by this entry.

a. Report Value. The Local Batch Elapsed Time Sum value, listed on the Model Report, represents the total elapsed time measured for Local Batch jobs (i.e., including Pre-Processing, System Processing, and Post Processing).

b. Form Entry. Enter the value for Local Batch Sum in the Value Column.

2. Step No. 2: Remote Batch "A" Sum

The System Section is prepared for the sequencing operation in Step Number 4 by this entry.

a. Report Value. The Remote Batch "A" Elapsed Time Sum, listed on the Model Report, represents the total amount of elapsed time spent by Remote Batch "A" jobs (i.e., including only System Processing).

b. Form Entry. Enter the value for Remote Batch "A" Sum in the Value Column.

3. Step No. 3: Remote Batch "B" Sum

The System Section of the Turnaround Scan Procedure Form is prepared for the sequencing operation in Step Number 4 by this entry.

a. Report Value. The Remote Batch "B" Elapsed Time Sum, listed on the Model Report, represents the elapsed time measured for Remote Batch "B" jobs (i.e., including System Processing and Post Processing).

b. Form Entry. Enter the value for Remote Batch "B" Sum in the Value Column.

4. Step No. 4: Sequence Major Model of Interest

Examine the three entries made in the Value Column of the System Section. Place a "1" next to the model entry with the largest elapsed time value. Place a "2" in the Sequence Column next to the model entry with the second largest value and a "3" next to the model entry with the smallest value.

Use the remainder of the Turnaround Time Model Scan procedures and forms to investigate the model selected at this step. Investigate the model with the "1" in its Sequence Slot first, the model with a "2" second, etc. Investigation of each model will require further executions of the procedures and extra copies of the Turnaround Time Model Scan Form.

D. SELECT PHASE

This procedure identifies the phase (i.e., Pre-Processing, System Processing, or Post Processing) that is to be examined during the scan. This procedure uses the Phase Section of the Turnaround Scan Procedure Form.

1. Step No. 1: Pre-Processing Percent Model

This entry prepares the Phase Section for the sequencing operation in Step Number 4.

a. Report Value. The Pre-Processing Percent Model value, listed on the Model Report, represents the total percentage of time devoted to Pre-Processing. Be sure to pick the value corresponding to the model picked above in Section III.C. Note that this value has meaning only when the Local Batch Model is being investigated.

b. Form Entry. Enter the value in the Value Column.

2. Step No. 2: System Processing Percent Model

This entry prepares the Phase Section for the sequencing operation in Step Number 4.

a. Report Value. The System Processing Percent Model value, listed on the Model Report, represents the percent of time devoted to System Processing. Note that all models incorporate this phase. Be sure to pick the value corresponding to the model picked above in Section III.C.

b. Form Entry. Enter the number in the Value Column.

3. Step No. 3: Post Processing Percent Model

This entry prepares the Phase Section for the sequencing operation in Step Number 4.

a. Report Value. The Post Processing Percent Model value, listed on the Model Report, represents the percentage of time devoted to Post Processing.

b. Form Entry. Enter the number in Value Column.

4. Step No. 4: Sequence Phase To Be Examined

This operation selects particular phases for analysis in the remaining steps of the Turnaround Scan.

Enter a "1" in the Sequence Column next to the entry with the largest value; enter a "2" in the Sequence Column of the entry with the second highest value and a "3" in the Sequence Column with the lowest entry.

If System Processing has the highest value, continue the current Turnaround Time Model Scan with the procedure steps in Section III.E, below.

If Pre-Processing has the largest value, investigate the queue points and processes involved in entering batch jobs into the system. Specific analysis points will depend upon the number and types of work stations, job staging queues, control points and tape or disk library processes at any particular site. Define processes for Pre-Processing and keep logs and collect data using the Batch Turnaround Time

Analysis System. The logs record for each job the times when each process of Pre-Processing has been entered. This data must be punched on cards and entered into the data reduction program as specified in the Batch Turnaround Time Analysis System User's Guide. The ending time of Pre-Processing will be recorded automatically for each job by the Batch Turnaround Time Analysis System. The processes contributing the most to Pre-Processing delays should be investigated for possible improvements. The method of investigation is left to the analyst.

If Post Processing has the highest value, investigate the queue points and processes involved in removing jobs from the system. Specific analysis points will depend upon the number and types of work stations, output decollation/bursting operations, and output staging queues at any particular site. Define processes, collect data, and investigate processes as with Pre-Processing (described above). The starting time for Post Processing will be recorded by the Batch Turnaround Time Analysis System.

E. SELECT PROCESSES OF PHASE TO BE INVESTIGATED

This procedure isolates those processes of the System Processing Phase that are to be investigated. This procedure uses the System Processing Section of the Turnaround Scan Procedure Form (see Figure III-2).

1. Step No. 1: System Input Percent Phase

This step prepares an entry for the sequencing operation in Step Number 8.

a. Report Value. The System Input and Scheduler Percent Phase value, listed on the System Report for the selected model, represents the percentage of System Processing Time devoted to the System Input and System Scheduling Process. Note that the same value is used in Step 2, below. Current GCOS trace instrumentation does not provide the means to distinguish between time spent in System Input and time spent in the System Scheduler. Future instrumentation will separate the two phases.

b. Form Entry. Enter the value in the Percent of System Processing Column.

2. Step No. 2: System Scheduling Percent Phase

This step prepares an entry for the sequencing operation in Step Number 8.

a. Report Value. The System Input and Scheduler Percent Phase value, listed on the System Report for the selected model, represents the percent of System Processing Time devoted to the System Input and System Scheduling Process. Note that the same value is used in Step 1, above. Current GCOS trace instrumentation does not provide the means to distinguish between time spent in System Input and time spent in the System Scheduler. Future instrumentation will separate the two phases.

b. Form Entry. Enter the value in the Percent of System Processing Column.

3. Step No. 3: Peripheral Allocation Percent Phase

This step prepares an entry for the sequencing operation in Step Number 8.

a. Report Value. The Peripheral Allocation Percent Phase value, listed on the System Report for the selected model, represents the percent of System Processing Time devoted to the Peripheral Allocation Process.

b. Form Entry. Enter the value the Percent of System Processing Column.

4. Step No. 4: Core Allocation Percent Phase

This step prepares an entry for the sequencing operation in Step Number 8.

a. Report Value. The Core Allocation Percent Phase value, listed on the System Report for the selected model, represents the percent of System Processing Time devoted to the Core Allocation Process.

b. Form Entry. Enter the value in the Percent of System Processing Column.

5. Step No. 5: Activity Execution Percent Phase

This step prepares an entry for the sequencing operation in Step Number 8.

a. Report Value. The Activity Execution Percent Phase value, listed in the System Report of the selected model, represents the percent of System Processing Time devoted to the Activity Execution Process.

b. Form Entry: Enter the value in the Percent of System Processing Column.

6. Step No. 6: Job Termination Percent Phase

This step prepares an entry for the sequencing operation in Step Number 8.

a. Report Value. The Activity Execution Percent Phase value, listed on the System Report of the selected model, represents the percent of System Processing Time devoted to the Job Termination Process.

b. Form Entry. Enter the value in the Percent of System Processing Column.

7. Step No. 7: System Output Percent Phase

This step prepares an entry for the sequencing operation in Step Number 8.

a. Report Values. The Ready for Output Percent Phase value, listed on the System Report of the selected model, represents the percent of System Processing Time devoted to waiting for output to start. The System Output Percent Phase value, listed just under it, represents the percent of System Processing Time devoted to the System Output Process after output starts for a job. See the User's Guide for the details of the distinction.

b. Form Entry. Sum the two values and enter the sum in the Percent of System Processing Column.

8. Step No. 8: Sequence Percent Phase Entries

This step is used to sequence the System Processing processes to be examined by the test procedures. Place a "1" in the Sequence Column of the process with the highest Percent of System Processing value. Enter a "2" next to the process with the second-highest value, etc. Choose the process marked "1" for further investigation using Section III.G. If this process has been recently investigated and little improvement is anticipated, an analyst may want to choose other processes for investigation, starting with the process marked "2".

9. Step No. 9: Activity Execution Detail

Continue with the analysis by conducting the following procedure if the Activity Execution Process has been selected for investigation. If the Activity Execution Process is not to be investigated, conduct the procedure in Section III.G.

F. OPTIONALLY DETERMINE EXECUTION STATE VALUES

This procedure provides detailed information about the jobs running under the selected model for the Activity Execution Process. This information can reduce the number of tests that must be conducted before an analyst reaches a tuning decision.

In order to perform this step it will be necessary to re-run the data collector and collect data for reports 3&4. (See GMF Users Manual) This option will result in a great increase in the amount of data collected. Obtain data for several time periods and repeat this section for each period collected, to insure that decisions are not made on the basis of one set of abnormal data. Consult appendix B for directions on the generation of Report 3.

This procedure uses the Activity Execution Section of the turnaround Scan Procedure Form (see Figure III-2).

Eleven ratios (see Table III-1) are produced by this procedure to describe the execution states displayed by the system. The dividend values for the eleven ratios are listed on the Activity Execution Model Report produced by the Batch Turnaround Time Analysis System. The ratio divisor values are determined from : (1) the test system configuration and (2) the emphasis placed on particular system states during the analysis. The following definitions apply to the divisors in the formulae in Table III-1:

(ET) = Execution Elapsed Time

(CPU) = Number of CPUs in the Test Configuration

(TAPE) = Number of Logical Tape Channels in the Test Configuration

(DISK) = Number of Logical Disk Channels in the Test Configuration

RATIO	RATIO FORMULA	DIVISOR FORMULA	DEFAULT VALUE FOR "n"
CPU Execution	CPU Execution Sum / n(ET)(CPU)	n(ET)(CPU)	1.0
Non-Dispatch	Non-Dispatch State Time / n(ET)(CPU)	n(ET)(CPU)	0.2
Swap/Compaction	Swap/Compaction Time / n(ET)(CPU)	n(ET)(CPU)	0.2
Tape Service	Tape I/O Service Time / n(ET)	n(ET)	1.0
Tape I/O Queue	Tape I/O Queue Time / n(ET)(TAPE)	n(ET)(TAPE)	1.0
Tape Device	Tape I/O Device Time / n(ET)(TAPE)	n(ET)(TAPE)	1.0
Tape Wait-Only	Tape I/O Wait-Only Time / n(ET)	n(ET)	0.1
Disk Service	Disk I/O Service Time / n(ET)(DISK)	n(ET)(DISK)	1.0
Disk I/O Queue	Disk I/O Queue Time / n(ET)(DISK)	n(ET)(DISK)	1.0
Disk Device	Disk I/O Device Time / n(ET)(DISK)	n(ET)(DISK)	1.0
Disk Wait-Only	Disk I/O Wait-Only Time / n(ET)	n(ET)	0.1

EXECUTION STATE FORMULAE

TABLE III-1

Before using the Execution State Formulae in the steps of this procedure, complete the optional elements of the formulae: (1) the configuration components and (2) the execution sequence determinator values (i.e., "n" for each formula). For example, if a dual processor configuration is being tested, substitute a "2" for the CPU parameter of the following divisors: (1) CPU Execution, (2) Non-Dispatch, and (3) Swap/Compaction.

The execution sequence determinator (i.e., "n" parameter in each divisor) provides a way of defining the relative "importance" of the eleven system states. A small amount of time in some of the states would be more "important" than a larger amount of time in some other states. For example, one expects to see very little "Tape I/O Wait-Only Time" (see the User's Guide for a definition of this value). Therefore, a modest value for this state could be significant while the same value of time for CPU execution would not be considered significant. The execution sequence determinators attempt to compensate for this problem by magnifying the amount of time reported for those states where a small value might be significant. Default values for the execution sequence determinators are listed in the right-hand column of Table III-1. Each site should carefully evaluate the determinators and modify their values if justified.

1. <u>Step No. 1: Activity Execution Elapsed Time</u>

This step prepares an entry for the sequencing operation in Step Number 31.

a. <u>Report Value</u>. The Execution Elapsed Time, listed on the System Report of the selected model, represents the amount of elapsed System Processing Time devoted to the Activity Execution Process.

b. <u>Form Entry</u>. Enter the value in the Value Column.

2. <u>Step No. 2: Calculate and Enter CPU Execution Divisor</u>

This step prepares an entry for the calculation of the CPU Execution Ratio in Step Number 20.

a. <u>Calculation</u>. Compute the CPU Execution Divisor using the appropriate divisor formula from Table III-1 and the Elapsed Time Value generated in Step Number 1. Provide an alternate value for the "n" factor in the formula if desired.

b. Form Entry. Enter the calculated value for the CPU Execution Divisor in the Value Column.

3. Step No. 3: Calculate and Enter Non-Dispatch Divisor

This step prepares an entry for the calculation of the Non-Dispatch Ratio in Step Number 21.

a. Calculation. Compute the Non-Dispatch Divisor using the Non-Dispatch Divisor Formula from Table III-1 and the Elapsed Time Value generated in Step Number 1. Note from Table III-1 that a default value for "n" of 0.2 is recommended. Change this value if desired.

b. Form Entry. Enter the calculated value for the Non-Dispatch Divisor in the Value Column.

4. Step No. 4: Calculate and Enter Swap/Compaction Divisor

This step prepares an entry for the calculation of the Swap/Compaction Ratio in Step Number 22.

a. Calculation: Compute the Swap/Compaction Divisor using the Swap/Compaction Divisor Formula from Table III-1 and the Elapsed Time Value generated in Step Number 1. Note that a value for "n" of 0.2 is recommended in the table. Change this value if desired.

b. Form Entry. Enter the value for the Swap/Compaction Divisor in the Value Column.

5. Step No. 5: Calculate and Enter I/O Service Divisor

This step prepares an entry for the calculation of the Tape Service Ratio in Step Number 23 and the Disk Service Ratio in Step Number 27. The same divisor is used in both steps.

a. Calculation. Compute the I/O Service Divisor using the Tape Service divisor formula from Table III-1 and in the Elapsed Time Value generated in Step Number 1.

b. Form Entry. Enter the value for the I/O Service Divisor in the Value Column.

6. Step No. 6: Calculate and Enter Tape I/O Queue and Device Divisor

This step prepares an entry for the calculation of the Tape Queue Ratio and the Tape Device Ratio in Step Number 24 and Step Number 25. The same divisor is used in both steps.

a. Calculation. Compute the value for the Tape Queue and Device Divisor using the Tape I/O Queue divisor formula from Table III-1 and the Elapsed Time Value generated in Step Number 1.

b. Form Entry. Enter the value in the Value Column.

7. Step No. 7: Calculate and Enter I/O Wait-Only Divisor

This step prepares an entry for the calculation of the Tape Wait-Only Ratio in Step Number 26 and the Disk Wait-Only Ratio in Step Number 30.

a. Calculation. Compute the value for the I/O Wait-Only Divisor using the Tape Wait-Only divisor formula from Table III-1 and the Elapsed Time Value generated in Step Number 1. Note that a default value of 0.1 is suggested for "n" in the table. Change this value if desired.

b. Form Entry. Enter the value in the Value Column.

8. Step No. 8: Calculate and Enter Disk I/O Queue and Device Divisor

This step prepares an entry for the calculation of the Disk Queue Ratio and the Disk Device Ratio in Step Number 28 and Step Number 29.

a. Calculation. Compute the value for the Disk I/O Queue and Device Divisor using the Disk I/O Queue divisor formula from Table III-1 and the Elapsed Time Value generated in Step Number 1.

b. Form Entry. Enter the value for the divisor in the Value Column.

9. Step No. 9: CPU Execution Time

This step prepares an entry for the calculation of the CPU Execution Ratio in Step Number 20.

a. Report Value. The CPU Execution Elapsed Time, listed on the Activity Execution Model Report, represents the sum of CPU busy time for all CPU's.

b. Form Entry. Enter the value for CPU Execution Time in the Value Column.

10. Step No. 10: Calculate and Enter Non-Dispatch State Time

This step prepares an entry for calculation of the Non-Dispatch Ratio in Step Number 21.

a. Report Value. The Monitor Session Time in minutes, listed in the upper right corner of the Activity Execution Model Report, represents the time during which the Batch Turnaround Time Analysis System collected data.

b. Calculation. Convert the Monitor Session Time to seconds (multiply by 60) and multiply it by the number of processors. Subtract from this product the CPU Execution Time just entered in the Value Column. The result will be a measure of CPU idleness.

c. Form Entry. Enter the result in the Value Column.

11. Step No. 11: Swap/Compaction Time

This step prepares an entry for calculation of the Swap/Compaction Ratio in Step Number 22.

a. Report Value. The Swap/Compaction Time value, listed on the Activity Execution Model Report, represents the amount of elapsed time devoted to the Swap/Compaction State of the Activity Execution Process.

b. Form Entry. Enter the Swap/Compaction Time value in the Value Column.

12. Step No. 12: Tape I/O Service Time

This step prepares an entry for calculation of the Tape Service Ratio in Step Number 23.

a. Report Value. The Tape Service Time, listed on the Activity Execution Model Report, represents the amount of time spent in the I/O Supervisor for Tape I/O.

b. Form Entry. Enter the value for Tape I/O Service Time in the Value Column.

13. Step No. 13: Tape I/O Queue Time

This step prepares an entry for calculation of the Tape I/O Queue Ratio in Step Number 24.

a. Report Value. The Tape Queue Time, listed on the Activity Execution Model Report, represents the amount of time that one or more entries were in I/O queues for tape I/O services.

b. Form Entry. Enter the value in the Value Column.

14. Step No. 14: Tape I/O Device Time

This step prepares an entry for calculation of the Tape Device Ratio in Step Number 25.

a. Report Value. The Tape Device Time, listed on the Activity Execution Model Report, represents the amount of elapsed time devoted to tape I/O operations outside the host processor (i.e., the physical I/O tape device time).

b. Form Entry. Enter the value in the Value Column.

15. Step No. 15: Tape I/O Wait-Only Time

This step prepares an entry for calculation of the Tape Wait-Only Ratio in Step Number 26.

a. Report Value. The Tape Wait-Only Time, listed on the Activity Execution Model Report, represents the amount of elapsed time when all host processors were in the Wait State (i.e., executing the DIS Instruction) and I/O was outstanding for one or more tapes (but no other type of device).

b. Form Entry. Enter the value in the Value Column.

16. Step No. 16: Disk I/O Service Time

This step prepares an entry for calculation of the Disk Service Ratio in Step Number 27.

a. Report Value. The IAS Service Time, listed on the Activity Execution Model Report, represents the elapsed time spent in the I/O Supervisor for disk service.

b. Form Entry. Enter the value in the Value Column.

17. Step No. 17: Disk I/O Queue Time

This step prepares an entry for calculation of the Disk Queue Ratio in Step Number 28.

a. Report Value. The IAS Queue Time, listed on the Activity Execution Model Report, represents the amount of elapsed time that one or more entries were in queues for disk I/O service.

b. Form Entry. Enter the value in the Value Column.

18. Step No. 18: Disk I/O Device Time

This step prepares an entry for calculation of the Disk Device Ratio in Step Number 29.

a. Report Value. The IAS Device Elapsed Time, listed on the Activity Execution Model Report, represents the amount of Activity Execution Elapsed Time devoted to disk device operations occurring outside the host system mainframe.

b. Form Entry. Enter the value in the Value Column.

19. Step No. 19: Disk I/O Wait-Only Time

This step prepares an entry for calculation of the Disk Wait-Only Ratio in Step Number 30.

a. Report Value. The IAS Wait-Only Time Value, listed on the Activity Execution Model Report, represents the amount of elapsed time when all host processors were in the Wait State (i.e., executing the DIS Instruction) and an I/O operation was outstanding for one or more disk devices (but no other type of device).

b. Form Entry. Enter the value in the Value Column.

20. Step No. 20: Calculate CPU Execution Ratio

This step calculates the CPU Execution Ratio, a value used to establish the turnaround time test sequence in Step Number 31.

a. Calculation. Divide the value for CPU Execution by the value for the CPU Divisor.

b. Form Entry. Enter the resulting quotient in the Calculation Column.

21. Step No. 21: Calculate Non-Dispatch Ratio

This step calculates the Non-Dispatch Ratio, a value used in the sequencing operation in Step Number 31.

a. Calculation. Divide the value for Non-Dispatch Time by the value for the Non-Dispatch Divisor.

b. Form Entry. Enter the resulting quotient in the Calculation Column.

22. Step No. 22: Calculate Swap/Compaction Ratio

This step calculates the Swap/Compaction Ratio, a value employed in the sequencing operation in Step Number 31.

a. Calculation. Divide the value for Swap/Compaction Time by the value for the Swap/Compaction Divisor.

b. Form Entry. Enter the resulting quotient in the Calculation Column.

23. Step No. 23: Calculate Tape Service Ratio

This step calculates the Tape Service Ratio, a value employed in the sequencing operation in Step Number 31.

a. Calculation. Divide the value for Tape I/O Service Time by the value for the I/O Service Divisor.

b. Form Entry. Enter the resulting quotient in the Calculation Column.

24. Step No. 24: Calculate Tape Queue Ratio

This step calculates the Tape Queue Ratio, a value employed in the sequencing operation in Step Number 31.

a. Calculation. Divide the value for Tape I/O Queue Time by the value for the Tape I/O Queue and Device Divisor.

b. Form Entry. Enter the resulting quotient in the Calculation Column.

25. Step No. 25: Calculate Tape Device Ratio

This step calculates the Tape Device Ratio, a value employed in the sequencing operation in Step Number 31.

a. Calculation. Divide the value for Tape I/O Device Time by the value for the Tape I/O Queue and Device Divisor.

b. Form Entry. Enter the resulting quotient in the Calculation Column.

26. Step No. 26: Calculate Tape Wait-Only Ratio

This step calculates the Tape Wait-Only Ratio, a value employed in the sequencing operation in Step Number 31.

a. Calculation. Divide the value for Tape I/O Wait-Only Time by the value for the Wait-Only Divisor.

b. Form Entry. Enter the resulting quotient in the Calculation Column.

27. Step No. 27: Calculate Disk Service Ratio

This step calculates the Disk Service Ratio, a value employed in the sequencing operation in Step Number 31.

a. Calculation. Divide the value for Disk I/O Service Time by the value for the I/O Service Divisor.

b. Form Entry. Enter the resulting quotient in the Calculation Column.

28. Step No. 28: Calculate Disk Queue Ratio

This step calculates the Disk Queue Ratio, a value employed in the sequencing operation in Step Number 31.

a. Calculation. Divide the value for Disk I/O Queue Time by the value for the Disk I/O Queue and Device Divisor.

b. Form Entry. Enter the resulting quotient in the Calculation Column.

29. Step No. 29: Calculate Disk Device Ratio

This step calculates the Disk Device Ratio, a value employed in the sequencing operation in Step Number 31.

a. Calculation. Divide the value for Disk I/O Device Time by the value for the Disk I/O Queue and Device Divisor.

b. Form Entry. Enter the resulting quotient in the Calculation Column.

30. Step No. 30: Calculate Disk Wait-Only Ratio

This step computes the Disk Wait-Only Ratio, a value employed in the sequencing operation in Step Number 31.

a. Calculation. Divide the value for Disk I/O Wait-Only Time by the value for the Wait-Only Divisor.

b. Form Entry. Enter the resulting quotient in the Calculation Column.

31. Step No. 31: Sequence Selected Entries

This step is used to establish the execution order of turnaround time elongation hypothesis confirmation tests. A Sequence Column, provided in the Activity Execution Section of the Turnaround Scan Procedure Form, is used for this step.

Sequence the Calculation Entries in descending order by the value entered. Enter a "1" in the Sequence Column of the entry with the highest calculated ratio; enter a "2" in the corresponding Sequence Column for the second highest ratio, etc.

Note the sequence of tests to be performed and proceed to a continuation of the scan process described below.

G. SELECT AND CONDUCT TESTS

The sequence of the tests to be conducted for each System Processing Process or Activity Execution State is defined in the Hypothesis Tests Column of either the System Processing or Activity Execution Section of the Turnaround Scan Procedure Form (see Figure III-2). As an example, note from Figure III-2 that the following tests are to be conducted (in the sequence indicated) for elongation of the System Input Process: (1) Seek Elongation Test, (2) Pathway Utilization Test, (3) Device Errors Test, (4) IOS Delays Test, and (5) Insufficient Devices Test. Each test checks one or more hypotheses (given in Tables III-2 through III-9). This sequence is suggested only. The site analyst may feel that the most likely hypotheses are given low priority in the test sequencing on the form. If this is the case, the sequencing should be modified so that the most likely hypotheses are tested first.

1. Step No. 1: Execute Tests

Note which processes (or states, if Activity Execution is picked) are to be investigated and the sequence of tests to be conducted. Modify the sequence of tests if circumstances justify it. Choose and execute a test and proceed to Step 2.

Some tests can be streamlined by keeping in mind the specific hypothesis being tested. Each test is designed to check all devices and/or programs for a particular problem; however, the specific hypothesis being tested may concern only one device or program. For example, the Insufficient Devices Test checks for adequate numbers of disk and tape devices. The analyst could be conducting the Insufficient Devices Test to evaluate the hypothesis that insufficient space exists for *J and J* files. If this were the case, the analyst should skip the part of the Insufficient Devices Test concerned with tapes, because the *J and J* files are disk only. While conducting one of the tests, the analyst should remember the specific hypotheses he is testing. Some analysts will also be able to devise a new test which tests more specifically for the hypotheses he is testing.

TEST	PROCESS STATE	HYPOTHESIS	NOTE
Seek Elongation	Read, NXTBF	1. Long Seeks to *J File 2. Long Seeks to J* File	
Pathway Utilization	Read, NXTBF	Long Queues to Devices And/Or Channels Containing J* And/Or *J Files	
Device Errors	Read, NXTBF	Excessive Errors on Disk Device Containing J* And/Or *J Files	
		Excessive Errors on Tape Drives	If Tape Input of Jobs
IOS Delays	Read, NXTBF	1. Inefficient Tape IOS* Code	If Tape Input of Jobs
		2. Inefficient Unit Record IOS* Code	If Card Reader Input of Jobs
		3. Inefficient Disk IOS* Code	
Insufficient Devices	Read, Start, NXTBF	1. Insufficient Space For J* File	
		2. Insufficient Space For *J File	
(Not Investigated)	Initialize	Insufficient Memory For GEIN	
	NXTBF	.CRDMP Set Too Often	
	.MGNAT Module	Inefficient Code	
	.MSCAN Module	Inefficient Code	
	Begin, Read	Excessive Errors on Card Reader	

*I/O Supervisor

ELONGATION HYPOTHESES FOR
SYSTEM INPUT PROCESS
GROUPED BY TEST

TABLE III-2

TEST	PROCESS STATE	HYPOTHESIS
Seek Elongation	Search For New Job Requests, Job Data Write to Disk	1. Long Seeks to *J File 2. Long Seeks to J* File
Pathway Utilization	Search For New Job Requests, Job Data Write to Disk	Long Queues to Devices And/Or Channels Containing J* And/Or *J Files
Device Errors	Search For New Job Requests, Job Data Write to Disk	Excessive Errors on Device(s) Containing J* And/Or *J Files
IOS Delays	Search For New Job Requests, Job Data Write to Disk	Inefficient Disk IOS* Code
Insufficient Devices	Search For New Job Requests, Job Data Write to Disk	1. Insufficient Space For J* File 2. Insufficient Space For *J File
(Not Investigated)	New Job Request Information Setup	1. .CRDSP Gate Shut Often 2. Insufficient Memory For RGEIN
	RQ Table Residence	Long RQ Table Entry Residence Time
	T1 Queue Residence	Long T1 Queue Entry Residence Time
	Job Data Write to Disk, Build .CRRGQ Entry	.CRRGQ Gate Shut Often
	.MGNAT Module	Inefficient Code
	.MSCAN Module	Inefficient Code

*I/O Supervisor

ELONGATION HYPOTHESES FOR
REMOTE SYSTEM INPUT PROCESS
GROUPED BY TEST

TABLE III-3

TEST	PROCESS STATE	HYPOTHESIS
Seek Elongation	Catalog Job, Update Catalog	Long Seeks to System Scheduler Catalog(s)
	Attempt Job Start	Long Seeks to SYSOUT Files
Pathway Utilization	Catalog Job, Update Catalog	Long Queues to Devices And/Or Channels Containing System Scheduler Catalog(s)
	Attempt Job Start	Long Queue to Devices And/Or Channels Containing SYSOUT Files
Device Errors	Catalog Job, Update Catalog	Excessive Errors on Devices And/Or Channels Containing System Scheduler Catalog(s)
	Attempt Job Start	Excessive Errors on Devices And/Or Channels Containing SYSOUT Files
IOS Delays	Catalog Job, Update Catalog, Attempt Job Start	Inefficient Disk IOS* Code
Insufficient Devices	Attempt Job Start	Insufficient Space For SYSOUT
(Not Investigated)	.CRRGQ Queue Residence	.CRRGQ Entry Residence Time Long
	Check .CRRGQ	Insufficient Core Available For Task Control Block Chain
	New Job Entering	1. MAX Parameter Setting Too Low 2. CLASS MAX Parameter Setting Too Low 3. MIN Parameter Setting Too High
	Catalog Residence	Job Catalog Entry Residence Time Long
	Attempt Job Start	1. Insufficient Buffer Space 2. Insufficient Program Numbers

*I/O Supervisor

ELONGATION HYPOTHESES FOR
SYSTEM SCHEDULING PROCESS
GROUPED BY TEST

TABLE III-4

TEST	PROCESS STATE	HYPOTHESIS
Insufficient Devices	Peripheral Allocation, PALC/CALC Damper Processing	1. Tapes Requested Not Available 2. Disk Space Requested Not Available
Memory Constraint	PALC/CALC Damper Processing	PALC/CALC Damper Set Often
Urgency Codes	New Activity	Urgency Code For Jobs Too Low
Seek Elongation	Peripheral Allocation	Long Seeks to PALC Scratch File
	Write SSA to *J File	Long Seeks to *J File
Device Errors	Peripheral Allocation	Excessive Errors on Device Containing PALC Scratch File
	Write SSA to *J File	Excessive Errors on Device Containing *J File
Pathway Utilization	Peripheral Allocation	Long Queues to Devices And/Or Channels Containing PALC Scratch File
	Write SSA to *J File	Long Queues to Devices And/Or Channels Containing *J File
IOS Delays	Peripheral Allocation	Inefficient Disk I/O Supervisor Code
	Write SSA to *J File	Inefficient Disk I/O Supervisor Code
(Not Investigated)	.CRJOB Queue Residence	.CRJOB Entry Residence Time Long

ELONGATION HYPOTHESES FOR
PERIPHERAL ALLOCATION PROCESS
GROUPED BY TEST

TABLE III-5

TEST	PROCESS STATE	HYPOTHESIS
Memory Constraint	Core Allocation	Amount of Memory Available Is Causing Elongation
Urgency Code	Priority Adjustment	Operator Changes to Urgency Codes Are Causing Elongation of Jobs
	Core Allocation	Urgency Codes of Other Jobs in System Are Causing Elongation
(Not Investigated)	.CRPOQ Queue Residence	.CRPOQ Entry Residence Time Too Long
	Core Allocation	"MXCNT" Parameter Causing Elongation Due To Excessive Core Compaction/Inadequate Core Compaction
	Start Job	Time to Zero Slave Core Too Long

ELONGATION HYPOTHESES FOR
CORE ALLOCATION PROCESS
GROUPED BY TEST

BLE III-6

TEST	PROCESS STATE	HYPOTHESIS
Seek Elongation	GECALL Process	Long Seeks to GECALL Library
	SSA Process, MME Process, Fault Process	Long Seeks to SSA Library
	Non-Dispatch State	Long Seeks to Activity's Files (Includes SYSOUT)
	Swap/Compaction State	Long Seeks to Swap File(s)
	Activity I/O Process	Long Seeks to Activity's Files (Includes SYSOUT)
Memory Constraint	Non-Dispatch State, Swap/Compaction State	1. Insufficient Memory 2. Wasted Memory
Urgency Codes	Non-Dispatch State, Swap/Compaction State	Urgencies of Other Activities Are Causing Elongation of Selected Jobs
"Few Activities In System"	Non-Dispatch State	Few (or no) Activities Is Cause For High Value
CPU Execution Characteristics	CPU Process	Inefficient Program Code
Device Errors	GECALL Process	Device Errors on Device That Contains GECALL Library
	SSA Process, MME Process, Fault Process	Device Errors on Device That Contains SSA Modules

ELONGATION HYPOTHESES FOR
ACTIVITY EXECUTION PROCESS
GROUPED BY TEST

TABLE III-7

TEST	PROCESS STATE	HYPOTHESIS
	Swap/Compaction State	Device Errors on Device That Contains Swap Files
	Activity I/O Process: Tape	Tape I/O Errors on Activity Files
	Activity I/O Process: Disk	Disk I/O Errors on Activity Files
Pathway Utilization	GECALL Process	Long Queues to Devices And/Or Channels Containing GECALL Library
	SSA Process, MME Process, Fault Process	Long Queues to Devices And/Or Channels Containing SSA Modules
	Non-Dispatch State	Long Queues to Devices And/Or Channels Containing Activity Files
	Swap/Compaction State	Long Queues to Devices And/Or Channels Containing Swap Files
	Activity I/O Process: Tape	Long Queues to Tape Channels
	Activity I/O Process: Disk	Long Queues to Disk Devices And/Or Channels

ELONGATION HYPOTHESES FOR
ACTIVITY EXECUTION PROCESS
GROUPED BY TEST

TABLE III-7 (Cont'd)

TEST	PROCESS STATE	HYPOTHESIS
IOS Delays	GECALL Process, SSA Process, MME Process, Fault Process, Non-Dispatch State, Swap/Compaction State, Activity I/O Process	Inefficient Disk I/O Supervisor Code
	Non-Dispatch State, Activity I/O Process	Inefficient Tape or Unit Record I/O Supervisor Code
(Not Investigated)	GECALL Process SSA Process, MME Process, Fault Process	Inefficient Called Code 1. Inefficient GCOS Module Code 2. High Module Push/Pop Overhead
	Activity I/O: Disk Activity I/O: Unit Record	Insufficient SYSOUT Buffers 1. Unit Record I/O Errors 2. Long Queues to Unit Record Channels
	Non-Dispatch State	Deadlock or Single-Thread Situations in Site Systems and Patches

ELONGATION HYPOTHESES FOR
ACTIVITY EXECUTION PROCESS
GROUPED BY TEST

TABLE III-7 (Cont'd)

TEST	PROCESS STATE	HYPOTHESIS
Seek Elongation	Activity Compression Processing, Activity Release and Accounting	Long Seeks to *J File
	Term Continuation Processing	Long Seeks to System Module Libraries
Pathway Utilization	Activity Compression Processing, Activity Release and Accounting	Long Queues to Devices And/Or Channels Containing *J File
	Term Continuation Processing	Long Queues to Devices And/Or Channels Containing System Module Libraries
Device Errors	Activity Compression Processing, Activity Release and Accounting	Device Errors on Devices That Contain *J File
	Term Continuation Processing	Device Errors on Devices with System Module Libraries
IOS Delays	Activity Compression Processing, Activity Release and Accounting	Inefficient Disk I/O Supervisor Code
(Not Investigated)	User Accounting Process	1. Inefficient User Code 2. If User I/O: a. Pathway Utilization b. Seek Elongation c. Device Errors d. IOS Delay

ELONGATION HYPOTHESES FOR
JOB TERMINATION PROCESS
GROUPED BY TEST

TABLE III-8

TEST	PROCESS STATE	HYPOTHESIS
Seek Elongation	SYSOUT Disperal Initialization	Long Seeks to SSA Library
	Start SYSOUT Processing, Buffer Input Processing	Long Seeks to SYSOUT Files
Pathway Utilization	SYSOUT Disperal Initialization	Long Queues to Devices And/Or Channels Containing SSA Modules
	Start SYSOUT Processing, Buffer Input processing	Long Queues to Devices And/Or Channels Containing SYSOUT Files
Device Errors	SYSOUT Dispersal Initialization	Device Errors on Device That Contains SSA Modules
	Start SYSOUT Processing, Buffer Input Processing	Device Errors on Device(s) That Contains SYSOUT Files
IOS Delays	SYSOUT Dispersal Initialization, Start SYSOUT Processing, Buffer Input Processing	Inefficient Disk I/O Supervisor Code
	Buffer Output Processing	Inefficient Unit Record I/O Supervisor Code
(Not Investigated)	.CRSYQ Queue Residence	.CRSYQ Entry Residence Time Long
	Buffer Input Processing	Lack of Sufficient SSA Buffers
	Buffer Output Processing	1. Long Queues to Channels Containing SYSOUT Output Devices 2. .MGEOT Serving Other Devices Than SYSOUT Devices

ELONGATION HYPOTHESES FOR
SYSTEM OUTPUT PROCESS
GROUPED BY TEST

TABLE III-9

Tables III-2 through III-9 (one table for each GCOS system process) show the specific hypotheses to be tested by each test, depending on the GCOS system process being investigated. Each time the analyst starts a test procedure, he should refer to the table for the GCOS system process he is investigating, and note the hypothesis or hypotheses to be tested by that test procedure.

2. Step No. 2: Decision to Continue

Once a test has been executed and tuning recommendations (if any) applied, an analyst must make the decision of whether to continue. This decision corresponds to the one discussed in Volume I, Section IV.F. If the tuning objective has been met, stop the analysis. If not, continue with Step Number 3, below.

3. Step No. 3: Continue the Analysis

The analysis may be continued in two ways.

a. Start at Turnaround Time Model Scan Procedure. If one or more tuning recommendations were implemented and their effect on system performance is not negligible, analysis should start again at the beginning of the Turnaround Time Model Scan Procedure. This will insure that decisions are not made on the basis of outdated data. The analyst may decide to rule certain phases/processes/states out of investigation because they were unsuccessfully investigated before. This should only be done with careful consideration that nothing has changed the chances for improvement in those phases/processes/states.

b. Select Another Test/State/Process/Phase. If no tuning recommendations were implemented or if they did not affect performance, the analyst may select another test, state, process, or phase. The process or state selected probably included several tests in its sequence of tests to execute. Refer to the proper row of the completed Turnaround Time Scan Procedure Form to find the next test in the sequence. If all the indicated tests have been executed or if the analyst has reasons to believe little is to be gained from the remaining tests, he may choose another state or process or phase for investigation. This is especially appropriate when the other state/process/phase was close in its ratio or elapsed time value to the state/process/phase just investigated. Start at the Select Phase, Select Processes of Phase, or Determine Execution State Values Section as appropriate.

IV. SEEK ELONGATION TEST

This section describes procedures for testing seek elongation hypotheses. The tests employ the Mass Store Monitor (MSM) as the primary data collection tool.

A. ANALYSIS SUMMARY

The Seek Elongation Test confirms the elongation hypothesis by calculating the Weighted Average Seek Length measured on all disk units in the configuration. The procedures isolate and identify both GCOS System files and user files on those units that exhibit a large Weighted Average Seek Length.

Procedures executed under the Seek Elongation Test include: (1) Confirm Seek Elongation, (2) Identify High-Use Extents, and (3) Identify Files. The steps of these procedures are charted in Figure IV-1. The Seek Elongation Test Form (see Figure IV-2) is provided as a guide to data collection for the procedures.

MSM reports used in the test procedures include: (1) Seek Movement Report, (2) Space Utilization Report, (3) System File Use Summary Report, and (4) File Summary Report.

B. CONFIRM SEEK ELONGATION

This procedure identifies the disk devices in the configuration with high Weighted Average Seek Lengths. The procedure employs the Seek Movement Section of the Seek Elongation Form.

1. Step No. 1: Device Address

This step prepares a form entry to identify the disk units in the configuration.

a. Report Value. A copy of the MSM Seek Movement Report is produced for each device attached to each disk subsystem in the configuration. The address of each configured device appears at the top of its Seek Movement Report.

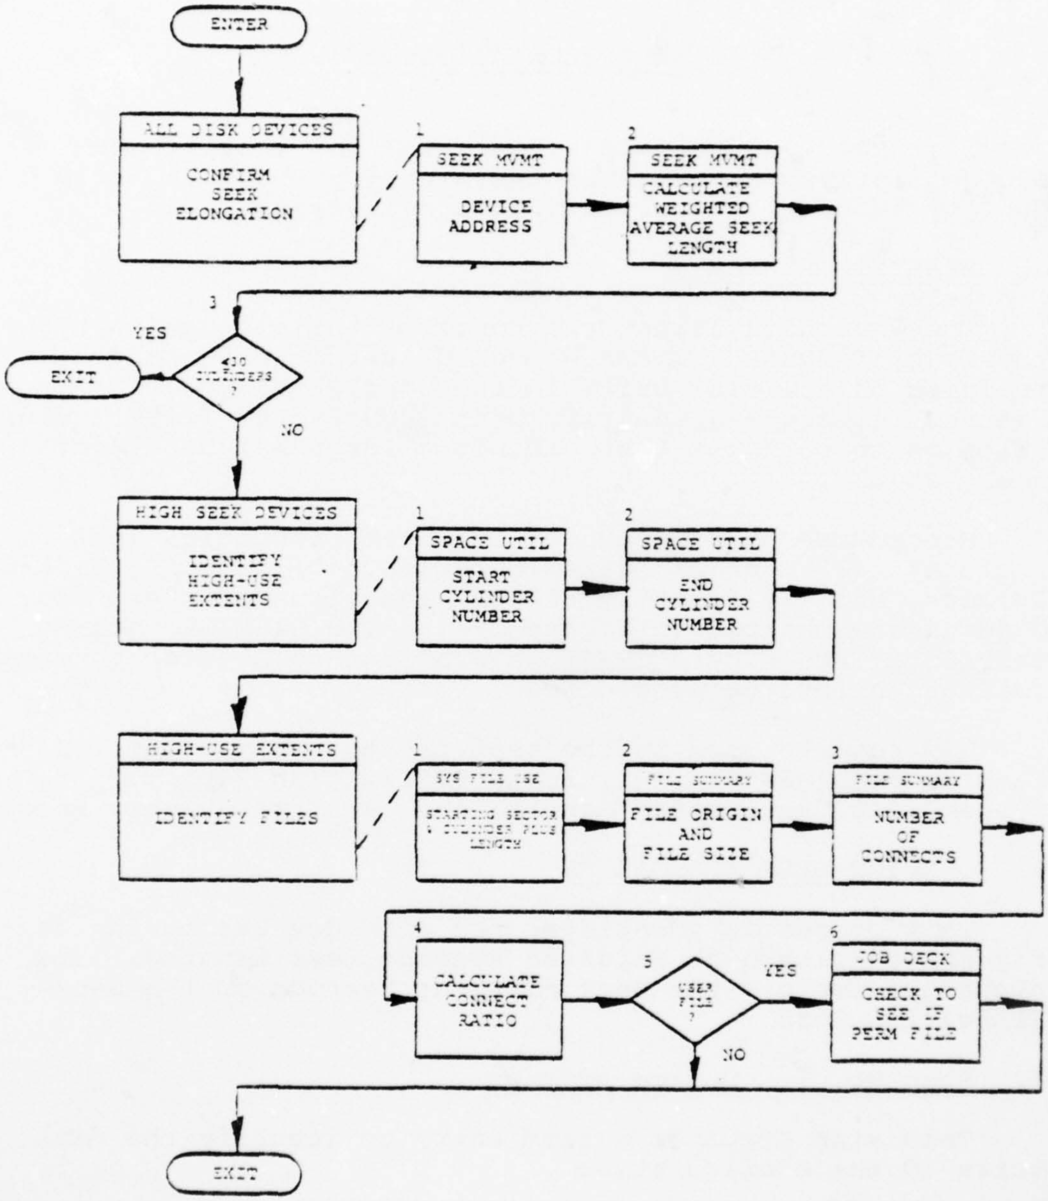

SEEK ELONGATION TEST PROCEDURES

FIGURE IV-1

SEEK ELONGATION	RUN		PAGE	

SEEK MOVEMENT	ADDRESS			WEIGHTED AVG SEEK LENGTH	✓	ADDRESS			WEIGHTED AVG SEEK LENGTH	✓	ADDRESS			WEIGHTED AVG SEEK LENGTH	✓
	ICM	CHNL	DEV			ICM	CHNL	DEV			ICM	CHNL	DEV		

SPACE UTILIZATION · ICCD	START	END	NAME	T	FC	CONNECTS	✓	NAME	T	FC	CONNECTS	✓	NAME	T	FC	CONNECTS	✓

SPACE UTILIZATION · ICCD	START	END	NAME	T	FC	CONNECTS	✓	NAME	T	FC	CONNECTS	✓	NAME	T	FC	CONNECTS	✓

SPACE UTILIZATION · ICCD	START	END	NAME	T	FC	CONNECTS	✓	NAME	T	FC	CONNECTS	✓	NAME	T	FC	CONNECTS	✓

SEEK ELONGATION TEST FORM

FIGURE IV-2

b. Form Entry. Enter the IOM Number, Channel Number (listed as PUB number), and Device Number listed at the top of each MSM Seek Movement Report in the slots provided in the Address Columns.

2. Step No. 2: Calculate Weighted Average Seek Length

This step calculates the value for weighted Average Seek Length on each configured disk device. The result of the calculation is used to determine if a seek elongation delay is exhibited by the test system.

a. Measures. The MSM Seek Movement Report displays a frequency distribution of the number of cylinders moved by seeks executed on the indicated disk unit.

b. Calculation. Compute the Weighted Average Seek Length for each disk in the system, using the Seek Movement Report for each device. Multiply the number in the Individual Probability Column by the larger of the two numbers in the Cylinders Moved Column. [Make this calculation only for those cylinder ranges that are plotted (i.e., exhibit at least one "x" at the right in the "Percent Probability of Occurence" graph).] Sum these products to form the Weighted Average Seek Length for the disk.

c. Form Entry. Enter the calculated value for Weighted Average Seek Length in the Weighted Average Seek Length Column for the corresponding device.

3. Step No. 3: Decision

Check (✓) the device for further analysis in the two procedures below if a high [>30 cylinders or ≥10% of the values are ≥100 cylinders] value for Weighted Average Seek Length is recorded for the device. Little seek elongation is indicated and the test may be exited at this point (return to the procedure that called for this test) if all low [≤30 cylinders] values are calculated. If any devices are checked, further data should be obtained to confirm that the problem occurs consistently. It does not necessarily have to occur consistently on the same devices.

C. IDENTIFY HIGH-USE EXTENTS

This procedure provides a range of addresses (i.e., an extent) on the active devices checked in B above. It uses the MSM Space Utilization Report and the Space Utilization

Section of the Seek Elongation Form. Enter the active unit's Device Address in the space marked "ICCDD" in the Space Utilization Section to identify the device.

1. Step No. 1: Start Cylinder Number

This step provides an entry value to identify the starting address of high-use areas on the disk units.

a. Report Values. Note the start of each extent on the active devices that has [≥5%] Individual Probability. Two or more adjacent extents, each of which has [≥5%] Individual Probability, can be considered one extent.

b. Form Entries. Enter the starting address (cylinder number) of each high-use extent in the Start Column.

2. Step No. 2: End Cylinder Number

This step provides an entry value identifying the ending address of a high-use area on the active disks. Enter the ending address of each high-use extent in the End Column of the Space Utilization Section.

D. *IDENTIFY FILES*

This procedure identifies heavily-used files on the active disk devices. The procedure uses the remaining columns of the Space Utilization Section: "Name", "T", "FC", "Connects", and "✓"

1. Step No. 1: System Files

This step identifies the GCOS System Libraries whose addresses are within the high-use extents identified in C above.

a. Report Value. The MSM System File Use Summary Report lists GCOS System Files that were initialized at System Startup. The device, starting cylinder number, and length for each GCOS System File are also listed. Consult the GCOS Startup File Map to determine the actual names of the GCOS System Files (e.g., SYSTEM FILE2 as identified on the MSM report is the second GCOS System File defined in the Startup Deck).

b. <u>Calculation</u>. For each listed GCOS System File, determine whether any part of the file falls within the high-use extents identified above.

Compare the starting address of the file (right side of Starting Sector/Cylinder Column) to the starting and ending addresses of each high-use extent identified on the same device. If the file does not start within any high-use extent, determine whether it overlaps a high-use extent on that device by adding the Length to the Starting Address to find the end of the file.

The length of the system file is listed on the MSM report in sectors and must be converted to cylinders to be added to the Starting Address. The bottom row of Table IV-1 gives the necessary conversion factors.

	DISK SUBSYSTEM			
PARAMETER	DSS181	DSS190	DSS191	MSU451
Sectors Per Track	18	31.0	40	40
Tracks Per Cylinder	20	19.0	19	19
Blocks Per Cylinder	72	117.8	152	152
Sectors Per Cylinder	360	589.0	760	760

DEFINITIONS FOR SECTOR-TO-CYLINDER VALUE CONVERSION

TABLE IV-1

c. <u>Form Entry</u>. List the following data for each GCOS System File identified as starting in or overlapping a high-use extent: (1) the file name (from the GCOS Startup File Map) in the Name Column, (2) an "S" (for System) in the Type ("T") Column, (3) a dash (-) in the File Code ("FC") Column (no file code), and (4) the number of accesses (from the Accesses Column of the System File Use Summary Report) in the Connects Column. List the data on the same row as the start/end of the high-use extent the file overlaps.

2. <u>Step No. 2: User Files</u>

This step identifies the user files and jobs that are causing the measured seek elongation.

a. Report Value. The MSM File Summary Report identifies the following information for each user file: (1) File Code, (2) Number of Connects, (3) File Size, (4) Allocated Device, and (5) File Origin. Each file is listed under the activity that used it.

Using the same technique described above in Step Number 1, identify the files that start in or overlap the identified high-use extents. Note that both the starting address and length are listed in cylinders; no conversion is needed. Ignore files that have [$\leq$10,000] connects. Reduce the number of connects in the previous sentence if all files in many extents are ignored.

Two unusual file codes will be frequently observed in the report: "00" and "--". The value of "00" for a file code identifies all disk accesses made without a GCOS Peripheral Allocation Table (PAT) entry (e.g., accesses made by GCOS as part of initialization). It also includes all Slave Service Areas (SSA) module calls and pushdowns. The "--" file code is constructed by the MSM data reduction program to account for all access made to SYSOUT files.

b. Form Entry. For the files that fall within high-use extents, enter the $IDENT Name in the Name Column on the same row (if possible) as the starting address of the extent. Enter a "U" (i.e., User) in the Type ("T") Column. Enter the File Code in the File Code ("FC") Column. Enter the number of connects to the reported file code in the Connects Column.

3. Step No. 3: Calculate Connect Ratio

This step calculates a Connect Ratio for each identified file. The ratios are employed in the decision in Step Number 4 to determine which files are candidates for investigation as high-use files. High activity has already been identified in the extent. This step merely determines which files are most responsible for that activity.

a. Calculations. For each device being investigated, total the values for the Number of Connects entered for each file on that device. Then, calculate a Connect Ratio for each entered file by dividing the Number of Connects for each file on the device by this total.

b. Form Entry. These calculations may be made on a separate piece of paper; they are not required to identify high-use files. Enter a "✓" for all files with a Connect Ratio [≥.1].

4. Step No. 4: Relocating Files

All files with a Connect Ratio value [≥.1] are candidates for relocation, eepecially those which consistently have high values over several periods. Unfortunately, GCOS does not allow an analyst to specify where within a device a file should be placed. Only thedevice name can be specified. This is done in three different ways, depending on the type of file: system, permanent, or temporary. User files can be classified by checking the file code entry on the file summary report. The type of file is indicated on the right side of the file code characters. Immediately to the right of the file code is the File Type (P = PERM, T = TEMPORARY). To the right of this is the access characteristic of the file (R=RANDOM and S=SEQUENTIAL). The next character defines the file as cataloged "C", or Non-Cataloged, "N", if it is a perm file. The final character defines the pack as fixed (F) or removable (R).

a. System Files. The device name is specified at Startup for each system file on the $FILDEF card for that file. High-use system files that are resident on devices with long seeks should be relocated to fast devices which are lightly used. If the relocation causes long seeks on the new device, some form of moving permanent and temporary files off that device should be tried.

b. Permanent Files. The device name for a permanent file can be specified at creation, whether through FILSYS or the ACCESS subsystem of Time Sharing. Files can be moved by changing their names, creating a new file with the old name, and moving the data. The new file can be created with a device name specification.

c. Temporary Files. The device name for a temporary file is specified in the second field of the $FILE card in the job control deck. Jobs which run frequently can have their $FILE cards changed. Other jobs can be controlled by policies governing user use of $FILE cards.

V. MEMORY CONSTRAINT TEST

This section describes the test for analyzing batch turnaround time elongation hypothesized as being caused by memory configuration or use. The procedures of this test use the Memory Utilization Monitor (MUM), the CPU monitor and the channel Monitor (PSM) as data collection tools.

A. ANALYSIS SUMMARY

The Memory Constraint Test analyzes the following memory usage metrics: (1) Average Elapsed Memory Wait Time, (2) Processor Idle Time, and (3) Number of I/O Entries Queued.

With this test, a Memory Wait Time Ratio is used to determine whether memory demand is high, with the rest of the test procedures used to determine whether additional jobs in memory at once would improve the system's utilization characteristics.

Procedures executed under the Memory Constraint Test include: (1) Determine If Memory Wait Time Is High, and (2) Analysis Entry. Figure V-1 charts the steps to be executed under the Memory Constraint Test procedures. The Memory Constraint Test Form (see Figure V-2) is provided as a guide to data collection for the procedures.

Other Batch Turnaround Time Analysis Tests referenced by this test are: (1) CPU Characteristics Test, (2) Urgency Codes Test, and (3) Pathway Utilization Test.

Table V-1 lists the PSM and MUM reports used in this test.

SYSTEM	REPORT
Channel Monitor	Channel and Device Reports
Memory Utilization Monitor	(1) Total Elapsed Time An Activity Was In Memory
	(2) Elapsed Wait Time For Memory Requests
CPU Monitor	(1) Processor Time Distribution Report

REPORTS USED IN MEMORY CONSTRAINT TEST PROCEDURES

TABLE V-1

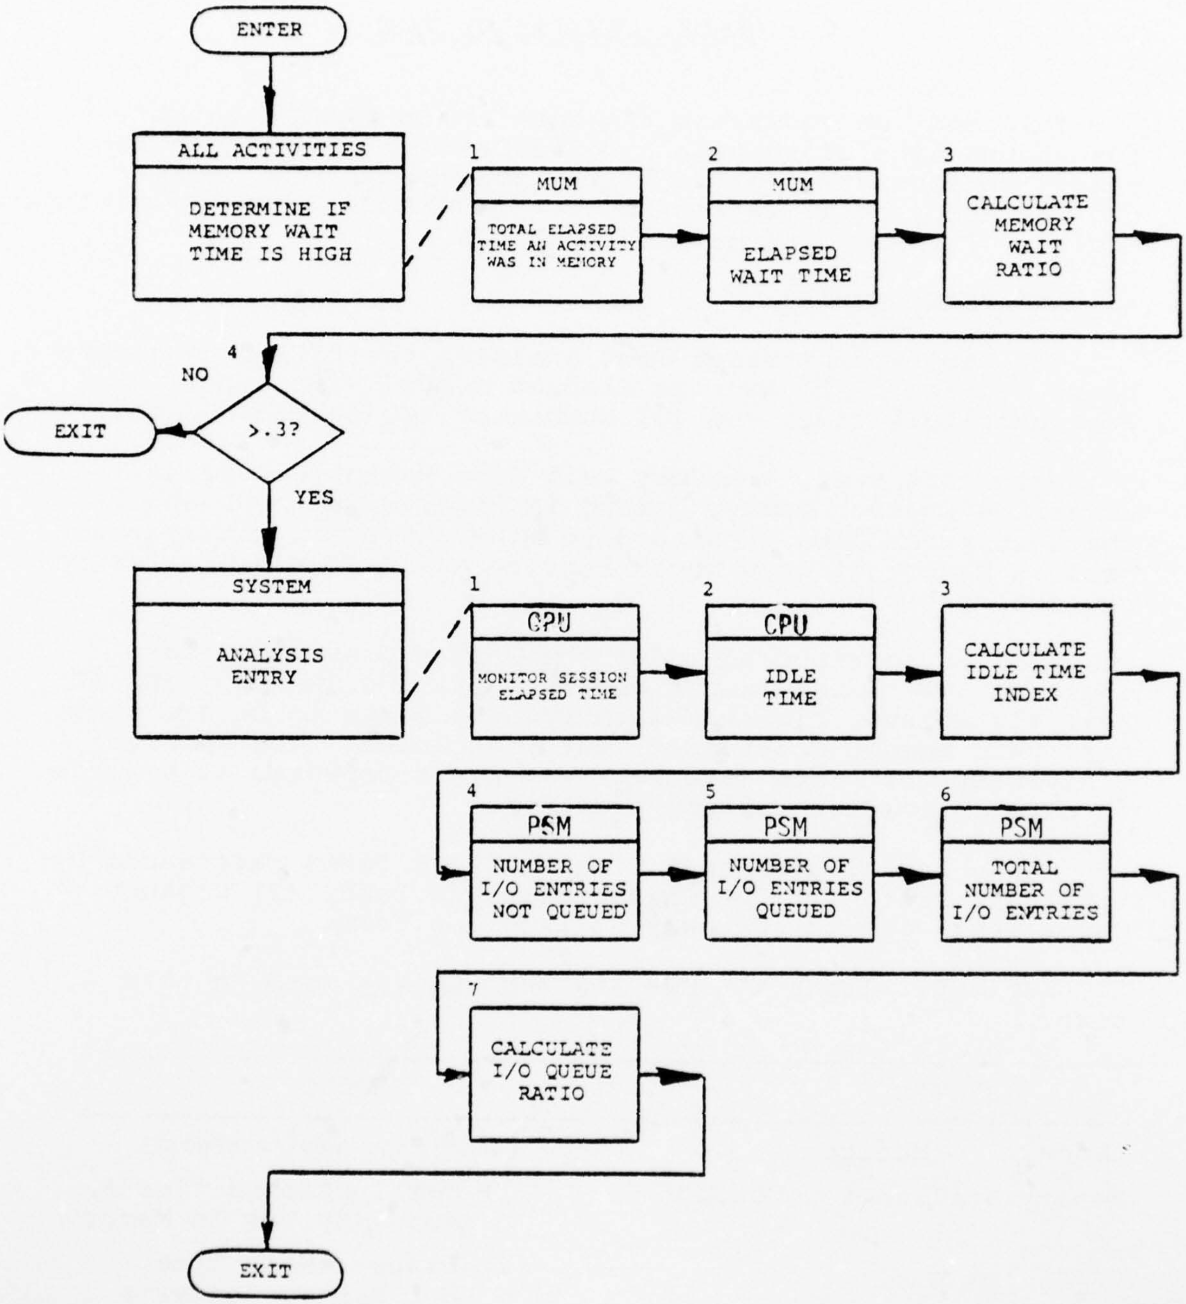

MEMORY CONSTRAINT TEST PROCEDURES

FIGURE V-1

MEMORY CONSTRAINT	RUN		PAGE	

MEMORY WAIT TIME CONFIRMATION				
TIME	TOTAL ELAPSED TIME	AVERAGE ELAPSED WAIT	MEMORY WAIT RATIO	✓

ANALYSIS ENTRY						
ELAPSED TIME	IDLE TIME	IDLE TIME INDEX	NO. I/O ENTRIES NOT QUEUED	NO. I/O ENTRIES QUEUED	TOTAL NO. I/O ENTRIES	I/O QUEUE RATIO

MEMORY CONSTRAINT TEST FORM

FIGURE V-2

B. DETERMINE IF MEMORY WAIT TIME IS HIGH

This procedure determines if Memory Wait Time is high enough to warrant analysis of memory demand and service. Data should be collected and this determination repeated for several days. Note that all data must be collected concurrently. Collect data on processor time and I/O queuing as well as memory usage, because synchronized processor and I/O queuing data are needed for later parts of the procedure.

1. Step No. 1: Total Elapsed Time An Activity Was In Memory

This value is used as the divisor in calculating the Memory Wait Ratio in Step Number 3.

a. Report Values. Determine the Representative Value for the frequency distribution in the MUM Total Elapsed Time An Activity Was In Memory Report.

b. Representative Values. The general rules for determining Representative Values of frequency distributions are presented in Section II.G.2.

c. Form Entry. Enter the Representative Value in the Total Elapsed Time Column of the Memory Wait Time Confirmation Section.

2. Step No. 2: Average Elapsed Wait Time

This value is used as the dividend when calculating the Memory Wait Ratio in Step Number 3.

a. Report Value. Determine the Representative Value for the distribution in the MUM Elapsed Wait Time for Memory Requests Report (see Section II.G.2.)

b. Form Entry. Enter the Representative Value in the Average Elapsed Wait Column of the Memory Wait Time Confirmation Section.

3. Step No. 3: Calculate Memory Wait Ratio

The ratio is calculated from the Memory Wait Time Confirmation Section entries made in the above two steps.

a. Computation. Divide the value for Elapsed Wait Time by the value for Total Elapsed Time.

b. Form Entry. Enter the quotient in the Memory Wait Ratio Column of the Memory Wait Time Confirmation Section. This represents the ratio of time waiting for memory to time in memory.

4. Step No. 4: Decision

The Memory Wait Time can be considered low if the resulting Memory Wait Ratio is [≤.3]; further analysis is not required. Return to the section that called for this test. If the resulting Memory Wait Ratio is [>.3], the Memory Wait Time can be considered high and the analyst should check (✓) the column to the right in the Memory Wait Time Confirmation Section. This determination should be repeated for several time periods (e.g., days). If Memory Wait Time is consistently high during the periods analyzed, proceed to the next section. Otherwise, exit the test (return to the section that called for this test).

C. ANALYSIS ENTRY

The purpose of this procedure is to direct the analysis toward investigation of system resources that are used by the workload. This procedure provides data that are interpreted by a decision logic table to direct the execution of three other Turnaround Time Tests: (1) CPU Execution Characteristics, (2) Urgency Codes, and (3) Pathway Utilization.

This procedure uses the Analysis Entry Section of the Memory Constraint Form.

1. Step No. 1: Monitor Session Elapsed Time

This step develops an entry used to calculate the Idle Time Index in Step Number 3.

a. Report Value. The Elapsed Time Thus Far value, listed on the CPU Utilization Report, represents the total elapsed time (in seconds) of the monitor session. Select the elapsed time value from the last CPU Utilization Report listed.

b. Form Entry. Enter the value in the Elapsed Time Column of the Analysis Entry Section.

2. Step No. 2: Idle Time

This step generates the dividend used to calculate the Idle Time Index in Step Number 3.

a. Report Value. The Idle Times values, listed on the CPU Utilization Report, represent the Processor Idle Time measured for each processor in the system. The values are displayed in 100th second units.

b. Calculation. Total the Idle Time values for processors configured. Convert the sum from 100th second units into seconds.

c. Form Entry. Enter the total for the configuration's Processor Idle Time in the Idle Time Column of the Analysis Entry Section.

3. Step No. 3: Calculate Idle Time Index

This step calculates the Idle Time Index used in the decision in the Analysis Entry Decision Logic Table.

a. Calculation. Divide the value for Idle Time by the value for Monitor Session Elapsed Time.

b. Form Entry. Enter the quotient in the Idle Time Index Column of the Analysis Entry Section.

4. Step No. 4: Number of I/O Entries Not Queued

This step prepares an entry for the calculation of the I/O Queue Ratio described in Step Number 7.

a. Report Value. The Channel and Device Free Report. Total all values appearing in this report. This total represents the number of I/O entries (i.e., requests for I/O service) that were immediately started by the GCOS I/O supervisor.

b. Form Entry. Enter the value in the Number I/O Entries Not Queued Column of the Analysis Entry Section.

5. Step No. 5: Number of I/O Entries Queued

This step prepares an entry for the calculation of the I/O Queue Ratio in Step Number 7.

AD-A059 938 COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON D C F/G 9/2
WORLDWIDE MILITARY COMMAND AND CONTROL SYSTEM (WWMCCS) H-6000 T--ETC(U)
FEB 78 B M WALLACK, G H GERO

UNCLASSIFIED CCTC-TM-180-78-VOL-2 SBIE-AD-E100 095 NL

2 OF 2
AD
A059938

END
DATE
FILMED
12-78
DDC

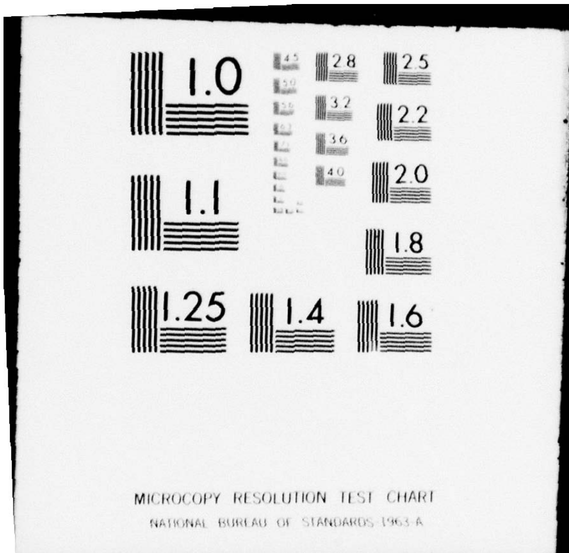
1.0
2.8
2.5
3.2
2.2
3.6
4.0
2.0
1.1
1.8
1.25
1.4
1.6
MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

a. Report Value. The Channel-Device Busy Report, Channel Busy-Device Free Report, and Channel Free-Device Busy Report. Total all values appearing in these reports. This total represents the number of I/O entries (i.e., requests for I/O service) that had to be queued by the GCOS I/O supervisor.

b. Form Entry. Enter the value in the Number I/O Entries Queued Column of the Analysis Entry Section.

6. Step No. 6: Total Number of I/O Entries

This step prepares an entry for the calculation of the I/O Queue Ratio in Step Number 7.

a. Calculation. Add the value in the Number I/O Entries Queued Column to the value in the Number I/O Entries Not Queued Column.

b. Form Entry. Enter the resulting sum in the Total Number I/O Entries Column of the Analysis Entry Section.

7. Step No. 7: Calculate I/O Queue Ratio

This step calculates the I/O Queue Ratio employed in the decision described in Paragraph D below.

a. Calculation. Divide the Number I/O Entries Queued by the Total Number of I/O Entries.

b. Form Entry. Enter the resulting quotient in the I/O Queue Ratio Column of the Analysis Entry Section.

D. FURTHER INVESTIGATION AND TUNING STEPS

This paragraph discusses interpretation of data collected by the Analysis Entry Procedure described in C above.

1. Form Entry Interpretation

Interpret the values entered on the Memory Constraint Test Form for the Idle Time Index and I/O Queue Ratio as either "HIGH" or "LOW" according to the following rules:

a. Idle Time Index. If the value computed for the Idle Time Index is [≤.2], consider it "LOW"; if [>.2], consider it "HIGH".

b. I/O Queue Ratio. If the value computed for the I/O Queue Ratio is [≤.4], consider it "LOW"; if [>.4], consider it "HIGH".

Make these determinations for each data period and ensure that a consistent pattern is observed. If the pattern is inconsistent, obtain professional help.

2. Decision Table

Table V-2 represents the decision logic for further analysis. Four decision rules are involved. In all cases, an analyst is directed to conduct further analysis by employing up to three tests: (1) CPU Characteristics Test, (2) Urgency Codes Test, and/or (3) Pathway Utilization Test. Of course, an analyst may call for any of the tests if extenuating circumstances require it.

	RULE			
	1	2	3	4
Idle Time Index	LOW	LOW	HIGH	HIGH
I/O Queue Ratio	LOW	HIGH	LOW	HIGH
Analyze CPU Execution Characteristics	X	X		
Analyze Urgency Codes Mix	X	X	X	
Analyze Pathway Utilization		X		X

ANALYSIS ENTRY DECISION LOGIC

TABLE V-2

3. Decision Rules and Tuning Steps

The correction of the conditions implied by the decision rules of Table V-2 depends upon information provided by the performance of other batch turnaround time analysis tests.

a. Rule No. 1: Low Idle Time Index and Low I/O Queue Ratio. This condition describes a system whose jobs are CPU-dominant and exhibit low I/O activity.

First, conduct the CPU Execution Characteristics Test (see Section VIII). If none of the tuning steps proposed under the CPU Execution Characteristics Test alleviate the condition, some improvement to turnaround time can probably be achieved by the following solutions (particularly if the Idle Time Index is [≥.1]):

(1) Parameter Solution. Propose reducing the size of GCOS system memory by decreasing the amount of core required by TSS and GCOS Hard Core.

(2) Programming Solution. Propose investigation of the jobs that were running at the time of the measurement session to determine if they actually required the memory that was requested. If not, reduce the amount requested on $LIMITS Card. Further analysis may be justified if this simple solution cannot be applied.

b. Rule No. 2: Low Idle Time Index and High I/O Queue Ratio. This condition describes a system that is approaching saturation. The system's CPU Utilization is high, I/O service requests have a high Queue Ratio, and Memory Wait Time is high.

(1) Scheduling Solution. Propose scheduling fewer jobs during periods that exhibit this condition. Adjust the site .MSCAN code to compensate. This solution involves classifying some jobs (the ones not scheduled) as not high enough in priority to require normal turnaround.

(2) Conduct Pathway Utilization Test. Remedies to the high I/O Queue Ratio condition can be obtained through scheduling, parameter changing, programming, and sizing solutions. Conduct the Pathway Utilization Test (see Section VII), proposing or implementing the solutions presented in that section before proceeding. If the solutions presented there are successful, re-enter the Turnaround Time Model Scan at the start if the Tuning Objective has not been met.

(3) Conduct CPU Execution Characteristics Test. Conduct the CPU Execution Characteristics Test (see Section VIII), proposing or executing the solutions presented in that section. If these solutions are successful, re-enter the Turnaround Time Model Scan at the start if the Tuning Objective has not been met.

(4) Conduct Urgency Codes Test. Perform the Urgency Codes Test (see Section XII) to examine the mix of Urgency Codes for the jobs executing in the system during the measurement period. Look for jobs that have high Urgency Codes or that request large areas of core. If possible, schedule these jobs so that they do not run at the same time or lower their Urgency Codes.

(5) Parameter Solution. This step proposes changes to GCOS system functions and only helps high priority jobs. Consider this solution only to get some critical jobs through the system. Determine, by examining the Startup Patch Deck, if the GCOS Urgency Throughput Dispatcher Option is employed. If not, determine whether the technique could be used to expedite execution of selected jobs.

c. Rule No. 3: High Idle Time Index and Low I/O Queue Ratio. This condition describes a system that could better use its resources if more jobs could reside in main memory at the same time.

(1) Scheduling Solution. Conduct the Urgency Codes Test (see Section XII). Look for jobs that request large areas of core using the GSEP Allocator/Termination Report. If possible, run these jobs at times when turnaround is less critical. If they have high urgency codes, reduce the urgency codes to a level closer to that of the other jobs running during the test.

(2) Parameter Solutions. This step proposes changes to GCOS system functions. Investigate reducing the size of GCOS System Memory by decreasing the amount of core allocated to TSS and to GCOS Hard Core.

(3) Programming Solution. Propose investigation of the jobs running at the time of the measurement session to determine if they actually required the memory requested and assigned to them. If the memory was not required, reduce the amount requested on the $LIMITS Card. Further analysis may be justified if this simple solution cannot be applied.

(4) Sizing Solution. If the condition continues, propose an increase in the amount of memory configured on the system.

d. Rule No. 4: High Idle Time Index And High I/O Queue Ratio. This condition describes a system whose jobs are I/O-dominant.

(1) Conduct Pathway Utilization Test. See Section VII for the procedures of this test. Propose or execute the solutions prescribed there. Return here only if nothing can be done.

(2) Scheduling Solution. Conduct the Urgency Codes Test (see Section XII) to examine the mix of urgency codes for the jobs executing during the measurement period. Look for jobs that have high urgency codes or that request large amounts of core. If possible, schedule these jobs so that they do not run at the same time or reduce their urgency codes.

(3) Parameter Solution. This step proposes changes to GCOS system functions. Investigate reducing the size of GCOS System Memory by decreasing the amount of core allocated to TSS and/or to GCOS Hard Core.

(4) Programming Solution. Propose an investigation of the jobs running at the time of the measurement session to determine if they actually required the memory requested and assigned to them. If the memory was not required, reduce the amount of core requested on the $LIMITS Card. Further analysis may be justified if this simple solution cannot be applied.

VI. DEVICE ERRORS TEST

This section describes the procedures for analyzing turnaround time elongation due to device errors. This test uses the HEALS II software monitor to gather data.

A. ANALYSIS SUMMARY

Procedure steps executed under the Device Errors Test include: (1) Determine Tape Handler Units In Error and (2) Determine Disk Units In Error. Figure VI-1 shows the procedure steps to be executed for this test. The Device Errors Form (see Figure VI-2) is used by this procedure for data collection.

HEALS II reports used in the Device Errors Test include: (1) Tape Unit Error Variance, (2) Tape Errors By Unit/Reel Numbers, and (3) MPC Statistics. The GCOS Console Log is employed in the Determine Disk Units in Error Test to identify disk pack numbers on which errors occurred.

B. DETERMINE TAPE HANDLER UNITS IN ERROR

The objective of this procedure is to identify the tape units and tape reels that are exhibiting a large number of I/O errors and are therefore contributing to batch turnaround time elongation.

This procedure uses the Tapes Section of the Device Errors Form.

1. Step No. 1: Tape Handler Address

This step identifies the tape handlers that reported at least one Tape Alert during the monitor session.

a. Report Value. The device address for which Tape Alerts are reported is listed on the Tape Unit Error Variance Report of the HEALS II System under "Handler."

b. Form Entry. Enter the tape handler addresses in the Handler Column in the Tapes Section. Enter these addresses in ascending sequence by IOM/Channel/Device Number.

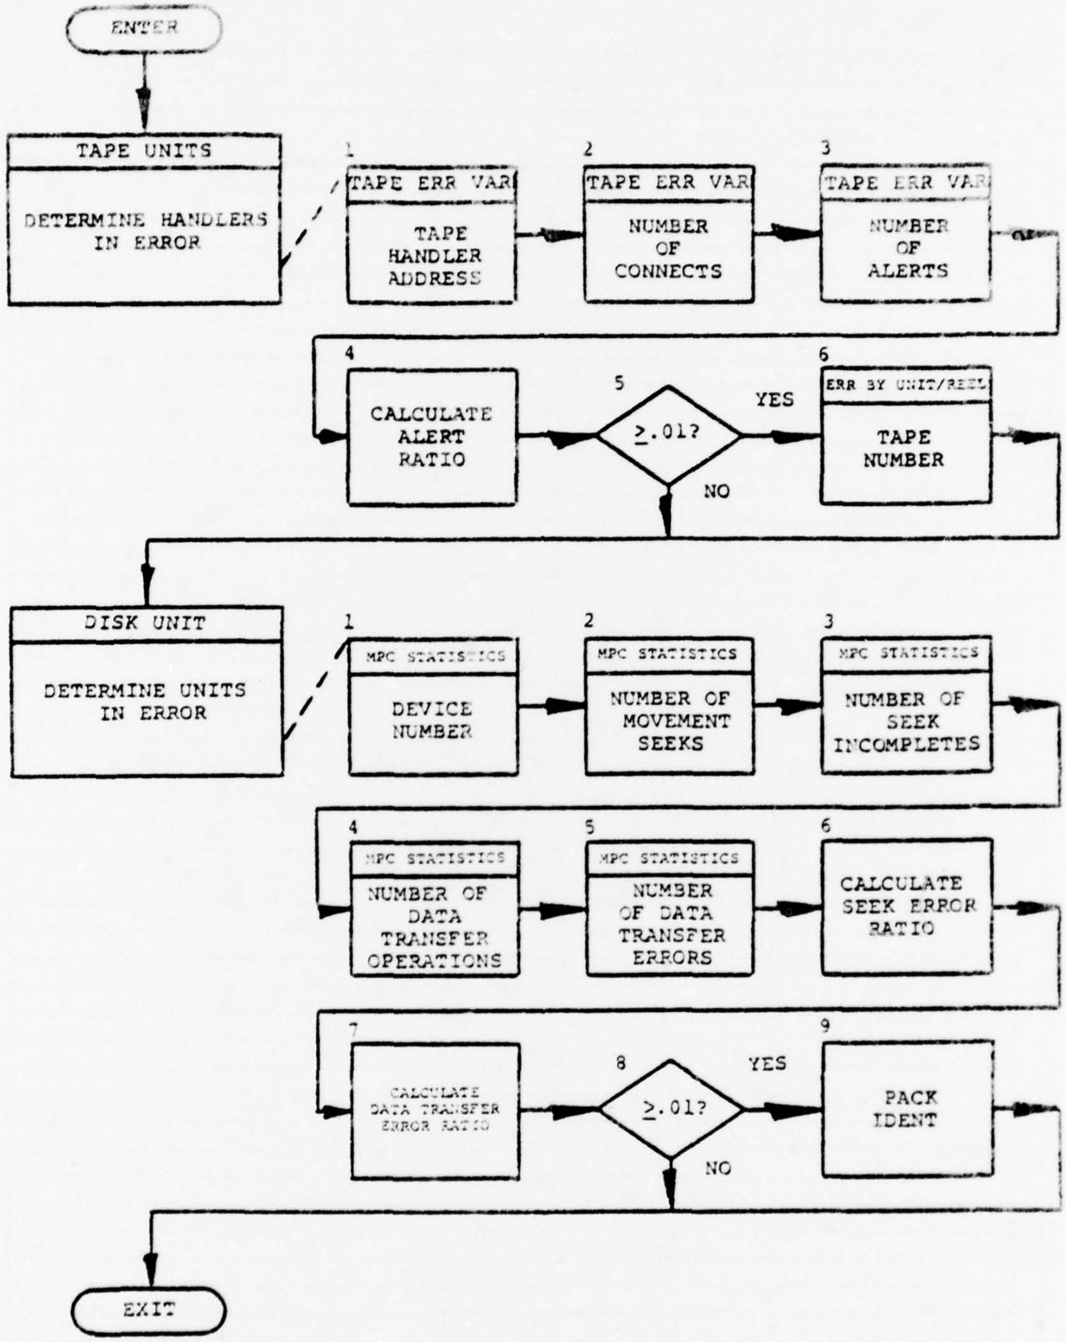

DEVICE ERRORS TEST PROCEDURES

FIGURE VI-1

DEVICE ERRORS	RUN		PAGE	

TAPES

HANDLER	CONNECTS	ALERTS	ALERT RATIO	TAPE NUMBER	✓

DISKS

DEVICE	SEEK MOVEMENT			DATA TRANSFER			PACK IDENT	✓
	SEEKS	ERRORS	RATIO	OPERATIONS	ERRORS	RATIO		

DEVICE ERRORS TEST FORM

FIGURE VI-2

2. Step No. 2: Number of Connects

This step prepares a value for the calculation of the Alert Ratio in Step Number 4.

a. Report Value. This value, listed on the Tape Unit Error Variance Report, is the total number of connects reported on the indicated tape device (including the detection of the Data Alert). The connect value is taken from the left column of the two-column section of the report labeled "Connect Values."

b. Form Entry. Enter the indicated value in the Connects Column of the Tape Section.

3. Step No. 3: Number of Alerts

This step provides a value used to calculate the Alert Ratio in Step Number 4.

a. Report Value. This value, listed on the Tape Unit Error Variance Report, is the total number of Data Alerts for the indicated tape unit. The value is taken from the left column of the two-column section of the report labelled "Alert Values."

b. Form Entry. Enter the indicated value in the Alerts Column of the Tapes Section.

4. Step No. 4: Calculate Alert Ratio

This step calculates an Alert Ratio that is used in the decision in Step Number 5.

a. Calculation. Divide the entered value for Alerts by the entered value for Connects for each of the tape handlers reported.

b. Form Entry. Enter the resulting quotients in the Alert Ratio Column of the Tapes Section.

5. Step No. 5: Decision

If the quotient generated in Step Number 4 is [$\geq$.01], mark (✓) the Tape Unit as a cause of elongation and execute Step Number 6; otherwise, continue the analysis by executing the procedure described in Section VI.C below.

The tape units checked should be scheduled for preventive maintenance unless most of the alerts occurred on the tape reels identified as defective in Step Number 6. Tape errors should continue to be monitored until no tape units are checked. The site may wish to continue monitoring indefinitely.

6. Step No. 6: Tape Number

This step identifies the tape reels on the tape units marked as elongators.

a. Report Value. Tape error numbers are listed by Tape Handler Address on the HEALS II Tape Errors By Unit Reel Numbers Report.

b. Form Entry. List tape error reel numbers in the Tape Number Column of the Tapes Section.

c. Decision. If a few reels are consistently reported as contributing tape alerts, they should be cleaned or their data should be converted to new tapes. Monitoring for reel errors should continue until few bad reels are indicated. The site may want to continue monitoring indefinitely.

C. DETERMINE DISK UNITS IN ERROR

This procedure identifies the disk units that are exhibiting a large number of I/O errors and thereby contributing to elongation of batch turnaround time. The Disks Section of the Device Errors Form is used in the procedure. The MPC Statistics Report of the HEALS II System is employed in the steps of this procedure.

1. Step No. 1: Device Number

This step prepares the Disk Section by identifying those devices configured on the system at the time of the monitor session.

a. Report Values. Each channel and disk device address is displayed on the MPC Statistics Report. The statistics are reported for each Logical Channel Address or Physical Channel Address for a disk device.

b. Form Entry. Enter the addresses for the devices reported (in ascending sequence by IOM/Channel/Device) in the Device Column of the Disks Section.

2. Step No. 2: Number of Movement Seeks

This value is used to calculate the Seek Error Ratio in Step Number 6.

a. Report Value. The value for number of movement seeks, listed on the MPC Statistics Report, represents the number of actuator movements recorded at the MPC Controller for the listed device.

b. Form Entry. Enter the value for each of the devices reported in the corresponding Seeks Column of the Seek Movement Section.

3. Step No. 3: Number of Seek Incompletes

This step prepares a value to be used to calculate the Seek Error Ratio in Step Number 6.

a. Report Value. The value for number of seek incompletes represents the total number of Seek Incompletes received from the indicated device.

b. Form Entry. Enter the number of seek incompletes next to the indicated device address in the Errors Column of the Seek Movement Section.

4. Step No. 4: Number of Data Transfer Operations

This step provides a value to calculate the Data Transfer Error Ratio in Step Number 7.

a. Report Values. Two values are summed to develop the entry for Data Transfer Operations: (1) number of data sectors written and (2) number of data sectors read. The former value is the total number of data sectors transferred to the indicated device; the latter value is the number of data sectors transferred from the device to memory.

b. Calculation. Sum the values for data sectors written and data sectors read for each of the indicated disk devices.

c. Form Entry. Enter the totals in the corresponding Operations Column of the Data Transfer Section.

5. Step No. 5: Number of Data Transfer Errors

This step provides a value that is employed to calculate the Data Transfer Error Ratio for the disk devices.

a. Report Value. The value for data check character alerts, listed on the MPC Statistics Report, represents the number of data errors detected on the indicated disk unit.

b. Form Entry. Enter the number of data transfer errors in the corresponding Errors Column in the Data Transfer Section.

6. Step No. 6: Calculate Seek Error Ratio

This step develops a ratio that is employed in the decision step described in Step Number 8.

a. Calculation. Divide the value for Number of Seek Incompletes by the value for Number of Movement Seeks for each of the disk devices in the configuration.

b. Form Entry. Enter the resulting quotients in the corresponding Ratio Column in the Seek Movement Section.

7. Step No. 7: Calculate Data Transfer Error Ratio

This step calculates the Data Transfer Error Ratio for each of the configured disk units. This ratio is employed in the decision in Step Number 8.

a. Calculation. Divide the values for Number of Data Transfer Errors by the values for Number of Data Transfer Operations for each of the configured disk units.

b. Form Entry. Enter the resulting quotients in the corresponding Ratio Column of the Data Transfer Section.

8. Step No. 8: Decision

If a quotient resulting from either of the above calculations is [$\geq$.01], the unit should be marked (✓) as a bottleneck. If it is checked, increase the effectiveness of preventive maintenance being performed on the disk device. An exception would occur if a particular disk pack was the major contributor to recorded errors. If this is the case, the disk pack should be cleaned or copied to a new pack. Error monitoring should continue until no units are checked.

If the quotients resulting from the above calaculations are [<.01], exit the procedure because no particular disk unit has been confirmed as an elongator of turnaround time. Return to the section that called for this test.

9. Step No. 9: Pack Identification

This step identifies the packs processed on the disk devices checked as bottlenecks.

a. Report Value. The Pack Serial Number is listed by the GCOS Media Message when it is assigned to each of the devices to which it is attached during the day. The format of the message is:

MEDIA -- S#SSSSS MNT i-cc-dd #nnnnn(filename) @tt.ttt

where,

S#SSSSS	--	SNUMB
MNT	--	Mount Action Indicator
i-cc-dd	--	Unit Address
#nnnnn	--	Pack Serial Number
@tt.ttt	--	Time Message Listed

Identify the message as coming from a disk mount by noting the Unit Address for Disks.

b. Console Log Scan. Scan the Console Log for the measurement period to identify pack serial numbers for those devices that have been checked as bottlenecks.

c. Operations Interview. This step identifies the serial numbers for the permanently mounted packs on the system during the measurement period. Determine from the site shift supervisor the numbers of packs that were permanently mounted on the disk units checked as bottlenecks.

d. Form Entry. Enter the pack serial numbers identified by Console Log Scan or by Operations Interview in the Pack Identification Column of the Disks Section.

e. Tuning Step. Disk packs whose pack serial numbers have been entered should be scheduled for copying to new packs unless maintenance on the disk drive removes the error indication.

VII. PATHWAY UTILIZATION TEST

This section describes procedures for analyzing turnaround time elongation caused by I/O Pathway Utilization. These procedures use the Mass Store Monitor (MSM) and the Channel Monitor (PSM) as data collection tools.

A. ANALYSIS SUMMARY

The Pathway Utilization Test uses direct measurement of the following types of disk I/O queues: (1) device queues, (2) channel queues, and (3) queues for both channel and device. Channel pathway queues for tapes are currently not measured directly. Unit record devices are not investigated with the test.

Pathway Utilization Test procedures include: (1) Analyze Tape Channel Queues, (not implemented) (2) Analyze Disk Device Queues, (3) Analyze Disk Channel Queues, (4) ANalyze Disk Device and Channel Queues, (5) Identify High-Use Extents, and (6) Identify Files. This section describes each procedure step(see Figure VII-1). The Pathway Utilization Test Form (see Figure VII-2) is provided to assist in data collection. No other Turnaround Time Analysis Tests are referenced by this test.

MSM reports used in the Pathway Utilization Test include: (1) Device Space Utilization Report, (2) System File Use Summary, and (3) File Summary.

A decision step (see Figure VII-1) is the first step to be executed in this test. If the pathway to be investigated is a disk channel, execute the procedure in Section VII.C.

B. ANALYZE TAPE CHANNEL QUEUES (Currently Not Implemented)

This procedure is used to determine if frequent tape channel I/O queues exist. The Tape Channel Queue Section of the Pathway Utilization Form is referenced by the steps in this procedure.

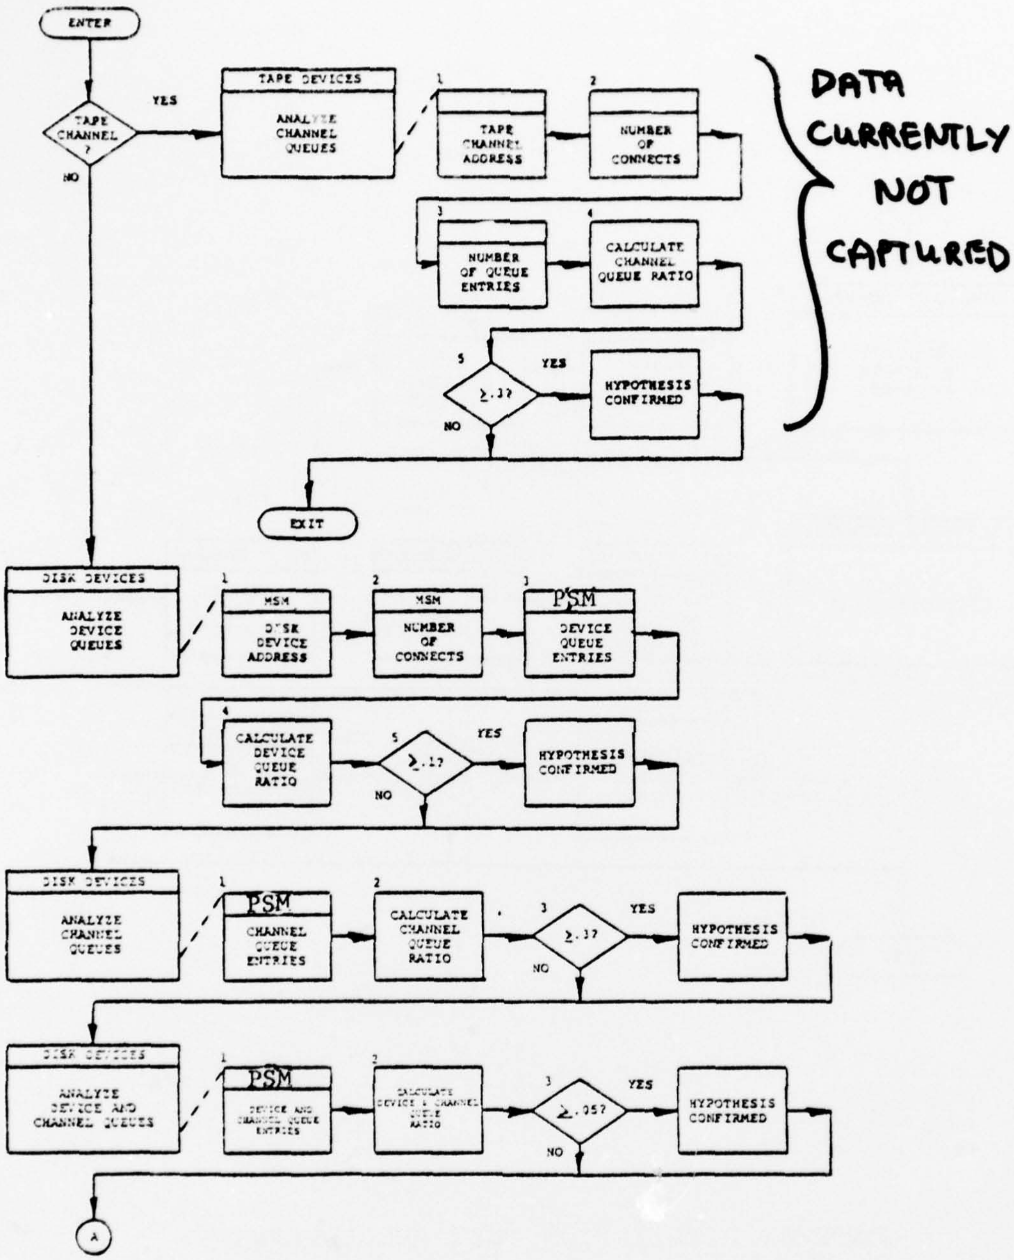

PATHWAY UTILIZATION TEST PROCEDURES

FIGURE VII-1

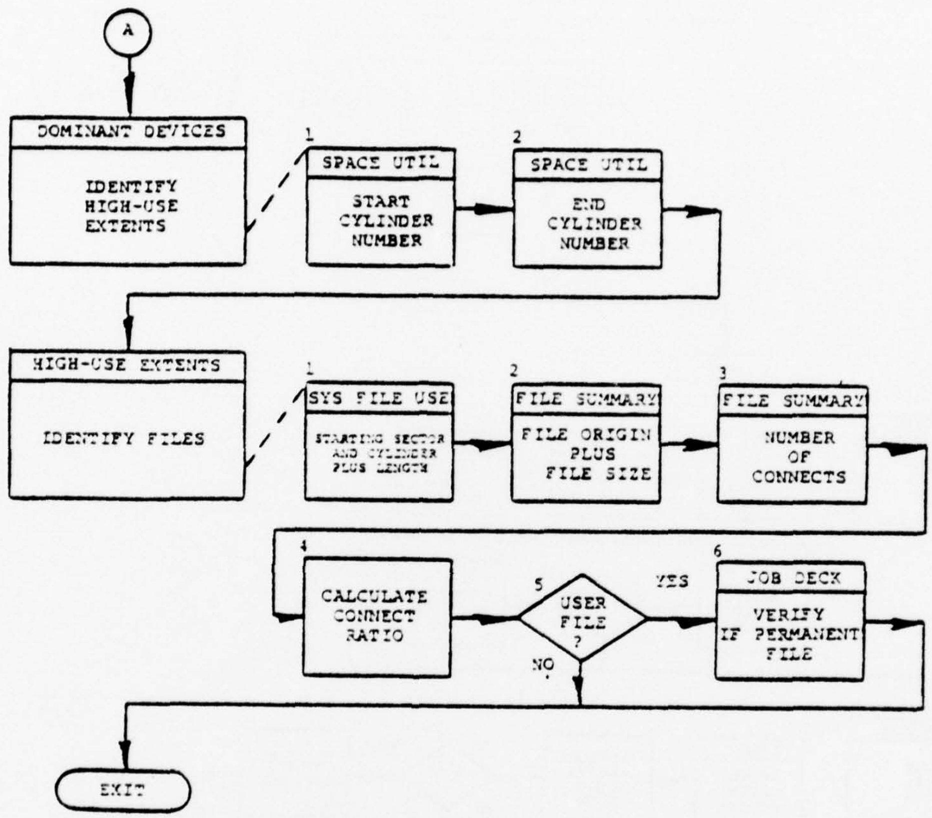

PATHWAY UTILIZATION TEST PROCEDURES

FIGURE VII-1 (Cont'd)

PATHWAY UTILIZATION	RUN		PAGE	

	ADDRESS		NUMBER OF CONNECTS	QUEUE		NOTE	✓
	IOM	CHNL		ENTRIES	RATIO		
TAPE CHANNEL QUEUE							

	ADDRESS			NUMBER OF CONNECTS	QUEUE DETERMINATION						NOTE	✓
					DEVICE		CHANNEL		DEVICE & CHNL			
	IOM	CHNL	DEV		NO.	RATIO	NO.	RATIO	NO.	RATIO		
DISK CHANNEL AND/OR DEVICE QUEUE												

	START	END	NAME	T	FC	CONN	✓	NAME	T	FC	CONN	✓	NAME	T	FC	CONN	✓
ICOD																	
SPACE UTILIZATION																	

	START	END	NAME	T	FC	CONN	✓	NAME	T	FC	CONN	✓	NAME	T	FC	CONN	✓
ICOD																	
SPACE UTILIZATION																	

PATHWAY UTILIZATION TEST FORM

FIGURE VII-2

1. Step No. 1: Tape Channel Address

This step is used to identify the tape channels in use.

a. Report Value. The tape channel addresses are listed on the MSM Tape Channel Queues Report.

b. Form Entry. Enter the tape channel addresses (IOM/Channel) in the Address Column.

2. Step No. 2: Number of Connects

This value is used to calculate the Tape Channel Queue Ratio in Step Number 4.

a. Report Value. This value, listed on the MSM Tape Channel Queues Report, is the number of I/O connects made to each tape channel in the configuration.

b. Form Entry. Enter the value in the Number of Connects Column of the Tape Channel Queue Section.

3. Step No. 3: Number of Entries Queued

This value is used in Step Number 4 to calculate the system's Tape Channel Queue Ratio.

a. Report Value. This value, listed on the MSM Tape Channel Queues Report, represents a count of the number of connects that had to be queued by IOS because the tape channel pathway was busy at the time of an I/O service request.

b. Form Entry. Enter the number in the Queue Entries Column for each channel.

4. Step No. 4: Calculation

This step is used to calculate the system's Tape Channel Queue Ratio.

a. Calculation. Divide the value for Tape Channel Queue Entries by the value for Number of Connects for each tape channel.

b. Form Entry. Enter the quotients in the corresponding Queue Ratio Columns.

5. Step No. 5: Decision

If a Tape Channel Queue Ratio is [>.3], consider the hypothesis confirmed and the pathway bottlenecked at the channel. Place a checkmark (✓) in the proper column. Repeat this procedure several times to ensure that the channel is consistently bottlenecked. Consider the hypothesis not confirmed if a Channel Queue Ratio is [<.3].

C. ANALYZE DISK DEVICE QUEUES

This procedure determines if frequent I/O service queues exist at the system's disk devices. The Disk Channel And/Or Device Queue Section of the Pathway Utilization Form is referenced by the steps in this procedure.

1. Step No. 1: Disk Device Address

This step identifies the disk devices that were allocated and in use on the system.

a. Report Value. The disk device addresses are listed on the Physical Device, Device IO Correlation Table of the MSM or PSM.

b. Form Entry. Enter the disk device addresses (IOM/Channel/Device) in the Address Column. Leave a blank row above the first device on each primary channel.

2. Step No. 2: Number of Connects

This value is used to calculate queue ratios in Step Number 4.

a. Report Value. The Device Utilization By Device Number Histogram (Report 1) of the MSM or PSM lists the number of I/O connects made to each device in the configuration. The number of connects is found under the column labeled "indiv number". The correlation between the device IO and the actual device is made by using the Physical Device, Device IO Correlation Table Report.

b. Form Entry. Enter the number in the Number of Connects Column of the Channel And/Or Device Queue Section.

3. Step No. 3: Device Queue Entries

This value is used in Step No. 4 to calculate the system's Device Queue Ratio..

a. Report Value. This value, listed on the PSM Channel Free Device Busy Report represents the number of connects that had to be queued by IOS because the device (and not the channel pathway) was busy at the time of I/O service request.

b. Form Entry. Enter the count for each device in the Device Number Column.

4. Step No. 4: Calculation

This step is used to calculate the System's Device Queue Ratio.

a. Calculation. Divide the Queue Determination Device Number value by the Number of Connects value for each disk in the configuration.

b. Form Entry. Enter the quotients in the Device Ratio Column.

5. Step No. 5: Decision

Consider the pathway bottlenecked at the device if a Device Queue Ratio is [≥.1]. Place a checkmark (✓) in the proper column next to each such device. Repeat this procedure several times to determine whether these devices are consistently bottlenecked. The bottlenecked devices may switch from day to day if the bottlenecking is due to temporary files. Consider the hypothesis not confirmed if a Device Queue Ratio is [≤.1].

D. ANALYZE DISK CHANNEL QUEUES

This procedure is used to determine if frequent I/O service queues exist at the system's disk channels. The Disk Channel And/Or Device Queue Section of the Pathway Utilization Form is referenced by the steps in this procedure.

1. Step No. 1: Channel Queue Entries

This value is used in Step No. 3 to calculate the system's Channel Queue Ratio.

a. Report Value. This value, listed for each primary channel on the PSM Channel Busy Device Free Report, represents a count of the number of connects that had to be queued by IOS because the channel pathway (and not the device) was busy at the time of I/O service request.

b. Form Entry. Enter the count for each channel in the Queue Determination Channel Number Column. Use the empty row next to that channel's first device.

2. Step No. 2: Number of Connects

This value is used in Step No. 3 to calculate the Channel Queue Ratio.

a. Report Value. In order to determine the number of I/O connects made across a channel configuration (i.e., the primary channel and all logical channels) it will be necessary to use the Channel Configuration REport of the MSM/PSM. For each primary channel configured the report lists all channels which are crossbarred to that channel. In addition, the number of connects to each channel (primary & logical is also given). The analyst should total all connects made to the primary channels and all channels which are crossbarred to it.

b. Form Entry. Enter the count in the Number of Connects Column for each channel.

3. Step No. 3: Calculation

This step is used to calculate the system's Channel Queue Ratio.

a. Calculation. Divide the values for Queue Determination Channel Number by the values for Number of Connects for each disk channel.

b. Form Entry. Enter the quotients in the corresponding Channel Ratio Columns.

4. Step No. 4: Decision

If a Channel Queue Ration is [≥.3], consider the hypothesis confirmed and the pathway bottlenecked at the channel, but not at the device. Place a checkmark (✓) in the proper column. Repeat this procedure several times to ensure that the channel is consistently bottlenecked. If a Channel Queue Ratio is [<.3], consider the hypothesis not confirmed.

E. ANALYZE DISK DEVICE AND CHANNEL QUEUES

This procedure determines if I/O service queues exist at the system's disk channels and devices. The Disk Channel And/Or Device Queue Section of the Pathway Utilization Form is referenced by the steps in this procedure.

1. Step No. 1: Device And Channel Queue Entries

This value is used in Step No. 2 to calculate the system's Device And Channel Queue Ratio.

a. Report Value. This value, listed on the PSM Channel, Device Busy Report, represents the number of connects resulting from I/O service requests that had to be queued by IOS because both the channel pathway and device were busy when the I/O request was made.

b. Form Entry. Enter the count for each channel in the Queue Determination Device And Channel Number Column.

2. Step No. 2: Calculation

This step is used to calculate the System's Device And Channel Queue Ratio.

a. Calculation. Divide the value for Device And Channel Queue Entries by the value for Number of Connects entered for each disk channel.

b. Form Entry. Enter the quotients in the corresponding Device And Channel Ratio Columns.

3. Step No. 3: Decision

If a Device and Channel Queue Ratio is [≥.05], consider the hypothesis confirmed and the pathway bottlenecked at both the device and channel. Place a checkmark (✓) next to each device on the effected channel(s). If a Device and Channel Queue Ratio is [<.05], consider the hypothesis not confirmed.

F. TUNING STEPS

This paragraph discusses tuning steps to be taken when hypotheses are confirmed by the Pathways Utilization Test.

1. Disk Tuning

Specific tuning steps are recommended depending upon which queue hypotheses are confirmed. Table VII-1 presents the decision logic that matches confirmations with specific turning actions.

a. Rule No. 1: No Device And/Or Channel Queue Confirmation. This condition describes a system that exhibits no bottlenecks at the disk device or channel. If the system consistently exhibits this condition, no constraint exists.

	RULE							
	1	2	3	4	5	6	7	8
Device Queue Confirmed	N	N	N	N	Y	Y	Y	Y
Channel Queue Confirmed	N	N	Y	Y	N	N	Y	Y
Channel & Device Queue Confirmed	N	Y	N	Y	N	Y	N	Y
Exit Procedure	X							
Condition Not Possible		X		X		X		
Separate Files Across Devices					X		X	
Separate Files Across Channels			X				X	
Separate Files Across Devices And Channels								X
Add Pathway Capacity			X				X	X
Check For Potential Seek Elongation					X		X	X

NOTE: N = No
Y = Yes
X = Take indicated action

PATHWAY UTILIZATION TEST
DECISION TABLE

TABLE VII-1

b. Rule No. 2: Channel And Device Queue Confirmation Only. This condition cannot exist. If it occurs, closely examine the bracketed decision values used and/or obtain professional help. Change the decision values if warranted.

c. Rule No. 3: Channel Queue Confirmation Only. This condition describes a system that is bottlenecked at the disk channel, but not at the device.

(1) Scheduling Solution. If possible, schedule the workload so that concurrently-running jobs spread their accesses equally across separate channel groupings.

(2) Parameter Solutions. These steps involve changes to GCOS system functions.

(a) Separate highly-used GCOS libraries and files across disk channels. Determine which libraries are highly-used by conducting the procedures described in Section VII.G and Section VII.H.

(b) Determine which GCOS Slave Service Area (SSA) modules should be assembled in GCOS Hard Core to eliminate both the delay and the channel capacity required to bring the modules into core. Conduct the procedure in Section VII.H.1.d to determine which modules to move.

(c) Expand I/O pathway capacity (if possible) by adding crossbar accessing routes via the $XBAR Card in the GCOS Startup Deck.

(3) Programming Solutions. These steps propose changes to jobs, programs, or files. Separate highly active user files across disk channels. Perform the procedures in Section VII.G and Section VII.H to determine which files are highly-used.

(a) Catalog active, permanent user files on disk packs that are accessed through separate channels. This can be done by creating new files on the proper devices, transferring the data, and renaming the files after the old files have been purged.

(b) Have the users change the $FILE cards for highly active user temporary files to specify devices on different channels (thereby separating the files). As an alternative, change highly active user temporary files to permanent files by cataloging enough space to house the largest temporary file. Next, follow the suggestion in paragraph (3) above for permanent file separation across channels.

(4) Sizing Solutions. These steps involve changes to the disk subsystem configuration. Propose these changes only after the above remedies have been attempted.

(a) If possible, increase the speed of the disk subsystem being used. Though the queue is not exhibited at the device, the speed of device service will affect the rate at which the channel is queued.

(b) If possible, increase the number of physical I/O channel paths to the devices. This may involve adding IOM's, MPC's and channels.

d. Rule No. 4: Channel And Channel/Device Queue Confirmation. This condition cannot exist. If it occurs, revise the bracketed decision values used and/or obtain professional help.

e. Rule No. 5: Device Queue Confirmation Only. This condition describes a system that is bottlenecked at the disk device but not at the channel.

(1) Scheduling Solution. If possible, schedule the workload so that concurrently running jobs do not access highly used files on the same disk devices.

(2) Parameter Solutions. These steps involve changing GCOS system functions.

(a) Separate highly-used GCOS libraries and files across disk devices. Execute the procedures in Section VII.G and Section VII.H to determine which libraries are highly-used.

(b) Execute the Seek Elongation Test (see Section IV) to confirm seek elongation on the devices checked.

(c) Determine which GCOS SSA modules should be assembled into GCOS Hard Core to eliminate both the delay and device time required to bring the modules into core. Conduct the procedure in Section VII.H.1.d to determine which modules to move.

(3) Programming Solutions. Separate highly used user files across disk devices. Execute the procedures in Sections VII.G and VII.H to determine which files are highly used. Catalog active, permanent, user files on separate disk packs. Have the user change the $FILE cards for highly active temporary files to specify relatively idle disk packs. As an alternative, change highly active user temporary files to permanent files by cataloging enough space to house the largest temporary file. Next, follow the suggestion above for permanent user file separation across devices.

(4) Sizing Solutions. These steps involve changes to the disk subsystem configuration. Propose these solutions only after the above changes have been attempted.

(a) If possible, increase the speed of the disk devices being used. The reduced latency time exhibited by the faster device should reduce disk queuing.

(b) If possible, increase the number of disk devices attached to the existing MPC. This may involve adding IOM's, MPC's, or channels. Catalog files across devices.

f. Rule No. 6: Device And Channel/Device Queue Confirmation. This condition cannot exist. If it does, revise the braketed decision values used and/or obtain professional help.

g. Rule No. 7: Device And Channel Queue Confirmation. This condition describes a system that is bottlenecked at both the disk device and channel. There is no confirmation, however, of queuing for both the device and channel at the same time.

(1) Scheduling. If possible, schedule the workload so that concurrently running jobs do not access files on the same disk subsystem. Also, schedule some jobs that make high use of disk files to other time periods.

(2) Parameter Solutions. These steps involve changes to GCOS system functions.

(a) Separate highly used GCOS libraries and files across disk subsystems. Execute the procedures in Section VII.G and Section VII.H to determine which libraries are highly used.

(b) Execute the Seek Elongation Test (see Section IV) to confirm seek elongation on the devices and channels checked.

(c) Determine which GCOS SSA modules should be assembled into GCOS Hard Core to eliminate the delay and channel/device capacity required to bring the modules into core. Conduct the procedure in Section VII.H.1.d to determine which modules to move.

(3) Programming Solutions. Separate highly used user files across subsystems. Execute the procedures in Sections VII.G and VII.H to determine which files are highly used. Catalog active, permanent, user files on separate disk subsystems. Have the users change the $FILE cards for the highly active temporary files to access devices on separate subsystems. As an alternate, change highly active user temporary files to permanent files by cataloging enough space to house the largest temporary file. Next, follow the suggestion above for permanent file separation across subsystems.

(4) Sizing Solution. This step is recommended only after the above remedies have been attempted. If possible, propose an increase in the number of disk subsystems configured. Recatalog user and system files/libraries on the new configuration.

h. Rule No. 8: Device, Channel, and Channel/Device Queue Confirmation. This condition describes a system that is simultaneously bottlenecked at both the disk device and the channel. Execute the tuning solutions described above for Rule Number 7.

2. Tape Tuning

High queuing on a tape device is a cause of delay only for the program that is using the tape. Other jobs waiting to run with tapes will be delayed because of too few tape devices (see Insufficient Devices Test in Section IX).

When high tape device queuing is confirmed, examine the job/activity use of tape files. If disk device/channel usage is not bottlenecked, it may be possible to transfer selected tape files to disk. Other approaches to change the program include: (1) an increased blocking factor for files, (2) a reduction and/or elimination of files or record types, (3) changes to file format, (4) changes to application design, etc.

High queuing on a tape channel is a cause of delay for all programs that use the queued channel(s).

a. Scheduling Solutions. If possible, schedule the workload so that concurrently run jobs do not access tapes via the same heavily queued channels.

b. Sizing Solution. This step is recommended only after the above remedy has been attempted. If possible, propose an increase in the number of channels attached to the tape subsystem.

G. IDENTIFY HIGH-USE EXTENTS

This procedure is used to identify each active extent on a bottlenecked disk device or channel in preparation for moving certain files to alleviate channel/device queuing. If all channels are bottlenecked, this procedure is optional. The MSM Space Utilization Report is referenced in this procedure.

The Space Utilization Section of the Pathway Utilization Form is used for data entry. Identify each device checked as a result of executing the procedures in Section VII.B through D. Enter the device address in the space marked "ICCDD" in the Space Utilization Section.

1. Step No. 1: Start Cylinder Number

This entry identifies the starting addresses of high-use extents on the bottlenecked disk units.

a. Report Values. For each bottlenecked device and each device on a bottlenecked channel, note the start of each extent that has [≥5%] Individual Probability. Adjacent extents, each of which has over [5%] Individual Probability, can be considered as one extent.

b. Form Entries. Enter the starting address (i.e., cylinder number) of each high-use extent into the Start Column.

2. Step No. 2: End Cylinder Number

This entry identifies the ending addresses of high-use extents on the bottlenecked disk units. Enter the ending address of each high-use extent in the End Column. Several clusters of activity may be listed with multiple Start-End entries.

H. IDENTIFY FILES

This procedure identifies heavily used libraries and other files on the high use extents.

The procedure references the remaining columns of the Space Utilization Section.

1. System Libraries

This step is used to identify GCOS system libraries that are recorded within the identified high-use extents.

a. Report Value. The MSM System File Use Summary Report lists GCOS libraries that were initialized at System Startup time. The report also lists the device, starting cylinder number, and length for each GCOS system library. Consult the GCOS Startup File Map to determine the specific names of the GCOS system libraries. (For example,SYSTEM FILE 2 is the second GCOS system file defined in the Startup Deck.)

b. Calculation. For each listed GCOS system library, determine whether any part of the library falls within a high-use extent identified above. First, compare the starting address of the library (right side of Starting Sector/Cylinder Column) to the starting and ending addresses of each high-use extent on the same device. If the library does not start within any high-use extent, determine whether it overlaps a high-use extent on that device by adding the length to the starting address to find the end of the file. Note that the length is given in sectors only and must be converted to cylinders to be added to the starting address. The bottom row of Table IV-1 lists the necessary conversion factors (i.e., divide the length in sectors by the appropriate number of sectors per cylinder to obtain the length in cylinders).

c. Form Entry. Identify each GCOS system library that starts in or overlaps a high-use extent in the same row as the starting and ending addresses of the extent. Enter the library name (from the Startup File Map) in the Name Column. Enter an "S" (System) in the Type ("T") Column, a dash (-) in the File Code ("FC") Column (no file code), and the number of accesses (from the Accesses Column of the System File Use Summary Report) in the Connects Column.

d. Individual Module Activity. The option of assembling highly used GCOS system code into GCOS Hard Core can be attempted if the appropriate modules can be identified. This step, using the MSM Individual Module Activity Report, identifies these active modules.

(1) Report Value. The MSM Individual Module Activity Report lists the activity on the GCOS libraries with the following column headings: (1) System File, (2) Module Name, (3) Type, (4) IOM-PUB-Device, (5) Sector In File, (6) Number of Accesses, and (7) Percent of Activity.

(2) Report Scan. Scan the MSM Individual Module Activity Report for code modules that exhibit [>10] Percent of Activity. If a module is consistently active, recommend that it be put into GCOS Hard Core. Note that if the module is not re-entrant, a core-to-core move will be made on each call. This still saves the disk I/O involved, but does not necessarily save any CPU time.

2. User Files

This step is used to identify the user files and jobs that are causing the queuing detected above in Sections VII.C through VII.E.

a. Report Value. The MSM File Summary Report lists the following data for each user file: (1) File Code, (2) Number of Connects, (3) File Size, (4) Allocated Device, and (5) File Origin. Each file is listed under the activity that used it, including the Number of Connects that the activity issued to it. Using the same technique described in 1, identify the files that start in or overlap identified high-use extents. Ignore files that have few [<1000] connects.

Two unusual file codes will be frequently observed: "00" and "--". The "00" file code identifies all disk accesses made without a Peripheral Allocation Table (PAT) entry (e.g., accesses made by GCOS as part of job initialization). It also includes all SSA module calls and SSA pushdown operations. The "--" file code is constructed by the MSM data reduction program for all accesses made to a SYSOUT file.

b. Form Entry. For files that fall within a high-use extent, enter the $IDENT Name in the Name Column on the same row as the starting address of the extent. Enter a "U" (User) in the Type Column. Enter the listed file code in the File Code Column and the number of connects in the Connects Column.

3. Calculate Connect Ratio

This step is used to calculate a Connect Ratio for each identified file.

a. Calculations. Add the number of connects for each file on each device being investigated. Divide the number of connects for each file on the device by the total for that device to develop a Connect Ratio for each file.

b. Form Entry. Enter a "✓" for all files with a Connect Ratio [≥.1].

4. Decision

All files with a Connect Ratio [≥.1] are candidates for relocation as high-use files. Refer to the Decision Table rules for instructions on how and where to relocate these files.

VIII. CPU EXECUTION CHARACTERISTICS

This section describes the procedures for determining the system's CPU execution characteristics and for identifying selected programs for detailed execution analysis. The Memory Utilization Monitor (MUM) and the GSEP Accounting Data Reduction System are used to obtain data for this analysis. A hardware monitor will be used if the optional Execution Activity Map is desired.

A. ANALYSIS SUMMARY

The CPU Execution Characteristics Test determines the general utilization level of the processor and then determines if the CPU is dominated by GCOS or user program execution. Two additional test procedures identify and investigate (i.e., create program execution maps of) selected user programs.

Procedures executed under the CPU Execution Characteristics Test (charted in Figure VIII-1) include: (1) Determine CPU Utilization, (2) Determine Dominant CPU User, (3) Isolate Jobs, and (4) Develop Execution Activity Map.

Reports used in the CPU Execution Characteristics Test include the MUM CPU Utilization Report and the GSEP Allocator/Termination Report.

The CPU Execution Characteristics Form (see Figure VIII-2) is used to assist in the analysis of CPU usage.

B. DETERMINE CPU UTILIZATION

This procedure determines the CPU utilization level for the configuration. A decision in Step Number 5 determines if the remaining CPU Execution Characteristics Test procedures are to be conducted. The CPU Utilization Section of the CPU Execution Characteristics form is used for this procedure.

1. Step No. 1: Monitor Session Elapsed Time

This step develops a value which is used to calculate Total CPU Utilization in Step Number 3.

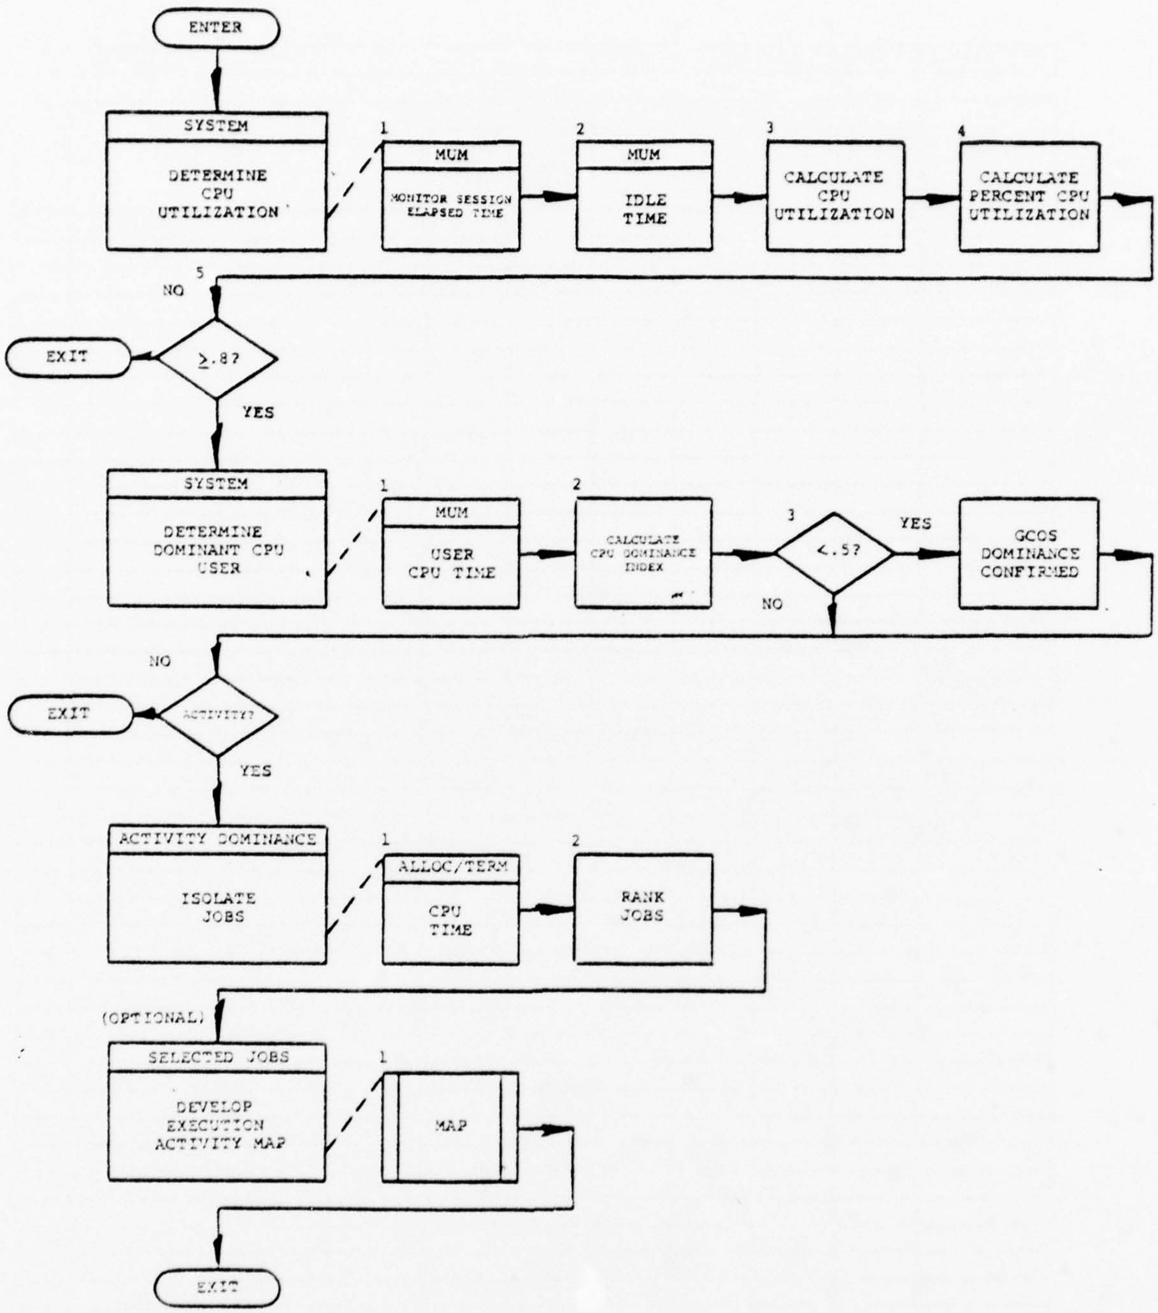

CPU EXECUTION CHARACTERISTICS TEST PROCEDURES

FIGURE VIII-1

CPU EXECUTION CHARACTERISTICS	RUN		PAGE	

TIME	CPU UTILIZATION				DOMINANT CPU USER		NOTE	✓
	MONITOR SESSION E/T	IDLE TIME	CPU UTIL TIME	PCT CPU UTIL	USER CPU TIME	CPU DOM INDEX		

CPU EXECUTION CHARACTERISTICS TEST FORM

FIGURE VIII-2

a. Report Value. The Elapsed Time Thus Far value, listed on the MUM CPU Utilization Report, represents the total elapsed time (in seconds) of the monitor session. Use the CPU Utilization Report with the largest value. If the period monitored was broken into several pieces by lost data conditions, add the pieces together.

b. Calculation. Multiply the value for elapsed time by the number of processors if more than one processor is used in the configuration being monitored. For example, if three processors are configured, multiply the elapsed time value by 3.

c. Form Entry. Enter the resulting number in the Monitor Session E/T Column of the CPU Utilization Section.

2. Step No. 2: Idle Time

This step generates a value for the calculation of CPU Utilization described in Step Number 3.

a. Report Value. Idle time is reported in the MUM CPU Utilization Report in 100th second units for each of the up to four processors of a configuration. Note that the value is reported in 100th second units and must be converted to seconds for use in Step Number 3.

b. Calculation. Total the values for the Idle Time of each processor in the configuration. Use the same CPU Utilization Report(s) used in Step Number 1 above.

c. Form Entry. Enter the sum (in seconds) in the Idle Time Column of the CPU Utilization Section.

3. Step No. 3: Calculate CPU Utilization

This step develops a value for CPU Utilization that is used in Step Number 4.

a. Calculation. Subtract the value for Idle Time from the value for Monitor Session Elapsed Time. The difference represents CPU Utilization for the configuration.

b. Form Entry. Enter the difference in the CPU Utilization Time Column of the CPU Utilization Section.

4. Step No. 4: Calculate Percent CPU Utilization

This step develops a value for Percent CPU Utilization that is used in Step Number 5.

a. Calculation. Divide the value for CPU utilization by the value for Monitor Session Elapsed Time.

b. Form Entry. Enter the resulting quotient in the Percent CPU Utilization Column of the CPU Utilization Section.

5. Step No. 5: Decision

Consider the processor as a potential bottleneck if the Percent CPU Utilization value is [≥.8]. Repeat this procedure several times to insure that the Percent CPU Utilization value is consistently [≥.8]. If Percent CPU Utilization is [<.8], consider the processor as not being a bottleneck. Exit the test (return to the section that called for this test) if the processor is not a bottleneck.

C. DETERMINE DOMINANT CPU USER

This procedure determines whether CPU execution on the measured system is dominated by GCOS code or by user code. Use the Dominant CPU User Section of the CPU Execution Characteristics Form for this procedure.

1. Step No. 1: User CPU Time

This step generates a value that is used to calculate a CPU Dominance Index in Step Number 2.

a. Report Value. The CPU Time Used Thus Far is displayed on the MUM CPU Utilization Report (in 100ths of a second) for each of a series of system and user programs. The value represents the CPU Time used by each user program during the monitor session.

b. Form Entry. Enter the value for User CPU Time (convert to seconds) in the User CPU Time Column of the Dominant CPU User Section.

2. Step No. 2: Calculate CPU Dominance Index

This step generates an index value used in the decision in Step No. 3.

a. Calculation. Divide the value for User CPU Time by the value for CPU Utilization Time.

b. Form Entry. Enter the resulting quotient in the CPU Dominance Index Column of the Dominant CPU User Section.

3. Step No. 3: Decision

User Dominance of the CPU is indicated by a CPU Dominance Index [≥50%]. GCOS dominance of the CPU is indicated if the value is [<50%].

If GCOS dominance of the CPU is indicated, the analyst should investigate means of reducing GCOS CPU utilization without seriously degrading service provided to user programs. Suggestion for this investigation are proposed in Section VIII.D.

D. TUNING STEPS

Three batch turnaround time elongation hypotheses are confirmed by the CPU Execution Characteristics Test: (1) High CPU Utilization, (2) GCOS Dominance of CPU Utilization, and (3) User Dominance of CPU Utilization. This section discusses steps to be taken to determine whether these conditions can be removed or reduced.

Take the steps proposed below for High CPU Utilization before those proposed for GCOS tuning or user program tuning.

1. High CPU Utilization

This indicator describes a system that exhibits [≥80%] CPU Utilization for the measurement period.

a. Scheduling Solutions. These steps involve changes to workload scheduling.

(1) If there are periods during the operating day when this condition is not indicated, and if the operating schedule permits, schedule CPU-dominant jobs to run during these periods. The CPU/IO

"dominance" of a job can be determined from an examination of the job's activity entries on the GSEP Allocator/Termination Report.

(2) Schedule I/O-dominant jobs during the high CPU periods to complement the CPU dominance of the workload (see also the use of GCOS I/O Priority Dispatching under Parameter Solutions below).

(3) Determine if a particular mix of urgency codes in effect during the high CPU periods is the source of the condition. Schedule the workload so that high urgency code jobs that are CPU-dominant are not run together. Execute the Urgency Codes Test (see Section XII) to determine if CPU-dominant jobs are executing with high urgency code values. Suggest that these jobs be run with urgency code values that are low enough to permit a balance of CPU and I/O dispatching. Use the GSEP Allocator/Termination Report to determine which jobs are CPU- and I/O-dominant. Try running without the Urgency Thruput Dispatcher Option. Note that this may tend to elongate the overall turnaround time of these jobs; determine whether this elongation is within site limits.

(4) Determine if an "unbalanced" use of GCOS Priority B Dispatching has been scheduled. Interview operations to determine which jobs (up to three permitted) have been granted the use of this dispatching option. It could be that too many CPU-dominant jobs are executing with the option enabled.

b. Parameter Solutions. This step involves changes to GCOS system functions.

(1) Recommend running with the GCOS I/O Priority Dispatcher Option enabled. This change will speed I/O-dominant jobs through execution.

(2) Interview field engineering to determine whether H-6000 memory is configured to run fully interleaved. Propose that the field engineer make the appropriate switch adjustments if it is not.

c. Sizing Solutions. These approaches involve changes to the processor configuration. Recommend these solutions only after those proposed above and those contained in Sections VIII.D.2 and VIII.D.3 below have been attempted. Note that these changes should be supported by benchmark testing to determine the exact number and type of processors to add.

(1) Recommend an increase in the number of processors configured at the current processor type (i.e., H-6060 or H-6080).

(2) Recommend an upgrade to the next processor type (i.e., H-6060 upgraded to H-6080).

2. GCOS Dominance of CPU Utilization

This indicator describes a system that exhibits [>80%] CPU Utilization and [>50%] of the CPU time is comprised of GCOS code execution. Because GCOS serves all jobs, solutions to GCOS dominance can be investigated examining both user programs and GCOS.

a. Program-Focused Solution. Examine job and activity use of GCOS functions. Determine how to reduce the number of passes required through particular GCOS service code. As an example, reduce a program's number of passes through IOS by increasing its I/O blocking factor. This general concept of reducing the need to execute code by changes in formats or design can apply to all user programs.

b. GCOS-Focused Solution. This is the most complex solution. Analyze the use of GCOS System Startup parameters to determine the functional areas within GCOS that are enabled at the site and might warrant investigation. Apply the program mapping procedure described in Section VIII.F to GCOS Hard Core and the GCOS Privileged Slave Programs to determine where in the code they spend the most CPU time. Examine these high-activity areas for potential parameter changes to reduce GCOS CPU utilization.

3. User Dominance of CPU Utilization

This indicator describes a system that exhibits [>80%] CPU Utilization and [≥50%] of the CPU time is comprised of user code execution. Analyze the programs for work that did not have to be done to accomplish the process defined for the program. Execute the selection and optional mapping procedures below under Sections VIII.E and VIII.F.

E. ISOLATE JOBS

This procedure provides a means of identifying the jobs and activities that are to be mapped in the procedure in Section VIII.F.

The GSEP Accounting Data Reduction System is used for this procedure, though other techniques or metrics (e.g., job elapsed time) can be used to select jobs for execution mapping.

1. Step No. 1: CPU Time Isolation

This step identifies user jobs or activities that exceed a CPU Time value defined by the analyst.

a. CPU Time Measure. CPU Active Time is displayed on the GSEP Allocator/Termination Report for each activity.

b. Establishing Isolation Threshold. Use a value of [300 seconds] as a lower limit. All jobs appearing in the GSEP Allocator/Termination Report with a CPU time of greater than [300 seconds] should be examined.

2. Step No. 2: Rank Jobs

Regardless of the isolation technique, the jobs or activities should be further examined in order to determine if they are candidates for mapping. Simply rank the [top ten] jobs or activities (i.e., those with the largest CPU Active Time) as candidates for mapping. Eliminate the jobs or activities that are run infrequently.

F. DEVELOP EXECUTION ACTIVITY MAP

This procedure produces a frequency distribution of code execution of each selected job or activity to determine what code in the program uses the most CPU time. Examination by a programmer may result in coding changes that reduce the use of the CPU. This procedure addresses the development of the program map; interpretation of the map will identify the areas of code that should be examined for changes to reduce CPU usage.

This procedure uses a hardware monitor to gather instruction execution data. Data is to be collected under a controlled experiment with the program forced into a fixed H-6000 memory area during its execution (i.e., neither swapped nor moved).

1. Monitor Setup

Detailed instructions for setting up the hardware monitor should be obtained from CCTC. (C702 Pentagon)

2. Special Considerations

The program must not move from its place in memory. Usually, the site will conduct the test while the system is idle. To guarantee that the program will not move whether the system is idle or not, patch the program's .STATE word using the PEEK, PATCH, and WRITE console verbs. Set bit 26 of this word to a 1 after the activity being investigated has begun. Verify using the STATS console verb that the program does not move during the test.

Usually probes will be attached to one processor only. Any program execution done by other processors will go unreported. Therefore the site must insure that only the processor with the probes attached executes the program. The last four bits of the .STATE word in the program's SSA afford one way to accomplish this. If the probes are on processor 0, setting bits 32, 33, and 34 will restrict program execution to this processor. If the probes are on processor 1, set bits 32, 33, and 35. Three of the bits are set and the remaining one determines which processor executes the program.

Both these considerations involve changing the .STATE word during execution. Since there is a possibility the system will crash, this test should not be attempted while critical work is going on.

To obtain the proper resolution, the hardware monitor must usually be set up to map only a certain segment of address space (i.e., memory). For this reason, the part of memory to be allocated to the job must usually be known before the experiment starts. One way to accomplish this is to run the job twice in an empty system--once to see where the Core Allocator will put it and once to actually collect the data.

3. Experiment Setup

Conduct this setup for each program being mapped.

a. Predict Area of Memory. Ascertain into what area of memory the program will be loaded. Execute the program to accomplish this, if necessary.

b. Monitor Setup. Set the hardware monitor software or plugboard logic to map the smallest area that completely includes the program.

c. Start Program. Start the program and wait until it reaches the activity to be mapped.

d. Start Monitor. Start collecting data once the program is loaded in the proper activity.

e. Set Bits. Set bit 26 and the proper three bits of bits 32-35 in the .STATE word. (Find the program's Lower Address Limit [LAL] and use the console PATCH and WRITE verbs.) Note that a PEEK must be done first so that the values of the other bits in the .STATE word are not changed.

f. Monitor Execution. Continue to collect data until the activity terminates. Then stop the monitor immediately. Check (using the STATS console verb) throughout the experiment to make sure the program remains in memory in the same place.

g. Instruction Activity Map. Print a map of the program's instruction activity by using the hardware monitor's data reduction software.

h. Instruction Activity. Examine the instruction activity map for areas in the program that exhibit [≥10%] of program activity. Direct a programmer to examine these areas of code for potential optimization.

IX. INSUFFICIENT DEVICES TEST

This section describes the test for batch turnaround time elongation caused by configuration of too few tapes or not enough disk space. The procedures of this test use the Memory Utilization Monitor (MUM), and the tape monitor for data collection.

A. ANALYSIS SUMMARY

The Insufficient Devices Test involves different analysis techniques for the two device types investigated. Tapes are analyzed by first examining Memory Wait Time to determine that Core Allocation is not the bottleneck and then analyzing measured tape delays. Disk space is analyzed by measuring the number of links refused a job each time a request is made for space.

Procedure steps (see Figure IX-1) executed under this test include: (1) Determine If Memory Wait Time Is High, (2) Determine If Tape-Caused Delay Is High, and (3) Determine If Disk Space-Caused Delay Is High. Each procedure is described in this section. This procedure uses the Insufficient Devices Test Form (see Figure IX-2) to assist in data collection. No other Turnaround Time Analysis Tests are referenced by this test.

MUM reports used in this test include: (1) Total Elapsed Time An Activity Was In Memory, (2) Elapsed Wait Time for Memory Requests, and (3) Tape Delay Report from the tape monitor. A fourth report, the MUM Disk Space Refusal Report, does not currently exist, but is hypothesized in order to complete the test description.

The first step (see Figure IX-1) in this test is a decision step. If disk space is to be investigated, the first procedure to execute is in Section IX-D. If tapes are to be investigated, start with the procedure in Section IX.B.

B. DETERMINE IF MEMORY WAIT TIME IS HIGH

This procedure determines if the system's Memory Wait Time is high. The setting of a software damper switch (i.e., the PALC-CALC Damper Switch) between the GCOS Peripheral Allocator (where tapes are assigned) and the Core Allocator (closer to the Activity Execution Process) could cause a

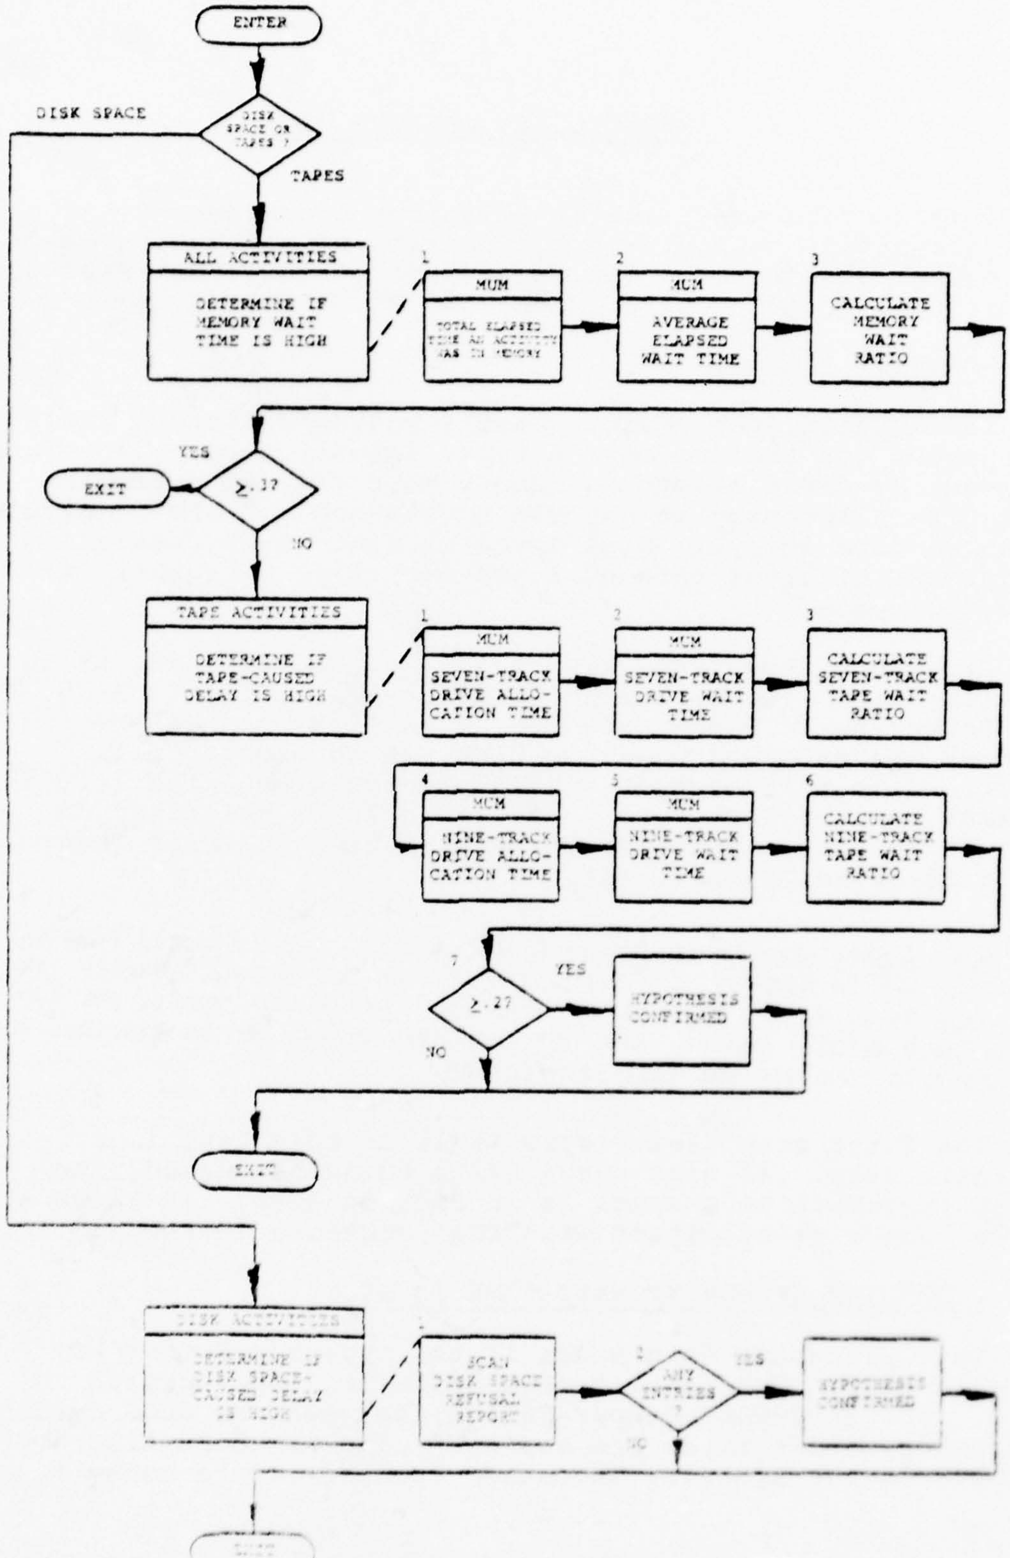

INSUFFICIENT DEVICES TEST PROCEDURES

FIGURE IX-1

INSUFFICIENT DEVICES	PAGE		RUN	

TIME	TOTAL ELAPSED TIME	AVERAGE ELAPSED WAIT TIME	MEMORY WAIT RATIO	SEVEN-TRACK DRIVES			NINE-TRACK DRIVES			✓
				ALLOC TIME	WAIT TIME	WAIT RATIO	ALLOC TIME	WAIT TIME	WAIT RATIO	

INSUFFICIENT DEVICES TEST FORM

FIGURE IX-2

delay that might be incorrectly attributed to a lack of available tapes. This procedure confirms that the delay is caused by memory requests.

This procedure uses the Insufficient Devices Test Form.

1. Step No. 1: Total Elapsed Time An Activity Was In Memory

This value is used in Step Number 3 to calculate the system's Memory Wait Ratio.

a. Report Value. Choose the Representative Value from the MUM Total Elapsed Time An Activity Was In Memory Report. Section II.G.2 gives instructions for choosing Representative Values.

b. Form Entry. Enter the Representative Value in the Total Elapsed Time Column.

2. Step No. 2: Average Elapsed Wait Time

This value is used in Step Number 3 to calculate the system's Memory Wait Ratio.

a. Report Value. Choose the Representative Value from the MUM Elapsed Wait Time For Memory Requests Report. Section II.G.2 gives instructions for choosing Representative Values.

b. Form Entry. Enter the Representative Value in the Average Elapsed Wait Time Column.

3. Step No. 3: Calculation

This step computes the system's Memory Wait Ratio for the decision in Step Number 4.

a. Calculation. Divide the value for Average Elapsed Wait Time by the value for Total Elapsed Time An Activity Was In Memory.

b. Form Entry. Enter the quotient in the Memory Wait Ratio Column.

4. Step No. 4: Decision

Consider the PALC-CALC Damper Switch elongation hypothesis confirmed if the Memory Wait Ratio [>.3]. This means that the Peripheral Allocation delay may be caused by a memory constraint. Investigate this hypothesis and remove its cause before continuing the tape-caused delay analysis. Execute the Memory Constraint Test (see Section V) to analyze memory demands and then continue with Section IX.C.

Consider the PALC-CALC Damper Switch elongation hypothesis not confirmed if the Memory Wait Ratio [<.3].

In either case, repeat the above procedure several times to insure that the decision consistently turns out the same way. If not, consider changing the decision value and/or obtaining professional help.

C. DETERMINE IF TAPE-CAUSED DELAY IS HIGH

This procedure determines if jobs are delayed due to wait for tape allocation.

1. Step No. 1: Seven-Track Drive Allocation Time

This value is used in Step Number 3 to compute the Seven-Track Tape Wait Ratio.

a. Report Value. The Time of Allocation for Seven-Track Drives, listed (in seconds) on the last page of the MUM Tape Delay Report, represents the total allocation time for seven-track tape drives.

b. Form Entry. Enter the value in the Allocation Time Column of the Seven-Track Drives Section.

2. Step No. 2: Seven-Track Drive Wait Time

This value is used in Step Number 3 to compute the Seven-Track Tape Wait Ratio.

a. Report Value. The Total Seven-Track Wait Time, listed (in seconds) on the last page of the MUM Tape Delay Report, represents program elongation time caused by seven-track tape allocation.

b. Form Entry. Enter the value in the Wait Time Column of the Seven-Track Drives Section.

3. Step No. 3: Calculate Seven-Track Tape Wait Ratio

This ratio is used in Step Number 7 to determine if jobs are elongated because of seven-track tape allocation.

a. Calculation. Divide the Seven-Track Drive Wait Time by the Seven-Track Drive Allocation Time.

b. Form Entry. Enter the quotient in the Wait Ratio Column of the Seven-Track Drives Section.

4. Step No. 4: Nine-Track Drive Allocation Time

This value is used in Step Number 6 to compute the Nine-Track Tape Wait Ratio.

a. Report Value. The Time of Allocation for Nine-Track Drives, listed (in seconds) on the last page of the MUM Tape Delay Report, represents the total allocation time for nine-track tape drives.

b. Form Entry. Enter the value in the Allocation Time Column of the Nine-Track Drives Section.

5. Step No. 5: Nine-Track Drive Wait Time

This value is used in Step Number 6 to compute the Nine-Track Tape Wait Ratio.

a. Report Value. The Total Nine-Track Wait Time, listed (in seconds) on the last page of the MUM Tape Delay Report, represents program elongation time caused by nine-track tape allocation.

b. Form Entry. Enter the value in the Wait Time Column of the Nine-Track Drives Section.

6. Step No. 6: Calculate Nine-Track Tape Wait Ratio

This ratio is used in Step Number 7 to determine if jobs are elongated because of nine-track tape allocation.

a. Calculation. Divide the Nine-Track Drive Wait Time by the Nine-Track Drive Allocation Time.

b. Form Entry. Enter the quotient in the Wait Ratio Column of the Nine-Track Drives Section.

7. Step No. 7: Decision

Consider the hypothesis of tape-caused delay confirmed if either the Seven-Track Tape Wait Ratio or the Nine-Track Tape Wait Ratio is [$\geq$.2]. Repeat the above procedure several times to insure that the hypothesis is consistently confirmed/denied. Check (✓) the entry and implement the tuning steps described under Section IX.E for tape-caused elongation.

Consider the hypothesis of tape-caused delay not confirmed if either tape wait ratio is [<.2]. In this case, return to the section that called for this test.

D. DETERMINE IF DISK SPACE-CAUSED DELAY IS HIGH

This procedure determines if jobs are being elongated because of the need to wait for disk space allocation. This delay occurs both before the jobs enter execution (i.e., at the processes listed in Table IX-1) and during the Execution Process. The execution delay can occur when SYSOUT space is exhausted by the jobs. Control of SYSOUT space is maintained by SYSOUT and not by the Peripheral Allocator.

The data used in this procedure are collected by GCOS Trace as entry number 53_8 (Mass Store Link Refusal). A report (Disk Space Refusal Report) lists each occurrence of a Mass Store Link Refusal from these trace entries. The report contains the following column headings: (1) Time of Occurrence, (2) Program Number, (3) Number of Llinks Requested, (4) Number of Llinks Remaining, and (5) Type (SYSOUT or other). The Insufficient Device Form is not used in this procedure.

1. Step No. 1: Scan Disk Space Refusal Report

Scan the Disk Space Refusal Report for entries listed during the period being investigated. Insufficient disk space can be so serious that even one refusal may be significant. Determine the source of the refusal (i.e., whether for SYSOUT space or for any other type).

2. Step No. 2: Decision

Consider the hypothesis confirmed if [any] entries appear on the report. Initiate the tuning solutions described below in Section IX.E.

If [no] entries appear on the report, the hypothesis is not confirmed and the test can be exited (return to the section that called for this test). Several days should probably be monitored to insure that entries never appear.

E. TUNING STEPS

This paragraph discusses tuning steps to take when delay is confirmed by the Insufficient Devices Test.

1. Insufficient Tapes Tuning Solutions

Apply the tuning steps proposed below when tape-caused delay is confirmed.

a. Parameter Solution. Scan the GCOS Console Log to determine if the total number of available tape units was reduced by operator action (i.e., RLSE Command). Verify that the release was required.

b. Scheduling Solutions. These steps involve changes in the workload schedule.

(1) Propose scheduling fewer tape jobs during the period measured in order to balance the use of tapes.

(2) Execute the Urgency Codes Test (see Section XII) to examine the urgency codes of the jobs executing during the measurement period. Look for jobs with high urgency codes that request a large number of tapes. If possible, schedule these jobs so that they do not conflict or reduce their urgency.

c. Programming Solutions. These solutions change jobs, programs, or files.

(1) Propose investigation of programs to determine if tape files that do not contain a large number of records can be transferred to temporary or permanent disk.

(2) Verify that the delayed programs require all of the requested tapes. If not, propose examination to reduce the number of drives allocated either to all runs of the job or to selected runs (because of a month-end file, for example) when all tapes are not needed.

d. Sizing. Propose an increase in the number of tape handlers configured on the system.

2. Insufficient Disk Space Tuning Solutions

Apply the tuning steps proposed below when disk space-caused delay has been confirmed. These steps are grouped into two areas that represent the source of disk-caused delays: (1) requests for SYSOUT space and (2) requests for other disk space.

a. SYSOUT Space Solutions. SYSOUT manages its own space allocation. Apply these steps to correct a lack of SYSOUT space.

(1) Parameter Solution. Increase the amount of disk space allocated to the SYSOUT files. This may impact other of the site's disk space requirements. SYSOUT file space allocation is made in the GCOS Startup Deck on the $FILDEF Card. All existing $FILDEF Card parameters for SYSOUT files must have the same size value specified.

(2) Scheduling Solution. If possible, schedule jobs so that a known conflict in SYSOUT space requirement does not develop. Verify (by consulting the GESEP Allocator/Termination Report's SYSOUT Record Count fields) that the concurrent scheduling of jobs with high SYSOUT requirements is not causing this conflict.

(3) Sizing Solution. This solution should be proposed only after the "Other Space Solutions" below are tried. If possible, propose an increase in the amount of disk space (i.e., available spindles) in the installation.

b. Other Space Solutions. Apply these steps to shortages of other disk space. The Scheduling, Programming, and/or the cold boot Paramater Solutions will probably be all that is needed to solve the problem of occasional refusals to very large requests.

(1) Scheduling Solutions. If possible, schedule jobs so that a known conflict in disk space requirements does not develop. Verify (by consulting the GESEP Allocator/Termination Report's SYSOUT Record Count fields) that the concurrent scheduling of jobs with large disk files is not causing this conflict.

(2) Programming Solutions. These steps apply to jobs, programs, or files.

(a) Propose examination of all permanent disk files cataloged to determine if they can be purged, copied to backup media, or put on removable packs.

(b) Propose examining programs to determine if they actually required the full amount of disk space requested. Suggest copying files to tape, if appropriate.

(3) Parameter Solutions.

(a) A cold boot of the system will consolidate disk space and may help alleviate disk space refusals, especially if it has been a month or more since the last cold boot.

(b) A reduction in the number of disk spindles configured as removable will add disk space. If the disk packs on one or more removable disk spindles are rarely changed, consider changing the spindles to permanent (fixed). The files on removable packs that are nearly always mounted could be changed to be permanently online. Then fewer removable spindles would be needed and the empty space on those packs would be available for system and user files.

(4) Sizing Solution. If possible, propose an increase in the amount of disk space (i.e., available disk spindles) in the configuration.

X. "FEW ACTIVITIES IN SYSTEM" TEST

This section describes the procedures for analyzing the operating periods when there are few activities in the system. These procedures use the Memory Utilization Monitor (MUM) to collect data.

A. ANALYSIS SUMMARY

The "Few Activities In System" test analyzes memory wait time to determine whether additional activities or jobs could have been scheduled in the system during the period being investigated. Figure X-1 charts the procedure steps executed under this test. The "Few Activities In System" Form (see Figure X-2) is provided with this procedure to assist in data collection. No other turnaround time analysis tests are referenced by this test.

The following MUM reports are used in the analysis: (1) Total Elapsed Time An Activity Was In Memory and (2) Elapsed Wait Time for Memory Requests.

This test is entered from the Turnaround Time Model Scan Procedure when the Activity Execution Process has been chosen for further investigation and the "Few Activities In System" hypothesis is to be tested.

B. DETERMINE IF MEMORY WAIT TIME IS HIGH

This procedure determines if Memory Wait Time is high in order to determine the type of solution to be proposed.

1. Step No. 1: Total Elapsed Time An Activity Was In Memory

This value is used as the divisor in calculating the Memory Wait Ratio in Step Number 3.

a. Report Values. Determine the Representative Value for the frequency distribution in the MUM Total Elapsed Time An Activity Was In Memory Report.

b. Representative Values. The general rules for determining Representative Values of frequency distributions are presented in Section II.G.2.

c. Form Entry. Enter the Representative Value in the Total Elapsed Time An Activity Was In Memory Column.

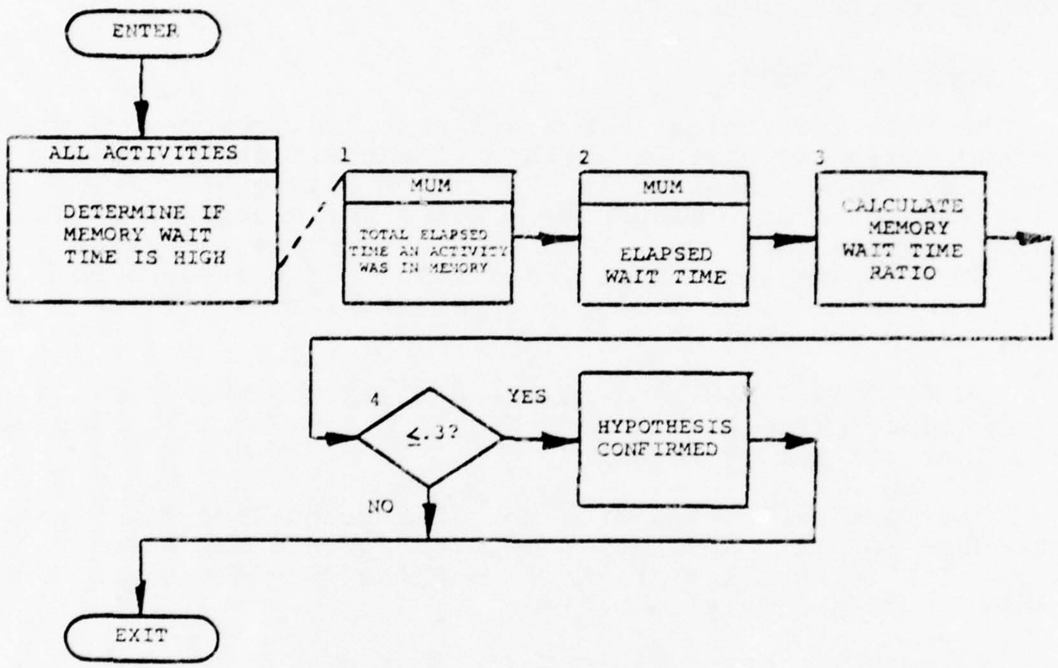

"FEW ACTIVITIES IN SYSTEM" TEST PROCEDURE

FIGURE X-1

"FEW ACTIVITIES IN SYSTEM"	RUN		PAGE	

TIME	TOTAL S/T AN ACTIVITY WAS IN MEMORY	AVERAGE ELAPSED WAIT TIME	MEMORY WAIT RATIO		TIME	TOTAL S/T AN ACTIVITY WAS IN MEMORY	AVERAGE ELAPSED WAIT TIME	MEMORY WAIT RATIO	

"FEW ACTIVITIES IN SYSTEM" TEST FORM

FIGURE X-2

2. Step No. 2: Average Elapsed Wait Time

This value is used as the dividend when calculating the Memory Wait Ratio in Step Number 3.

a. Report Value. Determine the Representative Value for the distribution in the MUM Elapsed Wait Time for Memory Requests Report. Rules for determining Representative Values are presented in Section II.G.2.

b. Form Entry. Enter the Representative Value in the Average Elapsed Wait Time Column.

3. Step No. 3: Calculate Memory Wait Time Ratio

The ratio is calculated from the entries made in the above two steps.

a. Computation. Divide the value for Elapsed Wait Time by the value for Total Elapsed Time An Activity Was In Memory.

b. Form Entry. Enter the quotient value in the Memory Wait Ratio Column. The resulting value represents the ratio of the elapsed time that an activity waited for memory to the elapsed time it used the memory.

4. Step No. 4: Decision

Consider the hypothesis confirmed if the Memory Wait Ratio is [<.3]. Repeat the above procedure several times to insure that the hypothesis is consistently confirmed. Follow the tuning steps proposed in Section X.C.

Consider the hypothesis not confirmed and the addition of more jobs to the system not the solution if the Memory Wait Ratio is [>.3]. Return to the section that called for this test.

C. TUNING STEPS

The hypothesis confirmed by the procedure in Section X.B was that more jobs (i.e., one or more) could have been scheduled in the system during the period undergoing analysis. These tuning steps assume that there was a backlog of jobs waiting to enter execution during the period. Otherwise, turnaround time would probably not have been a problem. If no jobs were backlogged, the originally specified site turnaround time requirement may be unrealistic.

1. Parameter Solution

a. System Scheduler Parameters. Examine the parameters for the System Scheduler on the $SSFILE Cards (see the GCOS Startup Listing) including: (1) Maximum Number Of Programs To Be Scheduled Concurrently, (2) Maximum Number Of Programs To Be Scheduled Concurrently Before Invoking Class Restriction, and (3) Maximum Number Of Programs To Be Scheduled Concurrently From Each Class. Adjust the entry value upwards if any entry appears low or constraining to jobs entering execution.

b. Sieve Parameters. Scan the GCOS Console Log for occurrences of the Sieve Message output during the period. This message indicates that the job demands of the listed SNUMB exceed the process time, core, or limit Sieve specified at the time of entry. The message is repeated every five minutes. If these messages occur frequently, consider raising the Sieve Parameters so that these jobs are not delayed.

2. Scheduling Solution

Try to schedule more jobs into the system during the period. The selection of jobs with a particular mix of resource requirements is beyond the scope of this Guide.

3. Memory Size Reduction

Consider reducing the size of memory if the Memory Wait Time Ratio is small enough. For example, a memory module (128K words) could be released from the system without reducing the number of jobs in memory if more than 128K is not used or if the CPU or I/O channels are bottlenecked without the last 128K of memory. Note that processor speed may be affected by the degree of interleaving possible with the new memory size. It may be increased (more interleaving possible) or decreased (less interleaving possible). Run the system with the new memory size and interleaving switch settings for a period to determine whether turnaround time is seriously impacted.

XI. IOS DELAYS TEST

This section describes the procedures for analyzing elongation hypotheses due to GCOS Input/Output Supervisor (IOS) delays. These procedures use the Turnaround Time Analysis System for data collection.

A. ANALYSIS SUMMARY

This test uses direct measurement of IOS delay time via GCOS Trace entry data that are collected and reduced by the Turnaround Time Analysis System. This test will probably be conducted infrequently at WWMCCS sites because IOS delay time should remain fixed for each program's use of a particular IOS device code module. The solutions proposed for this condition emphasize reducing the number of IOS requests made by particular programs.

Figure XI-1 charts the procedure steps executed under this test, including: (1) Determine If Tape I/O Service Causes Delay, (2) Determine If Disk I/O Service Causes Delay, and (3) Determine If Unit Record I/O Service Causes Delay. The IOS Delays Test Form (see Figure XI-2) is provided with this procedure to act as an assist in data collection. No other turnaround time analysis tests are referenced by this test.

The Activity Execution Model Report produced by the Turnaround Time Analysis System is the only report used in this test.

B. DETERMINE IF TAPE I/O SERVICE CAUSES DELAY

This procedure determines if tape I/O service contributes to batch turnaround time elongation.

Use the top section of the IOS Delays Form for this procedure. Note from Figure XI-2 that entries can be ordered on the form by time period.

1. Step No. 1: I/O Process Time--Tape Sum

This value is used in Step Number 3 to calculate a ratio of Tape I/O Service Time to Tape I/O Process Time.

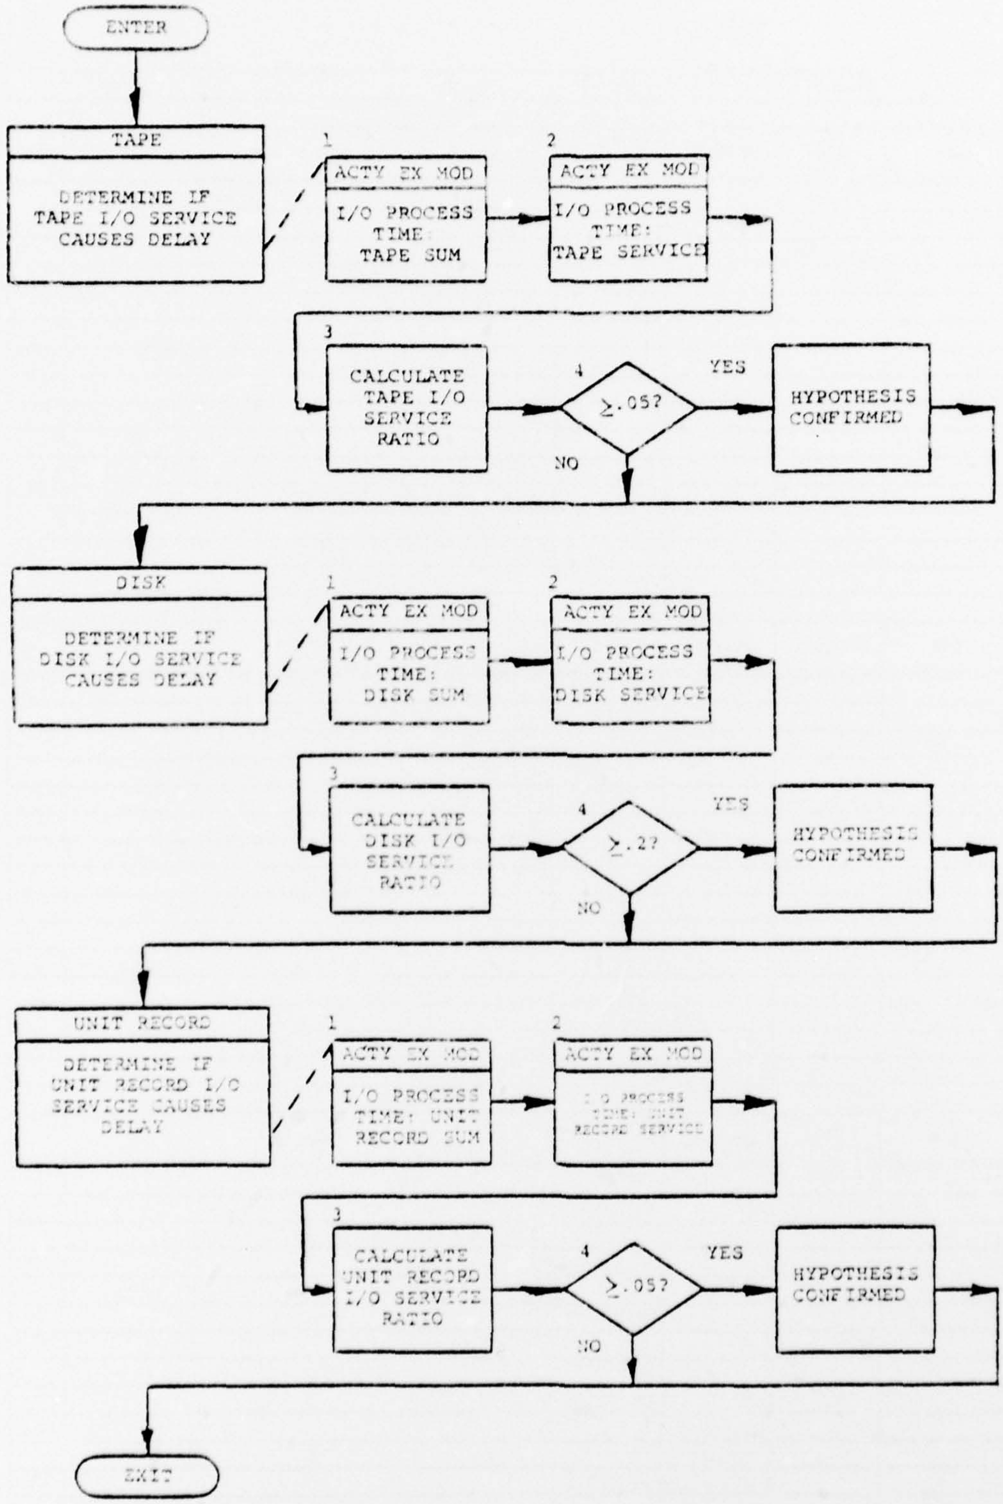

IOS DELAYS TEST PROCEDURES

FIGURE XI-1

IOS DELAYS	RUN		PAGE	

TAPE I/O SERVICE	TIME	SUM	SERVICE	SERVICE RATIO	N O T E	✓

DISK I/O SERVICE	TIME	SUM	SERVICE	SERVICE RATIO	N O T E	✓

UNIT RECORD (U/R) I/O SERVICE	TIME	SUM	SERVICE	SERVICE RATIO	N O T E	✓

IOS DELAYS TEST FORM

FIGURE XI-2

a. Report Value. The I/O Process Time--Tape Sum, listed on the Activity Execution Model Report, includes tape I/O operations occurring both inside and outside of the processor.

b. Form Entry. Enter the value in the Sum Column of the Tape I/O Service Section.

2. Step No. 2: I/O Process Time--Tape Service

This value is used in Step Number 3 to calculate a ratio of Tape I/O Service Time to Tape I/O Process Time.

a. Report Value. The Tape Service Elapsed Time, displayed on the Activity Execution Model Report, includes all in-processor tape operations, but excludes device queue residence time.

b. Form Entry. Enter the value in the Service Column of the Tape I/O Service Section.

3. Step No. 3: Calculate Tape I/O Service Ratio

This step calculates a Tape I/O Service Ratio. The ratio is used in the decision in Step Number 4.

a. Calculation. Divide the value for Tape I/O Service Time by the value for Tape I/O Process Time Sum.

b. Form Entry. Enter the quotient in the Service Ratio Column of the Tape I/O Service Section.

4. Step No. 4: Decision

Consider this hypothesis confirmed if the Tape I/O Service Ratio is [$\geq$.05]. Consider it not confirmed if the ratio [$<$.05]. Repeat the above procedure several times to insure that the hypothesis is consistently confirmed/denied.

If the hypothesis is confirmed and if the Disk I/O Service Ratio is low (relative to the Tape I/O Service Ratio), investigate the possible transfer of selected tape files to disk. Examine other approaches to reducing the number of passes through the Input/Output Supervisor by application programs, including: (1) increased blocking factors, (2) reduction and/or elimination of files or of record types, (3) changes to file format, and (4) change of application design. If these solutions do not help, discuss the condition with CCTC to attempt a solution through negotiation with the WWMCCS Contractor.

C. DETERMINE IF DISK I/O SERVICE CAUSES DELAY

This procedure determines if disk I/O service contributes to batch turnaround time elongation.

Use the center section of the IOS Delays Form for this procedure. Note from Figure XI-2 that entries can be ordered on the form by time period.

1. Step No. 1: I/O Process Time--Disk Sum

This value is used in Step Number 3 to calculate a ratio of Disk I/O Service Time to Disk I/O Process Time.

a. Report Value. The I/O Process Time--IAS Sum, displayed on the Activity Execution Model Report, is the Disk I/O Process Time Sum. The value includes disk I/O operations occurring both inside and outside of the processor.

b. Form Entry. Enter the value in the Sum Column of the Disk I/O Service Section.

2. Step No. 2: I/O Process Time--Disk Service

This value is used in Step Number 3 to calculate a ratio of Disk I/O Service Time to Disk I/O Process Time.

a. Report Value. The IAS Service Elapsed Time, listed on the Activity Execution Model Report, is the Disk I/O Service Time. The value includes all in-processor disk operations, but excludes device queue residence time.

b. Form Entry. Enter the value in the Service Column of the Disk I/O Service Section.

3. Step No. 3: Calculate Disk I/O Service Ratio

This step calculates the Disk I/O Service Ratio. The ratio is used in the decision in Step Number 4.

a. Calculation. Divide the value for Disk I/O Service Time by the value for Disk I/O Sum.

b. Form Entry. Enter the quotient in the Service Ratio Column of the Disk I/O Service Section.

4. Step No. 4: Decision

Consider the hypothesis confirmed if the Disk I/O Service Ratio is [≥.2]. Consider it not confirmed if the ratio is [<.2]. Repeat the above procedure several times to insure that the hypothesis is consistently confirmed/denied.

If the hypothesis is confirmed, examine programs that were executing during the test period to determine if any of the following solutions might be applied to reduce the number of passes through IOS: (1) increased blocking factors, (2) reduction and/or elimination of files or record types, (3) changes of file format, (4) conversion of disk records to tape format if appropriate, or (5) change of application design. If these solutions do not help, discuss the condition with CCTC to attempt a solution through negotiation with the WWMCCS Contractor.

D. DETERMINE IF UNIT RECORD I/O SERVICE CAUSES DELAY

This procedure determines if unit record I/O service contributes to batch turnaround time elongation.

Use the bottom section of the IOS Delays Test Form for this procedure. Note from Figure XI-2 that entries can be ordered on the form by time period.

1. Step No. 1: I/O Process Time--Unit Record Sum

This value is used in Step Number 3 to calculate a ratio of Unit Record I/O Service Time to Unit Record I/O Process Time.

a. Report Value. The I/O Process Time--Unit Record Sum, listed on the Activity Execution Model Report, includes all unit record I/O operations occurring both inside and outside of the processor.

b. Form Entry. Enter the value in the Sum Column of the Unit Record (U/R) I/O Service Section.

2. Step No. 2: I/O Process Time--Unit Record Service

This value is used in Step Number 3 to calculate a ratio of Unit Record I/O Service Time to Unit Record I/O Process Time.

a. Report Value. The Unit Record Service Elapsed Time, listed on the Activity Execution Model Report, includes all in-processor unit record operations, but excludes unit record queue residence time.

b. Form Entry. Enter the value in the Service Column of the Unit Record (U/R) I/O Service Section.

3. Step No. 3: Calculate Unit Record I/O Service Ratio

This step calculates the Unit Record I/O Service Ratio. The ratio is used in the decision in Step Number 4.

a. Calculation. Divided the value for Unit Record I/O Service Time by the value for Unit Record I/O Process Time Sum.

b. Form Entry. Enter the quotient in the Service Ratio Column of the Unit Record (U/R) I/O Service Section.

4. Step No. 4: Decision

Consider the hypothesis confirmed if the Unit Record I/O Service Ratio is [$\geq$.05]. Consider it not confirmed if the ratio is [<.05]. Repeat the above procedure several times to insure that the hypothesis is consistently confirmed/denied.

If the hypothesis is confirmed, examine application programs running at the time of experiment to determine if the requirement for reports can be reduced. If this solution does not help, discuss the condition with CCTC to attempt a solution through negotiation with the WWMCCS Contractor.

XII. URGENCY CODES TEST

This section describes procedures for analyzing batch turnaround time elongation attributed to job urgency code mix. These procedures use the GSEP Accounting Data Reduction System and the GCOS Console Log for data collection.

A. ANALYSIS SUMMARY

The Urgency Codes Test incorporates identification of jobs with particular urgency codes. The Urgency Codes Test procedures (see Figure XII-1) include: (1) Determine First and Last Urgency Codes and (2) Determine Evidence of Operator Intervention. This test references the Memory Constraint Test. Reports used in this test include the GSEP Allocator/Termination Report and the GCOS Console Log.

B. DETERMINE FIRST AND LAST URGENCY CODES

This procedure determines if potential job delays can be inferred from urgency code values.

1. Step No. 1: Scan For Potential Problem

The analyst can determine potential delays caused by urgency code value assignments by examining the initial and final urgency codes of jobs.

a. Urgency Code Values. GCOS (as of release WW6.3) serves jobs' core requests according to the following urgency code priorities:

(1) Urgency Code = 0,1. Activities with these urgency codes are swapped whenever their memory is needed by another job with an urgency code of 5 or greater.

(2) Urgency Code = 2-4. Activities with these urgency code values are lower in the core allocation service priority than normal jobs. Though not subject to swap for a job with urgency code ≤ 32, they will be swapped whenever their memory is needed by a job with urgency code >32.

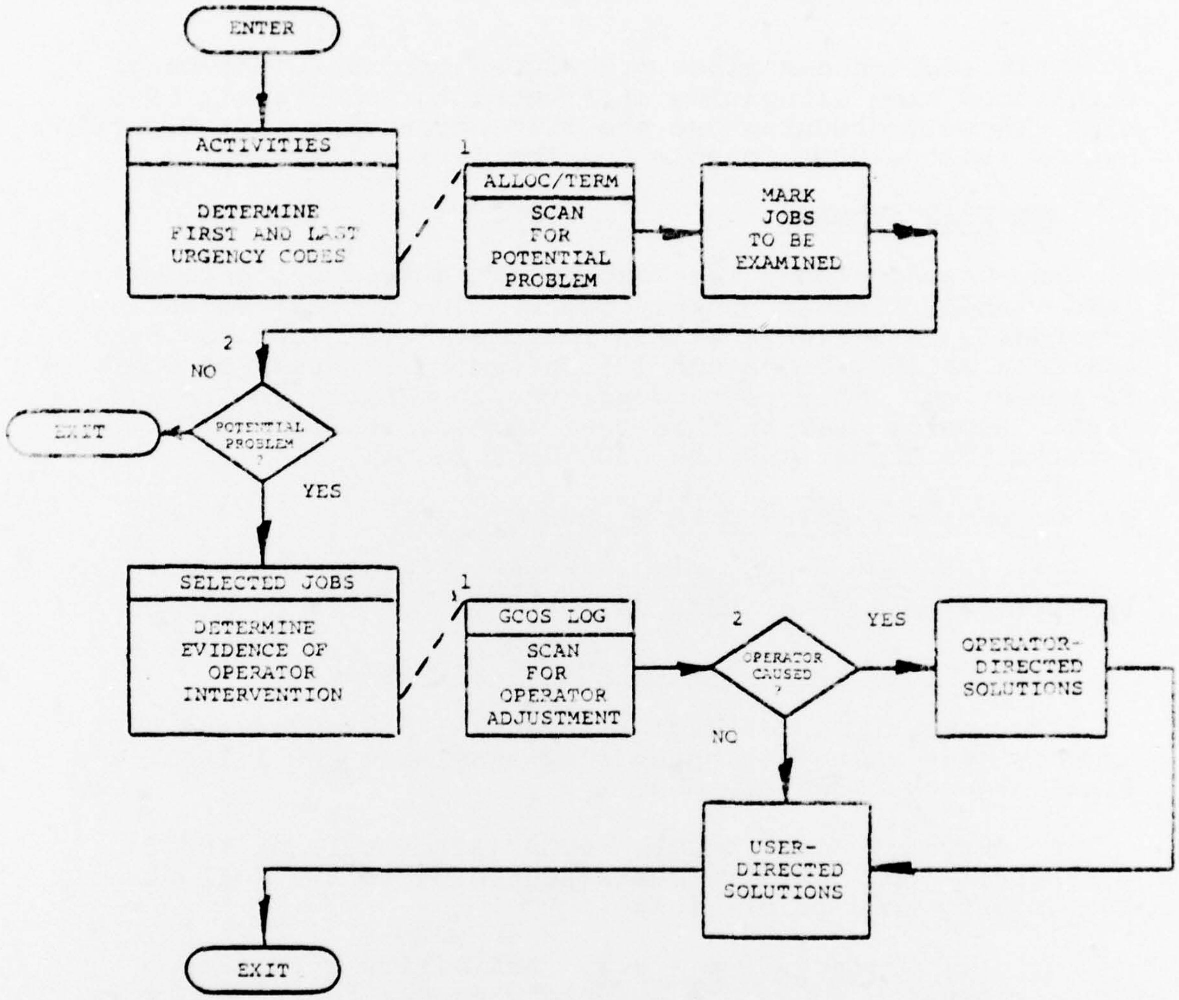

URGENCY CODES TEST PROCEDURES

FIGURE XII-1

(3) Urgency Code = 5-6. Activities with these urgency code values are the "normal" activities. These activities will have higher core service priority than those with urgency codes in the 2-4 range, but they will be swapped before any activities with urgency code ≥7.

(4) Urgency Code = 7-31. Activities in this range will have higher priority for memory allocation than those in lower ranges. They will not cause the swap of any activity with urgency code ≥2; they cannot be swapped until all activities with urgency codes ≤6 have been swapped.

(5) Urgency Code = 32-63. Perturbation of core access by the batch workload is caused by activities with urgency codes in this range. Continuously running activities such as the Transaction Processing Executive will frequently assume high urgency codes to force the system to provide memory space and frequent access to the processor (using the Urgency Throughput dispatcher option). Note that GCOS will cause the swap of jobs with lower urgency code values in order to find memory for jobs in this range.

b. Potential Urgency Code Isolation Conditions. This step investigates the following elongation hypotheses:

(1) Low Initial Code Value. This is the case when jobs that are experiencing bad turnaround time are run with urgency code values [≤4]. These are below the "normal" range described above. These jobs are affected by all other jobs in the system and will have to wait for access to both memory and other system resources.

(2) High Initial Code Value. In this case, some jobs are being run with initial urgency code values [≥7]. These jobs affect other jobs running with lower urgency codes.

(3) Change In Value. Turnaround time for any job can be affected by the changes made to its urgency code as the job executes. These changes, if made from a high initial value to a low one, will only delay the job in question. The change may delay

the other jobs in the mix if made from a low value to a higher one. The exact effect will depend upon the values of the other jobs.

(4) Other Conditions. The analyst may tailor the scan criteria to local installation urgency code thresholds other than those described above.

c. Report Value. The "IURG" and "CURG" values, listed on the GSEP Allocator/Termination Report, represent the urgency code assigned to the job when it entered the system (i.e., IURG) and when it exited the system (i.e., CURG).

d. Report Scan. This step is used to identify jobs with other than "normal" urgency code values.

(1) Look For Jobs With Low Initial Urgency Code Values. Scan the GSEP Allocator/Termination Report for jobs with low [<4] initial urgency code values. Identify them by (✓) on the report.

(2) Look For Jobs With High Initial Urgency Code Values. Scan the GSEP Allocator/Termination Report for jobs with high [≥7] initial urgency code values. Identify them by a (✓) on the report.

(3) Look For Jobs With Significant Changes In Urgency Code. Scan the GSEP Allocator/Termination Report for jobs with initial and final urgency code values [≥20] numbers apart. Identify them by a (✓) on the report.

(4) Look For Other Conditions. Scan the GSEP Allocator/Termination Report for jobs that exhibit other, locally-defined installation criteria that could elongate either selected jobs or impact other jobs in the mix.

2. Step No. 2: Decision

This step determines if potential urgency code problems are indicated.

a. Scan For Low Urgency Code Values. If the scan conducted under Step Number 1 results in identification of one or more jobs with low urgency code values, the next step is to confirm elongation of these jobs.

If these jobs experience long turnaround times, an urgency code problem has been confirmed. Continue the analysis at Section XII.C.

b. Scan For High/Normal Urgency Code Values. If the scan conducted under Step Number 1 results in the identification of one or more jobs with high urgency code values, the next step is to confirm elongation of the other jobs in the system at the time.

It may be impossible to determine if any of the turnaround time for other jobs is due to the jobs with high urgency codes. Consider an urgency code problem confirmed if jobs with high urgency codes are executing while other jobs are experiencing long turnaround times. Continue the analysis with the procedure in Section XII.C.

c. Scan For Significant Changes In Urgency Code. If the scan conducted under Step Number 1 results in identification of jobs with initial and final urgency code values [$\geq$20] numbers apart, the next step is to confirm turnaround time elongation: (1) of these jobs if the final value is smaller and (2) of other jobs running at the same time if the final value is larger. If the final value is smaller and these jobs experience long turnaround times, or if the final value is larger and other jobs running at the same time experience long turnaround times, consider the urgency code problem confirmed and proceed at Section XII.C.

d. Scan For Other Conditions. If the scan conducted under Step Number 1 results in identification of jobs that meet locally-defined elongation criteria, the next step is to confirm elongation of these jobs or the other jobs in the mix. Because this set of conditions is dependent upon locally-defined criteria, it cannot be further treated in this procedure.

e. Exit Procedure. Exit the procedure if none of the scans confirmed an urgency code problem. Return to the section that called for this test.

C. DETERMINE EVIDENCE OF OPERATOR INTERVENTION

This procedure determines if operator changes to urgency codes are the source of the urgency code problem(s) confirmed above. This procedure determines whether the analyst should direct the solution to operations and/or the user.

1. Step No. 1: Scan For Operator Adjustment

This step determines if the operator changed the urgency codes of the selected jobs as they executed.

a. Log Message Format. The URGC Message, listed on the GCOS Console Log, can be input by the operator to change the urgency code of a job. The message has the following format:

URGC sssss uu

where,

sssss = SNUMB of Job

uu = Urgency Value (01 through 60).

With the URGC Message, the operator sets the urgency code of job sssss to the decimal value, uu. If uu is not specified, the existing urgency code of the job is displayed, but not changed. For GCOS Version WW6.3, if uu = 00 is entered, the job's urgency code will be set to 01; if uu $\geq$60, the urgency code will be reset to 60.

b. Scan Start And Stop. When conducting the GCOS Console Log scan, begin at a point on the log that occurs immediately after the selected job(s) started processing and end at the point where they finished.

c. Scan Step. Scan the GCOS Console Log for operator changes to the urgency codes of jobs identified above in Section XII.B.2 as experiencing or causing long turnaround times.

(1) Low Urgency Code Values. Scan to verify whether the operator set the urgency code value to a low number after the job entered.

(2) High Urgency Code Values. Scan to verify if the operator set the urgency code value to a high value.

(3) Significant Urgency Code Changes. Scan to verify that the operator setting caused the significant change of urgency code value.

2. Step No. 2: Decision

Direct the solutions proposed in Section XII.D to operations if it is confirmed that operator changes to urgency codes are the source of job delay.

D. TUNING STEPS

This section discusses steps to be taken on the basis of Urgency Codes Test confirmations.

1. General Operator-Directed Solutions

Propose these steps to reduce operator changes to urgency code values while jobs are running. If possible, direct the installation to require operators to justify in writing any change to job urgency codes. Alternatively, direct the installation to formulate specific rules which determine when urgency codes should be changed from the console.

2. User-Directed Solutions

These solutions direct the user to change job urgency code values. Note that the implementation of these changes will require management approval.

a. Low Initial Urgency Code Solution. If it is confirmed that a low initial urgency codes problem exists, direct the user to employ a higher initial urgency code value.

b. High Initial Urgency Code Solutions. Direct operations and/or the user to reduce the urgency code value if it is confirmed that a high initial urgency codes problem exists. In many installations, the initial urgencies of jobs are governed by policies, written or unwritten. Reduction of initial urgencies of these jobs may require rethinking the policies to exclude some types of jobs from the privilege of high initial urgency. If this approach cannot be attempted, the addition of memory might alleviate the problem of other jobs in the system. Execute the Memory Constraint Test described in Section V to confirm the solution as appropriate.

c. Significant Urgency Code Changes. Usually, these changes will be due to operator changes or CALC Migration (incrementing the urgency code of the first activity in the Core Allocation Queue whenever it is passed over to give core to a lower urgency code activity). If operator changes are the cause of significant urgency code changes, attempt the General Operator-Directed Solutions above. No solution exists for CALC Migration; it may not be detrimental and the solution would involve a change to GCOS code.

Significant urgency code changes may occasionally have two other sources. An activity may change its own urgency code if it is allowed to execute in Master Mode. System programs (TSS, PALC, GEIN, etc.) do this frequently. Jobs with Master Mode (i.e., Privity) permission should not be allowed to do this unless it is necessary. Discuss the impact of this change with the user.

The second source is GCOS system programs. GEIN and PALC both change Urgency Codes for various reasons. Usually this change is to lower the urgency code; however, PALC will raise the urgency code to 38 if the activity is allocated unit record devices such as card readers, punches, or printers. No solution is proposed for these changes because it would involve modifications to GCOS code.