

SOFTWARE OPTIMIZATION
THROUGH RESEGMENTATION

by

MIKE CLARKSON

presented to;

HP-3000 USERS GROUP MEETING

in

MIAMI, FLORIDA

February 27, 1975

I. INTRODUCTION

During the past couple of years, we at Data Base Management Systems, Inc., have developed quite a few large, sophisticated software packages. These packages tax quite heavily the resources of the HP-3000 and tend to exercise the machine, both hardware and software, to its limits. In order to keep our program development from overwhelming the 3000, we have developed methods for efficiently maintaining our source code and for optimizing the execution of our programs. It is extremely important to have the most efficient software possible when dealing with a computer such as the 3000, which is very sophisticated, but yet has such small memory limitations. The 3000 can perform very well if the programmer will take time to optimize his software so that it complements the 3000's resource management. This paper will outline several areas for program optimization and will expound on the area of program resegmentation.

II. SOFTWARE OPTIMIZATION

The need for program optimization should be intuitively obvious, since it is very unlikely that a program will exist in its most optimal form when it is first written. Since most sophisticated programs will require optimizing before they will run at all, programmers will make a "first approximation" before the program is even run for the first time. From the initial run on, the programmer will continuously learn new short cuts as well as observe his program's running characteristics; thus, there is the continuing need for program optimization. There are many areas available that the programmer can attack in order to make his programs more efficient. Following are several examples:

Since MPE is a relatively unknown factor to most programmers, one should avoid calls to MPE as much as possible. If a call to MPE is unavoidable, one should attempt to call the lowest level intrinsic available in order to avoid unnecessary overhead, even if it means making the call in privileged mode. Privileged mode execution need not be avoided as long as one is very careful of each instruction that is executed while in privileged mode. The higher level, user mode MPE intrinsics have a lot of overhead determining if the user has the capability to perform the desired operations.

Another area for consideration is minimizing disc accesses.

If small scratch files are necessary, one should weigh the advantages of a local dynamic array or an "OWN" array as opposed to a disc file (keep in mind, however, that a larger stack could cause more disc accessing by memory management). Also, we have found that it is much more efficient to have our own buffer manager which uses direct access disc files with "no buffering" option, rather than depending on the file system to optimize our disc accesses.

Another area which is not in the scope of this paper, but is well worth mention, is modular, structured programming with clean paragraphing in order to minimize program maintenance time. ¹

Finally, but not exhaustively, is the area of program re-segmentation which is discussed in detail in subsequent sections.

III. RESEGMENTATION

We have found that one of the best and yet simple methods of program optimization is resegmentation. Since the memory size of the 3000 is limited to a maximum of 64K words and many SPL programs require more code than available memory, the programmer is able to group his code into segments by using the SPL command "\$CONTROL SEGMENT=name". This is called segmenting the code. One should attempt to gather into a segment, procedures which call each other since there is additional run time overhead involved when one procedure calls a second procedure that is in a different segment. If a relatively small procedure is called very often by some other procedure, and the small procedure is in a different segment and cannot be moved, then it is useful to make a copy of the small procedure and make it a subroutine inside the calling procedure. This eliminates both overhead of PCAL execution and segment swapping. We have found that an optimal code segment size is around 2000(octal) words. Reasonable limits are from 1400(octal) to 3000(octal) words. Procedures which are called infrequently such as initialization, termination, and error handling procedures should be grouped into larger segments. One should use as many segments as necessary to make his program run efficiently, but at the same time, should be frugal with segments, since there is a hardware limitation of 256 segments which can be active at any one time.

Not only is it advisable to segment one's code optimally, but at times one should also segment his data, the latter being slightly more complicated. For example, say a program needs a 4000 word buffer at various times during its execution, but it is undesirable to incorporate the buffer into the normal stack, then an extra data segment can be utilized. The intrinsic GETDSEG can be called, giving it the desired buffer size and it will return the DST number of the extra data segment. Then, during further execution whenever it is necessary to access the buffer, call the intrinsic EXCHANGEDB(DST#) and the DB register will then point to the extra data segment and it can be accessed just like normal DB storage. Remember that the regular storage normally accessed via the DB register and the extra data segment cannot be accessed simultaneously. In order to switch back to the regular DB storage, call EXCHANGEDB(0) and the DB register will be returned to its original value. Many extra data segments can be maintained in one program, but it is the responsibility of the program to keep track of each extra data segment's DST number. When the program is through with its processing, it is necessary to release each obtained extra data segment back to MPE by calling the intrinsic RELDSEG(DST#). The program must be in privileged mode in order to call any of these intrinsics. It is important to release all extra data segments back to MPE, or they will be lost until the system is COLDSTARTED (or COOLSTARTED).

Once a program is segmented, it can be run along with one of the available monitoring programs such as TRACE (described in subsequent sections) which will produce statistics describing segment swapping, et cetera. The program can then be resegmented by either recompiling after moving the "\$CONTROL SEGMENT" cards or by using the SEGMENTER command NEWSEG. The produced statistics will indicate which procedures call which other procedures in other segments, and using this information, the programmer can regroup procedures which call each other most often. The resegmenting, monitoring process can be performed repeatedly until the most optimal segmentation is achieved.

IV. AVAILABLE PROGRAMATIC ASSISTANCE

Presently, there are three packages available to aid the systems programmer in optimizing his software. They are (1) SAMPLER; (2) AUTOSEG; and (3) TRACE. SAMPLER is a software sampling system used for measuring the relative time spent executing various sections of code. SAMPLER requires an extra clock/TTY interface board to be installed in the system before it can run. The SAMPLER documentation in the appendix completely describes all necessary steps. Be sure to read the complete documentation before attempting to use SAMPLER.

AUTOSEG is a performance enhancement tool which provides for the automatic resegmentation of programs based on data gathered under actual program operation. AUTOSEG in its present form is not very useful because it is not quite smart enough to resegment a USL file any more optimally than could casually be done by a programmer. Sketchy documentation does exist in the appendix describing the three functions of AUTOSEG. AUTOSEG does prompt the user for the necessary information. TRACE uses the hardware "trace" facility to collect data pertaining to processes at the time of intersegment transfers caused by PCAL's and EXIT's. This package is very useful in helping the programmer resegment his

programs by hand. The following section describes TRACE in detail and additional information is located in the appendix.

V. TRACE

This section will discuss in detail the software monitoring package TRACE which analyzes procedure calls external to a segment, showing caller, callee, presence of the called segment and frequency. In order to use TRACE, a new version of MPE must be generated and the produced tape must be COLDSTARTED. The version of TRACE to be discussed runs under MPE32000B.00.09. It is important to note that TRACE periodically halts the machine and should therefore be run only when a program is to be TRACED and the regular operating system should be reloaded before any normal machine usage is resumed.

Following are the steps necessary to generate a TRACE cold load tape:

1. Restore all files from the distribution tape which belong in group/account "OUR.SYS."
2. Patch INITIAL to obtain a data segment for TRACE by changing an "IF FALSE THEN" to an "IF TRUE THEN".

The octal instruction to be changed is 25001 which is found around address 2244(octal). Enter the following command sequence:

```

:RUN PATCH
FILE=? INITIAL
? M, 2, 2244
25001, 0600
? E

```

3. Execute the following jobstream to generate the TRACE
version of MPE:

```

:JOB MANAGER.SYS,OUR
:PURGE ININ
:SPL M10M000B,, $NULL, S10S000B, NEWM10M
:SPL ININY,,, NEWM10M
:PREP $OLDPASS, ININ;CAP=PM
:SAVE ININ
:PURGE NEWM10M
:PURGE EXIN
:SPL M11M000B,, $NULL, S11S000B, NEWM11M
:SPL EXINX,,, NEWM11M
:SAVE $OLDPASS, EXIN
:PURGE NEWM11M
:EOJ
:JOB MANAGER.SYS
:FILE SYSTAPE;DEV=TAPE
:TELLOP MOUNT BLANK TAPE ON TAPE DRIVE UNIT
:SYSDUMP *SYSTAPE
YES
H8
NO
NO
NO
NO
NO
NO
NO
YES
ININ, ININ. OUR
EXIN, EXIN. OUR
NO
:EOD
:EOJ

```

<< BLANK CARD>>

<< SYSTEM PROGRAM CHANGES>>

<< BLANK CARD>>

<< BLANK CARD>>

4. Note that EXIN is a USL file and ININ is a PROG file.
5. The tape produced by SYSDUMP should be loaded with the COLDSTART option.

Once a TRACE cold start tape has been generated and a program is ready to be traced, perform the following steps to obtain printer output describing the program's segment activity:

1. COLDSTART the TRACE version of MPE.
2. Run the program to be TRACEd from a terminal, generating an LMAP on the line printer. If the CSTs are not allocated contiguously, abort the program and run it again until the CSTs are contiguous. Make sure the program does not terminate until after step # 3.
3. From the console :ALLOCATE the running program.
4. Allow running program to terminate.
5. From console :RUN TRACE.OUR.SYS - respond to prompts as follows:

? INIT 0

? TRACE %n/%m

where n is the first CST number on LMAP
and m is last CST number on LMAP

?RUN

If the system is equipped with an extra clock/TTY interface board, rather than entering 0 with INIT, enter the

decimal DRT number of the clock board (eg. ?INIT 12).

6. Hang scratch tape on tape unit #7 with write ring in.
7. From the terminal run the program to be TRACEd with valid input, output, etc.
8. When the program terminates on the terminal:

From console continue responding to TRACE prompts as follows:

```
?EOF
?STOP
?CLEAR %n/%m      [see above]
?EXIT
```

9. From console :RUN TRACERED.OUR.SYS
10. Respond to TRACERED prompt with:

```
?SEGMENT CALLER EXIT
```

11. Statistics will be printed on line printer.

(Note: If the machine halts during execution of the program being TRACEd, merely hit the RUN/HALT switch and the machine will continue executing. The halts are caused by such things as no write ring in tape or printer being used by another process. Also make sure that the program to be TRACEd is the only program running on the system [ie. besides TRACE.OUR.SYS]).

The next page is an example of one page of the output generated by TRACERED.

PAGE 1 HP3000 SEGMENT TRACE DATA REDUCTION PROGRAM

SEG ENTRY CNT.	ABSENCES	%ABS.	STI ENTRY CNT.	%SEGE TIMED ENT.	%SITE TIME/C	SDEVT/C	PIN SEG DELTP	CALLS	%SITE
121	564	73 12.9	001	564 100.0	432 76.6	.28	.13 026 144 00234	564 100.0	
122	2	1 50.0	001 002	1 50.0 1 50.0	0 0.0 0 0.0	.00 .00	.00 026 136 00402 .00 026 136 00375	1 100.0 1 100.0	
126	72	13 18.1	001 002 003 004	24 33.3 24 33.3 6 8.3 11 15.3	24 100.0 17 70.8 4 66.7 11 100.0	.16 .24 .30 .27	.01 026 130 02041 .00 026 130 02060 .28 026 133 00364 .04 026 133 00356 026 136 03337 026 133 00346 026 133 00331 026 133 00342	24 100.0 24 100.0 6 100.0 6 54.5 5 45.5 5 100.0 1 100.0 1 100.0	
127	45	40 88.9	001	45 100.0	4 8.9	1.39	.32 026 133 00227	45 100.0	
130	512	61 11.9	001 002	256 50.0 256 50.0	135 52.7 150 58.6	1.21 1.11	.21 026 133 00173 .15 026 133 00151	256 100.0 256 100.0	
131	2	2 100.0	006 010	1 50.0 1 50.0	0 0.0 0 0.0	.00 .00	.00 026 133 00250 .00 026 133 00252	1 100.0 1 100.0	
132	1	1 100.0	001	1 100.0	0 0.0	.00	.00 026 131 00642	1 100.0	
133	342	17 5.0	001	342 100.0	20 5.8	.47	.04 026 122 00767 026 134 00773 026 135 01111 026 136 03421	8 2.3 4 1.2 266 77.8 64 18.7	
134	1	1 100.0	001	1 100.0	0 0.0	.00	.00 026 135 00467	1 100.0	
135	165	30 18.2	001	119 72.1	115 96.6	.09	.00 026 133 00621 026 134 00735 026 136 03136 026 136 03215 026 133 00272 026 133 00216 026 134 00774	44 37.0 4 3.4 2 1.7 69 58.0 2 100.0 40 100.0 4 100.0	
136	15	5 33.3	006 010	1 6.7 5 33.3	0 0.0 5 100.0	.00 .09	.00 026 122 00222 .00 026 153 00066 026 157 00006 026 160 00044 026 161 00106 026 163 00124 026 122 00731	1 100.0 1 20.0 1 20.0 1 20.0 1 20.0 1 20.0 9 100.0	
137	1	1 100.0	001	1 100.0	0 0.0	.00	.00 026 135 01334	1 100.0	
140	35	31 88.6	001	35 100.0	0 0.0	.00	.00 026 133 00421	35 100.0	
141	20	12 60.0	001 002	1 5.0 19 95.0	1 100.0 9 47.4	.38 1.42	.00 026 147 00043 .44 026 147 00100 026 147 00206	1 100.0 7 36.8 12 63.2	

There are sixteen columns of values, four of which are non-zero because there was an extra clock/TTY interface board in the system when the program was TRACEd.

Following is a description of each of the sixteen columns of information on the TRACE reduction printout (the first eleven columns comprise the called segment statistics and the last five columns comprise the caller statistics):

1. three digit octal CST (code segment table) number of the segment
2. total PCAL entries to the segment
3. number of PCAL absences only (EXIT absences cannot be traced)
4. absences as a percentage of total entries
5. three digit octal STT (segment transfer table) entry number
6. total calls to that STT entry
7. STT entries as a percentage of the total entries
8. number of timed STT entries (total number of STT entries on which it was possible to gather timing information)
9. number of timed entries as a percentage of the total STT entry count
10. average time per STT call in milliseconds with all TRACE overhead removed (this time is the time spent in this segment less the time spent in other segments which were being traced and were called by this segment)

11. standard deviation of the time per call in milliseconds
12. three digit octal process identification number of caller
13. three digit octal CST number of caller
14. five digit octal delta P (ie. offset into calling segment to PCAL made to called segment)
15. total calls made to this STT by the caller from this delta P
16. number of calls as a percentage of the STT entries

An example of how to recognize potential optimizations is as follows: Notice that the segment at CST #140 was called 35 times which caused 31 disc accesses or 88.6% of the total. This is a very high percentage to have causing disc accesses. The only entry point in the segment which was called was STT #001 and it was called from only one place, which was 421(octal) words into the segment at CST #133. Now one would look at the LMAP describing these segments and determine which relative segments within the program were loaded at CST entries 140 and 133. Then looking at the PMAP of the traced program, one would analyze the present size of the segment corresponding to CST #133 to see if it has room to add the code of the RBM located at STT #001 of the segment corresponding to CST #140. If there is room in the segment, the RBM can be moved by using the SEGMENTER

command NEWSEG or by changing the appropriate "\$CONTROL SEGMENT" card in the SPL source and re-compiling the procedure. This one optimization would save 31 disc accesses. It is important to also determine if there are any other procedures in the CST #140 segment which make internal segment calls to the STT #001 procedure; since they could make more than 35 calls, the expected optimization could cause a degradation.

The timing information found on the TRACE reduction print-out is also useful. Through close observation of the timing statistics, one can learn where in the code a major portion of execution time is being spent. For example, two entry points (STT #001 and STT #002) in the segment loaded at CST #130 were each called 256 times. Note from column ten of the TRACE reduction printout that each call to these entry points averaged 1.21 milliseconds and 1.11 milliseconds respectively in duration. This time factor is somewhat high relative to most of the other TRACEd STT entries. Given this fact, one should analyze the code of the procedures which correspond to these STT entries and determine if any code optimization could be performed in order to help these procedures execute more quickly.

Following are some helpful hints for using TRACE that should be utilized. In order to get good timings for every RBM of the program,

the USL file should first be run through AUTOSEG's NICRSP processor to generate worst case segmentation so no internal segment PCALs are executed. This is done by (1) make a copy of the USL file;

(2) :RUN NICRSP.OUR.SYS and give it the name of the USL copy;

(3) then prepare the resegmented USL copy and run it through TRACE.

There are two things to keep in mind when using TRACE: (1) TRACE.

OUR.SYS runs in privileged mode so the user needs privileged and

account capabilities; (2) TRACE.OUR.SYS uses its own tape driver

for I/O so MPE must not be allowed to use any other tape drive on

the same controller that TRACE is using. An extra clock/TTY inter-

face board may be easily added to the system by assigning the extra

board an unused DRT number and plugging the board into the system.

MPE need not be reconfigured to recognize the new DRT.

VI. CONCLUSIONS

We hope that the reader has gained an appreciation for program optimization, especially through resegmentation. We at DBMS have enjoyed performance improvements of as much as 400% to 500%, using resegmentation alone. Resegmentation is an easy first pass to make before getting into the optimizations that require such things as changing existing code. The software monitoring packages described in this paper will be released on the contributions tape and are stored in the group/account "OUR.SYS". If any system anomalies occur during the use of TRACE, it should be assumed that they are attributable to that usage and, therefore, no problem reports should be submitted to Hewlett-Packard. The appendix contains a description of the contents of OUR.SYS which will be on the tape that is distributed to the users. If there are any questions concerning the software monitoring packages, please direct them to me at the following address:

Mike Clarkson, Vice President
Data Base Management Systems, Inc.
12100 N. E. 16th Avenue
North Miami, Florida 33161

REFERENCES

1. HP-3000 Users Group Proceedings, May 10, 1974, p. 69,
PROGRAM PERFORMANCE, by Stephen Sontz

APPENDIX

CONTENTS OF GROUP OUR.SYS

1.	SAMPLE	
2.	SAMPLES	<< SOURCE >>
3.	SAMPLING	
4.	SAMPLINS	<SOURCE>
5.	DRS	
6.	DRSS	<SOURCE>
7.	SYSGEN	<<OBSOLETE>>
8.	M10M000B	<<PATCH DECK>>
9.	M11M000B	<<PATCH DECK>>
10.	S10S000B	<<SOURCE>
11.	S11S000B	<SOURCE>
12.	EXINX	<<PATCH DECK>>
13.	ININX	<<PATCH DECK>>
14.	ININY	<<PATCH DECK>>
15.	TRACE	
16.	TRACES	<<SOURCE>
17.	TRACERED	
18.	TRREDS	<<SOURCE>
19.	NICRSP	
20.	NICRSS	<<SOURCE>
21.	ORSP	
22.	ORSS	<<SOURCE>
23.	ORTDRP	
24.	ORTDRS	<<SOURCE>

SAMPLER

SOFTWARE SAMPLING SYSTEM

March 7, 1973

Contents

	<u>Page</u>
A. Introduction	1
B. Hardware Requirements for Sampler	2
C. Sampler Operation and Format	3
D. Setting up the Sampler	4
E. The Data Reduction Program	6
F. A Case Study	8
G. Miscellaneous Comments	11

A. Introduction

The Software Sampling System is a useful tool for measuring the relative time spent executing various sections of code. The system consists of three parts:

I. The Sampler interrupts the CPU at a preselected frequency and records on tape the following information for any number of selected code segments:

1. Number of code segment interrupted.
2. Relative P in that code segment prior to the interrupt.
3. Approximate size of the stack.

All user interaction with the sampler is done via the Setup program.

II. There is a setup program with provision for initializing the sampler, selecting segments to be sampled and setting the sampling interval.

III. A Data Reduction Program is available to process the data tape. This program will provide the following histograms:

1. Relative time spent in each segment sampled.
2. Relative activity within selected segments.
3. Size of stack from DB to S for each sample.

A. Introduction

The Software Sampling System is a useful tool for measuring the relative time spent executing various sections of code. The system consists of three parts:

I. The Sampler interrupts the CPU at a preselected frequency and records on tape the following information for any number of selected code segments:

1. Number of code segment interrupted.
2. Relative P in that code segment prior to the interrupt.
3. Approximate size of the stack.

All user interaction with the sampler is done via the Setup program.

II. There is a setup program with provision for initializing the sampler, selecting segments to be sampled and setting the sampling interval.

III. A Data Reduction Program is available to process the data tape. This program will provide the following histograms:

1. Relative time spent in each segment sampled.
2. Relative activity within selected segments.
3. Size of stack from DS to S for each sample.

B. Hardware Requirements for Sampler

To run the Sampler in its present form requires the following hardware:

1. A 3000 System which will run MPE and for which the maximum DRT number has been set to 70 greater than required by any device.
2. An extra clock/TTY board with a known distinct device number.
3. A dedicated tape drive and controller.
4. An extra terminal from which to run the setup program.

C. Sampler Operation and Format

The sampler is an interrupt handler which resides in the upper DRT table together with its buffer. This interrupt handler is activated by timer interrupts from the extra clock/TTY interface board.

The sampler then traces stack markers backwards to determine which code segment was executing prior to the interrupt. This code segment number is used to index into an internal segment bit-table which determines whether this segment's data should be recorded. If the segment has been selected the contents of its status register is entered in the sampler's buffer together with the relative location (within that segment) of the next instruction to be executed. The relative value of the top of the stack is also entered and a user stack/interrupt control stack bit is set.

If the buffer is full interrupts are disabled and the buffer is written directly onto tape. If constant timing interval has been specified at setup time the count register of the timer is cleared just prior to exit. However, if the randomized jitter has been specified on the timing interval, a pseudo-random number between 0 and 255 is loaded into the count register. The switch register is used as one of the parameters of the random number generator to allow the operator to influence the sequence generated.

Tape records produced by the sampler are 128 words in length consisting of 16 logical records of 8 words each. The contents of each logical record is as follows:

- Word 0. Status word of interrupted procedure.
- Word 1. Relative location of next instruction to be executed within that procedure.
- Word 2. S-DB if on interrupt control stack and (DB-S) if on user stack.
Note that bit 0 indicates ICS/User Stack.
- Word 3. Unused at present.
- Word 4. Record type. Bit 0 is 1 to indicate that the record was produced by the sampler.
- Words 5, 6, 7. Unused at present.

D. Setting up the Sampler

Before starting-up MPE ensure that the extra clock/TTY interface board is inserted in the highest priority polled I/O slot. Coolstart MPE and reply "Y" to the question "ANY CHANGES?" When the question "HIGHEST DRT NUMBER = XX.?" is printed reply with a number which is 70 greater than the highest DRT used by the system. No other changes are required. Mount a tape fitted with a write ring and select unit 0. When the system is up log on and :RUN SETUP1. The machine replies "TRACER/SAMPLER" and prompts for commands with "?". The first command to be entered must be "\$LOAD" (as described below) to initialize the sampler after which sampler commands may be entered in any order.

Commands specifically affecting the sampler are strings of length less than 7 having "\$" as the first character. The second character determines the command, and following characters are optional. All parameters are in standard 3000 form ("%" prefix for octal numbers).

\$LOAD <timer device number>, <starting drt number>

This loads the sampler into the drt table starting at <starting drt number> (which must be at least one greater than the highest drt entry used by the system). The <timer device number> is the device number of the extra clock/TTY interface.

\$TIME <count> <quantum>

This command sets the time between completion of one sample and start of the next.

<quantum> is "U" = microseconds

"M" = milliseconds

or "S" = seconds

<count> is an integer less than 64K. The timer is set-up so that the precision of sampling interval is better than 0.1%. This command also halts the sampler if it was running.

\$RAND This command introduces a random jitter to the sampling interval. The purpose of this jitter is to avoid the possibility of the sampling intervals becoming synchronized with the sampled code. The switch

register is used as one of the parameters for the random number generator. Statistical tests have shown that %176523 is a good number to set in the switch register. Zero in the switch register resets the \$RAND command.

\$SET <segment range list>

This command specifies the segments to be sampled. <segment range list> is a list of ranges of segment numbers separated by commas, where a range is a segment number or first segment number/last segment number.*

\$CLEAR <segment range list>

This command inhibits sampling for all segments in the ranges of the <segment range list>.*

\$GO Starts the sampler. (At least a "\$LOAD" and a "\$TIME" command should have been issued previously.)

\$HALT Halts the sampler.

The following general commands are also useful.

EOF Write an end of file mark on the tape.

DEBUG Enter debug.

EXIT Terminate the setup program.

There are various self-explanatory error messages. However, at present there is no error checking on the parameter values of the "\$LOAD" command. Since the program runs in privileged mode care must be exercised when entering the \$LOAD command to avoid crashing MPE.

*<segment range list>:: = <segment range> | <segment range list>, <segment range>
<segment range>:: = <segment number> | <segment number>/<segment number>
<segment number>:: = any integer between 0 and 255

E. The Data Reduction Program

A program is available to process the data tape produced by the sampler. Output is in the form of histograms of segment usage and relative activity within each segment.

The data reduction program is activated by the following sequence of commands:

```
:FILE FTH09; DEV = LP ; CCTL  
:FILE FTH07; DEV = TAPE; REC = 128, 1, F, BINARY; NOLABEL  
:RUK -FTH TAPES
```

The program makes a first pass through the tape, processing all records, if any, produced by the old (now obsolete) trace routine and obtaining normalizing information about sampler records. At the end of the first pass the total number of records on the tape is printed on the line printer. Assuming all records were produced by the sampler "ENTERING DATA REDUCTION FOR SAMPLER" is printed on \$STDOUT, followed by a request for segments for which an internal histogram is desired. Enter these in the free format <segnum>, <p interval> followed by carriage return where <segnum> is the number of the segment and <p interval> is the "width" of the bars of the histogram. If <p interval> is zero the program chooses a bar width such that the histogram for that segment will fit on a single page. Of course, if <p interval> is 1 the histogram will show the number of samples which hit each individual instruction.

Prefix octal numbers with "%". If <segnum> is a valid number for a segment but no samples interrupted that segment the machine replies "BAD SEGMENT NUMBER". To terminate the list of segment number enter a <segnum> greater than 255.

A response of "Y" to the question "COMPLETE PLOT OF STACKSIZE?" will product a histogram of the value of S-DB for every sample. This histogram of stacksize consumes much time and paper and is usually of little interest.

A second pass through the tape is then made, after which the histograms requested earlier are printed on the line printer. The number of the segment

is shown at the top of each histogram. Each bar in an internal histogram represents the number of samples which hit the range of code locations shown at the left of the bar. The actual number of samples is shown at the right. Bar lengths are normalized separately for each segment to emphasize the high-usage sections of code within a segment.

The machine then prompts "SEGMENT HISTOGRAM?" on \$STDOUT. A reply "Y" will cause a histogram of segment usage to be printed on the line printer. This segment histogram shows the number of samples which hit each segment. For each segment which was sampled at least once a bar appears in the histogram showing at the left, the segment number and, at the right, the number of samples which hit that segment. Once again the bar lengths are normalized so that the longest bar spans the page.

The machine then prompts "EXIT?" on \$STDOUT. A reply "Y" terminates the program, whereas "N" rewinds the tape in preparation for repeating the second pass.

F. The Sampling System in Action: A Case Study

This section describes an example of the use of the software sampling system to improve the performance of a specific program, namely the data reduction program described in section E (henceforth referred to as DRS.)

DRS was written in Fortran and run originally on the HP 2100 and DOS-M. When transferred to HP 3000 Fortran and run under MPE execution time became distressingly large. In an attempt to improve this the sampler was run while DRS was executing. The following is a description of the procedure used.

The system was started up with the extra DRT space (maximum DRT set to 101) and with the extra clock/TTY interface board (device number %13 = 11) inserted in a polled I/O slot. Two terminals were required for the measurement, a running terminal and a sampling terminal. Since the only tape controller available would be required by the sampler it was not possible to run DRS with input from tape. The utility program FILECOPY was used to copy a data tape (from a previous sampling session) to a disc file called DATAS, and DRS was run using DATAS as the input file.

Before running DRS from the running terminal the following sequence of commands was entered-from the measuring terminal (system output underlined):

:RUN SETUP1

TRACER/SAMPLER

? \$LOAD %13, 31

? \$SET 0/255

? \$TIME 4 M

? \$RAND

(%176523 set in the switch register.)

? \$CLEAR %107

(%107 was segment number of the dispatcher.)

At this point the sampler is set up to interrupt at 4 millisecond intervals with random jitter (i.e. time between samples is randomly chosen in the range

3.75 to 4 milliseconds). All segments are set to be sampled except the dispatcher (thus no tape will be written while the system is paused waiting for input).

At this time DRS was started on the running terminal. As soon as DRS started executing the command

? \$G0

was entered on the measuring terminal to start the sampler. When DRS terminated, the sequence of commands

? \$HALT

? EOF

? EXIT

was entered on the measuring terminal to stop the sampler, write an end of file mark on the tape and then terminate the setup program.

The data tape was then manually rewound and DRS was run with the data tape as input to produce the histogram of segment activity shown in Figure 1. This histogram immediately indicates that the Formatter consumes more than twice the CPU time that the actual DRS code does. To determine exactly where in the formatter this time was spent histograms of activity within segment 132 were produced (during subsequent repetitions of the second pass). One such histogram is shown in Figure 2. Checking the high usage sections of code suggests that formatter time could be reduced considerably by changing some data definitions, and by modifying the statement which reads from tape (from a do-implied list to an array read), thereby reducing the number of calls to the formatter.

These changes were made to DRS and the above measurement procedure was repeated. The resulting histogram of segment usage is shown in Figure 3. (Note that in this measurement the dispatcher was sampled also.) This shows that the relative formatter time has been reduced to less than one quarter of its previous value. In fact total execution time of DRS was reduced 58%.

3.75 to 4 milliseconds). All segments are set to be sampled except the dispatcher (thus no tape will be written while the system is paused waiting for input).

At this time DRS was started on the running terminal. As soon as DRS started executing the command

? \$G0

was entered on the measuring terminal to start the sampler. When DRS terminated, the sequence of commands

? \$HALT

? EOF

? EXIT

was entered on the measuring terminal to stop the sampler, write an end of file mark on the tape and then terminate the setup program.

The data tape was then manually rewound and DRS was run with the data tape as input to produce the histogram of segment activity shown in Figure 1. This histogram immediately indicates that the Formatter consumes more than twice the CPU time that the actual DRS code does. To determine exactly where in the formatter this time was spent histograms of activity within segment 132 were produced (during subsequent repetitions of the second pass). One such histogram is shown in Figure 2. Checking the high usage sections of code suggests that formatter time could be reduced considerably by changing some data definitions, and by modifying the statement which reads from tape (from a do-implied list to an array read), thereby reducing the number of calls to the formatter.

These changes were made to DRS and the above measurement procedure was repeated. The resulting histogram of segment usage is shown in Figure 3. (Note that in this measurement the dispatcher was sampled also.) This shows that the relative formatter time has been reduced to less than one quarter of its previous value. In fact total execution time of DRS was reduced 58%.

SEGMENT HISTOGRAM

SEG. #	NUMBER
00000	2.00
00001	1979.00
00002	1.00
00003	302.00
00004	1.00
00005	976.00
00006	561.00
00007	207.00
00008	1.00
00009	1.00
00010	3.00
00011	1633.00
00012	1816.00
00013	1.00
00014	2.00
00015	4.00
00016	172.00
00017	1.00
00018	1363.00
00019	455.00
00020	53.00
00021	35.00
00022	89.00
00023	4.00
00024	1.00
00025	5065.00
00026	22425.00
00027	8754.00
00028	17.16

FORMATTER

Figure 1. Histogram of segment usage (excluding the dispatcher) for original version of DRS.

NUMBER
235.00
633.20
274.22
375.02
1392.03
27.00
13171.00
221.00
1.00
1.00
13742.03
765.00
21.00
26.00
34.00
3000.22
2354.00
155.00
3.00
6.00
0.00
0.00
0.00
0.00
55.00
3.00
1.00
10.00
0.00
160.00
319.00
244.00
11.00
71.00
27.00
19.00
5.00
5.00
6.00
1965.29
1221.00
146.00
0.00
32.00
42.00

Figure 2. Histogram of activity within Formatter(original-DRS).

1 SEGMENT HISTOGRAM

SEG #	NUMBER
000000 000023	1093.00
000000 000031	1339.00
000000 000042	3.00
000000 000055	5819.00
000000 000063	9120.00
000000 000067	4566.00
000000 000071	2280.00
000000 000101	3.00
000000 000125	11741.00
000000 000150	2.00
000000 000170	3095.00
000000 000171	1434.00
000000 000173	3.00
000000 000175	4.00
000000 000176	2.00
000000 000179	9393.00
000000 000181	15086.00
000000 000183	14.4%
000000 000184	23%

DISPATCHER

FORMATTER

DRS

Figure 3. Histogram of segment usage for DRS with improved formatter calls.

Although the formatter time was reduced it was still significant. To determine whether this could be reduced any further histograms were produced showing activity within the formatter. Figure 4 ~~shows~~^{shows} the default histogram with bar width of %43 (obtained by entering "%130, 0" when prompted for segment number of DRS). Clearly most of the formatter time is spent in the section of code between %430 and %472. To examine this in more detail a histogram was produced with unit bar width (by entering "%130, 1" when prompted for segment number) showing the relative usage of each instruction. Figure 5 ~~shows~~^{shows} a portion of this detailed histogram. It was interesting to note that the instructions in locations 467, 470 and 471 account for 75% of the formatter time and 10% of the total execution time! Examining the formatter code shows that these instructions comprise the inner loop of a FOR statement used to fill a buffer with blanks.

Rather than attempting to modify the formatter, a tape reading routine was written in SPL using file system instinsics and linked into DRS thereby bypassing the formatter. This version of DRS was more than three times faster than the original version. Its segment usage histogram is shown in Figure 6.

The sampling system could doubtlessly be used to improve the performance of DRS still further by concentrating now on the program itself and possibly the file system.

ADDRESS HISTOGRAM FOR SEGMENT 130

(FOR KATIA)

RANGE	NUMBER
00000 00042	.00
00043 00045	.60
00046 00050	.00
00051 00053	.00
00054 00055	.00
00056 00057	64.00
00058 00059	10.00
00060 00061	9.00
00062 00063	69.00
00064 00065	7190.00
00066 00067	162.00
00068 00069	1052.00
00070 00071	88.00
00072 00073	5.00
00074 00075	6.00
00076 00077	103.00
00078 00079	7.00
00080 00081	91.00
00082 00083	82.00
00084 00085	165.00
00086 00087	48.00
00088 00089	35.00
00090 00091	2.00
00092 00093	.00
00094 00095	.00
00096 00097	.00
00098 00099	.00
00100 00101	.00
00102 00103	.00
00104 00105	.00
00106 00107	.00
00108 00109	.00
00110 00111	.00
00112 00113	.00
00114 00115	.00
00116 00117	.00
00118 00119	.00
00120 00121	.00
00122 00123	.00
00124 00125	.00
00126 00127	.00
00128 00129	.00
00130 00131	.00
00132 00133	.00
00134 00135	.00
00136 00137	.00
00138 00139	.00
00140 00141	.00
00142 00143	.00
00144 00145	.00
00146 00147	.00
00148 00149	.00
00150 00151	.00
00152 00153	.00
00154 00155	.00
00156 00157	.00
00158 00159	.00
00160 00161	.00
00162 00163	.00
00164 00165	.00
00166 00167	.00
00168 00169	.00
00170 00171	.00
00172 00173	.00
00174 00175	.00
00176 00177	.00
00178 00179	.00
00180 00181	.00
00182 00183	.00
00184 00185	.00
00186 00187	.00
00188 00189	.00
00190 00191	.00
00192 00193	.00
00194 00195	.00
00196 00197	.00
00198 00199	.00
00200 00201	.00
00202 00203	.00
00204 00205	.00
00206 00207	.00
00208 00209	.00
00210 00211	.00
00212 00213	.00
00214 00215	.00
00216 00217	.00
00218 00219	.00
00220 00221	.00
00222 00223	.00
00224 00225	.00
00226 00227	.00
00228 00229	.00
00230 00231	.00
00232 00233	.00
00234 00235	.00
00236 00237	.00
00238 00239	.00
00240 00241	.00
00242 00243	.00
00244 00245	.00
00246 00247	.00
00248 00249	.00
00250 00251	.00
00252 00253	.00
00254 00255	.00
00256 00257	.00
00258 00259	.00
00260 00261	.00
00262 00263	.00
00264 00265	.00
00266 00267	.00
00268 00269	.00
00270 00271	.00
00272 00273	.00
00274 00275	.00
00276 00277	.00
00278 00279	.00
00280 00281	.00
00282 00283	.00
00284 00285	.00
00286 00287	.00
00288 00289	.00
00290 00291	.00
00292 00293	.00
00294 00295	.00
00296 00297	.00
00298 00299	.00
00300 00301	.00
00302 00303	.00
00304 00305	.00
00306 00307	.00
00308 00309	.00
00310 00311	.00
00312 00313	.00
00314 00315	.00
00316 00317	.00
00318 00319	.00
00320 00321	.00
00322 00323	.00
00324 00325	.00
00326 00327	.00
00328 00329	.00
00330 00331	.00
00332 00333	.00
00334 00335	.00
00336 00337	.00
00338 00339	.00
00340 00341	.00
00342 00343	.00
00344 00345	.00
00346 00347	.00
00348 00349	.00
00350 00351	.00
00352 00353	.00
00354 00355	.00
00356 00357	.00
00358 00359	.00
00360 00361	.00
00362 00363	.00
00364 00365	.00
00366 00367	.00
00368 00369	.00
00370 00371	.00
00372 00373	.00
00374 00375	.00
00376 00377	.00
00378 00379	.00
00380 00381	.00
00382 00383	.00
00384 00385	.00
00386 00387	.00
00388 00389	.00
00390 00391	.00
00392 00393	.00
00394 00395	.00
00396 00397	.00
00398 00399	.00
00400 00401	.00
00402 00403	.00
00404 00405	.00
00406 00407	.00
00408 00409	.00
00410 00411	.00
00412 00413	.00
00414 00415	.00
00416 00417	.00
00418 00419	.00
00420 00421	.00
00422 00423	.00
00424 00425	.00
00426 00427	.00
00428 00429	.00
00430 00431	.00
00432 00433	.00
00434 00435	.00
00436 00437	.00
00438 00439	.00
00440 00441	.00
00442 00443	.00
00444 00445	.00
00446 00447	.00
00448 00449	.00
00450 00451	.00
00452 00453	.00
00454 00455	.00
00456 00457	.00
00458 00459	.00
00460 00461	.00
00462 00463	.00
00464 00465	.00
00466 00467	.00
00468 00469	.00
00470 00471	.00
00472 00473	.00
00474 00475	.00
00476 00477	.00
00478 00479	.00
00480 00481	.00
00482 00483	.00
00484 00485	.00
00486 00487	.00
00488 00489	.00
00490 00491	.00
00492 00493	.00
00494 00495	.00
00496 00497	.00
00498 00499	.00
00500 00501	.00
00502 00503	.00
00504 00505	.00
00506 00507	.00
00508 00509	.00
00510 00511	.00
00512 00513	.00
00514 00515	.00
00516 00517	.00
00518 00519	.00
00520 00521	.00
00522 00523	.00
00524 00525	.00
00526 00527	.00
00528 00529	.00
00530 00531	.00
00532 00533	.00
00534 00535	.00
00536 00537	.00
00538 00539	.00
00540 00541	.00
00542 00543	.00
00544 00545	.00
00546 00547	.00
00548 00549	.00
00550 00551	.00
00552 00553	.00
00554 00555	.00
00556 00557	.00
00558 00559	.00
00560 00561	.00
00562 00563	.00
00564 00565	.00
00566 00567	.00
00568 00569	.00
00570 00571	.00
00572 00573	.00
00574 00575	.00
00576 00577	.00
00578 00579	.00
00580 00581	.00
00582 00583	.00
00584 00585	.00
00586 00587	.00
00588 00589	.00
00590 00591	.00
00592 00593	.00
00594 00595	.00
00596 00597	.00
00598 00599	.00
00600 00601	.00
00602 00603	.00
00604 00605	.00
00606 00607	.00
00608 00609	.00
00610 00611	.00
00612 00613	.00
00614 00615	.00
00616 00617	.00
00618 00619	.00
00620 00621	.00
00622 00623	.00
00624 00625	.00
00626 00627	.00
00628 00629	.00
00630 00631	.00
00632 00633	.00
00634 00635	.00
00636 00637	.00
00638 00639	.00
00640 00641	.00
00642 00643	.00
00644 00645	.00
00646 00647	.00
00648 00649	.00
00650 00651	.00
00652 00653	.00
00654 00655	.00
00656 00657	.00
00658 00659	.00
00660 00661	.00
00662 00663	.00
00664 00665	.00
00666 00667	.00
00668 00669	.00
00670 00671	.00
00672 00673	.00
00674 00675	.00
00676 00677	.00
00678 00679	.00
00680 00681	.00
00682 00683	.00
00684 00685	.00
00686 00687	.00
00688 00689	.00
00690 00691	.00
00692 00693	.00
00694 00695	.00
00696 00697	.00
00698 00699	.00
00700 00701	.00
00702 00703	.00
00704 00705	.00
00706 00707	.00
00708 00709	.00
00710 00711	.00
00712 00713	.00
00714 00715	.00
00716 00717	.00
00718 00719	.00
00720 00721	.00
00722 00723	.00
00724 00725	.00
00726 00727	.00
00728 00729	.00
00730 00731	.00
00732 00733	.00
00734 00735	.00
00736 00737	.00
00738 00739	.00
00740 00741	.00
00742 00743	.00
00744 00745	.00
00746 00747	.00
00748 00749	.00
00750 00751	.00
00752 00753	.00
00754 00755	.00
00756 00757	.00
00758 00759	.00
00760 00761	.00
00762 00763	.00
00764 00765	.00
00766 00767	.00
00768 00769	.00
00770 00771	.00
00772 00773	.00
00774 00775	.00
00776 00777	.00
00778 00779	.00
00780 00781	.00
00782 00783	.00
00784 00785	.00
00786 00787	.00
00788 00789	.00
00790 00791	.00
00792 00793	.00
00794 00795	.00
00796 00797	.00
00798 00799	.00
00800 00801	.00
00802 00803	.00
00804 00805	.00
00806 00807	.00
00808 00809	.00
00810 00811	.00
00812 00813	.00
00814 00815	.00
00816 00817	.00
00818 00819	.00
00820 00821	.00
00822 00823	.00
00824 00825	.00
00826 00827	.00
00828 00829	.00
00830 00831	.00
00832 00833	.00
00834 00835	.00
00836 00837	.00
00838 00839	.00
00840 00841	.00
00842 00843	.00
00844 00845	.00
00846 00847	.00
00848 00849	.00
00850 00851	.00
00852 00853	.00
00854 00855	.00
00856 00857	.00
00858 00859	.00
00860 00861	.00
00862 00863	.00
00864 00865	.00
00866 00867	.00
00868 00869	.00
00870 00871	.00
00872 00873	.00
00874 00875	.00
00876 00877	.00
00878 00879	.00
00880 00881	.00
00882 00883	.00
00884 00885	.00
00886 00887	.00
00888 00889	.00
00890 00891	.00
00892 00893	.00
00894 00895	.00
00896 00897	.00
00898 00899	.00
00900 00901	.00
00902 00903	.00
00904 00905	.00
00906 00907	.00
00908 00909	.00
00910 00911	.00
00912 00913	.00
00914 00915	.00
00916 009	

Address Histogram, 100 Segment Number 130

000000	0000
000001	0000
000002	2.00
000003	3.00
000004	2.00
000005	1.00
000006	2.00
000007	2.00
000008	.00
000009	.00
000010	2696.00
000011	635.00
000012	3788.00
000013	20.00
000014	10.00
000015	1.00
000016	1.00
000017	1.00
000018	1.00
000019	1.00
000020	1.00
000021	1.00
000022	1.00
000023	1.00
000024	1.00
000025	1.00
000026	1.00
000027	1.00
000028	1.00
000029	1.00
000030	1.00
000031	1.00
000032	1.00
000033	1.00
000034	1.00
000035	1.00
000036	1.00
000037	1.00
000038	1.00
000039	1.00
000040	1.00
000041	1.00
000042	1.00
000043	1.00
000044	1.00
000045	1.00
000046	1.00
000047	1.00
000048	1.00
000049	1.00
000050	1.00
000051	1.00
000052	1.00
000053	1.00
000054	1.00
000055	1.00
000056	1.00
000057	1.00
000058	1.00
000059	1.00
000060	1.00
000061	1.00
000062	1.00
000063	1.00
000064	1.00
000065	1.00
000066	1.00
000067	1.00
000068	1.00
000069	1.00
000070	1.00
000071	1.00
000072	1.00
000073	1.00
000074	1.00
000075	1.00
000076	1.00
000077	1.00
000078	1.00
000079	1.00
000080	1.00
000081	1.00
000082	1.00
000083	1.00
000084	1.00
000085	1.00
000086	1.00
000087	1.00
000088	1.00
000089	1.00
000090	1.00
000091	1.00
000092	1.00
000093	1.00
000094	1.00
000095	1.00
000096	1.00
000097	1.00
000098	1.00
000099	1.00
000100	1.00
000101	1.00
000102	1.00
000103	1.00
000104	1.00
000105	1.00
000106	1.00
000107	1.00
000108	1.00
000109	1.00
000110	1.00
000111	1.00
000112	1.00
000113	1.00
000114	1.00
000115	1.00
000116	1.00
000117	1.00
000118	1.00
000119	1.00
000120	1.00
000121	1.00
000122	1.00
000123	1.00
000124	1.00
000125	1.00
000126	1.00
000127	1.00
000128	1.00
000129	1.00
000130	1.00
000131	1.00
000132	1.00
000133	1.00
000134	1.00
000135	1.00
000136	1.00
000137	1.00
000138	1.00
000139	1.00
000140	1.00
000141	1.00
000142	1.00
000143	1.00
000144	1.00
000145	1.00
000146	1.00
000147	1.00
000148	1.00
000149	1.00
000150	1.00
000151	1.00
000152	1.00
000153	1.00
000154	1.00
000155	1.00
000156	1.00
000157	1.00
000158	1.00
000159	1.00
000160	1.00
000161	1.00
000162	1.00
000163	1.00
000164	1.00
000165	1.00
000166	1.00
000167	1.00
000168	1.00
000169	1.00
000170	1.00
000171	1.00
000172	1.00
000173	1.00
000174	1.00
000175	1.00
000176	1.00
000177	1.00
000178	1.00
000179	1.00
000180	1.00
000181	1.00
000182	1.00
000183	1.00
000184	1.00
000185	1.00
000186	1.00
000187	1.00
000188	1.00
000189	1.00
000190	1.00
000191	1.00
000192	1.00
000193	1.00
000194	1.00
000195	1.00
000196	1.00
000197	1.00
000198	1.00
000199	1.00
000200	1.00

7119
 1.00 > 10% of tot
 execution
 time.

Figure 5. Detail of Formatter Activity (compare Figure 4).

6. Miscellaneous Comments

It must be appreciated that the sampler does disturb the system. At least the timer interrupt will cause an extra dispatch if the CPU is operating on the user stack. If the sampling interval is made too small the sampler might be measuring the effects of its own perturbations. A sampling interval of greater than 3 milliseconds has been found to yield accurate measurements. However, when measuring events of short duration shorter sampling intervals may be necessary. In such cases it is important that the sampler does not interrupt the dispatch which it causes. To prevent this the sampling interval should never be less than 350 microseconds. If the sampler does interrupt its own dispatch the CPU will be locked up in an infinite loop. To terminate this loop the sampler can be manually disabled by setting bit 0 to 1 in the word at absolute location

4* <starting drt number> +3.

The sampling interval selected by means of the \$TIME command is the time between exit from the sampler and the subsequent interrupt returning to the sampler. Thus the time spent in the sampler itself is excluded and the only perturbation is the possible extra dispatch (approximately 300 microseconds). However, relative timing of other I/O operations is disturbed. When measuring programs with high I/O activity longer sampling intervals are recommended.

Overhead when sampling all segments at 3 millisecond intervals causes approximately 20% degradation in throughput.

In order to measure activity in segment zero (external interrupts) it is necessary to ensure that the extra clock/TTY interface is polled for highest priority. This will also guarantee the precision of the sampling interval. However, unless precise sampling intervals are specifically desired it is strongly recommended that the randomizing option (\$RAND) be used. When using constant sampling intervals the resultant histograms have been seen to exhibit spurious spikes due to synchronization of the sampler and the sampled program. The randomizing option has produced reliable data in all cases.

The sampling system is not particularly elegant and many improvements could be made in its operation. However it is a useful tool for determining which sections of code consume the most processing time. These results are sometimes surprising even to the programmer who wrote the code. Optimization of software can certainly be accelerated if it is possible to identify the "10% of code which consumes 90% of CPU time".

TO Measurement Distribution

DATE May 1, 1973
SUBJECT Modifications to The Software Sampling System

The original sampler used the switch register to provide a parameter to the random number generator for the randomized timing interval option. To free the switch register this parameter is now fixed and resides in the sampler's data area. A "\$FIX" command has been added to reset the "\$RAND" command. The default is a constant sampling interval; "\$RAND" introduces randomized jitter and "\$FIX" resets the sampling interval to be constant.

The magnitude of the randomized jitter has been doubled by using a random number between 0 and 511 to load the count register of the extra clock/TTY board. This requires that randomization should not be used for sampling intervals less than 600 microseconds. However, the greater randomness of the samples improves the validity of results. For most measurements a sampling interval of 11 milliseconds with randomized jitter will produce reliable results. In this case the sampling interval will vary randomly from 6 to 11 milliseconds ensuring that dispatches caused by the sampler are not sampled.

To conserve tape and reduce the overhead caused by the sampler the output record format has been modified. Two word records are now produced with a blocking factor of 64. The first word of each record contains the status register of the interrupted segment and the second word contains its P register. The data reduction program has been modified to handle the new format. Of course, stack size statistics are no longer available but could be resurrected if necessary.

The new sampler can be set up using the program SETUP9 and the associated data reduction program is TAPE9.

Overhead caused by sampling all segments with a randomized sampling interval of 11 milliseconds is 7%. The Central Limit Theorem predicts that 25,000 samples produce results accurate to 1% of total number of samples with 99.9% confidence. Indeed, experiments have shown that segment histograms produced by the sampler agree to this extent with measurements made using the SUM hardware monitor. These figures do not account for artifacts caused by disabling interrupts.

It should be noted that if interrupts are enabled just prior to an exit the sampler can only interrupt after the exit instruction has executed. This can produce spikes in the histograms of segments calling procedures which disable interrupts. The aberration will be avoided if there is at least one instruction between the enable and the exit.

In Sampler EOF now rewinds tape.

In Disk read 11 sec , `:FILE FTNOT, DEV=TAPE; REC=128, 1, F; NOLABEL`
↑

DATE May 14, 1974
SUBJECT Program Sampling

The 3000 System Section has a program, SAMPLER, which determines relative time spent in the various portions of a program. SAMPLER is very useful in optimizing program code by determining what sections of a program is most heavily used. Unfortunately SAMPLER requires hardware and software modifications to a 3000 system to run. A new version of SAMPLER called SAMPLE is now available which requires no hardware changes and only a minimal software change. Specifically about seven lines of SPL code must be added to procedure TIPC of EXIN. Some capabilities have been lost in SAMPLE that was available in SAMPLER. However the lost capabilities only affect users who are sampling system code segments. For user code segments SAMPLE is fully as capable as SAMPLER. SAMPLE is available for internal use from the 3000 System Section. A write up on SAMPLE has been enclosed for evaluation.

DL/kg

DL

INTRODUCTION

SAMPLER is a package of three programs for use in measuring the relative time spent within various portions of a program. The three programs are:

SAMPLE which the user runs to initiate the sampling of the program code,
SAMPLING which performs the actual sampling and
DRS which reduces the data generated by SAMPLING to a user readable form.

OPERATING INSTRUCTIONS

1. Load the program to be sampled and obtain the CST numbers of the segments to be sampled.
2. Run SAMPLE.
3. Run DRS to reduce the data tape created by SAMPLE.

SAMPLE writes the sample data on a file called SAMTAPE which is assumed to be a magnetic tape but can be a disk file. DRS reads the sample data from a file called DRSTAPE, also assumed to be a magnetic tape and lists the output on a file called DRSLIST, assumed to be a line printer.

MESSAGES

MACHINE ID? From SAMPLING. Type in name of 3000 that test is being run on.

SAMPLE ID? From SAMPLING. Type in identification for sample run.

SAMPLING INTERVAL = 50 MS? From SAMPLING. Type "Y" if a sampling interval of 50 milliseconds is to be used, else type the number of milliseconds between samples. The minimum interval is 10 milliseconds and the maximum interval is 1000 milliseconds. The shorter the interval the greater the number of samples that will be obtained in a given time span. The larger the number of samples, the more valid the results will be. However a short sampling interval will result in considerable degradation of system performance. At 10 milliseconds the system can be expected to run 70 percent slower, at 50 milliseconds degradation will be about 15 percent. A minimum of 7 percent degradation is to be expected.

CST #'S? From SAMPLING. Type in a list of CST numbers of segments to be sampled. The numbers are assumed to be octal and are typed in the following format

```
<CST list>::=<number range>|  
                <number range>,<CST list>  
<number range>::=<CST number>|  
<CST number 1>/<CST number 2>
```

<CST number 1> should be less than <CST number 2> and specifies that all segments with CST numbers between <CST number 1> and <CST number 2> inclusive is to be sampled. <CST list> can be continued onto a second line by typing an &.

Example:

231, 240/243, &
277

Indicates that segments 231, 240, 241, 242, 243 and 277 are to be sampled.

PRINT CSTAB(*)? From SAMPLING. Type "Y" to get a listing of CST numbers that will be sampled else hit carriage return.

TYPE 'STOP' TO STOP: From SAMPLE. Type "STOP" at any time to stop sampling.

UNABLE TO CREATE SAMPLING PROGRAM From SAMPLE.

Possible causes are

1. SAMPLE do not have PH capability.
2. User or account does not have AS priority.
3. SAMPLING exists in another group or account.

START = <CST number 1> > STOP = <CST number 2>.

From SAMPLING. Self-explanatory. If this error occurs, check CSTAB(*) afterwards.

*** BAD NUMBER *** From SAMPLING. A bad CST number was inputted. This message will be preceded by <CST list> up to and including the errant number but not beyond. Most likely cause is typing in decimal CST numbers or extraneous % preceding CST numbers. If this error occurs, check CSTAB(*) afterwards.

UNABLE TO OPEN SAMTAPE. From SAMPLING. Will be followed by file error information.

UNABLE TO WRITE HEADER From SAMPLING. Will be followed by file error information.

QUIT P=11 This is a MPE message. SAMPLING QUITs with P=11 when an error occurs on SAMTAPE.

DRS OUTPUT

DRS prints:

1. A header page which is self-explanatory.
2. A summary histogram showing relative time spent in each segment that was sampled.
3. Detail histograms of each segment that was sampled.

The summary histogram consists of:

1. The CST number of the segment.
2. The histogram.
3. The number of samples taken from the segment, percentage of total samples the numbers represents and cumulative percentages.

The user will be prompted for how detail he wants the detail histograms to be.
The prompt is

SEG # <CST number> (<number of samples>SAMPLES)?

Hit carriage return to omit the histogram of this segment. Type zero to get a one page histogram. DRS will scale the histogram to fit in one page. Type an integer n to get a histogram where the interval are n words wide. The detail histograms consists of:

1. Segment CST number.
2. Two PB relative addresses in octal specifying the beginning and end of each interval.
3. The histogram.
4. The number of samples in the interval.
5. The percentage the number is of the total number of samples for this segment.
6. Cumulate percentage.

DRS contains a restart facility and will ask the user if he wishes to restart.
Type "Y" or "N" as desired.

NOTES:

SAMPLE and SAMPLING require PH and PN capability.

SAMPLING requires AS priority.

AUTOMATIC PROGRAM RESEGMENTATION

DATE: February 1, 1974
SUBJECT: Automatic Program Resegmentation on the
HP 3000

A new performance enhancement tool has been developed. This tool provides for the resegmentation of programs based on data gathered under actual program operation. The resegmentation system will be discussed first, followed by a case study.

The resegmentation system consists of three programs, and utilizes the segment trace facility to gather the data necessary for resegmentation. See Figure 1 for an overall view of the operation. Only two items are needed to start the process. These are a USL file of the program requiring resegmentation, and an adequate test case to exercise the program when segment trace is performed.

The first program (NICRSP) takes the USL file and resegments it such that no procedure in a segment calls any other procedure in that segment. This elimination of internal segment calls is necessary for segment trace since internal procedure calls cannot be traced. The resultant USL file is then prepared into a program file. The program file is run and the segment trace data gathered. Then this USL file and the trace tape are input to program two (ORTDRP). This program reduces the trace data into a usable form for the third program and puts the data in a disk file. Finally, the resegmentation program is run (ORSP) using the file built by program two and the USL file. The output is the resegmented USL file.

It is a good idea to use a copy of the USL for resegmentation. Then the original can be kept in case another resegmentation is desired with a different

segment size. The maximum segment size is kept in DB+0 of both program one and three. This location may be modified by calling debug when the program is started. This may be done by starting the program at its secondary entry point DBG. Note that program's one and three require privileged mode and process handling capability. Another requirement is that when the program is traced, the segment numbers are assigned in a linear order. This may be verified by running the program with the LMAP option. This assignment order can be insured by starting immediately after cold load.

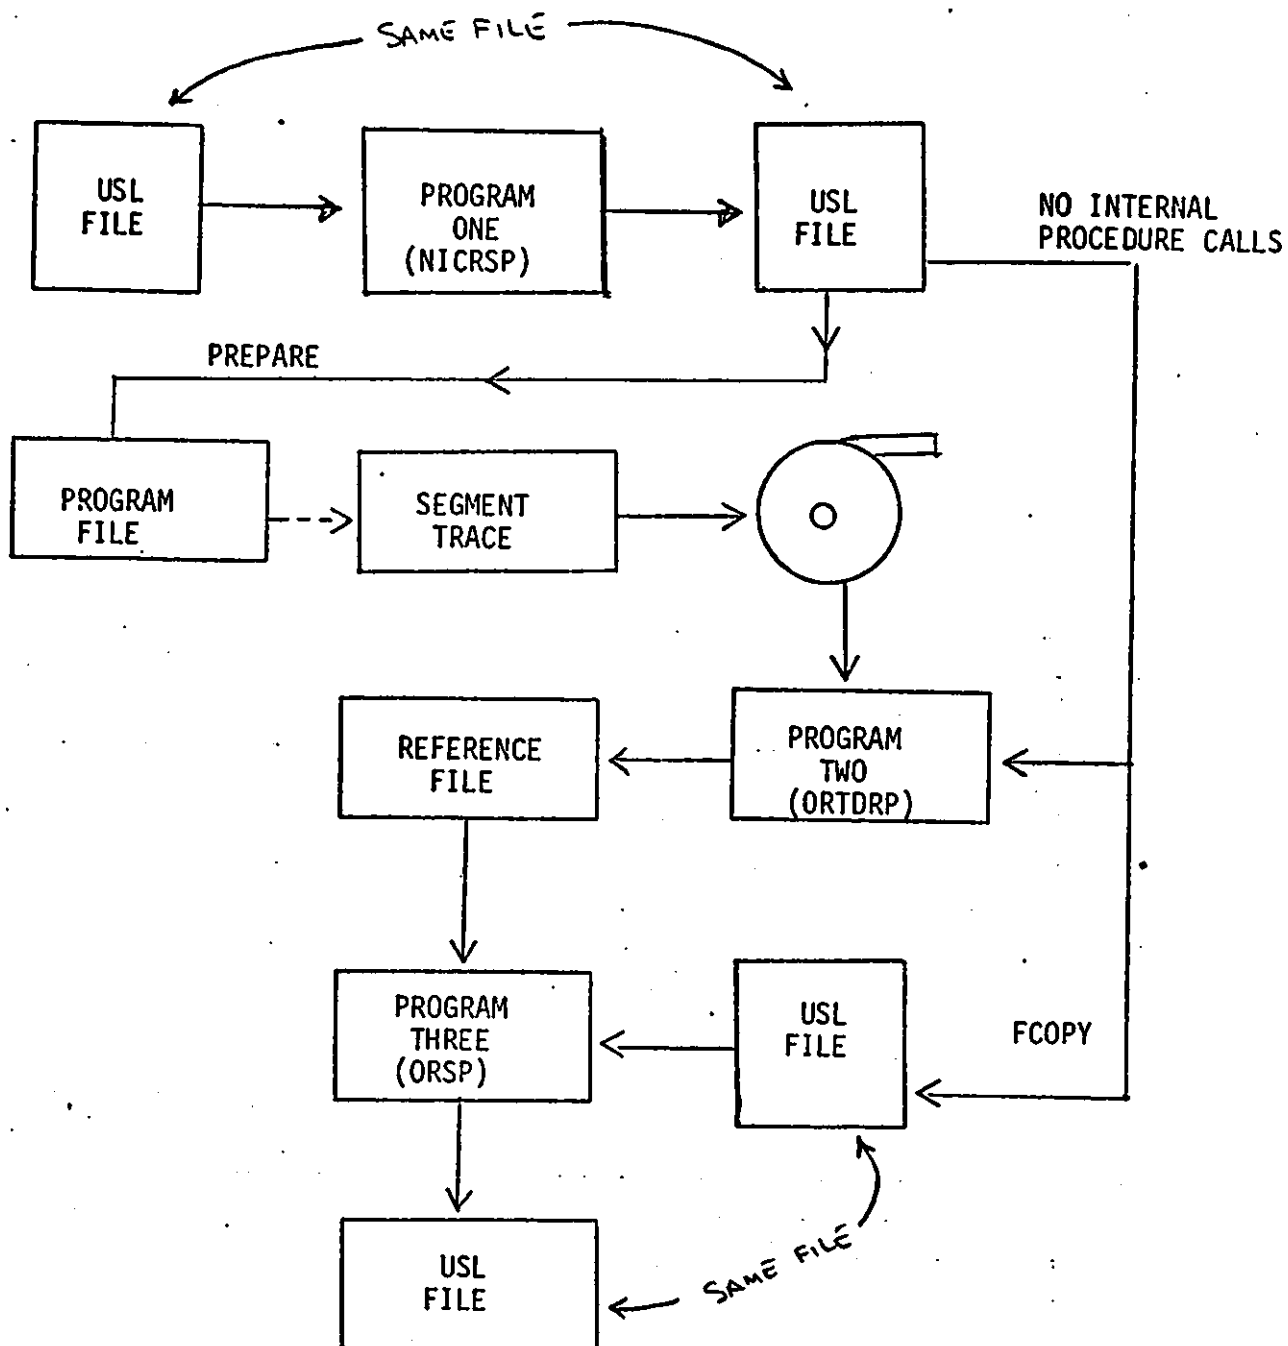


Figure 1
Resegmentation Procedure

A Case Study - COBOL

The COBOL compiler was chosen as a test case because it is one of the largest subsystems and its authors had already carefully segmented the code by hand. Thus, it would test the program's ability to resegment a large system, and the results could be directly compared to the original compiler. Four different tests were conducted. The first test consisted of resegmenting the compiler based on the results of tracing one compile. This same compile was then tested against the original compiler. The following three tests were conducted on a version of the compilation of nine different source files. Tests were made on three over the compile of nine different source files. Tests were made on three modes of operation; stand-alone, multiprogramming against itself (code sharing) and multiprogramming against another large subsystem. The results of each test are shown in tables one through four.

Each test was conducted on four different maximum segment sizes; one thousand through four thousand word segments in multiples of one thousand words.

PROGRAM	Number of Segments	Elapsed Time	Code Segment Faults
COBOL	28	147	587
4K Reseg.	20	206	1423
3K Reseg.	28	195	1376
2K Reseg.	39	162	949
1K Reseg.	56	132	451

Table 1 - One program traced & run alone

PROGRAM	Number of Segments	Elapsed Time	Code Segment Faults
COBOL	28	336	1069
4K Reseg.	23	484	2874
3K Reseg.	28	342	992
2K Reseg.	34	397	2078
1K Reseg.	52	325	859

Table 2 - Cobol run alone (4 Compiles)

PROGRAM	Elapsed Time	CPU Run Time	Code Segment Faults
COBOL	197	83	1915
4K Reseg.	178	77	1414
3K Reseg.	167	73	1450
2K Reseg.	165	71	1680
1K Reseg	153	71	1631

Table 3 - Cobol multiprogrammed against itself (one compile)

PROGRAM	Elapsed Time	CPU Run Time	Code Segment Faults
COBOL	206	85	1807
4K Reseg.	225	82	2023
3K Reseg.	189	80	1589
2K Reseg.	201	83	1857
1K Reseg.	225	94	2321

Table 4 - Cobol multiprogrammed against another large subsystem (one compile)

An examination of the tables shows that programmatic resegmentation can do about as well as a human. Different segment sizes yield much different results, so some experimentation must be done to pick a good segment size. For this experiment, the 3K segment size appears to be about the best of the four sizes used.

SEGMENT TRACE SYSTEM

Segment Trace System

The Segment Trace System (STS) uses the hardware "trace" facility to collect data pertaining to processes at the time of intersegment transfers caused by PCAL's and EXIT's. Sufficient information is available to gather statistics on both code and data of processes, or to examine an individual segment. Data are collected on unit 2 of the system magnetic tape unit (DRT 6) and reduced offline on the 3000 by a data reduction program. Because of the amount of code executed for each transfer traced and the non-overlapped tape I/O, tracing large numbers of segments can cause severe system performance degradation. The trace segment has been made an integral part of the operating system to reduce the problem of incorporating it after each MPE update and to make it available on every development system.

CONSTITUENTS OF STS

STS consists of the following three software components:

1. An MPE system containing a special version of the IIRII code segment and a 300 word data segment.
2. A program to control the tracing process (TCP).
3. A data reduction program.

STS is inert when no trace bits are set and should cause no noticeable loss of system performance.

USING STS

1. ~~COLD~~ load a copy of MPE containing the trace segment.
2. Ready a mag tape on unit 0.
3. Run the TCP program from a session. Use the command sequence (explained in detail on page 7).

INIT	clock dri
SET	set trace bits
RUN	tracing begins
STOP	tracing ceases
CLEAR	Clear trace bits
EOF	writes EOF ON TAPE
EXIT	terminates TCP

4. Rewind and dismount tape.
5. Process tape using reduction program.

STS - TRACE

Segment trace data are collected by a segment added to segment ININ in MPE. This segment is called from the absence trap segment (\$14) and the trace trap segment (\$16) when these segments have determined that tracing is to be done. The trace trap segment processes the Break and Control - Y features of a session, and these will pre-empt a segment trace. To aid understanding of the capabilities and limitations of trace, a description of the hardware trace feature follows.

A procedure call to a segment will cause a trap to the trace segment if bit 2 of the first word of that segment's CST entry is set to one and bit 0 is not set to one (absence).




Figure 1 CST Entry

Trace and absence traps due to PCAL's are identical in their effect on the stack. The result is two stack markers and an external label.

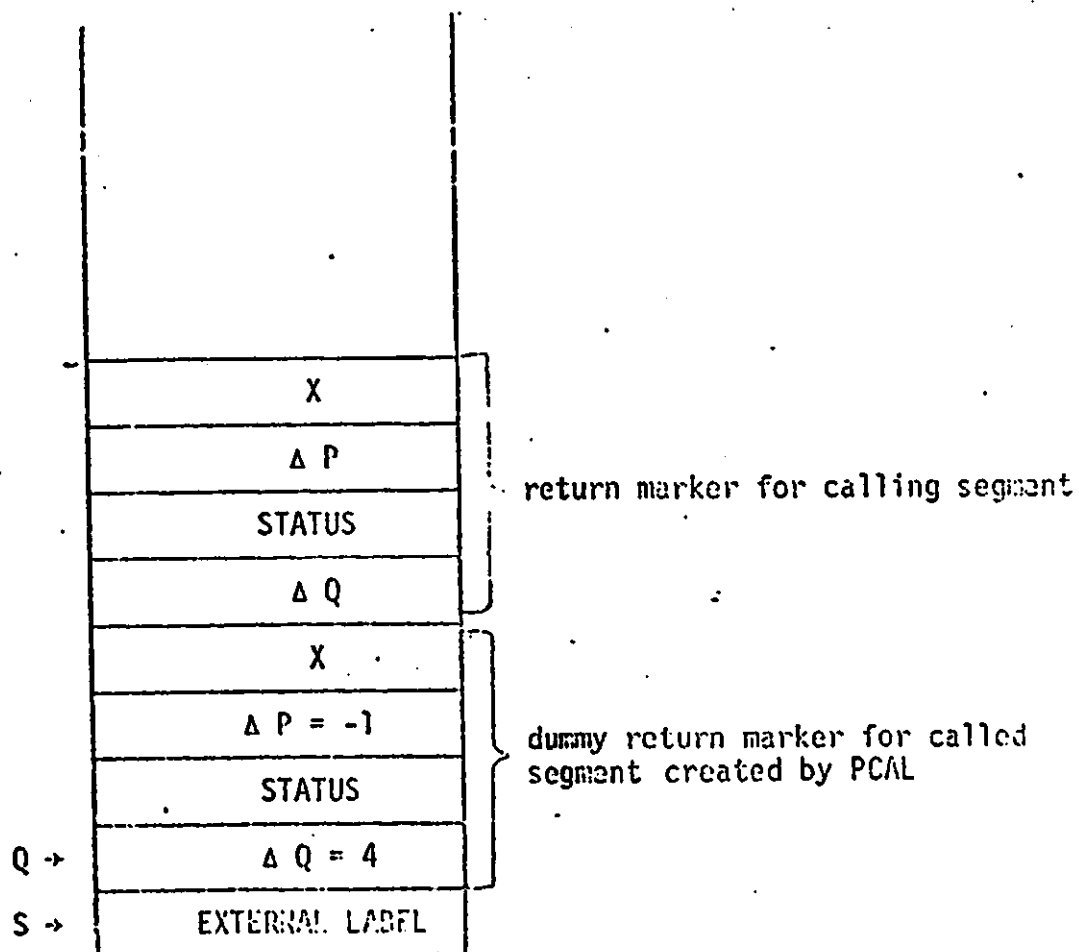


Figure 2 - Stack After PCAL Trace or Absence Trap

The external label is a copy of the one referenced by the PCAL in the calling segment. With this label, the trace routine can calculate the correct delta P in the dummy marker to enter the called segment.

An EXIT trace occurs if bit 0 of delta P in the return marker is a one and an EXIT instruction is executed using that marker. The marker is left on the stack and control is passed to the trace trap segment. The value N from the




Figure 3 - Stack After EXIT Trace Trap

EXIT N of the called segment is pushed on the stack so the trace trap processor can XEQ the correct EXIT N off the stack at the completion of the trace process. The trace routine must reset delta P to the correct value before the EXIT is executed. Currently, bit 1 of delta P determines if it is a trace trap (bit 1=1) or a break trap and the Trace Segment (\$16) takes the appropriate action.

The hardware trace facility has some limitations. PCAL's and EXIT's within a segment cannot be traced, nor can interrupts be traced. Another important exception is a PCAL of a segment to itself via an external label. The PCAL can be traced, but the corresponding EXIT cannot. The trace segment will suppress tracing such EXIT's.

The trace data collection segment works as follows. When the trace segment is called, the trace options word in the system global area (SYSPG + 245)(figure 4) is checked. If bit 1=0 then the trace segment will exit to the caller with condition code set to less than. This means that no tracing is being done. If bit 2 = 0, the trace request will exit without collecting any data. This bit is used to start and stop the physical collection of data. If bit 3 = 1 and the trace was from a PCAL, delta P of the return segment will be set so the corresponding EXIT can be traced if a segment is not calling itself via an external label. If the Clock DRT is non-zero, timing measurements will be made.

	0	1	2	3	4	5	8 - 15
WORD 0	R	I	A	T	O	C	PIN
1	LABEL (PCAL) or N (EXIT)						
2	CALLER STATUS						
3	CALLER DELTA P						
4	STACK DB (QI - 4)						
5	DB						
6	DL						
7	Q						
8	Z						
9	STACK DST <CPCB (2)>						
10	EXTRA DST <CPCB (3)>						
11	TIME						

where R = Record Type 0 = Trace Record

1 = INFORMATION RECORD

I = Interrupt Bit - on interrupt occurred between this record and the previous record.


A = Absence - the call for this record came from the absence trap.

T = extra clock being used.

O = timer overflow bit (18th bit).

C = timer carry bit (17th bit).

Figure 4 - Trace Record Format



where D = dispatch bit (set to one each time the dispatcher is executed)
 T = trace bit (set to one while tracing is active)
 A = arm bit (if this bit is a one, trace trap will be processed, else they will be ignored)
 E = exit bit (if this bit is a one, the trace processor will set the EXIT marker to invoke a trace trap each time a PCAL trace is processed)

Figure 5 - Trace Control Format in System Global Area

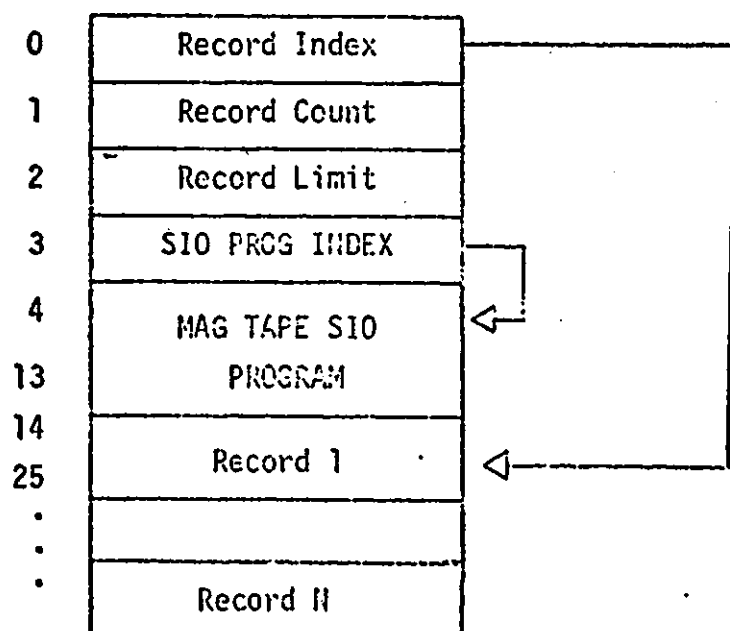


Figure 6 - Trace Data Segment Format

After the options check, the trace tape record is built. Word 0 contains 6 bits of status information and the process identification number as shown in Figure 4. Word 1 contains the label if a PCAL is traced or H if an EXIT H is traced. Words 2 and 3 contain the status and delta P as they appear in the calling segment stack marker. Word 4 contains the current process stack DB located in Q1 - 4. Words 5 through 8 contain the values of the process stack registers. Words 9 and 10 contain the stack DST and extra data segment DST located in the current process control block. If timing measurements are being made, word 11 will contain a current time stamp.

An extra clock board provides the time base for timing measurements. The clock is read at the beginning of each trace, and this value is loaded back into the clock at the completion of the trace, so that the trace overhead is removed from the measurement. The LR = CR and LR = CR overflow bits of the clock are used to extend the clock period to 18 bits instead of the 16 bits available in the clock counting register. At a ten microsecond counting interval, 1.9 seconds of time can elapse before measurement overflow occurs. This cannot happen since time of day updates occur once each second in MPE and timing analysis will not be done through interrupts.

Trace Control Program

The Trace Control Program (TCP) provides the user with a convenient means of setting and clearing trace bits, and initializing and controlling the trace operation. TCP must run in privileged mode so that user account capability is needed.

The commands to TCP may be input from a terminal (TCP prompts with a "?") or card reader, depending on the mode (session or job). Commands are entered one per line (card). Terminal input is terminated with a carriage return. Some examples of command syntax are shown in Figure 7.

A blank or comma may be used as a data item separator, but blanks are otherwise ignored. Numbers may be either octal or decimal, with a "%" character preceding an octal number. A "/" between two numbers indicates a range.

```
INIT 0
TRACE 45, %27/41 %102
RUN
STOP
CLEAR %20/130
EXIT
```

Figure 7 - TBC Command Examples

TBC features a comprehensive set of diagnostics. All but the following two abort the current command:

RAN CST NUMBER X (X <%20 or too large)

UNASSIGNED CST X (This CST not used)

If ~~there~~^{this} occurs within a <range>, execution of the <range> continues.

The TBC program itself is divided into three parts, which are the scanner, the interpreter, and the main program. The scanner picks "tokens" from the current line, leaving the ASCII representation in the byte array TOKEN, a type indicator in T, and the value, if numeric, in V. The main procedure of the scanner is NEXT. A token is defined to be a comma (",") a slash("/"), a carriage return, or a string of alphanumerics not including blanks, commas, slashes, or carriage returns.

The main program fetches new lines, requests the first token (by calling NEXT) and attempts to interpret it as a command. If this is successful, control passes to a procedure in the interpreter which executes the command.

One can add new commands simply by adding an equate for the "type" of the command, enlarging the case statement in the main program, and inserting an IF clause in procedure COMMAND to detect it. An actual procedure must be put in the interpreter section to perform the actual execution.

```

<command>:: = TRACE <t list> |
             CLEAR <t list> |
             SHOW <addr> |
             STORE <addr>, <value> |
             INIT <drt> |
             RUN |
             STOP |
             EOF |
             EXIT

<T list>:: = <t element>, <t list>
<t element>:: = <num>|<range>
<range>:: = <lower bound>|<upper bound>
<lower bound>:: = <num>
<upper bound>:: = <num>
               <addr>:: = <num>
               <value>:: = <num>
               <drt>:: = <num>
               <num>:: = Octal number  $\leq$  177777
                       (leading 0's ignored)

```

Figure 8 - TBC Command Syntax

TRACE/CLEAR <t list>

Trace bits are set/cleared in CST entries described by <t list>. If a <range> is specified, the trace bits in <lower bound> through <upper bound> inclusive are affected. <lower bound> must be \leq <upper bound>.

INIT <drt>

All necessary initialization is performed, <drt> is the DRT index of the extra clock board. This command must be issued before any RUN's.

RUN

Tracing begins.

STOP

Tracing ceases; it can be restarted with a RUN.

EOF

An end-of-file mark is written on the trace tape.

EXIT

TBS terminates (cannot be tracing when this command is executed).

Figure 9 - TBC Command Semantics

Data Reduction Program

The data reduction program will be capable of presenting the data either by process or by code segments. The segment statistics should be useful for analyzing library routines and other segments where only information about those particular segments are of interest. Process statistics can be used to examine both the code and data behavior of software sub-systems.

Segment statistics include both measurements of time and statistics concerning information about the segments which call the traced segment. Timing data includes the total time spent in the segment, the average time and standard deviation of all the calls to that segment. The calls will be broken down by STT numbers, calling segment numbers, and process numbers.

Process statistics will examine the minimum and maximum stack size; the number of times the stack is added to; the size of the user's own area, and how many times that size is changed. Extra data segment use can also be determined. The flow of program control can be shown by showing the segments used, who called them, and how many times they were called.

May 7, 1973

STS - TRACE Data Reduction Program

Segment Trace Data Reduction Program

A program is available to process the tape produced by the STS-TRACE program. (For a description of STS-TRACE, consult prior memo dated March 20, 1973.) Output is in the form of entry, caller, and timing statistics for each segment traced.

The data reduction program file name is TRACERED and is activated using the MPE Run command. The program may be run from batch or interactive access. If interactive access is used, the program prompts with a "?". Currently the program will recognize three commands, as shown:

SEGMENT - causes the program to process a trace tape and produce segment statistics. This command may be executed once only per line.

CALL - in addition to the segment statistics above, this option causes the processing of a trace tape to include the caller statistics, and the caller statistics to be included in the output. Note that this command must be

issued each time the segment command is issued if caller statistics are desired.

EXIT - causes the reduction program to terminate after other commands included in this line have been executed.

The commands may be input in any order. Any non-alpha character may be used for delimiting commands. Any alpha input other than the above will cause an error message to be emitted, but will otherwise be ignored. After input, the commands will be executed and, if no EXIT command was included, will request the next command.

All statistical output from the program is directed to the line printer. The following is an explanation of each column of data. Please refer to Fig. 1.

The first eleven columns comprise the called segment statistics. The first column is a three digit octal segment number. The second column is the total PCAL entries to that segment.

The third column is the occurrence of PCAL absences only. EXIT absences cannot be traced, and therefore cannot be displayed. The fourth column is absences as a percentage of the total entries. Column 5 is a three digit octal segment transfer table (STT) entry number. Column 6 is the total calls to that STT. Column 7 is STT entries as a percentage of the total entries. Column 8 is the number of timed STT entries. This is the total number of STT entries on which it was possible to gather timing information. Column 9 is the number of timed entries as a percentage of the total STT entry count. Column 10 is the average time per STT call in milliseconds with all trace overhead removed. Note that this time is the time spent in this segment less the time spent in other segments which were being traced and were called by this segment. Column 11 is the standard deviation of the time per call in milliseconds.

Columns twelve through sixteen comprise the caller statistics. Columns twelve through fourteen contain the caller ID, that is the process identification number, the segment number and delta P value of the calling segment. All three are represented in octal. Column fifteen is the total calls made to this segment by the caller, and Column sixteen is the number of calls as a percentage of the STT entries.

Segment Trace In Action - A Case Study

This section describes an example of the use of the software trace system. A calibration program has been written in order to calibrate the constants in the trace data reduction program which removes the tracing overhead from segment timings. All system output is underlined.

First, the trace control program is brought up and initialized.

:RUN TRACE

HP3000 TRACE CONTROL PROGRAM

? INIT 15

Next, the calibration program is started with the LMAP option to find out which segments the program would use. The segments used were %123 through %137. The calibration program requests the number of passes to make, and at this point, the following trace commands are issued.

? TRACE %123/%137

? RUN

Now trace is ready to operate, so 100 is issued to the calibration program. Tracing proceeds until the calibration program has run to completion. Next, the following commands are given to the trace control program.

? STOP

? EOF

? EXIT

This terminates the program execution. Note that the CLEAR command was not necessary since the CST entries were deallocated when the calibration program terminated.

After the trace operation is complete, the data is reduced. Two reductions were made with and without caller statistics are shown in Figures 2 and 3.

: RUN TRACERED

HP3000 TRACE TAPE REDUCTION PROGRAM

? SEGMENT

? SEGMENT CALL EXIT

Note that the measured time for an EXIT only is .02 milliseconds are shown in segments %131, %132, and %137; whereas a PCAL, EXIT pair is .05 milliseconds as shown in segments %124, %125, %133, and %134. These are approximate since an exit takes 21.7 microseconds, and both a PCAL and EXIT take 50.4 microseconds. They are not in exact agreement since the trace system only receives one clock count each 10 microseconds. Therefore the error could be as large as 10 microseconds for each segment traced.

SEG. ENTRY CNT.	ABSENCES.	2ABS.	STT ENTRY CNT.	ASLGE	TIMED ENT.	%STTE	TIME/C.	SOEVT/C.	PIN	SEG	DELTP	CALLS	%STTE
124	99	1	1.0	001	99	100.0	81	81.8	.05	.00	027 123 00041	99	100.0
125	99	1	1.0	001	99	100.0	85	85.9	.05	.00	027 124 00001	99	100.0
126	99	1	1.0	001	99	100.0	92	92.9	.02	.00	027 125 00001	99	100.0
127	99	1	1.0	001	99	100.0	95	96.0	.06	.00	027 126 00001	99	100.0
130	1	1	100.0	001	1	100.0	0	.0	.00	.00	027 123 00020	1	100.0
131	1	1	100.0	001	1	100.0	1	100.0	.02	.00	027 130 00001	1	100.0
132	1	1	100.0	001	1	100.0	1	100.0	.02	.00	027 123 00017	1	100.0
133	100	1	1.0	001	100	100.0	59	59.0	.05	.00	027 123 00030	100	100.0
134	100	1	1.0	001	100	100.0	61	61.0	.05	.00	027 133 00001	100	100.0
135	100	1	1.0	001	100	100.0	61	61.0	.08	.00	027 134 00001	100	100.0
136	200	1	.5	001	200	100.0	163	81.5	.08	.00	027 135 00001	100	50.0
											027 135 00002	100	50.0
137	400	1	.3	001	400	100.0	389	97.2	.02	.00	027 136 00001	200	50.0
											027 136 00002	200	50.0

Figure 1

SEG. ENTRY CNT. AUSENCES. XAUS. SIT ENTRY CNT. SSEG TIME/C. SSTE TIME/C. SDEVT/C. PIN SEG DELTP CALLS SSTE

124	99	1	1.0	001	99	100.0	81	81.8	.05	.00
125	99	1	1.0	001	99	100.0	85	85.9	.05	.00
126	99	1	1.0	001	99	100.0	92	92.9	.02	.00
127	99	1	1.0	001	99	100.0	95	96.0	.06	.00
130	1	1	100.0	001	1	100.0	0	.0	.00	.00
131	1	1	100.0	001	1	100.0	1	100.0	.02	.00
132	1	1	100.0	001	1	100.0	1	100.0	.02	.00
133	100	1	1.0	001	100	100.0	59	55.0	.05	.00
134	100	1	1.0	001	100	100.0	61	61.0	.05	.00
135	100	1	1.0	001	100	100.0	61	61.0	.08	.00
136	200	1	.5	001	200	100.0	163	01.5	.08	.00
137	400	1	.3	001	400	100.0	389	97.2	.02	.00

Figure 2.

SEG	ENTRY CNT.	ABSENCES	2AHS.	SIT	ENTRY CNT.	MSGE	TIMED ENT.	%SITE	TIME/C	SOEVT/C	PIN	SEG	DELTP	CALLS	%SITE
124	99	1	1.0	001	99	100.0	81	81.8	.05	.00	027	123	00041	99	100.0
125	99	1	1.0	001	99	100.0	85	85.9	.05	.00	027	124	00001	99	100.0
126	99	1	1.0	001	99	100.0	92	92.9	.02	.00	027	125	00001	99	100.0
127	99	1	1.0	001	99	100.0	95	94.0	.06	.00	027	126	00001	99	100.0
130	1	1	100.0	001	1	100.0	0	.0	.00	.00	027	123	00020	1	100.0
131	1	1	100.0	001	1	100.0	1	100.0	.02	.00	027	130	00001	1	100.0
132	1	1	100.0	001	1	100.0	1	100.0	.02	.00	027	123	00017	1	100.0
133	100	1	1.0	001	100	100.0	59	57.0	.05	.00	027	123	00030	100	100.0
134	100	1	1.0	001	100	100.0	61	61.0	.05	.00	027	133	00001	100	100.0
135	100	1	1.0	001	100	100.0	61	61.0	.08	.00	027	134	00001	100	100.0
136	200	1	.5	001	200	100.0	163	81.5	.08	.00	027	135	00001	100	50.0
											027	135	00002	100	50.0
137	400	1	.3	001	400	100.0	389	97.2	.02	.00	027	136	00001	200	50.0
											027	136	00002	200	50.0

Figure 3

May 15, 1973

Changes to the Segment Trace System

cc: Measurement Distribution

The following two changes have been made to the STS trace segment. The first involves the trace record format, and the record involves the operation of the trace code segment.

The following changes have been made to the output record format of the trace segment. Refer to Figure 1. The positions of DB and DL within the record have been reversed. Bite six and seven of the first word are now being used as indicators. Bite six indicates an end-of-data condition, that is, that this record and all subsequent records are invalid. Bit seven indicates that this record was not generated as the result of a trace or absence interrupt, but was generated as the result of a direct procedure all to the EXTRACE segment.

Word	Use								
0	1	2	3	4	5	6	7	8	15
R	I	A	T	O	C	L	E		
LABEL (PCAL) or H (EXIT)									
CALLER STATUS									
CALLER DELTAP									
STACK DB <QI-4>									
DL									
DB									
Q									
Z									
STACK DST <CPCB (2)>									
EXTRA DST <CPCB (3)>									
TIME									

where R = Record Type
I = Interrupt
A = Absence
T = Segment timing
O = Timer Overflow
C = Timer Carry
L = Last Record
E = External Call

The time segment now has the capability of recovering from a tape write error. If the error occurs, the program will stop with a HALT 1. The tape controller status will be displayed in RA so that the source of trouble can be determined. Press run, and the program will issue a backspace and gap to the tape drive. If this is successful, the operation will continue in a normal manner. If unsuccessful, the program will stop with a HALT 2. Pressing run will allow tracing to continue, but the tape will contain an error which will probably inhibit the data reduction program from reading the tape past that point. This means that it is best to terminate the tracing operation and restart with a different tape mounted.