



DECUS U.S. CHAPTER SIGs NEWSLETTERS

DATATRIEVE/4GL SIG	DTR
LARGE SYSTEMS SIG	LS
OFFICE AUTOMATION SIG	OA
PERSONAL COMPUTER SIG	PC
ARTIFICIAL INTELLIGENCE SIG	AI
BUSINESS APPLICATIONS SIG	BA
DAARC SIG	DAR
DATA MANAGEMENT SIG	DMS
EDUSIG	EDU
GRAPHICS SIG	GRA
HARDWARE MICRO SIG	HMS
IAS SIG	IAS
LANGUAGES AND TOOLS SIG	L&T
MUMPS SIG	MMP
NETWORKS SIG	NTW
RSTS SIG	RST
RSX SIG	RSX
RT SIG	RT
SITE MANAGEMENT & TRAINING SIG	SIT
UNISIG	UNI
VAX SYSTEMS SIG	VAX
LIBRARY INFORMATION SECTION	LIB
“HOW TO” SUBMIT AN ARTICLE . . . GROUP FORMS	HOW
SIG INFORMATION SECTION	SIC
QUESTIONNAIRE SECTION	QU

GENERAL TABLE OF CONTENTS

SECTIONS	PAGE NO.
DATATRIEVE-4GL SIG	
.Chairman's Corner	DTR-2
.High Performance DATATRIEVE MENUS	DTR-3
.Ask the Wombat Wizard	DTR-11
.Printing out Pending ALL-IN-1 Mail	DTR-15
.Ask the WOMBAT Wizard	QU-1
.DATATRIEVE/4GL SIG	QU-3
LARGE SYSTEMS SIG	
.Ask Dr. TOPS	LS-2
.Fall DECUS Symposium Preview	LS-4
OFFICE AUTOMATION SIG	
.From the Editor.	OA-1
.Some Things I Have Learned About Using the Message Router with ALL-IN-1.	OA-1
.OA SIG Events in Anaheim	OA-4
.Notes on Notes	OA-4
.OA SIG System Improvement Request Ballot	QU-15
PERSONAL COMPUTER SIG	
.A Rainbow Ballad	PC-1
.IBM PC Software Development on a DEC Rainbow	PC-7
.Digital Announces Phase Down for the Professional 380.	PC-10
.RSX on a PRO?.	PC-12
.Fall Symposium Happenings	PC-14
.PROgramming Quickie	PC-14
.PC Postscript	QU-17
AI SIG	
.Titles of AI Sessions Presented at Nashville, Spring '87.	AI-1
.Titles of AI Sessions Submitted for Anaheim, Fall '87.	AI-1
.Saturday Meeting at Symposium.	AI-2
.From the Symposia Rep:	AI-2
.Technical Sessions	AI-2
HMS SIG	
.From the Editor.	HMS-1
.Micronotes Update	HMS-2
.DECworld: A Hardware Bonanza.	HMS-3
.Traditional DEC Hardware Support	HMS-6
.Video Tape Drives.	HMS-11
.Hardware Submission Form	QU-5
IAS SIG	
.Contribution Guidelines.	IAS-i
.From The Editor's Terminal	IAS-1
.Ten Years Ago This Month.	IAS-2
.Fall Symposia Preview	IAS-3
.IAS Version 3.2C SPR's	IAS-4
.The Program of the Month Club	IAS-5
.A Round Tuit	IAS-8
.IAS WHIMS	QU-7

L&T SIG

.Ediotr's Notes	L&T-2
.L&T SIG - COBOL Working Group	L&T-3
.Languages and Tools Woods Meeting	L&T-5
.Languages and tools Working Group Statements	L&T-14
.Languages and Tools Masters Application	QU-9
.Languages and Tools Wishlist Questionnaire.	QU-11

MUMPS SIG

.GLOBAL ACCESS	MMP-1
--------------------------	-------

RSX SIG

.Food for Thought	RSX-1
.The Editor's Corner.	RSX-1
.The Bag of Tricks: MACRO-11	RSX-5
.Bulletin Board Notes	RSX-6
.Initializing DCL with Node Specific Prompts.	RSX-6
.PDP-11 Performance Benchmark.	RSX-8
.SIG Business Update	RSX-9
.VT220 Downline Loadable Character Sets	RSX-9
.Weird RSX Science.	RSX-18

RT SIG

.From the Editor.	RT-2
.A Symposium Paper? Who ME?	RT-3
.RT-11 Abstracts	RT-4
.RT-11 SIG Tapes on RX33	RT-17
.FPJ Chip Fix	RT-18
.Directory Search Program	RT-19
.SCURT Announcement.	RT-20
.A Simple Time Card System	RT-21
.The Flit Gun	RT-35
.RT-11 Wish List Survey	QU-23

SITE SIG

.Letter from the Editor	SIT-1
.Student Employees at a University Computing Facility	SIT-2

VAX SIG

.Undocumented System Space Watchpoint Utility	VAX-3
.Editor's Workfile	VAX-7
.Migration to MicroVMS from DEC Operating Systems	VAX-8
.Online Security Monitoring System	VAX-25
.VAX/VMS System Performance Management	VAX-35
.Converting Code for Parallel Processing	VAX-48
.INPUT/OUTPUT	VAX-62
.INPUT/OUTPUT Submission Form	QU-25
.System Improvement Request Submission Form	QU-27
.INPUT/OUTPUT Submission Form	QU-25
.System Improvement Request Submission Form	QU-27

LIBRARY

.New Library Programs Available	LIB-1
.Revisions to Library Programs	LIB-6
.DECUS Program Library Changes.	LIB-7

HOW TO SUBMIT TO A SPECIFIC SECTION OF THE NEWSLETTER

.Submitting Articles to Hard News HOW-1
.How to Submit Articles to the RSTS Newsletter HOW-2
.DECUS U.S. Chapter Subscription Service Order Form. HOW-3
.DECUS U.S. Chapter Application for Membership HOW-5

SIG INFORMATION SECTION

.Special Interest Committee Lists SIC-1

QUESTIONNAIRE SECTION

.Ask the WOMBAT Wizard Submission Form QU-1
.DATATRIEVE/4GL SIG PIR Submission Form QU-3
.HMS SIG Submission Form QU-5
.IAS WHIMS QU-7
.Languages & Tools SIG Masters Application. QU-9
.Languages & Tools SIG WISHLIST Questionnaire QU-11
.DATAGRAM QU-13
.Office Automation System Improvement Request Ballot. QU-15
.PC Postscript QU-17
.RSTS Newsletter Reader Survey QU-19
.RSTS/E Wishlist Form QU-21
.RT-11 Wish List Survey QU-23
.INPUT/OUTPUT Submission Form QU-25
.System Improvement Request Submission Form QU-27

The Wombat

EXAMINER

and 4GL
Dispatch

"Increases the Circulation of Anyone in America"

Volume 9

Number 3



DTR

Contributions

Submissions to this newsletter are constantly sought. A submission can be an article, a letter to the Wombat Wizard, a technical tip, or anything of interest to people using or considering the use of Datatrieve or any 4GL product. Submissions on magnetic media are preferred but almost any type will be considered.

Contributions for the newsletter can be sent to either of the following addresses:

Editor, DATATRIEVE Newsletter
c/o DECUS U.S. Chapter
219 Boston Post Road, BP02
Marlboro, MA 01752

Donald E. Stern, Jr
Warner Lambert Company
10 Webster Road
Milford, CT 06460

Table of Contents

DECUS U.S. Chapter SIG Newsletters, Volume 3, No. 3, Nov. 1987
Wombat Examiner and 4GL Dispatch, Volume 9, No. 3

DTR-2	-	Chairman's Corner
DTR-3	-	High Performance Datatrieve Menus
DTR-11	-	Ask the Wombat Wizard
DTR-15	-	Printing out Pending All-in-1 Mail

Chairman's Corner

Joe H. Gallagher, 4GL Solutions, Overland Park, KS

While you are contemplating what part of the turkey you will eat later this month, you had better consider how much further behind you will get if you don't attend the 1987 Fall DECUS Symposium in Anaheim in December.

With about 1000 hours of presentations on the latest developments in hardware and software (over 60 hours of which are in the area of DATATRIEVE and Fourth Generation Languages), a DECUS Symposia is one of the best training values that you could buy. There is no where else (not even DECWorld) where you can directly interact with Digital developers of operating systems, layered products, and hardware.

I have never attended a DECUS Symposia where I didn't learn at least one thing that more than paid for the trip; usually I learn several things.

This Symposium is also a time of special celebration. DATATRIEVE is 10 years old. Ten years ago this Fall, the first version of DATATRIEVE was announced. There will be a small party after Wombat Magic on December 10th to celebrate this tenth anniversary.

Again this symposium, we are encouraging new (and old) volunteers to assist the SIG during the Symposium. Those who wish to volunteer to a chair session or host the suite for two hours should plan to meet in the DTR/4GL SIG Suite at 5:30PM on Sunday, December 6. When you arrive, check the week-end addition of the Update.Daily for details and suite location. At Nashville, SIG volunteers were rewarded for their service with a very handsome wind breaker.

See you in Anaheim.

Developing Applications in DATATRIEVE: High performance Menus

Joe H. Gallagher, PH. D., 4GL Solutions, Overland Park, KS

Although DATATRIEVE was designed to be and is best suited as an on-line, interactive, query and report writing language, its lack of built-in menu capabilities has not stopped a very large number of users from writing big menu-driven applications in DATATRIEVE.

There are basically four ways to create menu-driven applications in DATATRIEVE:

1. using one very large "pre-compiled" procedure
2. using a DCL menu
3. using callable DATATRIEVE, and
4. using logicals and a pseudo-recursive procedure.

These four ways of creating menus in DATATRIEVE have been described in presentations by Chris Wool at the last three DECUS Symposia entitled "Writing Menu-Driven Systems in VAX-DATATRIEVE." A summary of the important features of each method has appeared in the DTR/4GL Session Notes at each Symposia. In addition, magic presentations by Bob Hoover,¹ Lorey Kimmel,² Gary Burton,³ Pat Scopelliti,⁴ Larry Jasmann,⁵ and articles by Diane Pinney using logicals with a pseudo-recursive procedure,⁶ by Barrie Gray using a very innovative approach to circumventing the poor performance of "pre-compiled" procedures,⁷ and by S. Begelman using callable DATATRIEVE⁸ illustrate examples of these techniques.

It is not the purpose of this article to compare and contract the features and capabilities of these four methods, but to describe an extension to one of these methods which, I believe, makes the technique far superior to all the others. The menu technique which I will fully describe below is of the type which uses logicals and a pseudo-recursive procedure, but it contains three important extensions to the basic technique. These three extensions are:

1. the control, prompt, security, and formatting information of the menu system is actually data in a DATATRIEVE domain;
2. the menu system can be subdivided into a very, very large system of sub-menus with essentially no performance penalty; (The method of accomplishing this concept was suggested to me by Tom Considine of Applied Video Systems at the Nashville Symposium.)
3. the pseudo-recursion itself is accomplished by the use of a logical which is only defined just before the pseudo-recursion is invoked.

Consider the following domain and record definitions:

```
DEFINE DOMAIN MENU USING MENU RECORD ON MENU$DATA:MENU.DAT;
DEFINE RECORD MENU_RECORD USING OPTIMIZE
01 MENU_REC.
    03 MENU_GROUP_KEY.      ! group key is primary key
    05 MENU_SYSTEM PIC XX.
    05 RECORD_TYPE PIC 9.   ! 0 - header, 1 - control
    05 OPTION PIC 99.      ! selection (and order)
    03 PROC_NAME PIC X(31). ! procedure name or blank if header
    03 RIGHTS_OWNER PIC X(10). ! tag of the owner of this menu function
    03 DESCRIPTION PIC X(80). ! header information or option
                                ! description
;
```

The fields in the domain MENU are described as follows:

MENU_SYSTEM: The values in the field MENU_SYSTEM are a two character abbreviation for the menu or sub-menu and are used to divide the menu system into easily manageable sub-menus. This field would have values such as "MN" for the main menu, "AR" for the accounts receivable sub-menu, "AP" for the accounts payable sub-menu, "SR" for the supervisor reports menu, etc. The information in this field (the values) is never seen by the user, but is used by the main menu procedure to determine which part or sub-part of the menu system is currently to be accessed.

RECORD_TYPE: There are two types of records in the MENU domain - header records and control records. RECORD_TYPE has the value "0" for header records and "1" for control records. Header records are used to create a menu header or title for each menu. Control records contain the control, security, and prompting information of each option on the menu.

OPTION: The field OPTION controls the order in which header records and control records are displayed on each menu. They are also the selection criterion for choosing an item from a menu.

MENU_GROUP_KEY: The group element, MENU_GROUP_KEY is the primary key of the indexed domain MENU. By using this group element as a primary key, it is possible to avoid sorting the records in order to display them on each menu. This technique of using a group element as the primary key plays a very important role in giving this implementation good performance characteristics.

PROC_NAME: The field PROC_NAME is blank for header records. For control records it contains the name of the DATATRIEVE procedure which is to be executed when this particular menu item is selected.

RIGHTS_OWNER: The field RIGHTS_OWNER contains a string which is the name of the rights identifier which must be possessed in order to be allowed to make this menu selection. The importance of this field will be clarified when we described the tables and procedures below.

DESCRIPTION: The field DESCRIPTION contains the text which is display on the menu for both header and control records. Note that for header records this string field may contain escape sequences which activate terminal characteristics such as double height/double width characters.

The listing of two typical records in the MENU domain might look like:

```
MENU SYSTEM : MN
RECORD TYPE : 0
OPTION      : 01
PROC NAME   :
RIGHTS OWNER : ALL
DESCRIPTION : <ESC>#3          Menu Data Management System
```

```
MENU SYSTEM : MN
RECORD TYPE : 1
OPTION      : 01
PROC NAME   : MOVE TO SUPERS_MENU
RIGHTS OWNER : SUPERVISOR
DESCRIPTION : Move to the Supervisor's Menu
```

The first record is a header record (RECORD_TYPE=0); it can be accessed by those who have the "ALL" rights identifier, and the DESCRIPTION field contains the first part (top part) of a two part entry which makes a double height, double width display of "Menu Data Management System". The second record is a control record (RECORD_TYPE=1); it is menu option 1 which activates the procedure MOVE_TO_SUPERS_MENU, and is described with "Move to the Supervisor's Menu". The MENU domain would contain one record for each menu option plus one or two header records for each menu and submenu.

In addition to the domain MENU, there are two domain tables which are used. These are as follows:

```
DEFINE TABLE MENU_SECURITY_TABLE FROM DOMAIN MENU USING
MENU_GROUP_KEY : RIGHTS_OWNER
END_TABLE
```

```
DEFINE TABLE MENU_TABLE FROM DOMAIN MENU USING
MENU_GROUP_KEY : PROC_NAME
END_TABLE
```

Before we can start the menu system, certain global variables and logicals must be properly set up. The LOGIN.COM file of the account or the SYLOGIN.COM file needs to contain some DCL command to establish the RIGHTS LIST for the user. This would be accomplished with some DCL that goes like:

```
$ . . .
$ assign/user rightsxxx.dat sys$output
$ show process/priv
$ open/read file rightsxxx.dat
$ rights := ""
$loop1:
$ read file r$
$ if (r$.eqs. "Process rights identifiers:") then goto loop2
$ goto loop1
$loop2:
$ read/end_of_file=eof file r$
$ rights = rights + r$
$ goto loop2
$eof:
$ close file
$ delete rightsxxx.dat;*
$ assign/process "'rights'" processrights
$ . . .
```

The net result of this DCL code is to create a process logical PROCESSRIGHTS which contains a list (separated by spaces) of the rights that the user's process possesses. The rights are granted to the user by the system manager with the GRANT/IDENTIFIER command within the AUTHORIZE utility. Determination of typical values for this logical might look like:

```
$ show logical processrights
"PROCESSRIGHTS" = " INTERACTIVE LOCAL ALL DATAENTRY" (LMN$PROCESS_TABLE)
$
```

The DATATRIEVE startup command file which is used to automatically activate the menu system is as follows:

```
! dtrstartup.com
DECLARE ESC PIC X(1).           ! This is going to hold the ESC character.
DECLARE CLEARIT PIC X(6)       ! This is going to hold the ESC sequence
    QUERY-HEADER IS -.         ! to clear a video screen with no header.
DECLARE BOTTOMIT PIC X(7)       ! This is going to hold the ESC sequence
    QUERY-HEADER IS -.         ! to position the cursor on line 22.
DECLARE RIGHTS_STRING PIC X(80). ! The rights list passed from DCL. If you
    ! have lots of rights you may need more
    ! than 80 characters in the global
    ! variable.
DECLARE MENU_TYPE PIC XX.      ! This is the variable which controls which
```

```

! menu or submenu we use. It is initially
! set by this startup file, but is changed
! by various procedures to move about.
DECLARE GROUP SELECTION COMPUTED BY          ! GROUP_SELECTION is used to
MENU_TYPE|"1"|FORMAT SELECTION USING 99    ! lookup entries in the tables
EDIT_STRING IS X(5).
DECLARE SELECTION PIC 99
VALID IF GROUP SELECTION IN MENU_TABLE AND
RIGHTS_STRING CONTAINING GROUP_SELECTION VIA MENU_SECURITY_TABLE||" " .
! The complex validation clause on SELECTION
! assures us that a menu selection entered
! is (first) a valid entry, and (second) an
! entry which this user is "privileged"
! (by holding rights list identifiers) to
! access the option.
DECLARE LOOP_PROCEDURE PIC X(8).           ! A string which contains the name of the
! procedure which is to be called pseudo-
! recursively. Usually this will be
! MAINLOOP, but it will change when it is
! time to exit and break the recursive
! loop.
ESC = "<ESC>"                             ! The escape character. Use GOLD 27 GOLD 3
! in the EDT editor to set this.
CLEARIT = ESC|"H"|ESC|"J"                 ! The escape sequence to home and clear
! screen on VT100 and VT200 type terminals
BOTTOMIT = ESC|"22;1H"                    ! The escape sequence to position the curser
! at the beginning of line 22.
:WORKING                                  ! A procedure which is given below that puts
! a blinking ". . . Working . . ." message
! on the screen to entertain the user while
! the rest of the startup command file is
! executed.
MENU_TYPE = "MN101" VIA MENU_TABLE         ! A dummy line which forces the menu-table
! to be loaded so that there will be no
! delay later.
MENU_TYPE = "MN101" VIA MENU_SECURITY_TABLE
! A dummy line which forces the
! menu-security-table to be loaded so that
! there will be no delay later.
MENU_TYPE = "MN"                           ! Initially point to the main menu.
RIGHTS_STRING = FN$TRANS_LOG("PROCESSRIGHTS")
! Move the rights list from the logical to
! the global variable. We will keep
! checking these rights and a global will
! be faster than a logical.
LOOP_PROCEDURE = "MAINLOOP"                ! Set the name of the procedure MAINLOOP in
! global variable.
FN$CREATE_LOG("MAINPROC",LOOP_PROCEDURE)
! The global variable is then moved to the
! logical.
READY MENU SHARED READ                     ! Shared ready is essential for a multi-user
! menu system
:MAINPROC                                  ! Execute the main loop procedure. The
! logical MAINPROC is translated into
! MAINLOOP and then executed.

```



```

EXIT                                ! This exit may or many not be executed
                                     ! depending on which type of "exit"
                                     ! procedure is activated instead of
                                     ! MAINLOOP.

```

This DATATRIEVE startup command file may appear to be unnecessarily complicated. The complexity of having both the global variable, LOOP_PROCEDURE, and the logical, MAILPROC, is, in fact, necessary to allow the menu system to be "gracefully" stopped and an exit to either the DATATRIEVE prompt or DCL be made. This, I hope, will become clear when we discuss the several exit procedures.

Now we have established enough of the environment that we can understand what is going on in the main pseudo-recursive procedure MAINLOOP. This procedure is as follows:

```

DEFINE PROCEDURE MAINLOOP
PRINT CLEARIT                        ! Clear the screen
PRINT DESCRIPTION(-) OF MENU WITH
  MENU_GROUP KEY STARTING WITH MENU_TYPE|"0" AND
  RIGHTS_STRING CONTAINING RIGHTS_OWNER||" "
                                     ! Print all the header lines for this menu.
PRINT SKIP 2                          ! Skip down a little bit
PRINT COL 10, OPTION(-) USING Z9, " - ",
  DESCRIPTION(-) USING T(55) OF MENU WITH
  MENU_GROUP KEY STARTING WITH MENU_TYPE|"1" AND
  RIGHTS_STRING CONTAINING RIGHTS_OWNER||" "
                                     ! Print all the menu option lines for this
                                     ! menu that this user is allowed to use.
PRINT BOTTOMIT                        ! Put cursor on the beginning of line 22.
SELECTION = *."selection from menu"
                                     ! Prompt user for selection. Remember the
                                     ! complex validation clause on selection
                                     ! assures us that the input is legal and
                                     ! that the user has the rights list
                                     ! privilege to choose the option
FN$CREATE_LOG("CURRPROC",GROUP_SELECTION VIA MENU_TABLE)
                                     ! Translate the users choice (by number)
                                     ! into the name of the procedure that is
                                     ! to be executed.
:CURRPROC                            ! Execute the procedure.
!
! Note that this procedure is not within a BEGIN-END loop. Therefore, any
! or all of the procedures so executed may contain DATATRIEVE commands and
! statements such as READY, DEFINE, DEFINEP, DELETE, FIND, and SELECT which
! can not be used in large "pre-compiled" procedure with one large
! BEGIN-END loop. This has the advantage that one does not have to ready
! every possible domain that one might need. You can wait to ready a domain
! when (and only when and if) you need it.
!
FN$DELETE_LOG("MAINPROC")             ! See discussion below.
FN$CREATE_LOG("MAINPROC",LOOP_PROCEDURE)
:MAINPROC                            ! See discussion below.
END-PROCEDURE

```

Because the execution of MAINPROC (which is usually MAINLOOP procedure) is the last statement in the procedure MAINLOOP, the procedure is not actually calling itself recursively. However, since the procedure does, in fact, include a logical reference which is resolved to itself, procedures of this type a called pseudo-recursive (they appear to be recursive but really are not).

The last three lines of the procedure (FN\$DELETE_LOG, FN\$CREATE_LOG, and execute MAINPROC) are at the heart of the pseudo-recursion. When the procedure MAINLOOP is invoked, the logical MAINPROC will be evaluated if it exists. Because the logical MAINPROC would not exist until after the FN\$CREATE_LOG has been performed, it forces the compiler in DATATRIEVE to re-evaluate the logical MAINPROC each time it is used. By this ruse, it is possible to have one or more procedures which are called in the menu system which change the contents of the global variable LOOP_PROCEDURE and thus affect a "graceful" breaking of the infinite pseudo-recursive loop.

The procedure to terminate the menu system and cause a normal exit from DATATRIEVE is as follows:

```
DEFINE PROCEDURE STOP_MENU
LOOP_PROCEDURE = "EXITPROC"
END-PROCEDURE
```

All this procedure does is to change the contents of the global variable LOOP_PROCEDURE. When the MAINLOOP procedure deletes and creates the logical MAINPROC, the procedure will now not point to MAINLOOP but EXITPROC. EXITPROC is as follows:

```
DEFINE PROCEDURE EXITPROC
PRINT CLEARIT
END-PROCEDURE
```

Essentially all the procedure EXITPROC need do is clear the screen since by not recursively calling MAINLOOP the loop is broken. Control then falls through to the last line of the startup file DTRSTARTUP.COM which is the DATATRIEVE command EXIT.

If one wants to break the loop and get to the DATATRIEVE prompt, one uses the procedure EXIT_PROCEDURE which is:

```
DEFINE PROCEDURE EXIT_PROCEDURE
FN$DELETE_LOG("MAINPROC")
SET ABORT
ABORT
END-PROCEDURE
```

Using the ABORT command will break all levels of the procedures and will deliver the user to the "DTR>" prompt of DATATRIEVE. Of course, one may not want a user in captured accounts to have access to this procedure.

Movement through the menu system is accomplished by very simple procedures which change the value of the global variable, MENU_TYPE. Two examples of such procedures are as follows:

```
!           A procedure to move to the supervisor's sub-menu.
DEFINE PROCEDURE MOVE_TO_SUPER_MENU
MENU_TYPE = "SU"           ! SU as the current menu
END-PROCEDURE
```

```
DEFINE PROCEDURE RETURN_TO_MAINMENU
MENU_TYPE = "MN"           ! set MN as the current menu again
END-PROCEDURE
```

There are several "utility" type procedures which are very useful to have available in the menu system. The first of these, WORKING, was used in the DTRSTARTUP.COM file to pacify the user while the startup command file was executing. This procedure is:

```
DEFINE PROCEDURE WORKING
DECLARE          BLINK          PIC X(4)
```

```

QUERY-HEADER IS -.
DECLARE      NOBLINK      PIC X(4)
QUERY-HEADER IS -.
BLINK       =             ESC|"[5m" ! The escape sequence for blink.
NOBLINK     =             ESC|"[0m" ! The escape sequence for no attributes.
PRINT CLEARIT, BOTTOMIT, "      ", BLINK,
      ". . . Working . . .", NOBLINK
END-PROCEDURE

```

Other useful procedures are:

```

DEFINE PROCEDURE NOT_YET
PRINT CLEARIT
PRINT "This procedure has not yet been implemented."
FN$DCL("$WAIT 00:00:02")
END-PROCEDURE

```

```

DEFINE PROCEDURE SPAWN_MAIL
PRINT CLEARIT
FN$DCL("$MAIL")
END-PROCEDURE

```

```

DEFINE PROCEDURE SPAWN_PHONE
FN$DCL("$PHONE")
END-PROCEDURE

```

The procedure NOT_YET is used as a dummy procedure to hold a place in procedure development. If one is going to write a procedure whose name is FOO, but you haven't yet written it, then you can put FOO in the menu system along with its rights and description, but let FOO point to NOT_YET. This is a kind of top-down implementation of procedures in this menu system. Note that NOT_YET contains a FN\$DCL to the WAIT DCL command. If you are handy with re-linking DATATRIEVE, you may want to use a user defined function like Don Stern's WL\$WAIT⁹ instead of FN\$DCL since it takes much less of the processor.

For menu system users who you wish to completely control, one can make the following command file as their login command file.

```

$!
$!      caplogin.com
$!
$on control_y then $goto done
$set control=(T,Y)
$assign [manager]dtrstartup.com dtr$startup
$assign/user tt: sys$input
$assign/user tt: sys$output
$dtr
$done:
$lo/full

```

Their account qualifiers should be modified to be

DISCTLY	disable control Y
DEFCLI	default command language interpreter
CAPTIVE	captive (can't change default disk)
LGICMD=[manager]caplogin.com	use captive login.com file
PRCLM=1	allow only 1 subprocess

and the protection of the files [MANAGER]DTRSTARTUP.COM and [MANAGER]CAPLOGIN.COM should be only read and execute (not write or delete). With such an environment, there is a pretty good chance that most users can not escape from the menu system nor can they do too much damage although it is still possible for them to use PHONE, MAIL, and the editor.

There is one more procedure which I find very helpful. I like the titles on my menus to be double height and double width. The following procedure is nice to create the two header records for a new sub-menu. The procedure is HEADER_LOAD.

```

DEFINE PROCEDURE HEADER_LOAD
!
!   header_load
!   A procedure to load header text for the menu system. This
!   procedure takes a text phrase, centers it, and loads it
!   into two lines of menu header (in double height, double width)
!   format. This procedure is not part of the operations of the menu
!   system but is used as a utility to support and initialize
!   new menus and sub-menus
!
READY MENU SHARED WRITE
DECLARE          BUFF          PIC X(42).
DECLARE          UPPERLINE     PIC X(3).
DECLARE          LOWERLINE     PIC X(3).
DECLARE          T_MENU_SYSTEM PIC XX.
DECLARE          T_OPTION      PIC 99.
DECLARE          T_RIGHTS_OWNER PIC X(10).
DECLARE          T_DESCRIPTION PIC X(80).
DECLARE          TWENTY_SPACES PIC X(20).
DECLARE          T_LENGTH      USAGE IS INTEGER.
!
TWENTY_SPACES = "          "
UPPERLINE = ESC|#3"
LOWERLINE = ESC|#4"
!
T_MENU_SYSTEM      = *."Menu system abbreviation"
T_RIGHTS_OWNER     = *."Rights owner"
BUFF               = *."header line of text (40 characters or less)"
T_LENGTH = ( 41 - FN$STR_LOC(BUFF," "))/2
STORE MENU USING BEGIN
    MENU_SYSTEM = T_MENU_SYSTEM
    RECORD_TYPE = 0
    OPTION = 1
    PROC_NAME = " "
    RIGHTS_OWNER = T_RIGHTS_OWNER
    DESCRIPTION = UPPERLINE|FN$STR_EXTRACT(TWENTY_SPACES,1,T_LENGTH)|BUFF
END
STORE MENU USING BEGIN
    MENU_SYSTEM = T_MENU_SYSTEM
    RECORD_TYPE = 0
    OPTION = 2
    PROC_NAME = " "
    RIGHTS_OWNER = T_RIGHTS_OWNER
    DESCRIPTION = LOWERLINE|FN$STR_EXTRACT(TWENTY_SPACES,1,T_LENGTH)|BUFF
END
END-PROCEDURE

```

What I have described is a complete, high-performance menu system. It appears very complex and difficult to initially set up. But the rewards of such a system pay off handsomely in maintainability of the system. Each procedure of the application is essentially independent of every other in the application. The greatest benefit of such a system occurs when the user needs a new function for the system. That function is written and tested. It is converted into a procedure. And then it is installed into the system by adding a single record to the menu domain. This enhancement, of course, can be done while multiple users are simultaneously accessing the menu system!

References

- 1 Bob Hoover's Magic using the "pre-compiled" technique, Fall 1986 Wombat Magic, DECUS U. S. Chapter SIG's Newsletters, Volume 2, Number 5, January 1987, pp. DTR-14 to 16.
- 2 Lorey Kimmel's Magic using logicals and a pseudo-recursive procedure, Spring 1987 Wombat Magic, DECUS U. S. Chapter SIG's Newsletters, Volume 2, Number 12, August 1987, pp. DTR-11 to 14.
- 3 Gary W. Burton's Magic using a variations of the pseudo-recursive procedure using indirect command files rather than executing a logical, Spring 1986 Wombat Magic, DECUS U. S. Chapter SIG's Newsletters, Volume 2, Number 1, September 1986, pp. DTR-41 to 43.
- 4 Pat Scopelliti's Magic using logicals and pseudo-recursive procedure, Fall 1985 Wombat Magic, DECUS U. S. Chapter SIG's Newsletters, Volume 1, Number 7, March 1986, pp. DTR-18 to 19
- 5 Larry Jasmann's Magic using callable DATATRIEVE, Spring 1984 Wombat Magic, The Wombat Examiner, Volume 5, Number 4, Fall 1984, pp. 17-21.
- 6 Diane Pinney, "DATATRIEVE Magic Using Logical Names", DECUS U. S. Chapter SIG's Newsletters, Volume 2, Number 4, December 1968, pp. DTR-18 to 20.
- 7 Barrie Gray, "So Who Needs a DATATRIEVE Compiler?", DECUS U. S. Chapter SIG's Newsletters, Volume 1, 12, August 1986, pp. DTR-32 to 42.
- 8 S. Begelman, "An Open, Menu-Driven, Graphic Report Generator using DTR and VAX-BASIC", The Wombat Examiner, Volume 6, Number 2, Winter/Spring 1985, pp. 30-35.
- 9 Donald E. Stern, Jr., "User Defined Functions (UDF's) for VAX DATATRIEVE", DECUS U. S. Chapter SIG Newsletters, Volume 2, Number 8, April 1987, pp. DTR-20 to 41.

Wombat Wizard

Philip A. Naecker, Consulting Software Engineer, Altadena, California

Care and Feeding of RMS Files

The subject of this month's column is the design and maintenance of RMS files, especially indexed files. This topic should be of extreme interest to you if you use Datatrieve or any other system or programming language that uses RMS. Since you are using RMS even if all you do is create listing files, it seems likely that you will want to read this article. Better yet, make copies and pass them around, including one appropriately marked in red for your system manager to read. This subject is too large and too important to be covered in a single column, so I hope to continue this discussion into the next couple of months.

Introduction to RMS

RMS is the Record Management System, and that's just what it does - manage records. A more complete description might be to call it the VAX file system, because RMS is also largely involved in the creation, deletion, and other operations on files of all types - whether or not they contain "records" as we Datatriever's know and love them. VAX RMS (sometimes called RMS-32) has grown up from RMS-11, and the two systems are still largely compatible. Files from one system can be operated on by the other - even

many indexed files. Much of what is discussed in this article will apply to both VAXes and 11's, but the specifics of commands, algorithms, syntax, FDL files, and such will be limited to the VAX implementation of RMS.

It would be redundant for me to provide a complete introduction to RMS -DEC has already done that with a pretty good manual, the Guide to File Operations. This book, which came out with V4.0 of VMS, provides a good introduction to RMS as well as providing great detailed information for those times you need it. A few real VMS hackers want considerably more detail about the file system and RMS internals (and maybe we'll see that in the V4 VMS Internals and Data Structures, whenever that comes out!), but for most mortals the Guide provides more than we'll ever need to know about RMS.

In case you didn't catch my drift, let me put it to you this way. If you want to be a real jock in data management on VMS, you will read and re-read the Guide to File Operations until you can recite it chapter and verse. It is very, VERY important that you understand what's in there.

"But, oh Great Wizard," you say, "I have no time to read that manual. I am working furiously on getting this new application to run. My boss says that it's much too slow and so I'm re-writing the whole thing."

Gotcha! If you've written a reasonably designed DTR application that is too slow, I'd be willing to bet that the reason it's too slow isn't the fault of DTR at all and doesn't need to be fixed by re-programming. It is probably a problem with the file operations of the system. Many otherwise perfectly good DTR applications suffer the label of "system hog" only because the system developer didn't properly design and tune the files. So perhaps the fastest way to improve the performance of your existing applications is to improve their file system performance.

To further whet your appetite, my own experience (and that of many other VMS software developers whom I respect) is that some simple file tuning and adjustment of RMS parameters will improve performance for a typical DTR application by something like 20 to 200 percent. I recently worked on an application (in COBOL, not DTR) where five minutes with the FDL editor improved run times (on a VAX 8530) from over four hours to under 15 minutes. In my consulting practice, nothing that I do quite generates the same look of awe from both managers and programmers as fifteen minutes of file tuning.

In the discussion that follows, we will assume that you are knowledgeable in the use of DTR with indexed files. You should understand the way that DTR selects keys, how to define files with keys, and the meaning of the DUP and CHANGE options in the DEFINE FILE command. If you don't, I suggest you first read the DTR Guides for more background on file definition.

Where to Begin?

The best place to begin is generally the beginning, and for a DTR application that is record and file design time. Most DTR applications are - at least the good ones - ISAM file applications, and the key in designing an indexed file (pardon the pun) is the correct selection of indexes, or keys. Here are a few of the more important decisions in selecting the indexes for ISAM files.

- Select keys that are unique whenever possible, especially for your primary key. For example, employee number is a good primary key for an employee file (there is only one employee with each number, so the key is unique) whereas department number is a poor choice for a primary key (since most departments have more than one employee, thus there will be multiple entries for any one key value in the file). A field such as SEX (the "M or F" kind, not the "Y or N" kind) would be a poor choice for any kind of key.
- Use the smallest key that does the job. Don't make a key for a larger field (or group field) than you need to use to select the records of interest.
- Use the appropriate data type. If you are dealing with numbers, use a numeric datatype.
- **NORMALIZE YOUR FILES.** Normalization will be the subject of a future WW column, but you can read about it in the DTR guides as well. Basically, the idea of normalization is to remove as much redundant information from a file as is possible.

- Don't create more indexes than you need. This is a biggy. As we will see later, maintaining indices is very, very time consuming for RMS. It typically requires many, many IO's to add or change an index, and if you have many indices the problem is obviously amplified. Very few applications will need more than two or three indices in any one file; if you have more than that, consider using a relational database instead of a file, create "derived databases" (see WW in the Wombat Examiner, December 1986), or at the very least minimize your updating of the file. (There are also some very clever tricks in this area provided by FDL, and we will cover those in next month's column. If you think you know what they are, write in Wiz and we'll include your suggestion in a subsequent column.)

File Definition Language

After your file and the indexes need to be laid out, you should next design the file characteristics. This is done using the functions of the File Definition Language (FDL) and the associated utilities. FDL is nothing more than an easily readable form of the information that is embedded in the file header and other RMS data structures called XAB's, or eXtended Attribute Blocks. These data structures contain information that RMS needs to know when it opens or creates a file. If you don't provide this information to RMS, it uses language- and system-defined defaults, but you can easily provide this information to RMS by using the FDL file.

(We will discuss more about system-defined defaults two months from now. That column will address the general topic of System Management and RMS, and we will discuss SYSGEN parameters that affect file system performance in general and Datatrieve performance in particular.)

The easiest way to use FDL is this:

1. Define the file using DTR. That is, define the record and define the domain for the file.
2. Use a DTR> DEFINE FILE to create a file with keys for the fields you want, including DUP and CHANGE options as required. Pay particular attention to your selection of a primary key.
3. Use the ANALYZE/RMS/FDL command to analyze the file you just created. For example:

```
$ ANA/RMS/FDL YACHTS
```

will create YACHTS.FDL.

4. You can now edit the FDL file using the command:

```
$ EDIT/FDL filename.FDL
```

command. At this point, you can't really "tune" the file, because you have no information about the data in the file, but you can add some additional entries to the FDL file if you are so inclined. We will talk more about these entries in next month's column. You certainly can view your file's attributes and you may wish to change any entries in the FDL file that you are certain you understand and that you want to change.

5. From now on, use the FDL file to create new copies of the file. You can do this in two ways:

```
DTR> DEFINE FILE FOR domain-name USING fdl-file
```

```
$ CREATE/FDL=fdl-file [filename]
```

Note that \$ CREATE will attempt to use the filename (including complete device and directory specification) from the FDL file in creating the file. It will also attempt to use the owner (UIC) information in the FDL, so if the original file was owned by someone different than the user running CREATE, you may want to either change the OWNER line of the FDL file or simply delete that line. (If you have CMKRNL privilege, the advantage of the OWNER line in the FDL file is that the file gets created with the "correct" owner, not with your UIC.)

By using the FDL every time you create the file, you will be sure to take advantage of any new attributes added to the FDL file as you go through the tuning process during the life cycle of a file.

The next step is to load the file with data. To be most useful, you should use "real" data so that the analysis of key duplicates is reasonably accurate. I generally try to perform this step after the application has been in use for a few weeks - enough for there to be a reasonable amount of data present but not enough time for performance to become a problem.

The FDL Editor

After you have loaded the file with real data, you can use the following command to obtain a new FDL file:

```
$ ANALYZE/FDL/STATISTICS[/OUTPUT=fdl-filename] filename
```

This command creates an FDL file with two parts - the FDL information like we found in the previous ANALYSIS command, and some Analysis of Key information that tells us about the use of RMS data structures in the file. This latter component will be used in our next step to optimize the file. We will do this using the command:

```
$ EDIT/FDL/SCRIPT=OPTIMIZE fdl-filename
```

Now, if you have modified the FDL file and done a CREATE/FDL or DEFINE FILE USING command, those modifications (which were embedded in the file header) will be found in the new FDL file. (There are some exceptions to this, mostly having to do with the use of the SET FILE command or changes in ownership.) If you want to use the FDL definitions you started with and the analysis information from the ANALYZE command, you can use the /ANALYSIS=filespec qualifier on the EDIT/FDL command. However, in general you should not need to worry about this and the simple EDIT/FDL/SCRIPT=OPTIMIZE command will do everything you need to do.

The FDL editor is more than just an editor that understands the format of the FDL language. It also knows how to use an analysis output file to tune a file. The OPTIMIZE script is just that - a script of questions. You provide answers to questions about the number of records in the file, the number of records you intend to add in between CONVERTs of the file (you DO convert your files regularly, don't you???), and similar items. The editor then goes off and calculates optimal settings for a number of FDL parameters. If you are not familiar with the meaning of a parameter, you should probably just use the settings that EDIT/FDL suggests. However, you can also change and parameters that you would like - just be careful and make sure you know what you are changing and how it will impact your applications.

The output from EDIT/FDL is another FDL file - this one without the analysis information. (By the way, you can type or use an editor on an FDL file - they are just text files. But it's usually easier to use EDIT/FDL on the file so that you avoid introducing errors into the file.) Use the new FDL file just as you used to use the old one - whenever you CREATE or DEFINE FILE. You can also use the FDL file with CONVERT to change the structure of a file. However, the two files must have the same record definition - only the RMS key information, buffering, bucket sizes, and the like can change using CONVERT. The record layout or datatypes cannot be changed by CONVERT.

In many cases, this is all you need to do to greatly improve ISAM file performance - use the ANALYZE/FDL and EDIT/FDL utilities. However, if the file changes or grows, you should go through the ANALYZE-EDIT cycle again to make sure the file design is still optimal for the data it is storing and the use you are making of the file. And if you want really GREAT performance from the file you will want to become a little more aggressive in the use of some of the more esoteric FDL parameters, especially options for buffering. We will cover these topics next month.

Printing out Pending ALL-IN-1 Mail (during Login, etc.) using Datatrieve.

B. Z. Lederman, ITT World Communications, New York, NY 10004-2464

I recently received a telephone call from a user seeking a way to find out during user login if there is any pending mail in ALL-IN-1. He stated that he had been told that you can have a program written for you by your local Software Services organization, but that he really wasn't in the position to spend the money. I was unable to find a public domain program which would do it, but it later occurred to me that it could be done fairly easily with Datatrieve. (I hope he sees this article!).

First, you need the record definition for the pending mail file used by ALL-IN-1:

```
DELETE AI1_PENDING_RECORD;
REDEFINE RECORD AI1_PENDING_RECORD OPTIMIZE
01 AI1_PENDING_REC.
!
! Read the ALL-IN-1 PENDING data file.
!
! B. Z. Lederman
!
    10 PENDING_KEY PIC X(65) EDIT_STRING T(32).
    10 C USAGE_BYTE EDIT_STRING SZZ9.
    10 A USAGE_WORD EDIT_STRING SZZ,ZZ9.
    10 B USAGE_WORD EDIT_STRING SZZ,ZZ9.
    10 D PIC XX.
    10 HIDE.
        20 FILLER PIC X(8).
    10 REAL REDEFINES HIDE.
        20 E PIC X(8).
    10 ENUM PIC 9(8) COMPUTED BY E EDIT_STRING ZZ,ZZZ,ZZ9.
    10 FILLER PIC X(1920).
;
```

You also have to define a domain which uses this record:

```
DELETE AI1_PENDING;
REDEFINE DOMAIN AI1_PENDING USING AI1_PENDING_RECORD ON OA$DATA:PENDING.DAT;
```

I have not quite figured out what AI1 is doing with all of these fields, but I am reasonably certain what two of them are. PENDING_KEY contains the name of the user queues (there are MAIL entries and message router queue entries), and ENUM contains the number of unread messages. I am using a COMPUTED BY field here because the data file contains the number with leading blanks, and I want to treat it as a numeric value. If all I ever wanted to do was print out the number, such as I am doing here, I could just use field E.

Once there is a domain which accesses this data, it's fairly easy to find out how many messages are waiting:

```
DELETE PRINT_PENDING;
REDEFINE PROCEDURE PRINT_PENDING
!
```

```

! Find out how much ALL-IN-1 Electronic Mail is Pending
!
! B. Z. Lederman
!
READY ALL PENDING SHARED
DECLARE KEY_FIELD PIC X(65).
!
! Get the pre-defined user name and make it into a retrieval key
! KEY_FIELD = "MAIL " | FN$TRANS_LOG ("USER_NAME")
!
! (note one blank space between MAIL and the user name.
!
! Now print out the proper record.
!
FOR ALL PENDING WITH PENDING_KEY = KEY_FIELD
    PRINT "Pending ALL-IN-1 Mail Messages = ", ENUM(-)
!
FINISH
END-PROCEDURE

```


The reason I'm translating a logical name here is because I'm assuming that the user would want to set up one procedure that could be used by everybody in the system (invoked from SYLOGIN.COM, for example), rather than hard-coding a dedicated version for each individual user. For this to work, you have to define a logical name which translated to the users' name: assuming that the UIC field translates to the users' name (as it usually does), you could do something like the following in DCL to retrieve the users' name, strip off the square brackets of the UIC format, define the logical name, and invoke the print procedure:

```


$ temp_name = F$USER()
$ end = F$LOCATE("]", temp_name)
$ end = end - 1
$ temp_name = F$EXTRACT(1, end, temp_name)
$ DEFINE user_name 'temp_name
$ DTR; :CDD$TOP.DTR$USERS$.ALLIN1.PRINT_PENDING

```

Note that this requires that you define DTR as \$SYS\$SYSTEM:DTR32 or something similar so you can put DTR commands in on the DCL command line. You would probably want to put the procedure PRINT_PENDING in some central directory to which everyone has read/execute access.



DIGITAL EQUIPMENT COMPUTER USERS SOCIETY



AT LARGE

The Newsletter of the Large Systems SIG

Contributions

Contributions and suggestions for this newsletter are constantly needed. Articles, letters, technical tips, or anything of interest to our SIG are greatly appreciated. The editor prefers submissions be made electronically, but magnetic tape and hard copy will be accepted.

Send your contributions to:

ARPAnet/CSnet: ctp@sally.utexas.edu
UUCP: ctp@ut-sally.uucp ({harvard,ihnp4}!ut-sally!ctp)
BITNET: CTP@UTADNX
CIS: 75226,3135

or if you must, use the U. S. Mails:

Clyde T. Poole
The University of Texas at Austin
Department of Computer Sciences
Taylor Hall 2.124
Austin, Texas 78712-1188

Table of Contents

Ask Dr. TOPS	LS-2
Fall DECUS Symposium Preview	LS-4

Ask Dr. TOPS

Dear Dr. Tops -

The last major upgrade increased the length of a file name from 9.3 to 39.39. What surprises await me in the next major upgrade?

Variable String Descriptor

Dear Character*(*) -

I am not allowed to tell you things that have not yet been announced. However, if you are familiar with the \$GETUAI and \$SETUAI system services, you will note that the fields for USERNAME and ACCOUNT return a length of 32 characters under VMS 4.6 (Now playing at your favorite delay center). A careful examination of STARLET.PAS for UAI\$ would prove interesting for the inquiring reader. If you are depending on physical locations of data items in ACCOUNT/FULL, you had better get out a sharp pencil and start figuring a way to find what you want in each new location.

Dr. Tops

Dear Dr. Tops -

I have heard a lot about Symmetric Multi-Processing. Could you tell me why VMS is making such a fuss over it?

Multi Tasker

Dear MT -

VMS has only had ASMP (ASymmetric Multi-Processing) since the 11/782 was announced years ago. TOPS-10 started out with ASMP on their 1099 systems. They quickly found out that the Master/Slave relationship could produce a situation where the pair ran at 95% of the power of a SINGLE CPU. The problem comes in when the task allocated to the second CPU wants to do anything but compute. This means normal timesharing I/O, system service calls, etc. When the second CPU saw the I/O request, it had to interrupt the primary CPU to re-schedule the task. This causes LOTS of overhead, context switching, cache purges and general thrashing. The simple solution (which wasn't so simple) was to give the "slave" a share of the I/O and overhead.

TOPS-10 pioneered SMP in DEC land. They were running SMP long before the 11/782 was dreamed up for the compute bound VAX community. The Master Slave scheme works well as long as all you are doing is computing. Today that is not the case for a large majority of the VAX user base. With an architecture designed for N processors, can you imagine the overhead of (N-1) slaves interrupting the single master? With the wrong load mix, the system would go nowhere! This is the reason the VMS crew is working so hard to make any CPU capable of performing overhead tasks.

To get down to the nitty-gritty of SMP, at least as far as the TOPS designers have made it, the system works like this... Each processor is assigned some private space in memory, and shares all the rest. Access is synchronized by a "spin lock", or a memory location that indicates if the resource is free or not. If the resource is free, the spin lock opens and the processor continues. If not, then the processor wanting the lock literally "spins". It sits in a VERY tight loop waiting

for the lock to free up. Now there may be cases where the lock could be in use for a long time. What would happen if the owner of a lock crashed? Well, this has been thought of and a way out has been found. Once passed the spin lock, the processor can perform whatever task it needed and then release the lock. For example, if a process page faults, the CPU will need to access that processes' page table for the location of the needed page. If the page is in memory, the memory allocation lock is spun for. If it is not free, the CPU can do something else, or wait for it. If the page isn't available, then the CPU goes through its usual routine of making an I/O request to get it. This means a bunch more spin locks and access to the common I/O queues.

Another example would be processor faults. In TOPS-10, if a processor encountered an error (like its registers had a parity error), then the processor would halt. The others would continue running, picking up the load from the failed processor. If the fault was common to all processors, like a parity error in one of the system page tables or common I/O database, then everyone was stopped. Usually just the currently active process hit the bit bucket. Can you say "High Availability"? Yes, I knew you could! Can you say "Concurrent Processing"? Can you say "Race Condition"? Can you say "Multi-Player, Real Time Klingon Bashing"? Aye, Captain!

The grand aim of SMP is to allow users to add CPU power when cycles are short without having to add another whole member of the cluster. Just like you add more memory, you could now add more CPU's. It is just another expandable resource.

Dr. Tops

Dear Dr. Tops -

What is, and why go through, a "major upgrade"? Is this going to be another 3.X to 4.X massacre?

Cheshire Cat

Dear Executioner Escapee -

There is no need to paint roses and assume the prone position for review. The sky isn't falling and things don't run backwards.

DIGITAL has learned from its earlier bludgeoning that continuity between releases is mandatory. The 4.X to 5.X transition should be much less of a shock than the last go round. There are the usual changes in a major release. Device drivers for third party equipment WILL NOT FUNCTION under 5.X since the underlying protocols have changed. This is to be expected with SMP. The changes to support SPIN LOCKS and multi-CPU access are not major. Simple device drivers could be repaired and operational after a few days of work. There is a large volume of documentation that needs to be read and UNDERSTOOD before attempting to hack a driver into being.

According to DEC's own documentation, a major upgrade brings significant enhancements in support into the O/S. SMP is clearly a step in the large systems direction. Expanding the definition of a cluster is also clearly a significant enhancement. Performance improvements are always welcome, but some do involve changing the way business is done. You cannot allow N CPU's to perform the same task at the same time. Some method must be introduced to arbitrate such access. Bigger machines + bigger disks + bigger clusters + more users + developers = necessity to change algorithms and data structures. VMS isn't RSX+++ on a little micro. It is VMS on an 8800 or 89xx with 100's of users and giga-bytes of disk. The old methods just couldn't handle the loads.

If you are running a plain vanilla DEC shop, with no foreign hardware, the upgrade should be simple. Make your BACKUPS (don't forget ANA/DISK/REP first), read the DOC and go! You

will want to INIT all (*ALL*) your tapes before you use them from now on. Pick a suitable set of labels for your saveset tapes (weeknn/yearnn/etc) so that the labels can be easily deciphered.

If you have a spare disk pack or drive, then install VMS on it and test fly to look for gotcha's before you turn it loose on your users. You will want to do a complete re-install of VMS. You can still @VMSINSTALL if you want, but that takes a long time, propagates DCLTABLE conflicts, and the like. This has the further benefit of making sure you run only those compilers / layered products you have licenses and current revisions of. No need to get into a legal hassle these days. Do you really want to drag around things left over from V3.1? What about those trojan horses left by the last security hole? Install the software, and make CHECKSUMS of all the .EXEs and other important files.

If you have multiple disk drives, you might consider multiple page/swap files split over controllers. More heads, more throughput. You can also use this opportunity to restructure the index files to be sized and positioned for maximum benefits. If you have 32K files on an RA81, you should allot at least that many headers at INIT time.

Dr. Tops

Fall DECUS Symposium Preview

Betsy Ramsey, American Mathematical Society, Providence, RI

The Large Systems SIG concerns itself with the issues and needs of users of Digital's large computer systems: large VAXs and DEC-10/20s in standalone, networked, clustered and super-computer environments. The SIG is presenting a number of sessions at the DECUS U.S. Chapter Fall Symposium in Anaheim this December that should be of interest to these users.

High-End System Planning and Management

On Monday and Tuesday, there are a number of sessions pertaining to the management of high-end systems, and planning for new high-end systems. Some sessions of particular interest are the following:

- VAX-8974/8978: Packaged VAXcluster Systems

This session describes two complete, large-scale, VAXcluster system offerings which are used as general-purpose systems for data processing, database management, information systems, and decision support applications. The VAXcluster systems are compared to the largest traditional mainframes in performance, capacity, and system availability in a computing environment which is characterized by multiple function, multiple application, and high capacity requirements.

- High-Performance On-line Transaction Processing

This session discusses Digital's approach to high performance On-line Transaction Processing (OLTP). The session includes a discussion of the functional components of a transaction processing system, transaction integrity, distributed transactions, high availability and fault tolerance.

- High-End VAX Configuring: Wires and Boxes

Configuring high-end VAX/VMS systems involves numerous details. While many individuals have a good knowledge of the overall concepts, information on some of the details which are

necessary to complete a "clean" configuration are sketchy. Most often these details involve two basic areas: cables (wires) and system options (boxes). This session covers, in detail, the configuration rules involved in each of these two areas; other areas are covered as time permits. A companion session "High-End VAX/VMS Configuring: Case Studies" presents sample full system configurations.

- Managing High-End Systems in a Multi-Vendor Environment

An increasing number of sites that use Digital's high-end systems are also supporting systems from a variety of other vendors. No single vendor seems to satisfy all needs, despite claims to the contrary. The support of systems (microcomputers through mainframes) in a multi-vendor environment is a challenge that presents new opportunities for creativity as well as new levels of complexity. Speakers at this session discuss their experiences in supporting and managing systems in the multi-vendor/multi-operating system environment.

DECSYSTEM-10/20 Sessions

On Thursday, the Large Systems SIG will present sessions of special interest to users of DEC-10 or DEC-20 computer systems, including the following.

- VMS Internals for DECsystem-10/20 System Programmers

DECsystem-10/20 system programmers have traditionally worked with a good understanding of the internals of their respective operating systems. This session is for those systems programmers, and is based on their existing knowledge of how an operating system functions. It presents details of key portions of the VMS operating system. This subject is quite broad, and is the topic of numerous training courses. The topics covered in this session are only a small subset of the entire picture, but are aimed at providing the most information in the time allotted.

- VMS Tape Handling for TOPS Users

This is a "hints and kinks" session for TOPS-10/20 users who are looking for ways to perform similar types of tape handling under VAX/VMS. Topics will include BACKUP techniques, reading and writing foreign tapes, and working around the lack of a mount queue.

- An Academic Conversion from a DECsystem-10 to a VAX 8650

The Catholic University of America has converted all of its academic computing from a DECsystem-10 Symmetric Multi-Processing (SMP) system to a VAXcluster. This paper describes the conversion of the academic computing; the planning and the work involved, the successes and the problems encountered.

Supercomputer Sessions

On Tuesday, the SIG will present a number of sessions on supercomputers: what they are, why you would want them, and how Digital systems are used in a supercomputing environment.

Internet and SMTP Sessions

On Wednesday, the Large Systems SIG will present a number of sessions oriented toward users of the Internet set of network protocols. In particular, the focus will be in the use of TCP/IP in a VAX/VMS environment and the implementation of electronic mail systems that implement the SMTP mail protocols under VMS.



**OFFICE
AUTOMATION**

IN THIS ISSUE...

From The Editor OA - 1

- Therese LeBlanc, T.M. LeBlanc & Assoc.

**Some Things I Have Learned
About Using the Message Router With ALL-IN-1. OA - 1-4**

- Bart Z. Lederman, ITT World Communications

OA SIG EVENTS IN ANAHEIM OA - 4

- OA SIG

Notes On Notes OA - -4-8

- Mark H. Hyde & C. J. Trayser, Digital Equipment Corp.

FROM THE EDITOR...

This month our featured technical article is from Bart Lederman about the Message Router and ALL-IN-1. We also have a great Notes column about using notes in Anaheim (for those of you who are attending Symposium).

The OA SIG will again be offering an outstanding line-up of session on Office Automation topics. Make sure to attend our Monday morning Roadmap session for an overview of the week's sessions and 'other' activities. If this will be your first Symposium, please let us know, we'll be happy to answer your questions and tell you more about what you can do to become involved in the SIG.

Regards,

Therese LeBlanc

OA Newsletter Editor
275 London Place
Wheeling, IL 60090
(312) 459-1784

SOME THINGS I HAVE LEARNED ABOUT USING THE MESSAGE ROUTER WITH ALL-IN-1

Bart Z. Lederman
ITT World Communications, New York, NY

Over the past 9 months or so, using the Message Router which is now supplied with ALL-IN-1, sending in and receiving answers to SPRs, working with TSC, and experimenting, I've picked up a few bits of information which may be useful. One (which was previously published here) was on automatically re-starting the message queues that move mail between the Message Router (MR) and AI1.

There are a number of options on when MR can run. The best one for most users appears to be that the Talker should run when activated by a message sent to it, and not as a batch job on schedule.

Message Router Cont'd

Bart Z. Lederman

I had to install MR so many times in the course of getting it to work that I got fed up editing the data file (mentioned in the release notes) to hand-apply the patches that DEC should have fixed before distributing the product. I finally created this EDT command file which can be used to make the changes with:

```
$ EDIT/COMMAND=FIXFILE.EDT MRINITDEF.DAT
(content of FIXFILE.EDT)
s/DEFDEC/DEFAULT DECNET/wh
s/ENBLOG/ENABLE LOGGING/wh
s/COMPULSORY LOG/COMPULSORY LOGGING/wh
s/LOGGER_BUFFER/INTERNAL_BUFFER_SIZE/wh
s/EXTEND_NO/INTERNAL_BUF_EXTEND_NO/wh
s/FILE_SIZE/FILE_SIZE IN_BLOCKS/wh
s/KEEP_NO/PURGE_KEEP_NO/wh
exit
```

Don't expect the IVP to work. Because the above mandatory patches, and the changes to account quotas (mentioned below) are not applied until after installation, the IVP will probably fail. If you want to take the time to re-install MR after you have made all of the corrections, and re-use your existing database (so you don't wipe out the patches), the IVP may eventually run. It's better to just send some mail messages from A11 to VMSmail and vice versa after installing all of the Message Router products. If the mail gets through it's just as good a check as the IVP: if it doesn't get through then something is wrong no matter what the IVP says.

When you invoke VMSINSTAL, it gives you a warning if DECnet is running. Many products can be installed while DECnet is up, but it's probably better not to install the Message Router while DECnet is running.

When you install a product on a cluster, the product usually goes into a node-specific account. If you are running a homogeneous cluster you can put the product on once in a root directory such as SYS\$SYSDEVICE:[MRMANAGER] and avoid duplicating all of the files for each node (you normally have only one copy of Message Router running per cluster anyway) and save disk space, etc.

There are some restrictions on password lengths which are not clear in the documentation. MR doesn't restrict password lengths, but due to a limitation within ALL-IN-1, the password on it's mailbox A1 must not exceed 9 characters. The password on the A1 mailbox in MR must be the same as in the ALL-IN-1 Messaging Manager (MM) menu under CMP. Similarly, the password on mailbox MRGATE (used to pass mail between the Message Router and the VMS Mail Gateway MRGATE) must not exceed 12 characters. I personally did a SET PASSWORD/GENERATE=8 on my account to have the system generate some long "nonsense" passwords for me to use on these mailboxes. (I copied down the passwords it generated, then rejected them so my account password didn't change.)

When MR is installed, the MRMANAGER account may end up with insufficient quotas for the product to run. For MR V2.0, the following quotas are recommended, and you may have to modify the account by hand to get them:

```
PRCLM = 60 BIOLM = 50 ASTLM = 60 ENQLM = 400
FILLM = 60 WSDEF = 1024 WSEXTENT = 8192 WSQUO = 2048
PGFLGQUO = 1600
```

The entire subject of running Message Router on a cluster is one which is not clear in the documentation (at least I didn't think so, and the answers to my SPRs state that DEC intends to clarify the documentation in the future). If you are not trying to do anything fancy, it turns out not to be too difficult. First, use the MBMANAGER routines (like MB\$START.COM and MB\$STOP.COM) to start and stop MR and MRGATE rather than the individual start and stop command files. If for some reason you must start MR individually then you may find, as I did, that there is some confusion about the RESTART qualifier. Basically, the command file should always be invoked as:

@SYS\$MANAGER:MRSTART.COM RESTART

The reason for not using it in a cluster is, apparently, to prevent a node which is being rebooted from starting the Message Router when it has been deliberately stopped on the other nodes. How you are supposed to know at boot time during SYSTARTUP that this is the case DEC doesn't say: apparently, they expect system managers to stand in front of the console and perform a conversational boot, whereas in my environment it is necessary for SYSTARTUP to run unattended. As I said, I find the best way out is to use MB\$START.COM, do not use MRSTART.COM or MRGATEST.COM, and hope for the best.

The defaults during installation are to put all batch jobs on the generic queue SYS\$BATCH. If you have a cluster, then the jobs will run on whatever node is handy: this is OK in a homogeneous cluster with MR on every node, but not in a heterogeneous cluster, or if you want to run MR on a specific node. MR\$:MRINITUTL.COM can be used to change the queue name, but it doesn't generate a new MRSTART.COM (which is needed even if you do use MB\$START.COM): if you don't take the defaults during installation you don't get MRTIDY.COM and the MR accounts fill up with unneeded log files. The solution is to take the defaults but edit the command files manually. There are no ill effects from editing MRSTART.COM, MRTIDY.COM, MRTALK_RUN.COM, and MRTALK_SCH.COM and changing the SUBMIT commands to use a specific queue.

I found the two questions asked during MR installation about access to mailboxes (access from remote systems, and access to SYSTEM mailboxes) to be a bit misleading. I was concerned about the implications of allowing remote access to privileged accounts, something we don't allow on our systems. The answer from DEC is that the privileges and /SYSTEM qualifier on mailboxes within MR are not the same as VMS privileges, and answering YES to the SYSTEM mailbox account does not give access to the VMS SYSTEM account. The reason for the second question is if AI1 (or other user agents) are not on the same node as MR but you still want them to be able to send and receive messages, then remote access must be allowed: you can conceivably set up a network where many nodes run AI1 or other types of mail user agents, but only one node runs the Message Router. If ALL-IN-1 and Message Router run on the same machine, then neither kind of remote access is needed.

If you have a definition in your (the system manager's) account that defines a symbol such as:

```
DEL*ETE := DELETED/LOG
```

or

```
DEL*ETE := DELETED/CONFIRM
```

you will see errors such as

```
%DCL_I_IGNQUAL, qualifiers appearing before this item were ignored  
\SYMBOL\
```

in MRTIDY.LOG and during other command procedures (like AUTOGEN). This is because these procedures use the DELETED/SYMBOL command to delete internal symbols. It's a good idea not to define a new symbol for DELETED in the system manager's account.

MRMAN currently has no good way to save the internal database (the list of mailboxes, etc.). You should save all of this information because you will need it for future installations, or in case the database becomes corrupted. One possibility is to execute a SHOW command in a batch job or during a SET HOST 0:./LOG=logfile to capture the output, something like this:

```
$ run sys$system:mrman
```

```
This is MRMAN V2.0
```

```
MRM> show */add_format
```

```
Add "A1" /Owner=ALLIN1 /System /Run="OA$LIB:OAMTMAIL.COM"
```

```
Add "DEFAULT_DECNET" /Owner=MRMANAGER /Network_node
```


etc. You can replace the "This is" line with:

```
$ RUN SYS$SYSTEM:MRMAN
```

to turn the listing into a command procedure to rebuild the mailboxes. NOTE: this will NOT capture the passwords which are present on the A1 and other mailboxes. You MUST be certain you have documented these passwords somewhere if you don't want to have to re-install all of the Message Router products and ALL-IN-1 in order to set new passwords.

I found that OAMTIMAIL.LOG and OAMTISEND.LOG (found in the ALLIN1 account) contained a lot of superfluous information because there are SET VERIFY commands in the command files which run on the batch queue to transfer mail between the Message Router and A11. DEC says this is deliberate: I find it a useless waste of disk space. If there are any errors, the messages come out in the log files, and are much easier to find if there are only the error and informational messages: therefore I have removed the SET VERIFY commands from OA\$DATA:OAMTIMAIL.COM and OA\$DATA:OAMTISEND.COM. (I should point out that the default on my system is for all batch jobs to be SET NOVERIFY.)

JOIN US IN ANAHEIM FOR OUTSTANDING SESSIONS !!

SUNDAY EVENING WELCOMING RECEPTION

Make sure you stop by our table Sunday evening during the welcoming reception. Say hello and meet the Committee members of the OA SIG. We'll have lots of good information on sessions, activities and maybe a flying pig or two.

SIR PROCESS CONTINUES

If you've been participating in the SIR process through the newsletter, make sure to attend the OA Wishlist session...DEC will be presenting their responses to our Top 10 SIR items and...you'll have the opportunity to personally give any other SIR items you've come up with.

QUESTION & ANSWER

Don't miss the Q&A session, you will have the opportunity to ask an panel of DEC experts questions about your OA software products. This is always an exciting and informative session!

AND MUCH, MUCH MORE

NOTES ON NOTES

— Discussions on VAX Notes, Volume 1, Number 10

Mark H. Hyde & C. J. "Buck" Trayser, Digital Equipment Corporation

Welcome back to Notes on Notes. This month, as opposed to our normal discussion of the VAX Notes product, we intend to use our space to give you a practical guide to using VAX Notes on the DECUS cluster - just in time to help you prepare for the Fall Symposium in Anaheim. We will include a general purpose review on how to use Notes for those of you who haven't been following our articles for the past year.

To start with we will need to give a little history as well as some background information for those new to the Symposium. For the past several symposia Digital has supplied a cluster on the exhibit floor to assist in demonstrating various products. VAX Notes has been one of the installed products on this cluster and its use has been beneficial (and popular) to both Digital and customers. The previous symposium in Nashville saw about 2 dozen conferences created with topics ranging from Artificial Intelligence to Workstations. Hundreds of people entered well over a thousand notes in a 5-day period!

There will be a large exhibit area where numerous demos, booths, and souvenirs are strategically positioned to make the best impact. Near the VAX family display will be the cluster terminals - dozens of them. From these terminals all the demos on the cluster will be available, as well as Notes.

Like last symposium, all of the attendees will have personal accounts in the system. The accounts will be set up with your first initial and last name, such as 'JSMITH'. If you have a common last name AND your first initial is not unique then they may put your middle initial in there as well, as in 'JQSMITH'. For security reasons it is best we don't discuss the password scheme here, but suffice it to say that it will be either very obvious OR the scheme will be posted right on the terminals as it was last year.

O. K., you got logged on and now a dollar prompt is staring you in the face. What next? Well, if you have been reading our articles from the beginning it would be obvious - type 'NOTES'. VAX Notes then sets up your notebook with the default values, which are the first things you will need to change. At the Notes prompt ('Notes>') type SHOW PROFILE. This will display your current profile settings on your screen. Then we start modifying. First, pick your favorite editor. The commands below outline the most popular choices:

```
Notes> SET PROFILE/EDIT=EDT      ! A decent TPU emulation of EDT
Notes> SET PROFILE/EDIT=EVE      ! The Notes default - a simple editor
Notes> SET PROFILE/EDIT=WPS      ! A pretty good WPS style TPU editor
Notes> SET PROFILE/EDIT=(CALL,EDT) ! The real EDT for the die-hards
```

As a suggestion, if you are familiar with EDT, we would encourage the first choice shown above since the TPU editors are more intimately integrated with Notes (split windows, message buffers, spelling checker, etc.).

The next item to modify is your personal name. This is probably best set to your first name followed by your company name. This should make you easier to identify since the combination of your personal name and your username name will give all of the pertinent information.

```
Notes> SET PROFILE/PERSONALNAME="John, Artificial Wombat Widget Co."
```

This will show up similar to:

```
=====
Note 1.0                The DECUS VAX Notes conference                1 reply
VAX::JSMITH "John, Artificial Wombat Widget Co." 26 lines 8-DEC-1987 09:19
-----
```

Don't make it too long, otherwise part of the personal name may disappear beneath the line count...

```
=====
Note 1.0                The DECUS VAX Notes conference                1 reply
VAX::JSMITH "John Smith - V.P. of Artificial Wo" 26 lines 8-DEC-1987 09:19
-----
```

Now that your profile is all set up, enter the DIR command to return to the listing of conferences. All you will probably see at this point is the SAMPLE_CONFERENCE. If you are brand new to VAX Notes we recommend that you read this conference all the way through first. To do this, press the SELECT key or the Keypad 7 key and then follow the instructions that you will find in the conference.

Now, once that is complete, we're off to find out what other conferences are around that we can read and participate in. Here's a preliminary list of the

conferences we expect to be available shortly after the exhibit floor is opened for public access:

DECUSERVE	Discussions on the DECUServe VAX Notes pilot
DATATRIEVE	Datatrieve/4th Generation Languages
VAX 897X	VAX 8974/8978 systems
CONVERSIONS	Conversions from HP, CDC, Burroughs to VAX
IBM-TO-VAX	Conversions from IBM to VAX
VPA	VAX Performance Advisor
VCS	VAXCluster Console System
TOPS	TOPS 10/20 systems
NOTES\$SAMPLE	Sample conference to teach Notes (Read-Only)
NOTES_ON_NOTES	Questions on how to use the VAX Notes utility
RAINBOW	Rainbow PC's
VAXMATE	VAXmate PC's
IBMPC	IBM PC's
PC-NETS	PC networks
RSX	Discussions on the various RSX operating systems and the products/tools/utilities that ship with them.
RSTS	Discussions on the RSTS/e operating system and the products/tools/utilities that ship with it including BASIC.
SYMPOSIUM	Comments on the Anaheim Symposium
TERMINALS	Terminals
TERMINAL_SERVERS	Terminal servers
PRINTERS	Printers laser and otherwise
VMS	Discussions on the VMS operating system and the products/tools/utilities that ship with it.
VMS_PERFORMANCE	System performance
ULTRIX	Ultrix

Of course there are many other potential topics, operating systems, languages, etc. However, the conferences listed above are the ones that we expect to be busy based on previous symposiums as well as other Notes services (such as Pageswapper and DECUServe). There will be an Introduction to VAX Notes session Monday night to review this information and answer questions. You can count on the latest information about the conferences to be available there.

To find the other conferences on the system type:

```
Notes > DIR/CONFERENCE
```

This command will list for you all the conferences that are found in the NOTES\$LIBRARY directory on the cluster, which is the standard location for conference files. Each entry in the display will look like:

```
>SYMPOSIUM
```

```
Title:      Comments on Nashville Symposium
Notice:     Enter your thoughts about the week here
Created:    28-APR-1987 08:38 22 topics, 98 notes Updated: 1-MAY-1987 12:12
```

This gives you relevant information about the subject matter of the conference, its size, etc. Note that an ">" character will point to the conference name. You can move this arrow up and down with the Keypad 2 and 5 keys. When the pointer is pointing to a conference that you would like to follow, press the SELECT key or Keypad 7 key and it will automatically be added to the list of conferences in your notebook. You will probably want to do this at least once every day to see what is new, or, there may be a 'catalog' conference set up where you can read about new conferences as they are created. Once you have selected all the conferences that you may be interested in type DIR to return to your notebook directory listing. Again, using your keypad 2 and 5 keys, move the ">" indicator to be positioned in front of the conference you want to read and press SELECT or keypad 7. A conference that has been added to your notebook is called an ENTRY in your notebook.

If you would like to request that a conference be created that covers a subject not already covered by an existing conference, send VMS mail to an account called MODERATOR with your request. However, make sure you browse the conference contents as well as conferences titles as often the topics are very diverse.

The notes in a conference are organized into topics and replies and are identified numerically in the format T.N where T is the topic number and N is the reply number. Topics are always of the format T.0 and are created with the WRITE command. Use WRITE to create a topic when you want to ask a new question or start a new chain of conversation. Replies are the notes that answer the question or continue the discussion found in the topic note and are created with the REPLY command. Replies to a topic are also numbered sequentially, thus 5.0 is a topic and it may have replies 5.1, 5.2, 5.3, etc. This can possibly be envisioned as we described it in an earlier newsletter...

VAX Notes conference (conceptual layout)

Topics	Replies----->								
1.0	.1								
2.0	.1	.2	.3	.4					
3.0	.1	.2	.3	.4	.5	.6	.7		
4.0	.1	.2	.3						
5.0	.1	.2	.3	.4	.5	.6	.7	.8	.9

Typing a DIR command when in a conference will show you the title of all of the topic notes in the conference. Use this to scan for discussions that interest you.

Although you may WRITE or REPLY to various topics, most people read at least 10 times the amount that they write, so reading the conference in a complete, efficient manner is important. There are several different commands that you may use to read the conferences, however, for simplicity we suggest the ENTER key. The ENTER key is the large key on the far right of the keypad located at the edge of the numeric keypad. Its function combines two of the more popular notes reading commands. Basically, after you open a conference Notes will display the next chronological note (topic or reply) in the conference that you have not yet read. You can press the RETURN key or the ENTER key to view the rest of the note as well as to see the following replies. The difference is that when you get to the last reply of a topic the RETURN key's function is finished. From here you would have to enter a command to find the NEXT UNSEEN note much like what was done when you entered the conference. However, the ENTER key, once you reach the end of a discussion, will automatically do this for you! The ENTER key is often referred to as One-Key-Noting, since it is the only key you have to press once you are in a conference to read it in its entirety. The last page in this article will be a 'cheat-sheet' that you can take with you to Anaheim as a quick reference guide.

Another command that will be most useful is the UPDATE command. Each time you logon to the system and start up Notes, issue the UPDATE command at the Notes> prompt. This will cause Notes to check all the conferences you are following and report to you, via the "Unseen" column on the notebook display, how many new notes there are for you to see in each conference.

If you find that you are interested in Notes and either want to learn the product or expand your skill with it, a PSS called "VAX Notes - from to Customization" has been scheduled and is described in the preliminary guide you should have already received. Maybe we will see you there!

Happy Noting :-)

MORE ON NOTES...

PF1 Shift key	PF2 Help	PF3	PF4 Numeric /comand
7 Select	8	9	- Last read
4	5 Back topic	6	, Next unseen
1 Back reply	2 Next topic	3 Next reply	Enter
0 Next screen		. Previou screen	One key reading

Getting started at Anaheim:

- o Login to system
- o Enter NOTES at DCL prompt
- o SHOW PROFILE at Notes prompt
- o SET PROFILE/PERSONAL = "my name"
- o SET PROFILE/EDITOR = EDT
- o DIR/CONF

Keypad 2 and 5 move arrow pointer ">"
Keypad 7 or SELECT key makes selection

The following table shows introductory commands and gives a brief description of each.

Command	Description
ADD ENTRY	Adds a conference to your Notebook.
CLOSE (or CTRL/Z)	Exit a conference, return to your Notebook.
DIRECTORY	Lists notes in a conference.
DIRECTORY/CONFERENCES	Lists available conferences.
DIRECTORY/NOTEBOOK	Displays the entries in your Notebook.
EXIT (or CTRL/Z)	Ends a VAX Notes session.
EXTRACT	Save a note to a text file.
HELP	Accesses online HELP about VAX Notes commands.
NOTES	Invokes VAX Notes from DCL level.
OPEN	Opens an existing conference in your Notebook.
PRINT	Prints notes.
READ	Displays a note.
REPLY	Enters a reply to a topic.
SHOW ERROR	Displays more information after an error.
UPDATE	Updates the "unseen" count in the Notebook.
WRITE	Enters a new topic.

PERSONAL COMPUTER SPECIAL INTEREST GROUP



NEWSLETTER

Table of Contents

A Rainbow Ballad	PC-1
IBM PC Software Development on a DEC Rainbow	PC-7
DIGITAL ANNOUNCES PHASE DOWN FOR THE PROFESSIONAL 380	PC-10
RSX on a PRO?	PC-12
Fall Symposium Happenings	PC-14
PROgramming Quickie	PC-14

A Rainbow Ballad

Submitted by Tom Warren, Stillwater, OK

INTRODUCTION

At the Dallas DECUS, I won a clock board. When I got home, I opened my Rainbow to install it and found, tucked away in the corner of the Mother Board, a piece of microfiche. What follows is a transcription of that material, secreted by a prophetic assembler at the factory.

THE BALLAD OF RAINBOW ROB

Being an Account of how the Son of Sir Kennen Sought Recognition for his true Heritage and Inheritance, and How Sir Kennen Villainously Betrayed Him, Leaving him to his Fate on the Field of Battle.

DEAR GENTLE READER

Please, pause and do not think
That I waste your time, or the printer's ink
By asking you to hear a tale
You may find trite, untrue, or stale.
I ask no boon except your ear,
And can but hope that you will hear
How Rob of Rainbow, fated knight of old,
Took on Sir Kennen, brave and bold.
It is a tale with moral deep
Of falseness, betrayal, in castle keep;
Of how young Rob was cast aside
After he had helped restore Sir Kennen's pride.
My verse may shock both friend and foe,
But Rainbow owners should certainly know;
Yet times suggest a memory bug
Because you act as though on a drug.
Pray listen, know and hear
How vital virtue will appear,
To confound an evil one
Whom faithful followers outwardly did shun.

I

When long ago, as men reckon time,
And greenery filled an undying world,
There sallied forth a great blue knight;
One whose flag, unfurled,
Proclaimed a saviors' role,
A damsel to console.

She was from Fair Trade and Business Deep,

And one whose fame spread far and wide;
She was distressed, did sorely weep,
Because she could not somehow abide
The press of Sir Data
And his ultimata.

He said, "Grow great with me,
And we will prosper forty-fold;
Great data processing leads to harmony
And that gives great increase in gold;
A promise this fair maid
Found that she could not evade.

She was engulfed by his suave lie;
Succumbed to his tall tale;
And found that she must always comply
No matter how much she'd wail;
Though she knew that it was wrong
She had no choice but to go along.

Her plight spread far; it spread wide;
Until two heard her cries of woe;
One, a Norse, Sir Epli, arrived at Yuletide,
And pledged relief from her foe.
But his way was too complex,
Especially in the user specs.

Then came Sir Thomas, son of Wat,
The blue knight brave and bold;
He easily seduced the maid with what he had brought,
A cure-all with a pledge, he told,
To give her what she'll need,
And to fend off Sir Data and his breed.

Sir Thomas advanced, challenging, brave,
Attacked to control, using his PC brand,
And enslave this scurrilous knave
Who made slave of her fair hand,
Locking him in Armok deep,
Controlled by DOS, that does not sleep.

Sir Data yielded to superior might,
Submitted, at least he seemed to,
Became controlled by that blue knight
Surrendered power — that is true —
At least for now, perhaps,
While laying many noble traps.

For Data knew Thomas' prideful flaw,
And bided time to hatch a plan,
While Thomas celebrated, inspiring awe
From all those under Data's scan,
Including Sir Epli's growing domain
And others outside his wide, wide reign.

Sir Thomas strutted, blew his own horn,
Proclaimed himself savior of all;
Demanded true loyalty, that all must be sworn
No matter what would befall.
But some, seeing a darker side,

Chose not to be a bride.

They saw that Thomas was not the best,
That boons he bestowed were bad;
They knew they'd better not protest,
Because he considers himself part of the Olympiad;
But their time they had to bide,
Praying for relief from outside.

When Sir Thomas sensed the growing split,
His wrath went far and near;
He would not stand for counterfeit;
No one would interfere.
The gloom touched our lady fair,
Who heaven sent a pleading prayer.

II

In a land outside Sir Thomas' reach
Lived a rebel from Thomas' own isle,
One Sir Kennen, a knight of subtle speech
Who would not abide Sir Thomas' guile,
Had set off on his own,
Building a castle of blue-gray stone.

Sir Kennen lived in Marbor Town,
Harboring thoughts to unseat his foe,
And claim a part of that renown
Which Thomas would not forgo.
Kennen'd give high quality,
Through his first son— PDP.

He gathered round him craftsmen fine,
To make resplendent his machine,
And wage unholy war divine
Against Sir Thomas, spiteful libertine,
Beating him soundly at his own game
Was Kennen's sole, obsessive aim.

Now, it chanced that Kennen heard a tale
Of suffering, pain, and dismay,
Involving Thomas and a certain female
In a land quite far away;
A land suppressed and sorely tried,
Under the sway of Thomas' awesome pride.

It was a land where people believed
That quality in goods would make a sale;
But now that practice no one believed
Because they learned they must avail
Themselves of service ever there
To fix mistakes found everywhere.

It seems Sir Thomas, once inside,
Fawned, pampered, and marketed his view;
Until the people no longer could decide
If service or that product blue
Solved problems that they had,
Including just when they should be glad.

Their queen, a stately blond without sons,

Stood tall and proud with royal bearing;
She was besieged by Thomas' minions;
Yet she denied him anything.
Still, on he pressed his loathsome suit,
Acting natural— an awful brute.

"Use mine," Sir Thomas pressed, "I set the standard;
Follow me and join my flock;
I know no innovation has occurred
Except as I unlock,
Because my way's the only one;
Those others (Kennen) have not begun."

"Good sir," our blushing maid replied,
"You say a deal of truth;
But you are known for tyrannicide;
Besides, you're a bit uncouth.
Your case we will not heed;
We all agree on what we need."

"What you need," his wrath they incurred;
"How can you tell what you need;
It's for me to say how you'll be served,
Because I know and you must heed.
My knowledge vast I'll use for you,
And not allow it for review!"

She hung her head, foresaw her fate,
And sent another message far and wide
To plead for champions, not sedate,
To rescue her, make her his bride,
And gain a measure of accolade
By saving this fair, helpless maid.

Sir Kennen heard her cry for help;
Resolved to go and there to shine;
Called forth his minions and his whelp,
His slaves from the Atlanta Hot Line,
As well as Rob, forgotten son,
At whose expense he'd have some fun.

This Rob, as those in Marbor knew,
Was sired as an after thought;
"But, here he is," Sir Kennen misconstrued;
"Don't worry; he'll come to naught.
Lock him away; don't tell a soul;
We've more than enough to gain control."

"So, onward we go and scatter debris;
Advance, attack, slay the base blue beast.
He sent his front line, his PDPs
And VAX-like creatures come forth from the East.
All day the battle raged and raged;
Because neither side would disengage.

Then, suddenly, Sir Thomas cried for Kennen to yield
And was about to slay him,
When Rob rode in, with lance and shield,
Forcing Thomas's hopes to dim.
Rob rescued Kennen there that day

Amidst all the horror and melee.

"My grudging thanks to you my (gulp) son,"
Kennen spit through teeth clinched tight.
You see he knew Rob was not the one
He favored with all his true might.
A mistake, a folly, a one-night stand,
Who helped a little, but really was not in demand.

Suddenly, Sir Thomas broke, made to advance;
He drew his sword, prepared to strike;
He aimed a blow, left naught to chance;
His target was the one we like,
Our Rob, resourceful youth,
Sired from a knight full of untruth.

Sir Kennen saw Sir Thomas advance;
He turned his back, rode on below,
Left Rob to die, his own fame enhance;
He'll rid himself of Rob; he'll simply go.
But Kennen reckoned not with Rob's wide fame,
Because a squire advanced, deflected Thomas' aim.

That squire Sir Kennen soundly cursed,
Refused to let him carry lance or shield;
That squire preferred Rob as master first,
And rejected Kennen; refused to yield.
That squire was one we've met before:
'Tis none other than whom Thomas would take for whore.

It was our maid who loved Rob from afar;
'Twas she who saved our hero bold;
She knew Sir Kennen was false— in peace and war;
But Rob was one whom she'd heard told
The truth about Sir Data's sway
And could ride forth to save the day.

III

The pattern was set that day of fate,
As Kennen went to save the damsel fair;
Poor Rob was banished from Kennen's estate—
Thrust on his own in deep despair,
Alone, or so Kennen thought was true,
Denying what was only Rob's due.

But Kennen little reckoned with Rob's spreading fame
Among a following they clearly recognized
Rob's truth and virtue; his acclaim,
They rallied to his side and loudly vocalized,
"Why abandon you a son so pure?"
" 'Tis simple, Kennen said: "He's mature."

"Mature," they cried. "He's just begun
To show how he can challenge blue's domain;
He's better bred; he's far from done.
So how can you disdain,
Cast him aside as if some punk;
He's far from being merely junk."

"You do not grasp my master plan;

The big picture of what I do.
 Rob's place is set; it's in Iran,
 With PDP and crew.
 So, leave me be; do not me bug.
 Disperse! Go home! Go wash a rug!"

The Group refused to let it die;
 So too should you of DECUS;
 Press hard for Rob; stand up; defy;
 This can't be the terminus.
 Rob must live on, develop and grow;
 Give him a chance— this you must know.

Sir Kennen slinked back to Marbor keep,
 Called forth a VAX-like creature;
 They plotted secrets, plans most deep;
 He knew Rob's group would fear to stir.
 He wanted them silent once and for all;
 Advanced a new son who appeared last fall.

That creature appeared, at once a paradox,
 Encased in garb strange and unearthly;
 He faltered, he fell, tripped on a pizza box.
 Sir Kennen clearly was in agony.
 This newest heir — at the expense of Rob—
 Threatened destruction from an unruly mob.

But Rob sprang forth; his half-brother to save;
 But he was brusquely pushed aside;
 Sir Kennen's anger, typical of the knave,
 Cut deeply into Rob's gentle pride.
 Rob wept to finally recognize,
 That Kennen did, indeed, him despise!

So, Rob withdrew to his faithful group.
 They cheered and rallied round.
 He said, "We need time to think, faithful troupe,
 Because we can Sir Kennen astound!"
 His followers cheered and lifted him on high;
 To Kennen they all yelled a last "Good-bye!"

"We need much help, if we're to live,"
 Rob told his faithful squire and throng.
 "We need support to be assertive,
 And hold our own where we belong."
 He then, with squire, rode off at midday,
 And still they ride as I end this lay.

POSTSCRIPT

My tale is done. I hope you'll hear
 About Rob and Kennen, his false sire.
 You are the ones who must interfere
 To change this scene most dire!

Stand tall, be bold, be proud of him;
 Press friends and strangers both
 To rally round the heir of Rainbow, now dim,
 And urge they support his growth.

Here ends the message in the computer. I don't know what happened to the author - whether he (or she) is now assembling VAXmates and preparing another message for some future owner contemplating the future of a mature product. Perhaps so. Knights have always been noted for deviousness - especially when confronted with damsels and unfavorite sons.

IBM PC Software Development on a DEC Rainbow

(c) 1987 by Kenneth Presley, President, CICA Inc., 4513 Flintlock Drive,
Louisville, KY 40216, (502)448-2719

I - INTRODUCTION

This article describes the use of a DEC RAINBOW for IBM PC software package development. Although not a common practice, it is certainly possible and cost-effective to use the DEC RAINBOW to develop software that is intended to run on IBM PCs and compatibles. Not only can the IBM PC target software be developed on a DEC RAINBOW, but it can be tested and packaged as well. Here at CICA, we have been in business for six years and have been engaged in DEC RAINBOW and DEC RAINBOW-based IBM PC software development for approximately two and a half years. The primary reasons for doing IBM PC development on a DEC RAINBOW are the following: (1) economic, i.e. not having to purchase a second development PC when one already owns a DEC RAINBOW, (2) the ease of dual machine development (developing versions of the product to run on both the IBM PC and the DEC RAINBOW), (3) the availability of the familiar and sometimes superior DEC RAINBOW software and RAINBOW MS/DOS environment, (4) the office space and logistical advantage of dealing with a single PC machine, (5) the pure satisfaction of using a superior machine like the DEC RAINBOW to develop products for the inferior but more widely-sold IBM PC.

II - THE SETUP HARDWARE & SOFTWARE TOOLS REQUIRED

The primary hardware addition to the DEC RAINBOW that is needed for easy IBM PC software development, testing and packaging is a double-sided double density IBM PC compatible diskette drive. We use the Suitable Solutions IDRIVE product, although there are also other products from companies such as HI-TECH materials (the RB-LINK) and DUNCAN MCDONALD. At CICA, we have found the IDRIVE IBM PC compatible diskette drive to be the most cost-effective, reliable, and easiest product to install and use. The IDRIVE (either the internal or external model) gives the software developer complete IBM PC MEDIA compatibility on the DEC RAINBOW. The IDRIVE gives the developer the ability to read or write to and from IBM PC compatible double-sided double density diskettes. This capability is very important for two reasons: (1) it allows the physical media transport to the DEC RAINBOW of various IBM PC software development packages, (2) it allows the creation of the final IBM PC software distribution diskettes on the DEC RAINBOW.

Additionally, on the DEC RAINBOW hardware side, it is almost a necessity to have a hard disk and the maximum amount of memory (896KB on a RAINBOW 100B and 832KB on a RAINBOW 100A).

There are several pieces of software that are needed to allow IBM PC development on the DEC RAINBOW. The first and foremost of these software packages is the CODE BLUE IBM emulator from INTERSECTING CONCEPTS. This splendid package gives the software developer the ability to run many IBM PC development tools and languages on the DEC RAINBOW. It also allows the IBM PC targetted software under development to be tested and run. It will NOT allow the use or development of IBM graphics-oriented packages. However, as you will see later, this is not really a serious handicap, as many fine pseudo-graphic packages will run and can also be developed on the DEC RAINBOW under the CODE BLUE emulator.

The second piece of software that will be needed for IBM PC software development on the DEC RAINBOW is an obvious one, a language compiler. Here at CICA, we use the IBM PC version of TURBO PASCAL (from BORLAND INTERNATIONAL). This highly professional and polished version of the pascal language is run under control of the above mentioned CODE BLUE IBM PC emulator. There are several other IBM PC language compilers that we have successfully run on the DEC RAINBOW (some with CODE BLUE & some without CODE BLUE!). Among these are several versions of MICROSOFT BASIC (from MICROSOFT INC.), and TURBO PROLOG (again from BORLAND INTERNATIONAL). There are most likely several IBM PC versions of 'C' compilers that will also run on the DEC RAINBOW, but we have not tested any of these. You should of course choose the language that you feel most comfortable with (although for serious IBM PC development you will probably use some form of 'C', PASCAL, or assembler).

A third piece of software that we find very useful, although not a necessity, is a public domain utility called DCOPY (from the pseudo-named JOERG GENIUS of Munich, West Germany!). This handy diskette utility has one very useful feature, the image copy feature. This feature allows the software developer to dump or backup the entire contents of a floppy diskette to a single image file on a hard disk. The utility supports almost all diskette types, including DEC RX50 quad density and IBM PC double and single sided diskettes. This feature is very useful for creating software distribution image files on hard disk and then using that image file to recreate distribution diskettes. This is particularly true when using the IDRIVE to create IBM PC distribution diskettes, as you can have only a single IDRIVE diskette drive on your DEC RAINBOW. Of course when using the IDRIVE to create IBM PC software distributions, you could also create your distribution diskettes by copying the various software components of your IBM-target package from the hard disk to the IDRIVE with a batch file. But that method is cumbersome and much slower, when compared with the DCOPY image file method.

III - PRODUCT DEVELOPMENT CONSIDERATIONS - There are several product development issues that must be considered when doing IBM PC software development on a DEC RAINBOW. These issues include: (1) general dual system compatibility, (2) video screen I/O methods including pseudo-graphics and window design, (3) keyboard compatibility and scanning methods, (4) miscellaneous issues such as memory differences between the two machines.

One of the first considerations in the product development process is whether a company wishes to have versions of the product to market for both the IBM PC and the DEC RAINBOW (after all, we ARE working on a DEC RAINBOW, so we should certainly consider having a version of the product that will run on it!). If the company decides that is desirable, (as we did at CICA), then the software developer should act accordingly and make high-level design decisions that make it easy to port the product to both machines. One of the first things that the developer should do to facilitate this dual-machine development is to write very modular code in a language that is available on both machines (such as TURBO-PASCAL). This modular development will allow the video I/O & keyboard routines that are appropriate for the particular product version to be inserted into the code rather easily.

The next and probably major consideration for IBM PC software development on the DEC RAINBOW is how to handle the video screen I/O. Such I/O is a major point of difference between the IBM PC and the DEC RAINBOW. If the target software will not be particularly video-oriented (such as an accounting or statistical package) then the developer can use the standard I/O statements available in the chosen language. This method certainly simplifies IBM PC software development and also greatly aids the IBM PC - DEC RAINBOW portability issue. However, screen I/O using standard language I/O statements is usually dreadfully slow and thus not suitable for exciting, video-oriented software (such as our own multi-window file viewing utility). In such cases the developer will usually want to write directly into the PC's video memory to achieve the desired screen update speed. This is certainly possible on both the IBM PC and the DEC RAINBOW but does require different software routines due to the different hardware nature of the DEC RAINBOW and the IBM PC (we did say to keep your code modular, remember!). Here is where the CODE BLUE IBM PC emulator again comes to the rescue. CODE BLUE will allow the use of DEC RAINBOW memory above 768K to be used to map the video-memory found in the IBM PC. Therefore it is possible to use IBM PC routines that write directly into IBM video-mapped memory and still be able to test programs containing these routines on the DEC RAINBOW under CODE BLUE.

There are also a few other relatively minor points of contention in video I/O development. One of these is the absence of an extended character ROM in the DEC RAINBOW. The IBM PC has an extended character ROM that will print characters such as the symbols for the four suits in a deck of cards. Since the RAINBOW has no equivalent hardware chip, the developer of IBM PC software on the RAINBOW must refrain from using the character codes that generate these characters, as there would be no way to test them. Another problem has been the fact that the IBM PC video monitor supports 25 lines whereas the DEC RAINBOW supports only 24. The solution in the past was simple, write IBM PC software that utilizes only 24 lines on the screen. Now however, the new version (V2.0) of CODE BLUE will support 25 line mode on the DEC RAINBOW, so one should be able to adequately test IBM PC software designed to utilize all 25 lines of the IBM PC video monitor.

Of course, we also mentioned previously the major limitation that IBM PC graphic software cannot really be tested or run on the DEC RAINBOW. However, what we like to call video pseudo-graphics can be tested and implemented in IBM PC software developed on the DEC RAINBOW. With the proper use of the ANSI line-drawing characters and other various character symbols, it is possible to create relatively good and visual screens, create windows, and otherwise simulate screen I/O that is normally done with graphics on the IBM PC (our packages developed at CICA, particularly our chess game program where line-drawing characters and other symbols are used to create the chess pieces, our menu system, and our multi-window file viewing utility all take advantage of such techniques).

The next major consideration for the IBM PC software developer on the DEC RAINBOW is the physical difference between the keyboards on the two machines. Also because of these physical differences in function keys and key layout, the keys are scanned differently from a software standpoint. Fortunately, the CODE BLUE emulator utility again comes to the rescue and relieves the software developer of many of the problems associated with these keyboard differences. CODE BLUE maps all of the IBM PC function keys (F1 - F10) onto various DEC RAINBOW function keys. (The new version V2.0 of CODE BLUE maps the IBM F1-F10 keys directly to the top-row and left-most F1-F10 keys on the DEC RAINBOW keyboard). That makes it possible to code the IBM PC targeted software package to perform with the standard IBM PC function keys, yet still allows the package to be debugged and tested on the RAINBOW with the use of CODE BLUE. There is one caveat however, using the previous versions of CODE BLUE (V1.x). The restriction is that the developer cannot code routines that use an IBM ALT-SHIFT key combination. That is because the DEC RAINBOW does not have a separate ALT key, it must be simulated by pressing the CTRL & SHIFT keys simultaneously on the RAINBOW. (You would in effect need to press a SHIFT & CTRL-SHIFT simultaneously on the DEC RAINBOW, which of course is not possible). The new version V2.0 of CODE BLUE solves this problem by mapping the IBM ALT key to the COMPOSE key on the DEC RAINBOW keyboard.

As we mentioned above, if product co-development is desired and the developer codes in a modular fashion, one keyboard routine can be constructed for the IBM PC version of the software product, and another routine can be constructed for the DEC RAINBOW version of the product, then easily integrated into the other software modules.

There are also a few other minor differences between the IBM PC and the DEC RAINBOW that must be taken into consideration during the development process. These include the amount of physical memory that is potentially available on the two machines, (640K for the IBM PC and 896K for the DEC RAINBOW), and differences in the number and makeup of the asynchronous serial ports on the two machines. Although not a factor here at CICA, if the developer is going to write memory drive software or communications software, these differences must be given close attention.

IV - CREATING FINAL PRODUCT DISTRIBUTIONS

The process of creating IBM PC diskette compatible software distributions on a DEC RAINBOW is rather simple with the aforementioned IDRIVE product. However, there are several software techniques that can make this process very fast and very automatic.

The first technique to fast and easy IBM PC distribution creation is the use of the also previously mentioned public domain DCOPY program. This program will allow the creation of a single image file containing an entire floppy's worth of files. We first create our master distribution floppy using normal MS/DOS COPY statements to an IBM diskette in the IDRIVE. Then we use DCOPY to make an image file of the distribution diskette onto our hard disk. We then use a batch file that also uses a memory (or RAM) drive to quickly make copies of the image file to blank IBM diskettes inserted into the IDRIVE. Below is the contents of such a batch file:

```
ECHO OFF
SETRAM J 400
COPY e:\util\dcopy.exe j:\zcopy.exe > NUL
ERASE j:*.*img > NUL
COPY f:\ibmimage\viewing.img j: > NUL
PATH j:\;e:\util;e:\
ECHO ON
:START
```



```
PAUSE VIEWING.BAT - Insert IBM formatted diskette into IDRIVE I:  
zcopy j:viewing.img i:  
dir/w i:  
GOTO START
```

The above batch file is just one example of creating IBM distribution floppies. Although there are probably many other methods, this one quickly clones IBM PC distributions while minimizing hard disk accesses and hard disk wear and tear.

V - CONCLUSIONS

As one can see from the preceding discussion, it is certainly possible to create and package IBM PC software on a DEC RAINBOW. I also believe that this article has shown that RAINBOW-based IBM PC software development is very economical and makes particular sense if a company is going to have a version of the target software package for both the IBM PC and the DEC RAINBOW. It has been our experience here at CICA that using a single machine for development has also improved productivity, primarily by reducing the learning curve that is required if one uses multiple machines and different development tools.

DIGITAL ANNOUNCES PHASE DOWN FOR THE PROFESSIONAL 380

Digital's Micro Systems Development announced on August 24, 1987 Phase Down for its Professional series of hardware and software. Digital now has follow-on products for the Professional Series in the MicroVAX 2000, VAXStation 2000, VAXmate, and MicroPDP-11/53. For this reason, the Professional Series will be retired during the coming year.

The retirement schedule, also announced, is designed to provide Digital's customers adequate migration planning time:

SYSTEMS

Last hardware system orders	March 25, 1988
Last hardware system shipments	July 1, 1988

This systems retirement schedule provides a 7-month lead time for final orders. It also provides a 10-month lead time from the phase down announcement to the final ship.

HARDWARE OPTIONS, APPLICATIONS SOFTWARE, AND P/OS

Last orders for hardware options, P/OS, and software applications	End of March, 1989
--	--------------------

Last shipments for hardware options, applications software, and P/OS	End of June, 1989.
---	--------------------

This options and software retirement schedule provides a 16-month lead time for final orders of options and software. It also provides a 19-month lead from the phase down announcement to the final ship.

SERVICE AND SUPPORT

Hardware support for systems and options will be provided by Digital Field Service for at least 7 years after end-of-life.

Software Product Services will be notifying PRO Telephone Advisory Service contract customers and Update Service customers that Software Product Services will continue to be available on these products until June, 1989.

PRO-BASED VAX CONSOLE SYSTEMS

The PRO 380-based VAX Console Systems WILL NOT BE RETIRED AT THIS TIME. A separate manufacturing effort will provide the means to continue using the PRO 380 as the console throughout the life of the VAX 8530/8550/8700/8800 products.

FOLLOW-ON PRODUCTS

MATCHING THE WIDE RANGE OF PRO APPLICATIONS

The follow-on Low End VAX systems and MicroPDP-11 products provide a range of functionality and performance that matches the range of Professional applications. For example:

- Graphics-based PRO applications can migrate to the VAXstation 2000 series of workstations, with high performance monochrome and color graphics, higher performance and higher disks capacity than the PRO, full DECnet and Ethernet support, and A-to-Z menu environment.
- Real-time PRO applications can migrate very cost-effectively to the MicroPDP-11/53 or VAXLAB or, for higher performance and disk capacity, to the MicroVAX 2000.
- PRO Office applications can migrate for higher performance, higher disk capacity and networking capability, and long-term growth to the MicroVAX 2000 or MicroVAX II series of products.

LONG-TERM GROWTH

Migration to the MicroVAX family of products offers long-term growth, the most sophisticated networking capability in the industry, and superior VAX/VMS-based applications and application development tools.

ENHANCEMENTS TO CURRENT PROFESSIONAL SYSTEMS AND SOFTWARE

NEW HARD DISK SUPPORT

Also on August 24, 1987, Digital announced the availability of the RCD32-AA and RCD53-A hard disk subsystems for the PRO 380.

The RCD32-AA is a 40 MB half-height disk subsystem incorporating the RD32 disk and its disk controller. It replaces the RCD52-A with improved price/performance and a lower total cost.

The RCD53-A is a 67 MB full-height subsystem incorporating the RD53 disk and its disk controller. It doubles the capacity of current hard disk offerings, with high performance and a competitive price.

P/OS V3.2 AND PRO/TOOL KIT V3.2 ANNOUNCED

Also on August 24, 1987, Digital announced P/OS V3.2 operating system and PRO/Tool Kit V3.2 for the Professional Series.

P/OS V3.2 provides support for the new hard disks. It also provides enhanced LA75 support to allow graphics to be printed on the LA75 in the correct aspect ratio. Previous versions of P/OS support the LA75 correctly only in text mode. Support for the LN03-PLUS in P/OS V3.2 also allows more complex images to be printed on the laser printer utilizing Print Services.

A new PDI driver is included in P/OS V3.2 to support the DECtouch monitor, the communications port, and the printer port.

A new utility, SYSTEM INSTALLATION AND CUSTOMIZATION, enables users of standalone or single-user systems to tailor the system to remove unneeded functions and recover disk space. This is not recommended for use in a server environment.

Another utility, DISK MAINTENANCE, allows the user to add bad blocks to the bad block file manually or to search the disk for lost files. It can be either installed or run as a standalone system application.

A GIDIS to SIXEL converter is included for those devices supported by P/OS V3.2 print services.

The Application Diskette Builder in PRO/Tool Kit V3.2 has been enhanced.

PRO/DECnet V2.1 ANNOUNCED

PRO/DECnet V2.1 supports P/OS V3.2 as well as the previous versions (V3.0 and V3.1). Known problems have been corrected. Password length has been increased from 6 to 8 characters. The Remote Terminal Utility allows the user to set the terminal characteristic to 8-bit. This allows Remote Terminal Utility to work with systems using international character sets.

TAPE BACKUP NOW AVAILABLE FROM A THIRD PARTY

Cipher Data Products is now offering a cartridge tape backup system for the PRO. It consists of a quarter-inch floppy tape drive with its own cabinet and power supply, a controller board, a complete software package, a chassis adapter and cabling. Further information can be obtained directly from the vendor at 415-964-2211.

PHASE DOWN SCHEDULE

DATE	ACTIVITY
8/24/87	Phase Down announced by Digital
3/26/88	Final HW Systems Orders
7/2/88	Final HW Systems Ships
3/31/89	Final Order of Options, P/OS, and Applications Software
7/1/89	Final Ship of Options, P/OS, and Applications Software
7/1/89	End of Software Support Contracts
7+ Years	Hardware Support Continues

RSX on a PRO?

by Gary Rice, PC SIG Newsletter Editor

For some time now rumors have circulated occasionally about PROs running the native RSX operating system. One reader told me that he actually saw it at a company he was visiting. Here is the REAL story straight from the person who DID it along with additional comments from his friends.

The Players:

Brian McCarthy - DEC Software Engineer (The guilty culprit)
Alan Frisbie - Anxious Fan
Bart Z. Ledernam - Even MORE Anxious Fan
Ed Cetron - Culinary Artist Extrordinaire
John Covert - DECie
Bill Tabor - Independent Observer

Their story begins . . .

Brian: Funny, I have no trouble with my PRO at all. But, then again, I'm running Micro/RSX V3.0.
Alan: Is there any way a user can obtain Micro/RSX for the PRO? A friend of mine would like very much to make the switch, and I might even buy a PRO if I could run real RSX on it.
Bart: It's comments like that that make the peasants storm the Bastille with pitchforks and torches! Seriously, there are many of us out here who would like to turn our PROs into something really useful with a "real" operating system.

John: What Brian doesn't tell you is that (unless things have changed) Micro/RSX V3.0 for the PRO doesn't support the PRO bitmap terminal. Brian has a VTxxx plugged into the printer/console port.

Bart: A lot would depend on how it was genned. If you look at an M-Plus kit you can see support for the PRO in conditional code. And if you look at the PRO Microfiche, you can see the drivers. Since P/OS is basically emasculated M-Plus, there is no reason why Micro-RSX couldn't be genned with the bitmap driver. The only "real" question is how much it would cost DEC to develop the kit (mostly packaging, SDC, and cataloging and literature costs, which I believe would cost more than actually genning the kit) and how much they could get in revenue from selling it. I think DEC would also receive a huge benefit in increasing the good will of the many customers who would appreciate the kit and would be better disposed towards buying things from DEC, but it's very hard to put a dollar value on that when trying to convince a business manager to support a product.

John: It isn't quite that simple; the bitmap driver doesn't support the normal CLI interface; that would have to be added and debugged before it could be used at all. In fact, isn't that really the main difference between P/OS and RSX: CLI support and the terminal interface to the CLI?

Brian: Before I'm swamped: NO, Nobody can have a copy of RSX for the PRO. Don't feel bad, I won't give it out inside DEC either. We did some work on the project and then abandoned it due to business reasons. It is running on one PRO-380. It is not finished and would take a goodly amount of work to make useful. And before John and Bart kill each other on speculation: The bitmap isn't the problem. I got the bitmap to work in the early days of the project. In fact, one of my favorite pastimes in that corner of the lab is the following:

```
>INS $PIPFSL/PAR=BITMAP/XHR=NO
>PIP NL:=[*]*.*;*
```

And then watch pip run on the bitmap display. You have to be careful where the cursor is or the blinking will kill PIP. Why /XHR=no? So that when you hit enough returns and the firmware scrolls PIP, only pip blows up. The system goes away if the header gets scrolled.

The major problem with the Bitmap is that part of TTDRV must be 32 word block aligned, so I couldn't easily get the driver to build in one pass, which doesn't help our system build procedure.

The major problems which precluded distribution are threefold:

- 1) One of the organs "emasculated" in P/OS was on-line reconfiguration and HRC. In Micro/RSX, HRC... does disk sizing (so does sav in HRSIZ) and I never finished the code to do disk sizing on the PRO.
- 2) The system does require a BCC08 and a terminal on the printer port to boot. I never did the SAVE work to get it to understand consoleless systems.
- 3) To be useful with RSX, you'd have to be able to use the Comm port, and use the real printer port instead of console emulation, and to use the mux board. I never did the port drivers for these.

There are also some small items like a SETUP task. Oh yeah, and the firmware doesn't fit in the installation system all that well. You need a Q-bus to CTI bus bridge to install the software on a PDP-11 and move it to a PRO.

Other than that, no problem. Well, XDT needs some work, and there never were any crash drivers done for the PRO disks. But other than that... Oh and Error logging, of course. But other than that... And I don't know if the firmware would work on a 380 with I/D space enabled.

Some guy like Alan Frisbie with a M+ 3.0 distribution kit and a P/OS source listing or fiche could finish it, I'm sure.

Alan: You are just trying to keep me from working on the VMS project, aren't you? :-) Besides, Jim McGlinchey just finished telling me about all the "wonderful" aspects of the PRO that I wasn't aware of. Now I can see why you compare it (unfavorably) to a bowling ball.

Ed: For those who don't know, I was in Spit Brook the day Brian 'announced' M+ 3.0 for the PRO. So blame me for the peasant rebellion not Brian

Bill: There is a version of RSX-11m running on the PRO. It was developed at Racal-Milgo by the D.E.C. resident software person. This software has been in the field for over a year. It is running a turn key application so I do not know if everything is there.

Brian: It is RSX-11M-PLUS, not RSX-11M. The specialist took the work we had done very early and got it to work for his applications. It has all of the restrictions I mentioned earlier, works only on the bounded configuration at Racal-Milgo, (It uses a terminal on the printer port and doesn't support the bitmap) and DEC has no intention of making it available elsewhere. It is also stuck (I believe) at an unsupported (or supported for Racal-Milgo under their specific contract) version of RSX-11M-PLUS.

Thus ends the saga of RSX on a PRO. It doesn't appear that RSX on a PRO will EVER be a reality, especially with the announcement of the PRO's retirement (see the article in this issue announcing that retirement).

Fall Symposium Happenings

by Gary Rice

For those of you who aren't going to the Fall Symposium, here is a list of the PRO sessions that you will be missing:

	<i>Time</i>	<i>Session #</i>	<i>Title</i>
Monday	1830	P015	RSX to POS Migration
	1915	P016	Writing a Synergy Application
	2100	P019	PRO Floppy DCL Bootstrap
	2200	P020	P300 Networked PRO Workstation
Tuesday	0900	P006	PRO Working Group Meeting
	1300	P021	PRO Programming Tips
Thursday	1230	P042	PRO to VAXststion Migration
Friday	1230	P024	PRO Public Domain Software
	1300	P018	PRO/SIGHT - An artist's view
	1400	P014	PRO/SIGHT - Image Display
	1500	P005	PRO Q&A Product Panel

I will be giving the Public Domain presentation. To date, I have collected over 130 RX50 diskettes full of goodies.

PROgramming Quickie

by Gary Rice

In this month's Quickie, I will give you a first look at the WIMP system directive. This feature is PRO specific. The letters of the directive stand for "What's In My Professional". The Executive Refernece manual where the directive is documented gives the calling sequence, but fails to mention what the numbers that are returned actually mean. Here is a program that has a portion of that missing documentation built in to it.

The program will tell you what options are present in your 350 or 380. The list is NOT complete, though. Perhaps you can provide the numbers for modules that I didn't have available for testing.

C EQUIP.FTN - This program can be used to "list" the contents of
C your PRO's option slots
C

C AUTHOR: Gary Rice
C

C CREATED: September 5, 1987
C

C REVISIONS: None
C

C INPUTS: None
C

C OUTPUTS: None
C

C NOTES: None
C

C*****

C

PROGRAM EQUIP

C

C

INTEGER*2 BUFFER(146)

C

C

CALL WIMP (12,BUFFER,146) ! Get the raw data

DO 120, I = 8, 23, 2

IF (BUFFER(I) .EQ. "1002) THEN
TYPE *, 'PRO 350 w/Extended Bitmap Option'
GOTO 120

END IF

IF (BUFFER(I) .EQ. "10050) THEN
TYPE *, 'PRO 380 w/Extended Bitmap Option'
GOTO 120

END IF

IF (BUFFER(I) .EQ. "2004) THEN
TYPE *, 'Includes RX50 diskette controller'
GOTO 120

END IF

IF (BUFFER(I) .EQ. "401) THEN
TYPE *, 'Includes RD50/52 controller'
GOTO 120

END IF

IF (BUFFER(I) .EQ. "41) THEN
TYPE *, 'Includes Telephone Management System'
GOTO 120

END IF

IF (BUFFER(I) .EQ. "46) THEN
TYPE *, 'Includes Real Time Interface Module'
GOTO 120

END IF

IF (BUFFER(I) .EQ. "34) THEN
TYPE *, 'Includes 256K memory expansion'
GOTO 120

END IF

```

        IF (BUFFER(I) .EQ. "42) THEN
            TYPE *, 'Includes Ethernet controller'
            GOTO 120
        END IF
        IF (BUFFER(I) .EQ. "43) THEN
            TYPE *, 'Includes C/PM card'
        END IF
120    CONTINUE
        END

```

The command file to link the previous program is as follows:


```

EQUIP/FP/CP=EQUIP
LB:[1,5]PROF77/LB
/
;
; EQUATE P/OS SYMBOLS TO LUNS
;
GBLDEF=TT$LUN:5
GBLDEF=WC$LUN:0
GBLDEF=MS$LUN:6
GBLDEF=HL$LUN:0
GBLDEF=MN$LUN:0
;
; DEFINE EVENT FLAG
;
GBLDEF=TT$EFN:1
;
; DEFINE CLUSTER SCHEME
;
CLSTR=PROF77,POSRES,RMSRES:RO
//

```


Send me your OWN PROgramming Quickie (on an RX50 please) and I will publish it here in this ongoing Newsletter segment. I will ALSO add it to the growing collection of Public Domain software available for the PRO.

My address can be found toward the back of these Newsletters in the "Steering Committee Lists" section.



ARTIFICIAL INTELLIGENCE

SPECIAL INTEREST GROUP



*"Sloths are so human in appearance--
and in some of their ways--
that inevitably one tends to judge them
by human standards."*

--Hermann Tirler, *A Sloth in the Family*

The AISIG strives to promote AI technology that the practioner can 'hang on to'... Join us in **Anaheim, CA, December 7-11, 1987** at the next DECUS semi-annual **Symposium**, where technical interchange is the name of the game!

➔ For more information about DECUS, the AISIG, or upcoming symposia, contact:
Digital Equipment Computer Users Society, 219 Boston Post Road (BP02), Marlboro, MA 01752

Titles of AI Sessions Presented at Nashville, Spring '87

Intro to Languages & Tools for AI Development
ART: Automated Reasoning Tool
Intro to PROLOG, Using PROLOG to Build Expert Systems
Using PROLOG for Advanced Application Development
Expert Systems Concepts Crash Course
Knowledge Representation Issues
An AI Project Case Study -- The First Six Months
VAX OPS5 Product Description, OPS5 Programming Techniques
Designing and Building Expert Systems in OPS5
OPS5 Control Constructs, Writing Efficient OPS5 Rules
Adding Meta-Rules to OPS5 -- A Proposed Extension, OPS5 Q & A
Semantic Problems in AI and Other Areas of Computing
Getting Started in AI -- Novice Q & A
AI Impact on Database Technology
Packaging and Configuring VAX Systems for AI
AI in a Distributed Network
Foundations of AI in Philosophy and Cognitive Science
Management Issues for AI
Corporate Strategy for AI and CIM
Intro to VAX LISP, VAX LISP for Advanced Users
How to Select a Project for Your First Expert System
How to Write Your First Expert System Proposal
AI on Personal Computers
What's an Inference Engine, Anyway?
Introduction to Machine Learning
Extending the VAX LISP Editor, Lisp Q & A
Neuron Data Nexpert Object
AI in the Public Domain, AI and Image Processing
THE GREAT AI TOOL DEBATE

Featured Speakers Included:

Earl Sacerdoti, Charles Weisbin, Charles Jorgensen
Mike Howard, Dennis O'Conner

Titles of AI Sessions Submitted for Anaheim, Fall '87

AI Novice Q & A
Survey Of AI Literature
AI At Digital
The Fifth Generation
Domain Dependent Shells
Expert Systems Project Management
AI In The Public Domain
Intro To Languages & Tools For AI Development
Expert Systems Concepts Crash Course
Federated Expert Systems
Some Simple Machine Learning Examples
Calling External Routines From VAX OPS5
Real-time Support In VAX OPS5
Designing & Building Systems In OPS5
Writing Efficient OPS5 Rules
Artificial Intelligence Applied To The Help Function
Selling AI To Management
AI And Image Processing
PROLOG Programming: A Crash Course
Natural Language Processing: A Crash Course
Incorporating Probabilities & Beliefs Into Exp Systems
Productivity Tools & The S/W Dev Life Cycle
Introduction To Natural Language Processing
Knowledge Representation Issues
THE GREAT AI TOOL DEBATE

Saturday Meeting at Symposium

The AI SIG will, as usual, have an open business meeting on the *Saturday before* the Anaheim Symposium. Topics will relate to the running of the SIG. The purpose in having the meeting at the symposium is to include as many interested people as choose to come.

Complete plans are not set at this deadline. Please come to the AI/L&T/UNISIG Suite at the symposium for additional information. The meeting will begin at 1:00 P.M. Saturday, December 5, 1987.

Cheryl Jalbert
AI SIG Chair

From the Symposia Rep:

The Anaheim Symposium promises to be exciting to Aiers for five full days! Below is a listing of the technical sessions each day (in case the Preliminary Program is tough for some to navigate). In brief:

Monday is dedicated to introductory and general interest topics.

Tuesday is OPS5 exclusively.

Wednesday Languages and Tools with some applications.

Thursday continues to languages and tools crescendo to the Great AI Tool Debate.

Technical sessions on Friday elaborate on advanced AI topics and conclude with a candid interchange with Digital.

We're very pleased that several very well-known speakers will join us again in Anaheim. Hope to see you there!

Technical Sessions

Monday

- (BA) AI Made Easy - Managers Guide
- AI as an Application Development Tool
- AI AI on a PC
- Intro to Natural Language
- Selling AI to Management
- Expert Systms Crash Course
- Knowledge Representation
- First Expert System Panel
- Natural Language
- AI in the Public Domain

Tuesday

- OPS5 Introduction
- Designing Systems
- Real-Time Support
- Control Constructs
- Writing Efficient Rules
- External Routines
- Programming Techniques
- Q&A/Wish List

AI-3

Wednesday

- AI Language and Tools
- Prolog
- AI Novice Q&A
- Federated Expert Systems
- VAX LISP Update
- ** Common LISP Object System
- NEXPERT Object System
- Image Processing
- Common LISP Symbolic Processing

Thursday

- ** AI and Database
- Fielded Expert Systems Survey
- Productivity Tools
- AI Tool Debate Preamble
- Machine Learning
- AI Applied to Help Function
- Expert System for Scheduling
- ** Domain-dependent Shells
- PHIGS and AI for Mechanical Engineering
- Factory Modeling
- The Great AI Tool Debate

Friday

- VAX Systems for AI
- Expert Systems Project Management
- What's an Inference Engine?
- MOBIUS - A Meta-expert System
- Belief Systems
- Digital Asks the Aiers

Note: (**) - Indicates Featured Speaker

AI-4

HARD

NEWS

The Newsletter of the DECUS Hardware/Micro SIG

CONTENTS OF THIS ISSUE

- o FROM THE EDITOR
- o MICRONOTES UPDATE
- o DECWORLD: A HARDWARE BONANZA
- o TRADITIONAL DEC HARDWARE SUPPORT
- o VIDEO TAPE DRIVES

FROM THE EDITOR -- Carmen Wiseman

DEADLINE FOR THE NEXT ISSUE: NOVEMBER 20, 1987

If good intentions were all it took to get this show on the road, **HARD NEWS** would come out every month! Well, I apologize for the hiatus, but real life sometimes gets in the way of newsletter production. Also, as my illustrious predecessor has pointed out, one overworked editor can't do every single issue without help. We need and welcome contributions--after all, it's your newsletter too. There's a form in the back of each volume of the combined newsletters that tells you the many ways you can get hardware-related information to me or to the HMS SIG chair, Bill Walker.

Unfortunately, the lecture isn't over yet, folks. Some readers have gotten the highly erroneous impression that because I happen to work for Digital Review, I am some kind of conduit to that publication. I've received calls about placing advertising in DR, requests for information about DEXPO West, and DR subscription inquiries. I have nothing to do with any of these areas, although I was able to make appropriate referrals. Still, it seems that people are confused about my position vis à vis DR and **HARD NEWS**, so I think it's time to set the record straight.

I am the editor of **HARD NEWS**, the DECUS Hardware/Micro SIG newsletter. That role is entirely separate from anything I do at Digital Review. Thus, if you have a question about the HMS newsletter, I'll be happy to take your call or answer your letter--but I simply cannot field queries about DR. Why? Because it's a conflict of interest, and because I'm not the person to answer the kinds of questions I've been receiving. But don't despair if you need DR information: you can always call (617) 375-4300 and explain your question to the person who answers the phone. He or she will make sure you're directed to someone who can help you.

And now for the good stuff. As most of you know, the winter DECUS symposium is in Anaheim, Calif., next month (December 6-11, to be precise), and we're going to try to get a roadmap in the December issue of **HARD NEWS** to steer you through the HMS sessions. DEC has a slew of new hardware (see the DECworld report in this issue), and there are sure to be plenty of interesting sessions on things like the MicroVAX 3500 family and the long-awaited TK70. Even if you can't make it to the symposium, no worries--we'll be reporting on it here in **HARD NEWS**, so you won't miss a thing.

Finally, it looks like our SIG finally has a logo, and if we can find an accommodating artist, you may see it hopping around the Anaheim symposium. I won't say any more than that until we have an actual graphic rendition, but stay tuned.

See you in Anaheim!

Carmen D. Wiseman
Editor

MICRONOTES UPDATE

Since last issue (August 1987), we've received a great many inquiries about MicroNotes. We've done some investigating, and here's what we've found out.

MicroNotes comes under the aegis of Digital's Channels Marketing Support Group, and Art Bigler is in charge of MicroNotes preparation. The last time I talked to Art, he said the group was trying to get another edition of MicroNotes (probably under a different title) ready in time for the Anaheim symposium. He also assured me that if you received MicroNotes in the past, you're probably still on the subscription list. If you're not sure that you're on the old list, however, you can call the Channels Marketing Group at (617) 467-7707 to check. You can also write to the following address for information:

Channels Marketing Support Group
Digital Equipment Corp.
3 Results Way
MS MR03-3/G/20
Marlboro, MA 01752

Thanks to Chris Munson of Ford Aerospace, we've also discovered that there are on-line MicroNotes (only the old editions, unfortunately) in an OEM database called RONNIE. RONNIE, part of Digital's Information Network for Complementary Selling, is also maintained by Channels Marketing. Accessing RONNIE is very easy: you just call your local TYMNET number and type RONNIE. But there's a catch: you need an account and a password to get into the database.

Once you have an account and are actually in RONNIE, you'll be greeted by a Digital Electronic Store-type interface. In a way, RONNIE is a higher-gear version of the Store for OEMs. For instance, if you're a prime OEM, you can take a class to use EXCEL, RONNIE's AI-based system configurator.

If you're a DEC OEM or work for a DEC OEM company, you should contact RONNIE system manager Jim Thacker ([603] 884-5171) or Jane Steele ([603] 884-5164) for information about obtaining a RONNIE account. As far as we know, accounts are not available to non-OEM DEC users, because RONNIE also contains sales updates, software ship dates, and other potentially sensitive information. But if you think you might qualify for an account, give Jim or Jane a call.

cdw

DECWORLD: A HARDWARE BONANZA

If you didn't attend DECworld 1987, by now you've probably heard all about it from someone who did. In either case, you'll be happy to know that this is one DECworld report that doesn't dwell ad nauseum on the wonders of the QE2 or how there wasn't a hotel room to be found within literally 100 miles of Boston. We're simply going to give you some capsule descriptions of the hardware and micro items introduced at DECworld. And there were a slew of 'em!

MICROVAX 3500

- o CVAX CPU (KA650 board, based on CMOS CPU chipset)
- o 12-slot Q-Bus backplane
- o 100 nsec clock rate (claims to be 3 to 4 times faster than μ VAX II, with twice the memory capacity)
- o System configuration includes 27-inch pedestal enclosure, CVAX CPU board, 16MB of ECC memory (expandable to 32MB), RA70 280MB 5 1/4-inch fixed-media disk drive with controller, TK70 296MB cartridge tape drive, Ethernet controller, diagnostics and doc, one-year on-site warranty, VMS or Ultrix-32 license, end-node DECnet, and VMS services for MS-DOS
- o Shipping by the end of this year; volume quantities available by early 1988

MICROVAX 3600

- o same CPU as 3500
- o System configuration includes 41-inch BA213 cabinet, KA650 CVAX CPU board, 32MB of ECC memory, RA82 622MB 14-inch fixed-media disk drive with controller, Ethernet controller, TK70 296MB cartridge tape drive, diagnostics and doc, one-year on-site warranty, VMS or Ultrix-32 license, end-node DECnet, and VMS services for MS-DOS
- o Expanded cabinet system includes two RA82s, a 1600/6250-bpi tape drive and 40-user VMS license upgrade
- o Shipping by the end of this year; volume quantities available by early 1988

VAXSTATION 3200 and 3500 WORKSTATIONS

- o KA650 CPU with GPX co-processor and FPU
- o The 3200 comes in BA23 pedestal cabinet (diskless or disk'ed) and has 8MB of memory, RD54 disk drive, 19-inch VR260 monochrome or color monitor, TK50, Ethernet controller, keyboard and mouse, and VMS or Ultrix-32 license
- o The 3500 comes in BA213 expansion cabinet and has 16MB to 32MB of memory, two RA70s, TK70, Ethernet controller, QDSS "Dragon" graphics subsystem, four or eight planes of video memory, 19-inch VR260 monochrome or color monitor, keyboard and mouse, and VMS or Ultrix-32 license
- o Shipping begins in December 1987

VAXSERVER 3500, 3600 AND 3602

- o KA650 CPU with FPU
- o The 3500 worksystem server comes in BA213 pedestal box, and has 16MB of memory, RA70, TK70, Ethernet controller, VMS or Ultrix-32 license (VMS license includes operating system, full-function DECnet and Local Area VAXcluster software; Ultrix license includes NFS), and one-year on-site warranty

- o The 3600 is basically the same as the 3500, but has an RA82 and comes in a different BA213 cabinet
- o The 3602 has two KA650 CPU boards with FPUs, two cabinets, 32MB of memory, dual-ported RA82, TK70, Ethernet controller and the same software and warranty as the other VAXservers

RA70 DISK DRIVE

- o Formatted capacity of 280MB
- o 5 1/4-inch form factor
- o Not OEMed--first 5 1/4-inch drive to be manufactured by DEC
- o Uses thin-film media with density of 30.4 million bits per in²
- o DSA/SDI
- o Bundled with MicroVAX 3500 and 3600 systems (not available separately yet)

RA82 DISK DRIVE

- o Formatted capacity of 622MB
- o 14-inch fixed-media drive
- o 170-bit ECC
- o Bad-block replacement
- o DSA/SDI
- o Dual-ported
- o Bundled with MicroVAX 3600, or available in expander cabinet with TU81-Plus tape drive or as optional storage in "three-pack" configuration (initial limited availability, in other words)
- o Ordering availability 90 days ARO

TK70 TAPE DRIVE

- o Formatted capacity of 296MB
- o Formatted transfer rate of 90KB/sec
- o 5 1/4-inch form factor
- o Uses CompactTape II cartridges; can read but not write TK50 CompactTape cartridges
- o Supposedly has better reliability and data integrity than TK50
- o Bundled with MicroVAX 3500 and 3600 systems; optional on MicroVAX II and MicroPDP-11/83 (no free upgrade, though)

TRADITIONAL DEC HARDWARE SUPPORT

EDITOR'S NOTE

The following comments originally appeared in a VAX Notes conference called "Traditional VAXen Support" that took place during the spring 1987 DECUS symposium in Nashville. We believe that the sentiments expressed in the conference are echo the way many of you feel, and that this material is definitely worth a public airing. Because there have been problems with material taken from bulletin boards in the past, however, we've done everything possible to "generalize" the notes and protect the noters' anonymity. That's why we have edited out any information that might be used to identify the noter--name, company, location, etc.

If you agree or disagree with these thoughts, drop **HARD NEWS** a line. Ongoing support of "extinct" DEC hardware is a problem that concerns almost all of us (especially if we own PDP-11s or old VAXes), and we'd like to know what you have to say about it.

INTRODUCTION

Digital is no longer actively marketing the following VAX products: VAX-11/725, 11/730, 11/780 and 11/782. Plans for the final manufacturing build of the 11/750, 11/751 and 11/785 processors have also been announced. In response to customer concerns about support for these "traditional" products, we are evaluating

strategies to address the issues. In this conference, we look to our customers for input to aid in designing a support policy for traditional VAXen.

INPUT

One of the main problems encountered by our company (which owns many older VAX-11/78X and VAX-11/750 processors) is that it is very difficult to throw away or even sell equipment that is not fully depreciated. Much of this equipment is on old five-year depreciation plans, and very little at the new three-year rates that only recently became available. Things like 785 upgrades bought an average of a year to a year and a half ago won't be fully depreciated for some time, even under the three-year depreciation scheme.

While 750s may not be very useful any more, 785s can still get a reasonable bit of work done, and should continue to be useful for many years to come--especially if you have older UNIBUS peripherals to support on them.

Perhaps what is needed are some real trade-in incentives à la the MASSBUS-to-RA81 trade-in program from a few years back. This was the only real cost-advantage trade-in program I ever saw DEC offer; when you worked out the numbers, you had no problem convincing management that the new equipment paid for itself in three years because of reduced maintenance costs.

A BUSINESS VIEW OF SUPPORT

The question of support is always difficult, but here's something to consider.

When a customer buys a system, he has to tell the government how long he plans to keep it (i.e., figure out the depreciation value for tax purposes). For VAXes and similar equipment, the system is generally depreciated over three to seven years. When a company drops a product, it should consider fully supporting that product for a minimum of the longest depreciation period. And field units should be supported for twice the longest depreciation period.

If this were the case, an engineer could look a business type in the eye when a system was purchased and say, "Don't worry--it's covered."

RESEARCH GRANT BUDGETS ARE A LIMIT

We have an 11/780 that was purchased with research grant funds. Given current funding, we see no prospect for an upgrade in the foreseeable future. We have not upgraded to the 64KB or 256KB chip memory yet for the same reason, although we have maintained software currency. Therefore, we would like to continue using our 780 for some time to come.

We would hope that VMS development will continue to support this hardware, which has done very well for us. That would mean continued support for UNIBUS tapes and MASSBUS disks as well.

I have no suggestion as to how this continued support should be packaged, except that I hope it will not be much more expensive than current packages--again, because of research budget limitations.

A significant increase in the cost of VMS support for old hardware in an attempt to force users to upgrade their systems could be counterproductive. This might end up in a loss of all support. I'm fairly certain this would be the case for our system.

For hardware support, I expect gradual increases will come, simply for lack of spares, but this should be several years in the future because of the number of 780s still in use.

DON'T FEED THE BEAR

Adding fuel to the fire...

I agree that the major concern with the use of older equipment revolves around depreciation schedules. This is a problem at our company as well. But I'd like to call attention to another issue: support for the current installation base (depreciation notwithstanding). At a small company, it's not practical to discard usable and functioning equipment just because it has no proven value on the books. Given the cost of new equipment and the proven quality (i.e., maintenance record) of the 700 series, we see no reason to move to a new machine. If it works, don't fix it!

Granted, this is not in DEC's best interests--sales are the life blood of the company. However, DEC has also prided itself on good customer relations and customer support. Though we are not currently looking for new equipment, the time may come when we need to. At that point, a migration to the new DEC machines will occur, provided we feel that DEC is still the best game in town. The taste left in the mouths of management will depend quite heavily on the recent involvement of the vendor, and there will be little the lowly user can do to change that.

This is even more of an issue considering the rapid rate of turnover within the industry. The decision to upgrade the 8200/8300s to 8250/8350s painfully demonstrated this problem. It is hard to convince someone to buy a product that will be obsolete in less than a year. It is very tempting for management to get stuck in a "wait and see" mode that could last forever. This has scared management and made them wary. They are starting to actually question our view that DEC is the only choice.

When asked why we should continue to choose DEC and choose it now, we as users need ammunition. The VAX series is that ammunition--full compatibility up and down the line.

If DEC discontinues support for a sector of that community, there will be two major consequences: it will cost us a sorely needed bargaining chip, and it will scare management even more. I don't pretend to understand how the powers that be make decisions, but I don't know that change really upsets them. You have to ease them into it. Two whammies too close together may cause them to head for the hills.

This is not to say that DEC must continue to support the line forever. I propose that the phase-out of support coincide with the removal of those machines from the marketplace. DEC must know how big its installed base is. It would not be difficult to survey the expected lifespan of those sites. The reduction in level of support should correspond to the reduction in the number of sites. However, no pressure should be applied to sites to retire equipment or migrate to new equipment. And DEC must assure that sites do not feel they have been abandoned.

DON'T FORGET MASSBUS SUPPORT

Don't forget about the MASSBUS devices! I know that DEC doesn't like them any more and wants us to get rid of them, but the TU7X tape drives are the only ones that can be more than 24 feet away from a CPU that is not in a cluster. Our computer room is set up with an operations area and a systems area, which are separated by a wall. Until we replaced the UNIBUS TU80 with a MASSBUS TU78, we could not keep the VAX in the computer room with the IBM machines.

WE STILL NEED OUR 750

I work for a federal government agency and thus cannot replace equipment that fast. Would you believe we are still using a PDP-11/20? Our 750 is only two and a half years old, and will have to last us for several more years. I certainly hope that maintenance and full software support would be available for this machine for years to come and at a fair price.

There are questions that we need to be answered before we replace the 750 with, say, a BI machine, such as what to do with our TU78 (see above!). Also, since we are a research laboratory doing more and more data acquisition work all the time, we are concerned about moving to a closed-architecture machine. The currently available BI hardware is much too limited, and I don't see it improving sufficiently until the BI bus is opened to one and all.

BUY IT BACK AND SELL IT TO THE CRUSHER!

Most people solve problems by buying computers. Many times, buying a computer causes the largest problem: hardware becomes outdated. Clearly those of us who understand where the technology has been and where it might be going should also understand that a five-year system lifespan may be unreasonable. (That's not to say that I don't understand the need for budgets and five-year plans.)

Maybe DEC needs to see that when major changes occur in the architecture of its "compatible" systems, it should offer some way that you can refuse and not lose in getting new technology. Face it--eventually the world will see the light and own a DEC system, and then the only market DEC will have will be its installed base. Maybe a "good" program to replace or supplant old 7XX with new 8XXX systems will evolve when that happens.

We need a way to replace 730s and 750s with 8250s. Licenses are only a very small part of the issue. How about a buy-back of the nonusable items (e.g., non-MSCP RA disks, CPU boxes with associated boards, etc.)?

SUPPORT IS PART OF THE PRODUCT

The lifespan is part of the product, too.

One of the big problems in making purchasing decisions is that many of the most important factors are intangible and fall under the heading of "policy." What Digital does about upgrades and support can determine whether the useful life of a machine is five or 15 years. By any reasonable accounting, this is a much more important factor than the cycle time or the number of bits.

When we choose to buy from Digital, we guess what it will do in the future based on what it has done in the past. We perceive DEC as a company that traditionally does the things that ensures a long, useful life for its machines--not just durable construction, but things like field service having spares for old machines, old machines being able to run new software and participate in new networks, and so on.

We now tend to buy DEC terminals, memory and disks rather than third-party add-ons. Our rule of thumb is that we won't go to a third party just to get 30 percent more performance or a 30 percent lower price. We go to third parties only for functionality we can't get from Digital. About 15 years with DEC and third parties have convinced us that the third-party stuff is often excellent, but consistently has a shorter lifespan of useful support than DEC stuff.

DEC only wins this kind of credibility by doing good things consistently, over and over again, for a long time. But such credibility is also fragile. If we get the idea that support for old equipment is just something that the "old Digital" used to do, the whole thing is blown.

We damn well expect Digital to support our 780 and our two 750s. We see the 780 as a fine, reliable workhorse that has old technology but a modern architecture. We expect to use it for at least another five years. And we expect DEC to help us keep this machine going and not pressure us into treating it like disposable tissue.

VIDEO TAPE DRIVES

Here's an idea whose time has definitely come: system backup on 8mm videocassettes!


As you might know, 8mm is the latest video technology. It's the video analog to the old Super8 movie film (which, like the PDP-11, is still flourishing, thank you), and its advantage is that it's small but holds a lot of information--much more, for instance, than a TK50 or even a TK70.

At least one company is now making 5 1/4-inch helical-scan 8mm tape subsystems for VAXes and MicroVAXes, both Q-Bus and UNIBUS. In helical-scan recording (the same technique used in regular video recording), rotating read/write heads write very narrow tracks in a diagonal pattern on the tape.

Error-correcting 8mm subsystems on the market claim to have a formatted capacity as high as 2.3GB and to run at an effective head-to-tape speed of 150 ips, even though actual tape movement is less than one-half inch per second.

We think it makes eminent good sense to use video technology (especially the ubiquitous half-inch VHS format) in the quest for a better backup medium, but it's all so new we don't know anyone who's used a video drive or subsystem. How about you? Think this is a good or a putrid idea? Should DEC abandon the TK50/70 and jump on the video bandwagon?

cdw



THE DEVIAS LETTER

Contribution guidelines.....	i
From the Editor's Terminal.....	1
Ten Years Ago This Month.....	2
Fall Symposia Preview.....	3
IAS Version 3.2C SPR's.....	4
The Program of the Month Club.....	5
A Round Tuit.....	8

CONTRIBUTION GUIDELINES

Contributions for the newsletter should be sent to:

Frank R. Borger
Michael Reese Medical Center
Department of Radiation Therapy
Lake Shore Drive at 31st St
Chicago, IL 60616

Contributions of letters, articles, important SPR's etc will be accepted in any form, (including notes jotted in pencil on gravy-stained tablecloths.) Contributions will be much more graciously accepted in one of the following formats:


1. Non machine readable sources, (SPR's etc,) should be reasonably dark to insure good photocopying. Text whatever should be the equivalent of 66 lines at 6 lpi, with 4-line top margin, 5-line bottom margin, left-margin 10, right margin 74 at 10cpi. If using a DEC LN03 for output, use left-margin 8. right margin 72.
2. Machine readable sources may be submitted on 9-track Mag-tape, (800,1600, or 6250 BPI,) DEC-tape II, DecMate floppies, or whatever. We're not fussy, we'll even accept paper tape or cards. Preferred format is DOS or BRU for tapes, Files-11 for DEC-tape II.
3. 1200 baud dial-up modems are available on our IAS system and our VAX, with KERMIT servers available. Give the editor a call at (312)-791-2515 (preferably later in the day,) to obtain access information, etc.
4. If long distance dialout is not possible on your system, we'll be willing to call your system and do the work, (unless you want to transfer the entire manual set at 300 baud.)

Any media sent to us will be promptly returned.

ASK THE DEVIAS WIZARD

If you have a problem you would like to submit to the Devias wizzard, write a letter or fill out a copy of a standard SPR and send it to the Editor at the above address. Answers to problems from members (or anyone) should also be sent to the Editor.

From the Editor's Terminal



First the good news. The IAS V3.2 Update D field test has arrived at our institution. Now the bad news. Until we get our spare RM02 drive fixed I can't create a copy of our system disk to play with at my heart's content, so I will not be testing the update. (Anyone foolish enough to try a field test update on their running system deserves having Murphy do his best to them.) The drive works beautifully on the tester, but flunks when connected to the system. Hans has been given top priority to getting the drive back on line, so by next week (when I return from a DECUS meeting out west,) we should have our backup/development system back to full strength, and we will play with the field test. We should be able to have a full report ready for the next newsletter.

Also in this issue, another SPR from Ted Smith, and a preview of the Fall Symposia from Lynda Roenicke.

This months program of the month is an example of our work to add Digital Command Language syntax to our MCR oriented system. Its a program which supports the DCL "COPY" command. (A slight variation also emulates the "RENAME" command. Now that we have users trained on VAXes its more important for us to support a similar syntax across our VMS, RT11 and IAS systems. IBM eat your heart out.

And finally, we are including something that everybody needs. Its the ultimate solution to all those projects that you haven't finished because you said, "I'll finish that when I get around to it."

Ten Years Ago This Month

The November 1977 Multi-Tasker contained:

Phil Cannon reported on the Fall Symposium Tape Copy facility procedures. In essence, if you submitted something you brought two tapes, one for your submission, and a blank tape to receive the resultant symposium tape. Talk about programs hot off the tape drive!

An SPR reported that two undocumented error messages: "COMMAND INPUT SOURCE ERROR" and "TTnn OPEN FAILURE ON COMMAND INPUT OR OUTPUT FILE" occurred due to a user logging out on TTnn after having changed his terminal speed, (and not changed it back.) The SPR also reported that the terminal in question ended up hung. DEC software support reported "unable to reproduce the problem."

Jim Downward reported that an August Software Dispatch patch to the Mag tape ACP, (MTAACP,) seemed to be crashing his system. The symptom he saw was that after each DISMOUNT, vector location 0 increased.

Jim also included a correction to a typo in the July 77 Multi-tasker,

Incorrect Code	Correct Code
MOV HLPNAM,INDEX	MOV #HLPNAM,INDEX
MOV BUF,R0	MOV #BUF,R0

(If your editor had a dollar for every time he lost a "#" sign by running Macro source through RUNOFF, he could retire today.)

And finally, a Fortran IV SPR reported that the following test program:

```
PROGRAM TEST
IMPLICIT INTEGER (A-Z)
STOP
END
```

Produced the following error message:

```
NON-STANDARD STATEMENT ORDERING
```

Talk about the matchbook school of computer compiler writing...

Fall Symposia Preview

Date: 20-Aug-1987 05:38pm EST
From: Lynda Roenicke
Dept: Symp. Committee

Subject: Anaheim DEC 1987

FYI a peek of IAS activities at next symposia

MONDAY	09:00am - IAS Opening Session 10:00am - IAS Product Panel 8:00pm - Auto Dial Out 8:30pm - IAS Device Drivers
TUESDAY	9:00am - User Written CLI
WEDNESDAY	1:00pm - RSX-11D Working Group 2:00pm - IAS+ Working Group 4:00pm - AN/GYQ-21 Handlers/Lims 5:00pm - AN/GYQ-21 Users Forum
THURSDAY	10:00am - IAS Sig Plan 10:30am - IAS Sig Lib
FRIDAY	09:00am - IAS User Forum 10:00am - Parsing without TPARS 12:00am - IAS Closing

The schedule looks pretty good! See you in Anaheim. We will also have a campground. Some sessions will be in the campground.

The Program of the Month Club

From: Ted Smith
Subject: Another SPR...

I have discovered another "bug" but cannot determine if the problem is in the Fortran-77 V5.0 compiler or Linker (TKB).

Create the Fortran program A.FTN:

```
INTEGER*2 TABLE
I = 1
I = TABLE(I)
END
```

A.FTN compiles without error. Linker produces one *Diag* undefined symbols segment A. The undefined symbol is TABLE. This appears to be correct since the Fortran compiler is assuming that TABLE is a INTEGER*2 function subroutine.

Create the Fortran program B.FTN:

```
INTEGER*2 ITABLE(10)
CALL C(ITABLE, I)
END

SUBROUTINE C(ITABLE, I)
I = 1
I = ITABLE(I)
RETURN
END
```

B.FTN both compiles and links without error. However, when B.TSK is executed, the following fatal error is produced:

```
JOB10 -- Exiting due to ERROR 3
Odd Address Trap (SST0)
at PC=001376
IN "C" AT OR AFTER 00002
FROM ".MAIN." AT OR AFTER 00002
```

Questions (problems):

1. In B.FTN; ITABLE is passed to subroutine C. In subroutine C, the Fortran compiler should know that ITABLE is a Integer variable since ITABLE is passed as an argument and is not specified as an array. Therefore, why did the Fortran compiler NOT generate a Warning or ERROR message for the statement
I = ITABLE(I)
in subroutine C?
2. In B.FTN; If the Fortran compiler assumed ITABLE to be a function in subroutine C then why did Linker not print a *DIAG* undefined symbol in segment as it did when linking A.FTN?

I have submitted an SPR and informed the Telephone Support center

At Michael Reese, we support full DCL syntax on our IAS system, but retain MCR as the standard CLI, (mainly to support a specialized login system we use, and to reduce the overhead of PDS.) Rather than using the PDS or VMS system of switching between DCL and MCR modes, (or of preceding MCR mode commands with "MCR ",) we wanted our system to be transparent. We felt that the CLI should be intelligent enough to know that QUE is an MCR command, while PRINT is a DCL command. To do this, we

1. Modified MCR to pass the command line to a catchall task if the requested task, (...DIR, ...PRI, etc.) is not installed.
2. Wrote a catchall task that accepts DCL syntax commands and parses them to their MCR equivalents and then evokes PIP, QUE or whatever to process the request, just as PDS does.

Although this catchall task works fine for one argument commands, (Directory, Print, Type, etc.) it was not designed to handle two argument commands such as COPY or RENAME, (which also prompt for the two file specifications if they are left out of the command line.) For cases such as COPY, we ended up writing simple programs to emulate the DCL commands. Note that our catchall task also is intelligent enough to do temporary installs of CUSPs such as our COPY task, so that the system task directory does not become inordinately large and node consuming.

The following program is an example. It emulates the DCL COPY command under MCR. The task build command file is:

```
cop,copy/-sp=copy
lb:[1,1]exec.stb/ss
/
task=...COP
//
```

Finally note that simple variations on the above are possible. To create a RENAME command emulator, make the following changes to the COPY program.

1. Change "copy" to "rename" in .TITLE, .END and "copy:"
2. Insert the following 3 lines in TWO locations before the lines that read:
"movb #'=(r2)+" ;put in the equal sign

movb #'/(r2)+ ;put in "/re
movb #'R,(r2)+
movb #'E,(r2)+


```

.TITLE copy
.IDENT      /MRH001/
;
;
;   program to emulate the VMS COPY command
;
;   FRANK BORGER
;   MICHAEL REESE MEDICAL CENTER
;   August, 1987
;
;   .MCALL      dir$,gmcr$,qiow$,exit$$,spwn$
;
;GET COMMAND LINE
;
copy:  dir$      #getmcr          ;get mcr command line
      mov       @$dsw,r1        ;get length of line
      mov       #getmcr+g.mcrb,r0;point to buffer
      mov       r0,r4          ;create end of line pointer
      add       r1,r4
      add       #3,r0          ;point past "COP"
1$:    cmpb     (r0),#40        ;found a space ?
      beq       2$            ;br if yes
      inc       r0
      cmp       r0,r4          ;end of the command line ?
      bge      from          ;yes, prompt for names
      br       1$            ;check again
2$:    cmpb     (r0),#40        ;now skip spaces
      bne      3$
      inc       r0
      cmp       r0,r4          ;end of the line
      bge      from
3$:    mov       #temp,r2      ;point to temp buffer
      call     coptil         ;copy to comma or space
      cmp       r0,r4          ;anything left ?
      bge      to            ;if not, prompt for output
      mov       #tobuf,r2     ;point to final buffer
      call     coptil
      movb     #'=(r2)+       ;put in the equal sign
      mov       #temp,r0      ;point to the from name
      call     coptil
      sub      #spacmd,r2     ;figure length
      mov       r2,spapip+s.pwcl
      call     uppcas
      dir$     #spapip        ;evoke pip
      exit$$    ;and exit
;
;do from and/or to prompts and reads
;
from:  dir$      #fropro        ;read the from stuff
to:    dir$      #topro         ;read the to stuff
      mov       #tobuf,r2     ;point to to buffer
      add       rlen,r2       ;add length of line
      movb     #'=(r2)+       ;put in the equal sign
      mov       #temp,r0      ;point to from name

```

```

      call     coptil
      sub      #spacmd,r2     ;figure length
      mov       r2,spapip+s.pwcl
      call     uppcas
      dir$     #spapip        ;evoke pip
      exit$$
;
;coptil routine, to copy from (r0) to (r2) until get terminator
;
coptil: movb     (r0)+,(r2)+   ;copy a character
      movb     (r0),r5        ;get next character
      beq     done           ;terminate on null
      cmpb     r5,#40        ;a space
      beq     done
      cmpb     r5,#15        ;a carret ?
      bne     coptil
done:  inc      r0            ;bump r0 past terminator
      rts     pc
;
;uppcas, convert to uppercase for pip
;
uppcas: mov      #spacmd,r0    ;point to line
1$:    cmpb     (r0),#100     ;lower case ?
      ble     2$            ;no
      bic     #40,r0
2$:    inc      r0
      sob     r2,1$         ;do whole line
      rts     pc
;
;get mcr stuff
;
getmcr: gmcr$
;
;from and too prompts
;
fropro: qiow$   io.rpr,5,5,,rstat,,<temp,80.,,fromas,6>
fromas: .byte   12
      .ascii  /From:/
topro:  qiow$   io.rpr,5,5,,rstat,,<tobuf,80.,,toas,6>
toas:   .byte   12
      .ascii  /To: /
rstat:  .word   0            ;read status
rlen:   .word   0            ;read size
;
temp:   .blkb   40.         ;temp buffer
;
;spawn stuff
;
spapip: spwn$   ...PIP,,,,4,,,spacmd,3
spacmd: .ascii  /PIP /
tobuf:  .blkb   80.         ;room for rest of command line
spasta: .word   0            ;spawn status
      .end   copy

```


Leverage




TABLE OF CONTENTS

Editor's Notes	2
L & T SIG — Cobol Working Group <i>Bill Leroy, The Software House, Inc</i>	3
Languages and Tools Woods Meeting <i>Sam Whidden, American Mathematical Society</i>	5
Languages And Tools Working Group Statements <i>Joe Pollizzi, Space Telescope Science Institute</i>	14

EDITOR'S NOTES

This issue of *Leverage* contains only a short selection of articles, largely due to your poor, overworked, editor having to spend a week in Finland, just prior to the publication deadlines. Nonetheless, I believe we have an interesting collection of information here.

First is the monthly column by Bill Leroy regarding the COBOL working group. My thanks to Bill for his usual timely submission.

Next comes the second half of the Languages and Tools Woods Meeting report. This was cut by the DECUS staff at the last minute from the prior issue, due to space considerations. Also cut from that issue due to space restrictions, by the way, were the L & T symposium abstracts for Anaheim. I know many of you have requested them; I'm sorry they weren't able to be printed.

Finally, Joe Pollizzi has submitted a summary of the Working Group statements. This is to give you information regarding the goals and activities of the new Working Groups. The SIG leadership has high hopes that these groups will provide the technical leadership in the SIG. If you look these over, you may find folks interested in the same concerns you are.

That's all for this month. See you in Anaheim!

Al Folsom
Leverage Editor

To all users of the COBOL language:

This month is kind of light, since I only received one letter. However, we at TSH are documenting the use of COBOL shareable libraries, and will include this in the next issue. It sure is handy to re-compile a sub-routine, and not have to also re-compile / re-link the fifty COBOL programs that use it. We are also producing form letters from WPS-Plus in "batch" mode using variable data from COBOL programs, but this is more of a submission to the OA SIG.

Included this month is a response to last month's article by George Mayer on "Suppression of VFC Record Formats". The response is in the form of a letter and a sub-routine from James Bair of Grinnell Infosystems, Inc. We have generalized and added some comments to his command file. What works is his -- what doesn't is ours.

Membership in the DECUS L&T COBOL Working Group is open to **ALL users of "anybodys" COBOL on any DEC computer**, including Rainbows, VAX-Mates, PDP-11s, VAXs and DEC-10s and 20s, and clones. If you would like to be on our mailing list, please send me your name, address, & telephone number.

And finally, please send any articles, questions, hints, kinks, wish list items, bugs, "puff" sheets on yourself to be showcased in a future issue, new members, or whatever -- on pieces of paper, or if a long article, in VAX/VMS format, on 1600 bpi magnetic tape, or TK-50 magnetic tape -- to me at:

Bill Leroy (404) 231-1484
The Software House, Inc.
P. O. Box 52661
Atlanta, GA 30355-0661

(OR)

Bill Leroy (404) 231-1484
The Software House, Inc.
2964 Peachtree Road, N.W. #300
Atlanta, GA 30305-2120

Suppression of VFC Record Formats

August 31, 1987

Dear Bill:

The COBOL newsletter seems to be a good way to address common COBOL problems. We have been using VAX COBOL since early 1983 as our primary language to develop insurance applications. We have found that in most cases VAX COBOL is flexible enough to meet our needs. We do use System Service and run-time library calls from our programs to handle many situations that are impossible (or very difficult) in COBOL.

Alan Cline, from our technical Services Systems Programming area, has a solution to point 3 "Suppression of VFC Record Formats" from a list of questions submitted by George Mayer in the last issue of the newsletter. I hope the attached solution will be helpful.

s/James F. Bair

Manager, Technical Support
Grinnell Info Systems, Inc. (515) 236-6121
P. O. Box 790
Grinnell IA 50112-0795

```

$! VFC_EDIT.COM
$!
$! The following is an example of how to convert a VFC formatted file to a sequential file for editing
$! and then converting it back to a VFC file.
$!
$! To use this procedure type @VFC_EDIT FILE.NAME
$!
$! and the output will be your file name with "_VFC" appended to the end.
$!
$! example: $@VFC_EDIT SALES.LIS
$!
$! produces an output file called: SALES.LIS_VFC
$!
$! First you need to create an empty sequential file using the DCL CREATE command.
$!
$! CREATE VFC_TEMP.SEQ
$!
$! Next use the following CONVERT command to transfer all the records in the VFC file to the new empty
$! file, including the two byte fixed length control header, using the attributes of the new empty file.
$!
$! CONVERT/NOCREATE 'P1/FIXED_CONTROL VFC_TEMP.SEQ
$!
$! Now go into an editor and extract needed portions or change the contents of the file VFC_TEMP.SEQ.
$! When you are editing the file, make sure NOT to change the first two characters of each line, which
$! contains the two byte VFC control header.
$!
$! ASSIGN/USER SYS$COMMAND SYS$INPUT
$!
$! EDIT VFC_TEMP.SEQ
$!
$! After you exit your favorite editor, you can convert the sequential file back to a VFC file
$! by using the following commands:
$!
$! 1) Create an FDL file from your original VFC file.
$!
$! ANALYZE/RMS/FDL/OUTPUT=VFC_TEMP.FDL 'P1
$!
$! 2) Use CONVERT again to transfer the records from the edited sequential file to the new VFC file.
$!
$! CONVERT/CREATE/FDL=VFC_TEMP.FDL VFC_TEMP.SEQ 'P1'_VFC/FIXED_CONTROL
$! SET FILE/TRUNCATE 'P1'_VFC
$!
$! 3) Prove to yourself what you just did, and get rid of your temporary files.
$!
$! DIR/FULL 'P1'_VFC
$!
$! DELETE/NOCONFIRM VFC_TEMP.FDL;*,VFC_TEMP.SEQ;*
$! EXIT
    
```

LANGUAGES & TOOLS WOODS MEETING

5. The relationship between PSSes and the technical content of symposium sessions was discussed. It was suggested that the seminar speakers be encouraged to present sessions based upon their seminars; the general feeling was that presenting PSS contents as sessions would not detract from PSS attendance and would result in some good sessions.
6. Some discussion was had about a PSS being a deliverable for a working group. The consensus was that a Seminar would be a good deliverable for some WGs, but that it should not be required of any.

During this discussion, Sam explained that the organization of DECUS and that of the SIG are very similar; that is, as there are technical groups (working groups) and service groups (scheduling, newsletter, session notes, etc.) within the SIG, so there are technical groups (SIGS) and support groups (CommComm, Symposium Committee, Seminars Unit, etc.) that serve somewhat analogous purposes in the larger organization. At both levels, there is a degree of crossover between groups that sometimes causes things to work in a less than perfect manner. Members of our SIG who serve on DECUS committees should help their committees to work smoothly.

LIBRARY COMMITTEE --- Discussion Leader, Tony Scandora, Library Committee Rep; Recorder Bob van Keuren.

We are developing a new tape from the Nashville Symposium. Most submissions for it arrived after Nashville and many came from Glenn Everhart.

We got a lot of material from UNISIG - many generic tools. One is a new version of GNU EMACS. We are working on getting materials from the Free Software Foundation. Right now a legal question concerning software distribution is being worked through by Richard Stallman (developer of EMACS and one of the founders of the Free Software Foundation). We are waiting because we might have to change something on the VMS tape. The Free Software Foundation is trying to put together a free version of UNIX, and also public domain C code for software tools. We are trying to get these things on the SIG tape. Also a C compiler and GNU Chess.

On the SIG tape is micro EMACS. EMACS was first written on the 10's in TECO. Micro Emacs was written later without the LISP extensions. There is also a micro GNU EMACS from Texas. Mark of the Unicorn has a more word-processing-like product. On the tape also is a C compiler for the VAX and a lot of TPU as well. One person mentioned a patch to EDT that jacks the priority up to five while EDT is updating the screen, and then down to four afterwards. We are interested in having the Large Computer Group public domain library (largely TOPS-to-VMS conversion aids). We found here at the meeting one volunteer who has it.

For convenience, the new tape will be a complete replacement for the previous L&T tape (San Francisco). As for timing, Tony will start distributing the Nashville tape in August; he will talk to NLC Tape Distribution about it. TeX, and a lot of other good stuff, will be on the new tape. Our tape will eventually be available in the DECUS Library.

L&T-5

Tony presented a report on the recent meeting of the Library Committee, which he felt would benefit from improved organization.

Tony reported that the Public Domain Software Working Group, of which he is Chair, will act as a Library Committee for the SIG and assume responsibility for the the SIG tape. It was suggested that we need a well-indexed tape for the SIG and the tape library. One person has volunteered to edit and organize the TPU extensions. We need a LaTeX styles editor.

The question came up of non-VMS tapes. Only one PDP-11 tape submission has come in--an updated version of DECUS C. There is no 36-bit stuff, either. A lot of material is VMS-specific. We would like to set up other tapes if we can get submissions.

PDP-11 ISSUES WORKING GROUP --- Discussion Leader, Sam Whidden; Recorder, Steve Jackson.

Since Bill Tabor, the PDP-11 Layered Products Working Group Chairman was unable to attend the L & T woods meeting, Sam Whidden, who also was present at the PDP-11 Issues Working Group meeting, reported on that meeting. The following is an expansion of Sam's report:

Jeff Killeen, Chair of the SIG Council's Product Directions Committee, was the convenor of the PDP-11 Issues Working Group, and Ed Cetron, RSX SIG, its chair. The Group contained representatives from all interested SIGS involved with PDP-11 hardware or software. The SIG Chairs attended (for this meeting only), along with the designated SIG reps.

At the opening of the meeting's first day, Sunday, Killeen guided the Group toward a productive operational philosophy. He pointed out that any group involved in product planning cannot simply complain; it must have positive suggestions for achievable goals. It's essential, as well, to communicate concerns to people able to address them; it does little good to berate a language developer for flaws in Digital's long-term parallel-processor strategy.

If the Group suggests a solution to a problem, it must avoid the temptation to tell Digital how to implement that solution; execution has to be left up to DEC. Better, frame concerns as well-defined problems that DEC might be able to solve, rather than as specific solutions.

DECUS members must remember that the general user represents only one facet of DEC's marketplace, albeit an important one. Digital has some very large accounts at the corporate level whose influence over Digital's market strategy has to be great. In fact, for the most part, DECUS members' influence is felt most by DEC at the tactical, rather than the larger strategic, level.

With these guidelines in mind, the group developed a series of specific points which were presented to officials of the MicroSystems Development group (MSD) on Monday. In general, the Working Group stressed the fact that PDP-11's remain a viable---even unique---solution to an important range of computing problems which are often not well-handled by VAXs (e.g., real-time, low-cost, or OEM

L&T-6

turnkey applications). Frequently, a user's only good alternative may be another manufacturer's equipment when PDP-11s are not offered as a solution by DEC.

MSD officials were responsive to the Group's concerns and firm in their assertion that PDP-11s would be developed and marketed for the foreseeable future, citing, among other things, statements by Ken Olsen to that effect. The Working Group noted that users' perception of DEC's withdrawal of PDP-11 support was nearly as destabilizing as the reality would be. It agreed with MSD that the Digital sales force was knowledgeable on the subject of PDP-11s, but pointed out that members of the sales force had greater motivation to offer customers VAX solutions than lower-priced Eleven solutions. There was concern in the Group, too, that Field Service representatives tended to be better trained to maintain VAXes than PDP-11s. The Group was clear its feeling that Elevens need not be available forever, as long as any successor solution offers the same functionality, simplicity, cost-effectiveness, and reliability.

The meeting concluded with an agreement to continue discussions. Through the RSX SIG, the Group will sponsor two sessions at Anaheim (probably scheduled back-to-back) in which users will share in creating input on PDP-11 issues and in interacting with Digital on its responses.

LANGUAGES & TOOLS WORKING GROUPS --- Discussion Leader, Joe Pollizzi, Working Groups Coordinator; Recorder, Gerald Lester.

The following working group chairs gave a statement of purpose of their working groups (see statement of purposes from Joe Pollizzi): PL/I, SCAN, C, BASIC, Project Management, Technical Prod. of Doc., Low Level Languages, Configuration Management, Software Metrics, Large Software Systems, COBOL, TECO, APL, DIBOL, Pascal, Modula II, and Public Domain Software. The remaining working groups were not represented, so no report was given for them.

In addition to the statement mentioned above, the SCAN working group reported that SPR response time was too long for its area. It was noted that the C working group needs involvement with the PDP-11 and Unix SIG's C working groups. Joe noted that the working group chairs must communicate with their members not only at Symposia but between Symposia. Earl Cory, L&T Symposium Coordinator, stated that he gets a copy of all the audio tapes for L&T and will send to the working groups the tape(s) that pertain to their areas. It was decided that the COBOL Generator will be the responsibility of the COBOL Working Group. It was also noted that a handout should be in the L&T packet which identifies the responsibility of each Working Group.

MASTERS PROGRAM --- Discussion Leader, Dena Shelton, Masters Coordinator; Recorder, Mark Hyde.

The segment opened with a discussion of methods for classifying Masters' subject areas---in particular, how should someone who classifies himself as a Master in 'BASIC' and in 'VAX BASIC' be designated? Should these be two separate categories? It was decided to list only 'BASIC' as a category, and to

mention any operating systems noted by the Master as comments in his listing.

The question arose of including listings for subjects outside the areas covered by L&T, such as DCL, for example. Sam voiced concern that L&T not overstep its area, although he and Dena noted that some subjects were listed because some Masters had felt that their knowledge of a field within L&T's sphere was inextricably bound to that of another outside it. They also pointed out that one subject---MUMPS---had been included at the request of the responsible SIG. They noted that the VAX SIG had plans to generate its own Masters Directory, and that perhaps eventually this service could graduate to become a DECUS, rather than a SIG, product.

Sam stated that he was having trouble formatting the Masters Directory in its present form, with all of each Master's areas of expertise listed under each category in which the Master appears. After some discussion, his suggestion was agreed to that the full listing of a Master's areas would appear just in the name & address portion of the directory, along with an expanded set of self-descriptive comments by the Master (comments intended to help the user select an appropriate Master). Sam plans to revise the Directory format as soon as he has time.

Any changes in the listings should be sent back to Dena ASAP. Dena needs all the help she can get with recruiting Masters. WG chairs could help by recruiting within their groups. The consensus was that most, but perhaps not all, WG Chairs should themselves be Masters. Dena has received at least one "thank-you" letter from a satisfied user, which she will publish. Any that WG chairs receive should probably be published, too.

Dena also needs feedback and suggestions about the program. Can we set criteria for what constitutes a Master? Must Masters be DECUS members? (The floor consensus was no, that they just need to be comfortable with their subject.) In answer to a question on monitoring the service, Dena responded that there would be no general review; she would intervene in cases of repeated problems. To evaluate the program, Dena needs feedback on services rendered by individual Masters; she will note any complaints. So far, there has been no negative feedback. If any negative comments do occur, Dena will discuss the situation with the Master involved. A suggestion that the feedback mechanism be formalized was adopted, and a feedback area will be added to the Newsletter Masters form.

The segment closed with a floor suggestion to get feedback on the frequency with which Masters are contacted.

SPEAKERS PROGRAM--Open Position; Discussion Leader, Sam Whidden; Recorder, Mark Hyde.

The SIG has in the past helped find LUG speakers, and should be continue to send speakers to LUG meetings. Terry Medlin suggests that we could also be sending speakers to trade shows, to advertise both and DECUS.

There is definitely positive interest in this program on the part of DECUS Leadership. They feel that DECUS has a need to do a better job of marketing itself. Further, there do not appear to be any by-law restrictions on it.

There was a discussion of which shows to send speakers to. Emphasis was on DEC-related shows like DEXPO, but also included COMDEX.

Next came a discussion of subjects for speakers to address. The question was raised of whether "DECUS promotion" (or at least description) was a valid topic, or should speakers stick strictly to technical subjects. Terry should collect more ideas for Speaker opportunities (e. g. conferences in addition to trade shows). The discussion ended with the conclusion that the question still bears further thought.

In the case of LUGs, we should co-ordinate with the LUG regions. Where possible, Steering Committee members should make themselves available to LUGs as speakers.

The segment ended with the conclusion that the idea of the Speakers Program needs someone who is interested in developing it further.

WISHLIST PROGRAM --- Discussion Leader, Shava Nerad, Wishlist Coordinator; Recorder, Bruce Mebust.

o In Anaheim, the Wishlist will be constructed from wishlists generated at Working Group meetings and in the Campground an Clinic, plus wishes expressed at the general L&T Wishlist session. Wishlist forms will be available in the Anaheim Folder, at the campground at Anaheim and in Session Chairs' material. Often in sessions, wishes are expressed informally and never get on the Wishlist. Session Chairs could be on the lookout for such wishes.

o The Wishlist will be organized by working group and voted on separately by category to avoid burying high-priority wishes of smaller groups under the long lists of big ones. All wishes will go on the main SIG Wishlist in addition to individual WG lists for DEC developers.

o The current Wishlist is being reformatted and will go to the Newsletter in time to appear in the September issue as a ballot, with vote-responses to Shava and prioritized list to Celeste by November. The vote summary and developers responses will be presented at the session in Anaheim and in the February 88 newsletter.

o Can there be an on-line "Wishes" conference on VAXnotes at Anaheim? How would we organize/publicize this?

o Normally, at each Symposium DEC will give its responses to the current Wishlist items. The timing of the 6-month Wishlist cycle must therefore allow for response time from DEC, the Wishlist going to DEC at least two months before Symposium.

o After each symposium, the Wishlist responses from DEC will be published, along with the new wishlist items that were generated at the Symposium or via the Newsletter Wishlist Form. Users will vote on the latter so DEC can be given a priority of items before the next symposium.

o We should include UNIX software tool wishes. The cross fertilization of ideas would be valuable.

o New product wishes are better handled initially in product forums and forwarded to the Wishlist later.

o Wishes related to DECUS products should be handled differently (but no discussion of particulars occurred.)

o Too many wishes are related to better documentation.

THE LANGUAGES & TOOLS CLINIC --- Discussion Leader, George Scott, Clinic Director; Recorder Wayne Sewell.

Due to the huge success of the Clinic at the last symposium, the Clinic has been expanded to three full hours for Anaheim. This amount of time is too long for any of the doctors (developer or master) to be expected to remain available for the entire Clinic. The Clinic will begin with relatively light staffing and will build up to a climax.

The tentative schedule provided by George is a first draft and is not locked in. The staffing in the draft schedule is based on the number of ailments diagnosed in each disease (language/tool) by the doctors of the Nashville Clinic. There was no data on the commercial languages because they were not part of L&T at the time. George just made a guess for those.

Some of the topics require staffing throughout the entire Clinic because of the multitude of questions asked at Nashville. These include CMS, MMS, C, and FORTRAN. Dena Shelton wanted BASIC added to that list, based on the results of similar Clinics held by the Commercial Languages SIG. She also pointed out that DIBOL and RPG were not listed at all. George was not aware that they that they should be, and will add them. Dena felt that DIBOL required staffing for the full three hours also, while RPG only needed half an hour. It was decided that Scan would be covered in the time slot for compilers as well as the one for editors and text formatters.

There was some confusion about how to handle the SMG and other RTL questions. These may be the providence of the VAX SIG. Of course, a doctor who is asked a question about a non-L&T product is free to answer it if he can.

George requested feedback on the tentative schedule, both during and after the meeting. The working groups chairs should contact him if any hole in the coverage is seen.

The developers are to be scheduled for one hour each. Because of the limited number of developers sent to the symposium, each individual developer may cover more than one topic. It may be necessary to set up appointments.

Space in the campground shouldn't be a problem at Anaheim. Sometimes people can cluster around while someone else is being helped and can get their own problem solved in parallel. Sometimes one of these peripheral people can even provide the answer to the question.

The Clinic can be held early in the week---it is not related to anything else. The working groups should contact the Masters in their

topic and ask them to be there during the slot for that topic. The final schedule for the Clinic should be coordinated with the sessions in competing time slots. The availability of both developers and Masters for a particular subject depends on the lack of sessions in the same subject at the same time.

STEERING COMMITTEE MANUAL --- Discussion Leader Terry Medlin, Symposium Logistics Coordinator; Recorder, Kathy Hornbach.

After Nashville, Terry asked whether a SIG Manual would be a good idea. L&T is a big SIG, with lots of positions and many new members from CL. It's possible that new people in the SIG could be flabbergasted by its size and structure - intimidated and scared away.

Terry called on Mike Terrazas, a newcomer to the SIG in Nashville and a new Working Group Chair.

- Mike responded that he had a lot better understanding of the SIG after one day of the Woods Meeting
- He asked, Who can I contact via DCS? (he didn't know or remember which SIG members he could contact, and noted general DCS problems)
- He felt the need for a job description and organization description
- He suggested a DECUS Jargon Dictionary

Some of these are DECUS leadership issues, but some of it is SIG-specific.

People really need to understand entire DECUS structure, even more than SIG structure. There was an LDEC session that had this info last time - but it went only to newcomers - should go out to/all leaders.

Sam is working on job descriptions, and has been, but its a huge job and positions proliferate. How about trying to get people (again) to submit descriptions for their own job? Terry did a really good, and detailed, job on the PSS coordinator's and campground coordinator's job descriptions. Working Group Chairs, especially new ones, could write out their job descriptions to help them understand their job.

Shava suggested having a volunteer who writes well interview people and write up their job descriptions - maybe for the newsletter.

Steering Committee members please tell Sam on DCS what they think should go in this manual. (One suggestion - give hints for newcomers on the dress code for Woods Meetings).

COUNTERPART'S REPORT --- Discussion, Celeste LaRock, Digital counterpart; Recorder, Ed Woodward.

Working With DEC:

There was an interest on the part of Celeste to understand how Digital can help

L&T-11

both the Languages and Tools SIG as well as the individual Working Groups which are the technical thrust of the SIG. To achieve a more cohesive relation with the SIG and the Working Groups, it is her intent and that of Digital to develop the roles of the development counterparts assigned to the SIG and the various Working Groups. These are Leslie Klein, Jim Totten and Joe Mulvey. Celeste would like to provide the development counterparts with copies of all materials distributed at the L&T Woods Meeting.

Limitations of the Counterparts:

The primary limitation on the use of counterparts is the fact that the number of Working Groups exceeds the number of counterparts available. It will be necessary to understand this situation and be willing to share the available counterpart resources.

Digital's Plan for DEC Symposium Attendees:

A meeting with all DEC developers who will be attending Symposium will be held prior to DECUS Symposium to discuss their activities and roles at the Symposium. She will make sure that all are up-to-date on Digital products and DECUS activities.

Agenda of Activities at the Spit Brook Road Facility:

We were informed of the agenda of activities scheduled for tomorrow (Monday) at DEC's Spit Brook facility. Included will be discussions with the Development Counterparts of Working Group goals and what each WG expects from Digital.

DECUS Symposium Equipment Requirements:

Celeste solicited feedback on the supply of equipment that was available in the Nashville Campground. The general response was that the equipment provided by Digital was good, but the service was not as good. In addition, because of hotel complications, serious power problems were experienced. It was decided that the SIG would have the same equipment requirements at the upcoming Anaheim Symposium.

Responsiveness of Digital:

Celeste wanted to find out if Digital was providing the services needed by the L&T SIG. There was general agreement that Digital was satisfying the needs of the SIG.

WRAPUP --- Sam Whidden.

The Sunday session was excellent, accomplishing everything expected of it. I believe the many new Working Group Chairs learned much about how the SIG functions and about their valuable and important place in it. We spent considerable time on the issues and goals each WG has established and reached a good understanding of the SIG's expectations of the Working Groups: deliverables such as Symposium Sessions, Wishlists, and others, as well as the Working Group's role as a focus for DECUS members sharing their concerns.

L&T-12

Many SIG issues were considered and discussed. The meeting was kept nicely on time by our timekeeper, Dave Ream, preventing drift and ensuring that agenda topics were covered. Presenters were well-prepared, and the recorders did an outstanding job of reporting each segment in a manner permitting this edited report to be prepared in little time. Thanks go to every participant.

**Language and Tools Working Group Statements
for Spring 1987.**

edited on September 14, 1987

Joseph Pollizzi

In order to provide better technical coverage of the diverse language and tool products supported by the Languages and Tools SIG of DECUS, L&T adopted the creation of technical "Working Groups". There are no formal requirements placed on any working group other than it demonstrate its service to the DECUS community by defining "deliverables" - something of concrete value that the membership of DECUS can take advantage of. Deliverables can vary from the collection of and presentation of sessions at a Symposia, to submitting articles to the newsletter, to making submissions to the SIG Tape (as examples).

The following is a compendium of the summary reports for each of L&T technical Working Groups. If you are interested in any of these topic areas and wish to participate, we heartily encourage you to contact the Working Group's Chair.

If on the other hand, you have recognized the need for a technical group that is not represented by one of the following, and you are interested in forming your own working group, or if you have any other questions about L&T Working Groups, please contact the L&T Working Group coordinator at:

Joseph A. Pollizzi, III
Space Telescope Science Institute
3700 San Martin Drive
Baltimore, MD 21218
(301) 338-4901
ARPA: POLLIZZI@STSCI.ARPA

1 APL Working Group

We had three sessions on APL at the Nashville Symposium:

- An update from DEC on the latest version of VAX-APL.
- An introductory session: "What is APL?"
- An Open Working Group Meeting.

Attendance was good at the sessions: in fact, higher than expected (over 40 people attended the introductory session) and discussions were quite lively.

We started a name and address list to help people interested in the working group to keep in touch and to share information. The list has about 15 people at present.

Several of those that attended the sessions promised to contact others in their companies to give sessions at future symposia and to make submissions to the newsletter. We are interested in increasing the number of sessions at Symposia, especially on applications using APL. Among other issues discussed at the sessions were ANSI standards, the syntax for nested arrays, external calls, prototyping in APL, music applications, object-oriented programming, and APL for scientific applications.

Working Group Chair Address:

Bob van Keuren
UserWare International, Inc.
2235 Meyers Ave.
Escondido, CA 92025
(619) 745-6006

2 Basic Working Group

The goals of the Basic Working Group are: To be the DECUS focal point for the language BASIC. To encourage the development and growth of BASIC users on DEC computers. To help guide DEC in the development and compatibility of BASIC on all operating systems. To participate in the current ANSI standards effort.

Stated Working Group Goals:

- Symposium working group, clinic and wishlist sessions.
- Encourage user participant sessions and guide DEC in user desired sessions.
- Solicit newsletter articles and L & T volunteers.
- Participate in ANSI X3J2 committee.

Working Group Chair Address:

Stephen C. Jackson
7260 University Avenue NE, Suite 105
Minneapolis, MN 55432
Work: (612) 571-8430
Home: (612) 571-1499

3 C Working Group Statement

The function of the "C" WG should perform the following:

- Follow the ANSI "C" standard
 - How does the standard effect the current users of the language ?
 - What does the standard mean ?
 - How can the "C" WG effect the standard (ie What inputs can can the WG input to the standard) ?
- Provide help to new users of the language.
- Determine under what application environment the "C" language is is used and how the WG can promote the usage of the language.
 - Is portability an issue ?
 - VMS unique requirements, (i.e. writing of system type software such as executive system service and/or device handler in C).
- Since this is the first time this WG will be meeting , request input and help from the current user's of the language and the software tools environment. How can the tools better aid the development of the application program environment of "C" programs.

Stated Working Group Goals:

- Symposia sessions in using VAX C and System Services.
- I hope to submit a tape to the L&T for using global section within VAXcluster, were the global section is mapped for read/write access within the VAXcluster.

Working Group Chair Address:

James A. Maves (Jim)
Sr. Staff Eng.
Eaton Corp
Information Management System Division (IMSD)
31717 La Tiuienda Dr.
Westlake Village, Ca. 91359
Phone work: (818) 706-5395
Phone home: (805) 498-9826

4 COBOL Working Group

Over the last three decades, COBOL and the tools that can be interfaced to COBOL have undergone many changes. It is the goal of the COBOL Working Group to serve as a resource for the dispensing of up-to-date information to COBOL programmers and their management. In addition to presenting information in these areas, the COBOL WG is also interested in helping the professional programmer become more productive. To this end, the Working Group is anxious to hear about COBOL related concerns, problems, and interests which could be presented as a session at a future DECUS Symposia and/or relayed to DEC.

Stated Working Group Goals:

For Anaheim Symposium:

- Check into possibility of having "hands-on" use of the COBOL Generator and other tools present at the Symposium.
- Provide a COBOL sessions "Road Map" for the Symposia if warranted.
- Publicize significant COBOL-oriented sessions in the *UPDATE.DAILY*.

Long Term Goals:

- Look into developing a COBOL interest database.
- Develop an ongoing contact with the COBOL ANSI Standards representatives.
- To seek out, or develop future symposia sessions that are both sensitive and timely to the needs of COBOL users.

Working Group Chair Address:

Edward W. Woodward
Science Applications International Corp.
10210 Campus Point Drive, M/S #24
San Diego, CA 92121
(619) 546-3758

5 DIBOL Working group

The Goals of the DIBOL Working Group are to provide a forum for users to exchange information about the uses DIBOL in business and other application areas, and to have access to the Digital developers for questions and the voicing of concerns.

Working Group Chair Address:

Bruce Mebust
VICOM
9713 Valley View Road
Eden Prairie, MN 55344
(612) 944-7135

6 Large SW Systems WG Information

Large Software Systems addresses the special problems of generation and maintenance of the larger software systems that tend to take over a whole computer single-handedly. Configuration management, testing needs, life cycle maintenance, and applicability of standards are all considered.

There will be an open meeting of the working group at the Fall 1987 Symposium in Anaheim to determine the direction and activities of the group.

In addition, the working group is sponsoring a user session on the use of the VAXSet products in a Large Software System environment.

Working Group Chair Address:

George L. Scott
 Computer Sciences Corporation
 P.O. Box N
 Moorestown, NJ 08057
 (609) 234-1100
 Temporary phone (215) 657-4500

7 Low Level Languages Working Group

This working group address the use of "low level" languages on Digital Equipment Corporation computer and the concerns and needs of their users. The working group will also be concerned with how and why these languages are used and what tools are needed to support the efficient use of these languages.

Stated Working Group Goals:

At the Fall 87 symposium the working group will hold its first organizational meeting at which I hope to better define the scope of the working group. Out of this meeting I will attempt to write a paper detailing:

1. What we mean by a low level language.
2. How they are used.
3. Why they are used.
4. What features would improve them.
5. What tools are needed to support their use.

Working Group Chair Address:

Gerald W. Lester
 2901 Houma Blvd., Suite 5
 Metairie, LA 70006
 (504)-889-2784

8 Modula-2 Working Group Report

The Modula-2 Working Group has been formed for the purpose of promoting the Modula-2 language. The W/G intends to do this by supplying the DECUS membership with articles discussing the features and capabilities of Modula-2, by publishing lists of available compilers, and working to find a way to distribute Modula-2 compilers on the SIG tapes. We also expect to track the standards efforts that are taking place, and submit articles about these efforts to the newsletter. Our long term goal is to popularize Modula-2 to the extent that DEC will see it a viable language and offer compilers.

The Modula-2 working group is actively soliciting articles, SIG tape submissions, and people who would be willing to participate in and contribute to the goals of the working group, if you have a contribution, or wish to participate in the working group, please call or write.

Working Group Chair Address:

Jack Davis
 NAP Consumer Electronics Corp.
 Videowriter Business Unit
 1111 Northshore Drive, Building 2
 Knoxville, Tn 37919
 (615) 558 5206 (8-4:30 eastern)

9 Pascal Working Group

The initial goals of the Pascal W/G are:

1. Organize sessions (including tutorials, panels, and clinics) and BOFs at Symposia pertaining to Pascal.
2. Solicit newsletter and Proceedings articles on Pascal.
3. Interface with Pascal development staff at Digital.
4. Coordinate with Masters Program to make effective use of Pascal Masters for user-assistance, clinics, field-testing, product reference, etc., and to recruit new Pascal Masters for the program.
5. Report on current status of Pascal standardization effort.
6. Solicit wish list items for Pascal to be incorporated into the main L & T list.
7. Solicit Pascal-related submissions to SIG tape, such as environment, include files and utilities and sample programs written in Pascal.

Stated Working Group Goals:

In realization of the first goal, the working group came up with some ideas for sessions at Anaheim. In addition to the three I am already proposing (software engineering in Pascal: WEB, which isn't strictly Pascal, but uses it as a byproduct: and the working group meeting, of course), we thought of a couple of topics.

1. A panel of users discussing exotic applications using Pascal.
2. A presentation explaining use of VAX Pascal attributes, one of the most confusing areas of VAX Pascal and the largest departure from standard Pascal.

Future goals to strive for after the group is more established:

1. Pre-Symposium Seminars on Pascal programming.
2. Direct involvement in standardization effort. (Dependent on overall DECUS policy on standards).

Working Group Chair Address:

Wayne Sewell, Chair
Software Engineering Specialist
E-Systems, Garland Div, MS 53730
Box 660023
Dallas TX 75266-0023
(214)272-0515. ext 3553
CIS — 70037,1755
BIX — wsewell
MCI Mail — esewell

10 PL/I Working Group Report

Release 3.0 has been distributed per announcement at Spring 1987 Symposium in Nashville. User comments were very favorable. The only user request not seemingly satisfied was for a method of creating indexed sequential files without having to use FDL\$CREATE.

A new standard is out and will be obtained for review.

Nothing is known about what is next on the horizon from DEC.

Immediate goals are to solicit comments from masters and prepare for working group session at Anaheim in December.

Working Group Chair Address:

David K. Ream
Lexi-Comp
26173 Tallwood Dr.
N. Olmsted, OH 44070
(216)777-0095
(216)468-0744 (for messages)

11 PDP-11 Languages & Tools Working Group

The CL and L&T sigs have both chosen me to be the working group chair for their PDP-11 Working Groups. At this time I am soliciting input from all persons interested in assisting me with this working group.

During discussions at the last CL woods meeting, the initial issues for this working group were brought out:

1. The Coexistence of VAX and the PDP-11,
2. Compatibility of PDP-11 Software with VAX software,
3. Support of PDP-11 peripheral hardware on the VAX,
4. Continued Development of software on the PDP-11.

This is just the beginning of this working group, and the above items are not listed in any order of importance. As this working group progresses, we will continue to add to the list as new issues are uncovered, as well as embellish upon the currently defined issues.

Working Group Chair Address:

Bill Tabor
W. I. Tabor, Inc.
11652 N. W. 30th Street
Coral Springs, FL 33065
(305) 755-7895

12 SCAN Working Group Report

Trying to develop a large list of SCAN users and sessions for symposia is a high priority.

Many enhancements could be made to SCAN to allow access to other VAX system features from within the language as well as adding more built-in functions, tokens, and "parsing" capabilities.

SPR response is basically non-existent. TSC support, except for simple syntax situations, is available only from one software specialist who is not even on a language team. (A DEC developer at Spit Brook claimed all this was due to the transition of SCAN support to another group.)

Working Group Chair Address:

David K. Ream
Lexi-Comp
26173 Tallwood Dr.
N. Olmsted, OH 44070
(216)777-0095
(216)468-0744 (for messages)

13 Software Metrics

To share ideas, recommend solutions to DEC, and raise questions regarding the measurement of development productivity and maintenance effort for software projects. Emphasis will be placed on the search for methods and tools for obtaining meaningful metrics.

Stated Working Group Goals:

- DECUS Sessions (the WG meeting & a paper by Anne Duncan have both been submitted for Anaheim)
- A bibliography on software metrics (to be published in the SIG newsletter)
- Software metrics book reviews (also in SIG newsletter)
- Other things as suggested by WG members.

Working Group Chair Address:

Michael S. Terrazas
LDS Church - 27th Floor
50 E. North Temple
Salt Lake City, UT 84150-0001
(801) 531-3246

14 TECO Working Group

Description: To promote and facilitate the use of TECO across all DEC product lines and operating systems. In addition, to provide a central clearinghouse for TECO implementations, documentation, and macros.

Stated Working Group Goals:

- Sessions - Intro to TECO (F87 & S88)
VMS Native TECO (S88? VMS 5.0 Time frame)
TECO BOF/WG Meeting (F87 & S88)
- Tape - Machine readable documentation
Implementations
Macro library
- Other - Update of 1978 TECO Reference Card
Update of 1980 TECO Manual(?)

Working Group Chair Address:

Mark J. Hyde
Advanced Computing Services
209 Ardsley Drive
DeWitt, NY 13214
(315) 446-7223

The following Working Groups were unable to provide the editor with a brief summary of their activities by press time. In their place, the editor submits the following brief summaries on their group chair's behalf

15 Fortran Working Group

This Working Group is concerned with the usage, development and compatibility of the various DEC implementations of the FORTRAN Language. Of special interest to this group is impact and resolution of the recommendations from the ANSI Committee's new FORTRAN Standard.

Working Group Chair Address:

Scott Krusemark
8473 Daisywood Ave Nw
North Canton, OH 44720
(261) 499-6251

16 Technical Production of Documentation

This working group is interested in the various aspects in producing quality documentation. These issues cross from the use of various editing facilities, markup vs wysiwyg processing products, and typesetting systems. The advent of low-cost "desk-top" publishing systems, the availability of numerous documentation systems, and the rapidly falling costs of quality printing devices has made this area especially dynamic and exciting.

Working Group Chair Address:

Howard Holcombe
RCA
Front & Cooper St.
Camden, NJ 08055
(609) 338-4946

17 Public Domain Software Working Group

In the "Good-Ole-Days" of computer science, new and interesting developments in the field were generally made available to the public at large. Many of the original implementations of compilers, editors, even complete systems were distributed in this manner. Well such products still do exist in this day and age, and the goal of this working group is to discover such items, report on them, and solicit/encourage others to make their usefull inhouse tools available to all (were permissible). It should also be noted, that the chair of this working group is also (quite naturally) the L&T SIG's Tape Librarian. In effect, the end product of this working group is the L&T Tape.

Working Group Chair Address:

Tony Scandora
Argonne National Laboratory
CMT 205
Argonne, Illinois 60439
(312) 972-7541

18 Project Management Working Group

With the cost of software development ever increasing with respect to the total cost of a computing system, much interest has been given to the better overall management of such system development. Capitalizing on this interest, several "management methodologies" and tools that support these methodologies have been produced. Member of this working group pay particular attention to these developments, analyzing the applicabilty of these techniques in their own situations, and to make recommendations on their needs as they perceive them in project management.

Clearly this working group crosses over in interest with those of several other working groups (Notably the Configuration Management, Large System's Maintenance, Software Metrics, and Tools Integration groups). However, the primary focus of this group is on the management of such projects, and the development/use of tools in assisting with project management.

Working Group Chair Address:

Dena Shelton
Cullinet Software Inc.
2860 Zanker RD. Suite 206
San Jose, CA 95134
(408) 434-6636

19 Tools Integration Working Group

This working group is primarily concerned with the interaction (or lack there-off) of the various tools products used in producing software. This group is concerned with the overall tool interconnectivity from product definition (editors, project management/tracking systems, documentation systems), to product design (compilers, building and configuration tools), to product assurance (test tools, debuggers). The group is also interested with the compatibility of using various languages together.

Working Group Chair Address:

Jay Wiley
Bectel Power Corp.
12400 East Imperial Highway
Norwalk, CA 90650
(213) 807-4016


20 RPG Working Group

The RPG is another language working group. Its primary concern is advancing the usage of and the ongoing development of RPG.

Working Group Chair Address:

Charles Williamson
Hargray Telephone Co.
P.O. Box 5519
Hilton Head Is., SC 29938
(803) 686-1204

GLOBAL ACCESS



"If you don't want it in the paper, don't let it happen."



MMP

GLOBAL ACCESS

Volume 1, No. 3

November, 1987

MUMPS means you never have to say you're sorting.

\$VIEW(Editor)

The most pervasive hassle at a DECUS Symposium is schedule conflicts. The knottiest conflict confronts everybody first thing Monday morning. That's right, the Roadmap sessions. You always want to get to at least four of them, and they're ALL at 9:00 AM. On top of this, you want to get there early enough to collect the Roadmap handout, which is a necessary adjunct to the Sessions-at-a-Glance because it is organized by subject rather than by time. Thus, you spend the first hour rushing around from Roadmap to Roadmap. I believe that it would take only a few minor tweaks to eliminate most of this problem. Accordingly, I submit to the Symposium Committee the following suggestions:

- o Stagger scheduling of Roadmap sessions. As of Nashville, the MUMPS Roadmap is not at 9:00 AM; some other SIGs are even starting their streams on Tuesday.
- o Combine Roadmap sessions. I seem to remember hearing that two SIGs will have a combined Roadmap in Anaheim, but I can't verify this. I am sure that many of the smaller SIGs (including MUMPS) could share Roadmaps just as we share Campgrounds.
- o Print more handouts, and leave copies of all of them out in a central location (the Registration area?) where attendees can get them if they miss the session. Long range, perhaps the format could be standardized, and the handouts consolidated into a pamphlet to be given to attendees (but let's NOT let it turn into a megillah like Session Notes, with its own committee and raft of political/money grief!).

In my opinion, these constitute minor changes that would give a major return on investment.

\$DATA

As promised last month, here goes with "Sessions You Won't Want to Miss in Anaheim:"

General Interest

MUMPS Roadmap	M001	Mo	10:30AM
Early History of MUMPS	M018	Mo	10:00PM
A Voice Response Inquiry System Using DECTalk	M020	Tu	4:00PM
MUMPS Dump Busters	M015	Th	7:00PM
MUMPS System Management Panel	M016	Fr	9:00AM
MUMPS Development Committee Report	M005	Fr	3:00PM
MUMPS Future Planning Session	M007	Fr	4:00PM

Tutorials

Nearly Everything You Didn't Know Enough To Ask About MUMPS	M011	Mo	9:00PM
MUMPS Functions and String Handling	M002	Tu	10:00AM
How to Tame a Wild Programming Staff	M004	Tu	5:00PM
Relational Vectors Proposal for MUMPS	M013	We	10:00AM

Public Domain Software

Public Domain Trauma Registry	M022	Tu	12:00PM
Overview of Public Domain VA Software	M008	We	12:00PM
Using VA File Manager as a Relational DBMS	M019	We	4:00PM
Changes to VA Kernel for VAX DSM	M021	We	5:00PM

DSM Specific

DSM Product Panel	M017	Mo	12:00PM
DSM and Networking	M027	Tu	11:00AM
MUMPS Utilities and Their Usage	M014	Fr	10:00AM
DSM-11 Performance and Tuning Workshop	M026	Th	9:00PM
Setting up a DSM-11 \$ZCALL	M009	Th	12:00PM
VAX DSM System Overview	M024	Mo	1:00PM
DASL--DSM Application Software Library, A VAX-Based Application Generator	M023	We	11:00AM
Calling Other Languages from VAX DSM Using the \$ZCALL Function	M028	Th	11:00AM
VAX DSM Performance and Tuning	M025	Th	8:00PM

\$HOROLOG

November 20	Submission deadline for Jan. newsletter
Dec. 7-11	Fall '87 Symposium; Anaheim, CA
Feb. 8-12, 1988	Canadian '88 Symposium; Toronto
May 16-20, 1988	Spring '88 Symposium; Cincinnati, OH
June 13-17, 1988	MUG '88 Conference; New Orleans

\$ORDER("Author")

I have observed (very painfully, as each instance screeches across my ears like fingernails down a blackboard) that the word *author* is coming into increasing use as a verb. This is typified by constructs such as "Representative Gramm and Senators Rudman and Hollings *authored* the Deficit Reduction Act." In fact, several years ago, even DEC released a package called the "Courseware [hmm, but that's another whole article...--Ed.] *Authoring System*." This trend ignores the fact that, in light of the rich history of English, it is very common for a "perpetrator" noun and its corresponding verb to have sprung from different roots. For example, you don't take a car to the garage to have it *mechanicked*; you have it *repaired*. Similarly, you certainly do not go to the hospital to be *doctored*; I am sure you would much rather be *cured*. Furthermore, while my dictionaries are inconclusive on the point, my "internal compass" is not: *author* is not a verb. English is replete with verbs which express the creation of manuscripts: *to compose*, *to draft*, *to formulate*, and even the humble *to write*. A visit to the local thesaurus will elicit many more. It thus seems to me unnecessary to force *author* into such a graceless, unnatural role.

\$NEXT

Did you know that the MUMPS SIG is currently under serious SIG Council scrutiny concerning its continued existence? Stay tuned in January for developments in this heart-rending story.

\$NEXT(\$ORDER)="Ship"

\$RANDOM

Oscar Kelly Allen, the Louisiana governor who succeeded (and was owned by) Huey "The Kingfish" Long, was described thusly: "He was so agreeable that if a leaf blew in the window across his desk, he would sign it."

RSX

MULTI-TASKER



RSX

The RSX Multi-Tasker

November, 1987

"Stand Up, Ed"

Fine Realtime Commentary Since 1975

Table of Contents

Food for Thought	RSX-1
The Editor's Corner	RSX-1
Productivity and Relative Value	RSX-2
Submitting Articles to the Multi-Tasker	RSX-4
And That's The Way Things Are	RSX-4
The Bag of Tricks: MACRO-11	RSX-5
Bulletin Board Notes	RSX-6
Initializing DCL with Node Specific Prompts	RSX-6
PDP-11 Performance Benchmark	RSX-8
SIG Business Update	RSX-9
VT220 Downline Loadable Character Sets	RSX-9
Weird RSX Science	RSX-18

Opinions expressed in the editorial section of the Multi-Tasker are those of the Editor. They do not represent the official position of the RSX SIG or that of DECUS leadership in general.

Food for Thought

"Another difference between murder and war is when and where they take place. Murder takes place in dark alleys and behind closed doors and when nobody is looking. War takes place on battlefields, out in the open, with newspaper reporters and photographers and even television cameras to record it. Obviously there is something sneaky about murder, as contrasted with war."

- *Richard Armour*
It All Started with Stones and Clubs

The Editor's Corner

Bruce R. Mitchell

Have we got problems up here at the Multi-Tasker offices. What with the polar bears migrating down into Minnesota already, the shower delivering ice cubes again, and the snowmobile not working so I can't get to work, what else can go wrong? Uff da! (= Scandinavian "Oy!")

Copy for the September issue got to the DECUS offices just one day too late to be printed in the September issue. OK, says I, we'll print September and October combined in one big issue.

The October issue hit the DECUS offices one day later than expected, but still in time. Just one little problem, though - everybody else wanted to publish in that issue. So the October issue was trimmed. September made it into the October issue, at least. I hope so - I have no way of knowing at this writing.

So now - if nothing else goes wrong - I hope you are now

reading the combined October and November issues. Well, it's all good reading material, and worth waiting for.

Got your symposium registration in yet? There will be more RSX activity than you can shake a stick at in Anaheim. Shucks, there'll be more RSX activity than you can shake a stick at anywhere. Just a few of the sessions to be presented include:

- o Advanced Device Driver Techniques
- o Digital's Lemons
- o Justin L. Hewser's Big Book of System Optimization
- o Overlaying: Fast, Quick and Easy
- o RSX to VMS Communication Methods
- o Running M-Plus on 18-Bit Addressing PDP-11s
- o VAX Coprocessor/RSX Update

plus the usual RSX sponsored fun and games, including the time-hallowed visit to the VMS Magic session. (*Anything we do more than twice is a time-hallowed tradition. -- The Editor.*)

The Spring 1987 Proceedings recently arrived by dogsled at the well-insulated Multi-Tasker offices. After reading the preface a few times, the Editor would like to serve up a few reactionary comments for your delectation.

----- Productivity and Relative Value -----

Productivity is a big buzzword these days. Extremely clever MBA managers, as well as a slew of writers who have never heard of Frederick W. Taylor and couldn't define a "therblig" if their lives depended on it, have suddenly awakened to the fact that Productivity is a Good Thing. And it should be Encouraged.

Well! Somebody above line supervisor level finally realized that making product and getting it out the door is the name of the game.

On a seemingly unrelated topic, I now hear that the constituency of the "New DECUS" is not that of the "Old DECUS", and that the "New DECUS" has different needs and different directions to take. The "Old DECUS" membership of technical, laboratory and industrial programmers has lost its mandate, and - I am so much as told - has no business in the new Society, composed of middle managers and supervisors.

For a moment, let's discount the results of the recent DECUS Board election as a rebuttal. Instead, let me examine this "interesting" idea from the standpoint of productivity.

I shall apply the concept of relative value in this situation. As Robert Heinlein says, given dough and fresh apples, one can make a pie. A professional chef, given the same

materials, can create a superior confection of higher value. Conversely, incompetent use can transform the already valuable raw materials into an inedible mess.

This is all to say that labor is not a uniform product. Some types of labor are more valuable than other types of labor.

It should be obvious that "productive" people are valuable to a company. Likewise, it should be obvious that superior productivity allows a company to compete more effectively in the marketplace. Companies which do not compete effectively in the marketplace do not survive. Therefore, from a company's point of view, the value of an employee to the company is proportional to the "productivity" of that employee.

Given two employees of a company - a production programmer and a middle manager - who contributes more? The production programmer writes programs which control machines. The machines take raw materials, add energy and time, and create valuable products. Without the production programmer, those products either would not exist or would be available in lesser quantities, costing more. Basic microeconomics.

The paper-shuffling MBA manager, on the other hand, contributes little to the production capacity of the company. Whether the manager is there or not has small effect on production. Take away the manager, the product still goes out the door. The corollary is - or should be - obvious.

It may be argued that management is necessary to sell the product, balance the books, and run the show. A limited amount of management is necessary. Good enough; I'll argue it from that point of view as well.

Assume that the company has sufficient programmer and management capacity to run at its current level. What is the incremental effect? Adding one production programmer generally has a positive effect on cash flow. Adding one manager doesn't.

On that basis, which is more valuable to the company? Is it the MBA who thinks a CNC mill is used for grinding his \$5 morning cup of Jamaica Blue Mountain coffee? Or is it the production programmer who forces those greasy, recalcitrant, stupid robots to weld auto frames?

Assume that attending DECUS symposia makes one more productive in one's job. Which person should be reading the newsletters and attending symposia?

The reason for Japan, Inc.'s success? They produce. A lot. They do it cheaply and well. And much of that is true because they keep management out of the way of production.

There are plenty of nice professional management organizations in the good 'ol US of A where managers can go to drink a nice lunch in a nice air-conditioned meeting room and listen to other managers talk about management. Let them stay there. DECUS is our turf. I mean the "Old DECUS", the obsolete old farts who program computers for a living and think MBO is just another damn TLA and a form to fill out.

DECUS was once a productive society of productive people. And we're going to make it one again.

As it is said: "Push. Pull. Or get out of the way."

Because there's a storm rising.

----- Submitting Articles to the Multi-Tasker -----

Please submit machine readable media. RX01/2, RX50 and 9 channel 800/1600 BPI magtape are OK. Almost any medium I can't read I can get converted; don't let that stop you. All RSX volume formats are acceptable, but I can't read VMS Backup or ODS-2 stuff.

You can also submit articles through the RSX bulletin board system at (612) 777-7664. Send them via MAIL to username MULTITASKER.

Submissions which aren't machine readable may take longer to get into print. The editor is lazy and types mass quantities only once a month when progress reports are due.

If you preformat a submission in RUNOFF, please set page size 58,80; left margin 10; right margin 75; and, when changing margins, use incremental changes rather than absolute. The editor blesses you for the consideration.

Send your articles and other submissions to the luxurious Multi-Tasker offices:

Bruce R. Mitchell
Machine Intelligence and Industrial Magic
PO Box 816
Byron, MN 55920
(507) 775-6268

----- And That's The Way Things Are -----

... this month in Pool Lowbegone, where the drink in the SIG suite is strong, the RSX SIG's Fall Symposium presentation to VMS will be handsome, and the tendency of RSX users toward productive work is above average.

The Bag of Tricks: MACRO-11

Richard Neitzel
312 Laveta Pass
Golden, CO 80401

In this month's article we examine the problem of counting bits set in a word.

There are times when one needs to know how many bits in a word are set or clear. For example, one may be checking a disk's bitmap to see if there is enough free space left. The numeric value of the word is useless to us, and checking each individual bit is clumsy and time consuming.

There is a simple algorithm that gives the correct value. It is based on the fact that the logical AND of a number and that number less one clears the lowest set bit in the number. Using the ANDed result as the next value and repeating the process until zero is reached yields the number of set bits.

The implementation in Macro is simple, as follows:

```

;      Input:  R0 - Integer to be bit counted
;
;      Output: R2 - Bit set count
;
;      Register dispositions:  R0, R1, R2 destroyed

... code precedes ...

CLR      R2                ; Zero the bit set count

1$:      MOV      R0, R1    ; Get value to decrement
         DEC      R1        ; Minus one
         BIT      R1, R0    ; AND them together
         INC      R2        ; Increment bit set count
         TST      R0        ; Zero?
         BNE     1$        ; If not, go try it again

... code continues ...

```

And it's quite equally simple when done in Fortran:

```

200      BIT = 0                ! Zero the bit counter
         TSTBLK = IAND (TSTBLK,TSTBLK-1) ! .and. the 2 values
         BIT = BIT + 1          ! Increment bit set count
         IF (TSTBLK .NE. 0) GOTO 200    ! If result not 0 do more

```

Obviously, if the initial value may be zero, one should test for this before starting to avoid incorrect results.

Bulletin Board Notes

The BBS lost it big in mid-August; the RK711 disk controller went down in flames, taking the system disks with it. At this writing the BBS management is acquiring a new controller, and hopes that it will be up and available in early October.

As of mid-September, RSX MAIL and Kermit are available. Conferencing is unavailable but is a high priority.

A Vadic 3451P modem was installed on the system in early September, so the line now supports Vadic as well as Bell protocol.

The system still needs hardware. Anything. The biggest needs right now are a couple 80 Mb SMD winchesters and a 2400 baud modem. But anything and everything goes, so pack up all your disused treasure and ship it off to the BBS management c/o your friendly Multi-Tasker editor at the address above.

The BBS number: 1-612-SPR-PONG / 1-612-777-7664. This line is autobaud 110 - 1200 baud. To request an account, log in with account name ACCOUNT, password REQUEST.

Initializing DCL with Node Specific Prompts

Wayne Steffen
Cap Gemini America
134 Cedar Avenue
Stirling, NJ 07890

We are using DECserver 200s to communicate with several PDP-11s. I am the system manager, so I pop from one PDP to another quite often. I would often become confused as to which system I was on, and enter the right command on the wrong system at the wrong time, such as dismantling the active source disk while the development staff was rebuilding the application system.

I found some help by changing the DCL prompt. Those who have learned MCR and continue to refuse to upgrade may leave by the rear exits at this time (most of the programmers at this site).

I moved the code in STARTUP.COMD that initializes DCL to occur after the DECnet startup code. I also added the /DPR (default prompt) and /CPR (control-C prompt) switches to my DCL initialization command. The argument to these prompt switches is built from the INDirect special symbol <NETNOD>, which receives the executor node name:

```
/DPR="<15><12>'<NETNOD>' $ "  
/CPR="<15><12>'<NETNOD>' DCL $ "
```

Here is the command sequence as it appears in STARTUP.COMD after the DECnet NCP SET SYSTEM command:

```
.ENABLE SUBSTITUTION  
;  
; Enable DCL CLI  
;  
.SETS NOD <NETNOD>  
CLI /INIT=DCL/CPR="<15><12>'NOD' DCL $ "/DPR="<15><12>'NOD' $ "
```

and here is the command line as it is actually executed on node "DLCC" at system startup time:

```
CLI /INIT=DCL/CPR="<15><12>DLCC DCL $ "/DPR="<15><12>DLCC $ "
```

When the system is released for use after STARTUP, DCL prompts with the standard sequence:

```
DLCC $
```

and, in response to a control-C, prompts:

```
DLCC DCL $
```

with a space following the \$ character in both cases.

The next step is to set your CLI to DCL. The MCR command SET /DCL=TI: does this for your terminal, TI: . DCL as the default CLI can be specified in the account file as well, then the user's terminal is automatically set to DCL when the user logs on.

Now, by typing a carriage return or control-C, I can see which node I am on before I do something someone may regret.

SIG Business Update

In a conference held via the DCS computer system, the RSX SIG Executive Committee, meeting *ad hoc*, created the position of SIG Vice-Chair.

There is no precedent for creation of such a position. It was decided that creation of the position required approval of a majority of elected positions on the SIG Executive Committee. In votes taken via DCS and tallied 2 September, the following affirmative votes were cast:

27-Aug 14:34 EST, Robert Uleski; 28-Aug 12:14 EST, Bruce Mitchell; 29-Aug 18:54 EST, Denny Walthers; 30-Aug 23:14 EST, Ed Cetron; 2-Sep 16:19 EST, Gary Maxwell. No dissenting votes were cast.

The SIG By-laws will be amended to reflect creation of this position and the responsibilities pertaining to it.

Reported by acting recorder for the SIG Executive Committee, Bruce R. Mitchell.

PDP-11 Performance Benchmark

Richard Neitzel
312 Laveta Pass
Golden, CO 80401

Some interesting benchmark numbers from the July 87 issue of BYTE caught my eye recently. I decided that I would try their tests on my 11/73.

The results show (at least to me) what a good engine that CPU is. All the tests were coded as closely as possible to BYTE's C code with the Fortran-77 compiler. I didn't bother with all the tests, only those that I could steal time from productive work to write.

Here are the results, with the BYTE systems included. All times are in seconds with precision to three places.

Test	Recursive Fibonacci	Float	Sieve of Eratosthenes	Savage
Iterations	100	10000	100	25000
AT, 8 MHz with FPU	950 121	116 9.7	26.7 25.3	1100 38.3
Deskpro 386	3.1	5.4	6.0	35.1
Mac SE 16 MHz w/FPU	264 71.6	230 4.2	64.7 14.9	1880 8.8
Arete 1100	70.2	2.9	12.8	24.8
11/73	5.0	3.5	58.2	106

With minor tweaks (/TR:NONE, PRI=125) and an 11/83 or /84, I believe it could run at least twice as fast. Anybody care to try it? -- The Editor.

VT220 Downline Loadable Character Sets

K. F. Uhland
Scientific Micro Systems
339 North Bernardo Avenue
Mountain View, CA 94039-7777

The documentation set shipped with a VT220 has a phenomenal dearth of information with respect to programming the special features of the terminal, particularly in the area of defining a Downline Loadable Character Set (DLCS). The only programming information provided at all is the VT220 Programmer Pocket Guide (DEC P/N EK-VT220-HR-001), hereafter referred to as the Guide.

On page 19 of the Guide are the escape sequences for defining a DLCS as G0, G1, G2 or G3, and for then invoking one of these into GL (Graphics Left) or GR (Graphics Right).

On pages 30 and 31 are the escape sequences for downline loading a character set. It is referred to as a "DRCS character set", but nowhere is "DRCS" defined. The (corrected) information on page 30 is reproduced below:

Down-Line Loading Characters (DRCS)

You can down-line load your DRCS character set using the following DECDLD device control string format:

DCS Pfn;Pcn;Pe;Pcms;Pw;Pt{DscsSxbp1;Sxbp2;...;Sxbpn ST

Parameter descriptions are as follows:

DECDLD Parameter Characters

<u>Parameter</u>	<u>Name</u>	<u>Description</u>
Pfn	Font number	0 and 1.
Pcn	Starting character number	Selects starting character in DRCS font buffer to be loaded.
Pe	Erase control	0 = erase all characters in this DRCS set 1 = erase only the characters that are being reloaded 2 = erase all characters in all DRCS sets (this font buffer number and other font buffer numbers)
Pcms	Character matrix size	0 = device default (7 x 10) 1 = (not used) 2 = 5 x 10 3 = 6 x 10 4 = 7 x 10
Pw	Width attribute	0 = device default (80 columns) 1 = 80 column 2 = 132 column
Pt	Text / full cell	0 = device default (text) 1 = text 2 = full-cell (not used)

Dscs defines the character set "name" for the soft font, and is used in the SCS (select character set) escape sequence.

Sxbp1;Sxbp2;...;Sxbpn are sixel bit patterns (1 to 94 patterns) for the characters, separated by semicolons. Each sixel bit pattern has the form:

S...S/...S

where the first S...S represents the upper columns (sixels) of the DRCS character, the slash advances the sixel pattern to the

lower columns of the DRCS character, and the second S...S represents the lower columns (sixels) of the DRCS.

Obviously there are some rather large holes in this description. First of all, the description of parameter *Pcms* is defective; *Pcms* is missing from the first column and there is an extra blank line in the description. (Ed. note: Corrected in text above.)

Second, the phrase "starting character number" for parameter *Pcn* leaves a number of unanswered questions. What number corresponds to which character? Is this character number decimal, hexadecimal or what?

Third, and perhaps most important, just what the devil are "sixels", what is meant by "upper columns" and "lower columns" of a DRCS character, and how are all these things related?

Finally, there are a few "minor details" like those spaces shown in the command format - are they or are they not to be included in the actual command sent to the terminal?

The Guide answers none of these questions. True, the following paragraph does appear on page 1 (which isn't numbered, by the way):

This Pocket Guide provides a summary of the information contained in the VT220 Programmer's Reference Manual EK=VT220-RM which can be ordered from Digital. It is provided for use by people with a knowledge of computer programming to access the VT220 features.

However, the first entry in the Table of Contents is for page 2. One must accidentally stumble over this paragraph to learn even of the existence of the Programmer's Reference Manual (hereafter referred to as the Manual). The neophyte programmer would indeed be hard pressed to learn the terminal's secrets without access to the Manual; even with it, not much more can be ascertained. One would require either tremendous insight or recourse to trial and error methods.

The purpose of this article is to try to demystify the downline loading of characters and to provide additional information not available from DEC that this author discovered while experimenting with the terminal.

The Manual defines "DRCS" as "Dynamically Redefinable Character Set". The DRCS is composed of those characters whose decimal ASCII codes fall into the range 33 through 126, octal 41 through 176, or in other words, the printable characters!

through ~.

These 94 characters are sequentially numbered 1 through 94 in the *Pcn* parameter of the DEC/DLD command. Thus, to downline load any of these characters, you subtract 32 from the ASCII code (both in decimal) and use the result* as the *Pcn* parameter. The first sixel bit pattern, *Sxbp1*, loads this character. Each subsequent sixel bit pattern *Sxbp2* through *Sxbpn* loads characters *Pcn+1* through *Pcn+n-1*. More about sixels later.

VT220 characters occupy a 7 column by 10 row matrix. This matrix looks like the one shown in Figure 1. Characters are composed by "turning on" cells in the matrix, which produces bright spots in the corresponding locations on the screen.




Figure 1. Basic Matrix




Figure 2. Pattern for Sigma

Let us define a special character, the capital Greek letter sigma. The layout for this character is shown in figure 2 above. The cells filled with "X" represent bright spots in the character pattern.

To load this pattern into the VT220 via the DEC/DLD command, it must be converted to a series of ASCII characters. The first

* As with all ANSI escape sequences, this parameter is the ASCII string representing the result, not the character whose ASCII code is the result. In other words, if we want to load a character for decimal ASCII code 92, the *Pcn* parameter is the two character ASCII string "60".

step in the conversion is to add two dummy rows at the bottom and then divide the 7 x 12 result into an upper half and a lower half.

We also must sequentially number the rows in each group from top to bottom, starting with zero. See Figure 3 below:




Figure 3. Sixel Generation Matrix

The resulting character cell is divided into two individual character cells for the next step in the process.

We now take each six-row column, rotate it clockwise by 90 degrees, and substitute a "1" for each illuminated dot, and a "0" for each dark one. The row numbers we assigned in the previous step are used to represent bit positions.

The result is a set of fourteen six-bit binary numbers corresponding to illuminated and dark dots in the character cell. (By now you should be getting an idea of where the term "sixel" comes from.)

```

  5  4  3  2  1  0  5  4  3  2  1  0
+---+---+---+---+---+---+
| | | | | X | | 0 0 0 0 1 0
+---+---+---+---+---+---+
| | | | X X | | 0 0 0 1 1 0
+---+---+---+---+---+---+
| X | | X | | X | | 1 0 1 0 1 0
+---+---+---+---+---+---+
| | X | | | X | | 0 1 0 0 1 0
+---+---+---+---+---+---+
| | | | | X | | 0 0 0 0 1 0
+---+---+---+---+---+---+
| | | | X X | | 0 0 0 1 1 0
+---+---+---+---+---+---+

```

Figure 4. Sixel Generation Matrix, Upper Half

```

  5  4  3  2  1  0  5  4  3  2  1  0
+---+---+---+---+---+---+
| | | | | X | | 0 0 0 0 1 0
+---+---+---+---+---+---+
| | | | X X | | 0 0 0 1 1 0
+---+---+---+---+---+---+
| X | | X | | X | | 1 0 1 0 1 0
+---+---+---+---+---+---+
| | X | | | X | | 0 1 0 0 1 0
+---+---+---+---+---+---+
| | | | | X | | 0 0 0 0 1 0
+---+---+---+---+---+---+
| | | | X X | | 0 0 0 1 1 0
+---+---+---+---+---+---+

```

Figure 6. Sixel Generation Matrix, Lower Half

At this point the process is almost done. The next step is to add a binary 111111, octal 77, to each of these six-bit numbers. The result of each addition is interpreted as an ASCII code.

```

000010 + 111111 = 1000001 = A
000110 + 111111 = 1000101 = E
101010 + 111111 = 1101001 = i
010010 + 111111 = 1010001 = Q
000010 + 111111 = 1000001 = A
000010 + 111111 = 1000001 = A
000110 + 111111 = 1000101 = E
000010 + 111111 = 1000001 = A
000011 + 111111 = 1000011 = B
000010 + 111111 = 1000001 = A
000010 + 111111 = 1000001 = A
000010 + 111111 = 1000001 = A
000011 + 111111 = 1000011 = B

```

Figure 7. ASCII Lower Half

Figure 8. ASCII Upper Half

The final step is inserting these ASCII characters into the DECtLD command string. Assume that we want to use ASCII code 92, octal 134, to display sigma, and that the name to be assigned to this character set is "\$G". (Refer to page 19 of the [Guide](#) for more information on character set names.)

The resulting escape sequence is:

```

DCS 0;60;1;0;1;0;$G AEiQAAe/ABAAAAB ST
| | | | | | | | | | +--- ASCII 156 or ESC \
| | | | | | | | +----- Sxbsp1: Sixel bit pattern
| | | | | | | | +----- Dscs: "Name" of font
| | | | | | | | +----- Pt: Device default
| | | | | | | | +----- Pw: 80 column
| | | | | | | | +----- Pcms: Matrix size (default)
| | | | | | | | +----- Pe: Erase only this char
| | | | | | | | +----- Pcn: Starting char number 60
| | | | | | | | +----- Pfn: Font 0
+----- ASCII 144 or ESC P

```

Figure 9. Font Definition Escape Sequence

It should be noted that although the command string is shown with spaces embedded for the sake of clarity, spaces should not be included in the actual command. In commands where a space character is required, "sp" is used to indicate the location. This applies to all command strings shown in this article and in the [Guide](#).

If it were also desired to define ASCII 93, octal 135, "]" in the character set, simply append a semicolon and the sixel definition of the new character to the *Sxbsp1* segment of the command string as follows. What will be displayed by *Sxbsp2* is left as an exercise for the reader:

```
DCS 0;60;1;0;1;0;$G AEiQAAe/ABAAAAB;wCAAqAC/?@AANIH ST
```

To display the new characters, the information on page 9 of the [Guide](#) is used. The character set name (\$G) is designated as

character set G0 and the appropriate ASCII codes are sent to the terminal.

By way of example, the command:

ESC (\$G\]

designates \$G as the G0 character set (ESC \$G) and then displays the two characters just defined (\]).

To return to the normal ASCII character set, use the command:

ESC (B

That's basically all there is to it. As for the *Pfn* parameter, the Manual says: "The VT220 has only one DRCS font buffer. This parameter has two valid values: 0 and 1."

Why there are two valid values for the parameter when there is only one font buffer, or why there is even a parameter at all, is never explained. Presumably this is reserved for compatibility with future terminals. The author has tried both values; they appear to behave identically.

The *Pcms* and *Pw* parameters are similar; various combinations of the specified values have been tried, including *Pcms* = 1 (allegedly not used). No difference in the display was noted.

The *Pe* parameter does appear to do just exactly what the Guide says it does, though exhaustive testing was not attempted.

The Manual provides the following information about the *Pt* parameter: "Full cell fonts can individually address all pixels in a cell, while text fonts, in general, may not be able to address all pixels individually."

Setting this parameter to 2 (full cell) seems to disable loading a character; the other two values appear to be interchangeable.

As most of us who have been around DEC equipment and software know, DEC has a curious habit of including "undocumented" features (probably bugs they couldn't get rid of) in their products. The VT100 is loaded with them; try playing with the LEDs or send it a control-Z character sometime.

The VT220 is no exception. The Manual sort of hints that there might be something more to the DECDLD command than meets the eye. It says that the "DRCS cell size" is an 8 x 10 array, but the "normal terminal character cell size" is only 7 x 10. It

even includes an illustration showing the difference.

If an eighth column is defined for each character, the terminal accepts the additional information, squirrels it away somewhere in its innards, and then does strange things when that character is requested for display. Extrapolating from known information, one might expect that this additional column represents the gap between two adjacent characters and that dots corresponding to 1s should be illuminated.

This does indeed happen, but in addition this dot pattern is inclusive ORed with the first column of the following character. For example, the command:

DCS 0;60;1;0;1;0{\$G AEIQAAEO/ABAAAAB? ST ESC (\$G\ ESC (B#

produces the following display:




Figure 10. Eighth Bit Setting Effect

One final note: If a particular character is redefined while that character is displayed on the screen, every occurrence of that character is replaced by the new character as soon as the new definition is loaded into the terminal font memory.

Weird RSX Science

Alan Frisbie
Flying Disk Systems
4759 Round Top Drive
Los Angeles, CA 90065

Transparent spooling (or, as some implementers refer to it, "translucent spooling") allows the user to do some apparently strange things with the system.

When a file is opened on a spooled device, what actually happens is that a workfile is created. When the file is closed, the file is automatically spooled to the device and (normally) deleted after spooling.

For example, if TT1: is spooled, and one opens TT1:filename, writes it and closes it, one copy of the file is printed on TT1: and the file disappears.

This leads me to an interesting interaction between EDT and spooling. Suppose that I want to extract internal documentation from a source file and print it on a spooled printer. The source file is FLZOPN.MAC, and most of it is code I don't care to see. I could, of course, do:

```
EDT TEMP.TMP=FLZOPN.MAC
```

wipe all but the documentation, exit the editor, and then enter


```
PRINT TT1:=TEMP.TMP/DE
```

but that seems like a waste of a perfectly good command line, and a long and complex one to type in at that.

It's so much simpler to just edit the printer directly. I merely type in the magic command line:

```
EDT TT1:=FLZOPN.MAC
```


which has the desired effect of letting me edit the source in a temporary file, print the resulting temporary file, and delete it after printing - all in one line.



THE
mini-
tasker

DECUS

RT-11 SIG NEWSLETTER



RT-11 MINITASKER

November, 1987

A Symposium Paper? Who Me?

by
Milton Campbell

From the Editor:

Those of you who are sick-and-tired of my endless pleas for Minitasker submissions may skip ahead to the Table of Contents. The rest of you, listen up!! As of the writing of this paragraph, my "December Minitasker" folder is empty. You have fair warning, if I don't get some stuff to publish, the next issue may be filled with "My Favorite TECO Macros." So if you don't want to be subjected to sixteen pages of line noise, send your submissions to:

John M. Crowell
RT-11 Newsletter Editor
Multiware, Inc.
2121-B Second St. Suite 107
Davis, CA 95616

CONTENTS

2. From the Editor	RT-2
3. A Symposium Paper? Who ME?	RT-3
4. RT-11 Abstracts	RT-4
5. RT-11 SIG Tapes on RX33	RT-17
6. FPJ chip fix	RT-18
7. Directory Search Program	RT-19
8. SCURT Announcement	RT-20
9. A Simple Time Card System	RT-21
10. The Flit Gun	RT-35

I realize that you will be reading this before the Symposium in Anaheim this December, but it is not too early to think about presenting a session at the Spring 1988 Symposium (May 16 through May 20) in Cincinnati. The deadline for abstract submission ("Calls for Participation" or "CFPs" in DECUS-speak) is December 14, 1987. That is the Monday after the Fall Symposium.

There are a number of advantages to being a speaker a DECUS Symposium. The most tangible is probably that it makes it easier to justify the cost of attending the Symposium. If you work for a large organization, speaking at a technical conference often greases the skids of the approval process. Speaking is also a way to announce your expertise in a subject. This has benefits for companies large and small, as well as for individuals. Other obvious advantages are being allowed into the Sunday reception before the crowd and the green speaker ribbon to wear on your Symposium badge.

One of the hard to quantify, but very important, reasons for going to a symposium is meeting people with similar interests and problems. By being a speaker you enhance this effect by making some of your interests known to a larger audience. This informal part of the conference probably provides the most long term benefits of attending, so increasing these interactions can be well worth the effort of preparing a talk.

Note that the submission of an abstract does not imply that you already have your talk complete. Except for sessions that are repeats from previous symposia, most talks are probably completed in the last few weeks before the symposium. The abstract does mean that you have an idea of what to talk about and that you can describe it.

Once an abstract gets to DECUS, it is entered into a computer where the scheduling process begins. The various SIG Symposia Coordinators review the abstracts and make sure that they are assigned to the right SIG and that they do not duplicate some other abstract. Eventually, the time comes to assign the sessions to actual times and rooms. Before I do this for the RT-11 sessions, I call each of the speakers to make sure they still plan to give the session. Up until you say "yes" to this phone call, you can drop the session without any problems. For Cincinnati, I will be calling speakers around the third week in January. This is the actual "commitment" deadline.

The upshot of all this is that you have up to mid-January to decide not to give a session, but only if you have submitted an abstract by the December deadline.

There is always the question of what to talk about. The best answer is often to talk about some aspect of what you do on your job. Note that sessions do not need to address some deep technical aspect of RT-11 (or whatever the subject). There is a wide diversity of interests in DECUS attendees and we need sessions to match. If you would like to give a session but are having trouble finding a topic, give me a call at (213)318-2206. I can help you select something from your job, or I can suggest some other ideas.

RT-11 ABSTRACTS

submitted by
Milton Campbell
RT-11 SIG Symposium Coordinator

The following are the session abstracts for the sessions sponsored by the RT-11 SIG at the Fall 1987 Symposium in Anaheim, California.

RT001 RT-11 USERS SPEAK OUT

Time: Thursday Room: Avila
9:00-11:00 p.m. Anaheim Hilton

Chair: John M. Crowell
Multiware, Inc.

Abstract:

A panel of alleged RT-11 experts conducts a program of fun, history, war stories, and technical information which is unavailable from any other source for all RT-11 users. Audience participation is encouraged and expected.

The panel is prepared to answer questions from the floor to the best of their ability. Scheduled presentations may include slide shows of historic RT-11 installations, lessons on how to knot flat ribbon cables, minimal keystroke methods of destroying your system disk, magical chants to revive down systems, and field circus approved locations on the computer which improve system performance when direct force is applied.

The highlight of the evening may be the semi-annual awards presentation for most sensitive response to a user question.

All in all, this is a fun and informative evening. The only rule is that you cannot ask a question the panel cannot answer (of course, you cannot know this until you ask the question). Proper dress is requested; costumes are optional.

RT002 ADEP: SOFTWARE FOR REAL TIME DATA ACQUISITION AND CONTROL

Time: Tuesday Room: Palos Verdes
9:00-10:00 a.m. Anaheim Hilton

Spkrs: Ellen Bachmann Chair: Bruce Sidlinger
Monsanto Research Corp. Sidlinger Computer Corp.
Jim Lindesmith
Monsanto Research Corp.

Abstract:

This session describes the content, use, and rationale of the Applications Development Package (ADEP). ADEP is a generalized set of software tools which aids in the efficient implementation of laboratory data acquisition and control systems. It consists of an integrated package of handlers, utilities and standards for data acquisition, control, listing and graphics.

ADEP is divided into three parts. Section I consists of the basic utilities and software to initialize data files, control the data acquisition process, acquire data and place it in a given file. The software includes an RT-11 device handler which provides the control, sequencing and data acquisition functionality, and a set of basic support utilities for building event queues, channel lists, file headers and archival data files.

Section II is a utility to access and list the acquisition parameters and data in a report format. Section III includes the basic utilities and software to define graphic output requirements, create XY plots and output plots to a number of graphic devices. In general, the ADEP routines are used either as is, or as the functional core of a more complex, customized system.

RT003 BASIC-PLUS/RT V3.1 AND FUTURES

Time: Monday Room: Santa Monica
10:00-10:30 a.m. Anaheim Hilton
Spkr: Peter Hosford Chair: John Davis
Digital Equipment Corp. Naval Ship Research Center

Abstract:

This session presents plans for future releases of BASIC-PLUS/RT-11. Special attention is given to features which are included in the upcoming release, V3.1. The most prominent of these is "Language Extension", a feature which allows users to define new keywords and statements in their copies of the BASIC-PLUS language, and to access user-written MACRO subroutines from BASIC-PLUS.

Features being considered for releases beyond V3.1 are also discussed. Among these are mechanisms for in-place editing (editing BASIC-PLUS programs without leaving the BASIC environment) and for passing data across a CHAIN.

Audience questions and feedback are welcome.

RT005 TSX-PLUS MAGIC

Time: Wednesday 5:00-6:00 p.m. **Room:** El Capitan Anaheim Hilton
Chair: Tim Clarke
Omnex Corporation

Abstract:

A panel of experienced TSX-PLUS users, system managers and system programmers are on hand to assist users in making more effective use of their TSX-PLUS systems. Some brief presentations of special techniques, utilities, handlers, command files, and programs may be made by panel members, but most of the session is oriented toward audience questions, problems, solutions, and wishlist items.

All TSX-PLUS users are encouraged to attend. This is your chance to get an answer to that elusive problem, to learn how others have made their systems better, and to share the knowledge you have gained while using TSX-PLUS.

RT006 CONFIGURING A TSX-PLUS SYSTEM

Time: Wednesday 3:00-3:45 p.m. **Room:** El Capitan Anaheim Hilton
Spkr: Greg Adams
GABA **Chair:** Carl Linden

Abstract:

TSX-PLUS is easy to install, but there is a difference between "installing" and "configuring" a system. Because there is no documented standard for configuring TSX-PLUS, environments tend to be either nightmarish for the manager or over restrictive for the user.

This session discusses methods of configuring TSX-PLUS to provide a balance between restriction and usability while at the same time improving performance and generally helping to keep things tidy.

RT010 RUNNING TSX-PLUS ON A Q-BUS CO-PROCESSOR

Time: Wednesday 3:45-4:30 p.m. **Room:** El Capitan Anaheim Hilton
Spkr: Milton Campbell
Talisman Systems **Chair:** Ned Rhodes
Software Systems Group

Abstract:

The availability of several J-11 based co-processors for the Q-bus has opened up some interesting possibilities. This session discusses some of the considerations in getting TSX-PLUS running on the MDB JFEP-11 co-processor. Included in the presentation will be a discussion of how to use the co-processor with a TSX-PLUS host, some of the problems involved and what can be done about them. The use of a TSX-PLUS and co-processor combination with a MicroVAX/VMS host will also be discussed.

RT012 AN OVERLAY METHOD TO MAXIMIZE PROGRAM EXPANDABILITY

Time: Tuesday 10:00-11:30 a.m. **Room:** Palos Verdes Anaheim Hilton
Spkrs: Marshall Lajoie
LTV Missiles & Electronics **Chair:** Robert Roddy
Naval Ship Research Center
Jim H. Nicholson
LTV Missiles & Electronics

Abstract:

This session presents a software structure and overlay method allowing virtually unlimited expansion by allowing a program segment to overlay the segment invoking it. The program example is a test control program used in acceptance testing systems containing an embedded processor. Primary objectives were: simple user interface, large expansion capability, and ease of maintenance while limited to 64kb addressing. Developed under RSX-11M, written in FORTRAN-77 and MACRO-11 for a 248kb RT-11 environment, the method is applicable to both operating systems. Topics include: Program structure, Overlay structure and methods, Maintenance, DMA I/O, A/D & D/A conversion, Interrupts, Linearization of non-linear transducer data, Extended memory usage under RT-11XM, Example program features.

RT013 RUNNING RT-11 PROGRAMS ON KXT11-C AND KXJ11-C I/O PROCESSORS

Time: Tuesday 11:30-12:30 p.m. **Room:** Palos Verdes Anaheim Hilton
Spkr: John M. Crowell
Multiware, Inc. **Chair:** Laura DeChellis-Barry
MDB Systems, Inc.

Abstract:

This session presents a method for running RT-11 programs on the KXT11-C and KXJ11-C I/O processors (IOP). Programs which do not require any system services can run unattended on the IOP. A program which downloads a .SAV image onto the IOP and then jump starts the IOP is described. The host processor may be either a Q-bus PDP-11 running RT-11 or a MicroVAX running MicroVMS. Problems of avoiding the IOP firmware's stack and distinguishing the characteristics of the two types of IOPs are discussed as well as suggested methods for communication between a program running on an IOP and programs running on the host processor.

RT014 RT-11 FORTRAN AND SYSLIB INTERACTIONS

Time: Monday 3:00-4:00 p.m. Room: Malibu Anaheim Hilton
Spkr: Ned W. Rhodes Software Systems Group Chair: John Davis Naval Ship Research Center

Abstract:

This session introduces the use of the RT-11 System Library and how it can be used with FORTRAN programs. Example programs are given which demonstrate the PEEK, POKE, ISLEEP and TIME/DATE routines. In addition, the use of the library string routines and the RT-11 I/O routines are shown.

RT015 RT-11 TO VMS MIGRATION

Time: Thursday 7:00-8:00 p.m. Room: Avila Anaheim Hilton
Spkr: Ned W. Rhodes Software Systems Group Chair: Shal Farley Cheshire Engineering Corp.

Abstract:

Many RT-11 users have the opportunity to develop code on VAX/VMS systems, but many die-hard RT-11 users are resisting because they perceive that VMS does not allow them to do the same types of things that they are used to under RT-11. This session shows RT-11 programmers how to develop programs under VMS using the concepts that were learned under RT-11. Some of the key differences and similarities between the two operating systems are shown.

RT017 FORTRAN-77/RT PROGRAMMING STYLE

Time: Wednesday 1:00-2:00 p.m. Room: El Capitan Anaheim Hilton
Spkr: Robert Walraven Multiware, Inc. Chair: Ralston Barnard Sandia Labs

Abstract:

Programming style can affect the readability and maintainability of your code. This session does not tell you what style you should use (because there are no absolute answers), but does present techniques that can be effective for FORTRAN-77. Ample time is planned for questions, answers, and comments.

RT018 THE RT-11 MARKET IN EUROPE

Time: Monday 6:00-6:30 p.m. Room: Malibu Anaheim Hilton
Spkr: Robert Walraven Multiware, Inc. Chair: Jim Crapuchettes Omnex Corp.

Abstract:

European countries provide a large customer base for RT-11. This session discusses how RT-11 (and other) products are typically marketed in Europe, what the RT-11 customer base is like, and what DECUS in Europe is like. Time is given at the end for questions, answers, and comments.

RT019 TIMING STUDIES ON RT-11 FILE I/O

Time: Thursday 5:00-6:00 p.m. Room: Palisades Anaheim Hilton
Spkr: Selina Tourjee Multiware, Inc. Chair: Robert Peckham Computer Programming Services

Abstract:

This session describes file I/O timing studies conducted for RT-11. The speeds of SYSLIB I/O functions and FORTRAN read/write statements are compared for a variety of file-structured devices and a selection of CPUs. The effect of the type of I/O on real-time applications are discussed, and examples of programming style are presented.

RT020 **RT-11 LANGUAGE AND LAYERED PRODUCT PANEL**
Time: Monday Room: Santa Monica
 11:30-12:00 noon Anaheim Hilton
Spkr: PDP-11 Languages Chair: Gary Sallee
 Digital Equipment Corp. Sallee Software

Abstract:

This session presents the status of Digital-supported languages and layered products under the RT-11 operating system.

RT021 **RT-11 PRODUCT PANEL**
Time: Monday Room: Santa Monica
 10:30-11:30 a.m. Anaheim Hilton
Spkr: RT-11 Engineer Chair: Gary Sallee
 Digital Equipment Corp. Sallee Software

Abstract:

This session presents an overview of RT-11 Engineering, the future direction of RT-11 and how Digital's licensing policies apply to RT-11.

RT022 **RT-11 HANDLERS IN XM**
Time: Monday Room: Malibu
 6:30-7:30 p.m. Anaheim Hilton
Spkr: RT-11 Engineer Chair: John Rose
 Digital Equipment Corp. Omnex Corp.

Abstract:

This session presents the advanced techniques that are used by RT-11 Engineering to reduce low memory usage of large handlers in an extended memory (XM) environment. The following topics are covered:

- o the use of GLOBAL REGIONS to hold code and data;
- o the use of \$REL MACRO; and
- o the use of non-memory resident one-time code

Many of the techniques presented are applicable to single job (SJ) and foreground background (FB).

RT026 **RT-11 AND BIG DISKS**
Time: Friday Room: Palos Verdes
 9:00-10:00 a.m. Anaheim Hilton
Spkr: RT-11 Engineer Chair: John M. Crowell
 Digital Equipment Corp. Multiware, Inc.

Abstract:

With the advent of the Mass Storage Control Protocol (MSCP) controllers, RT-11 supports large capacity disks. Due to some limitations within the operating system, programs cannot directly access all the space on these large disks. In order to access the full disk capacity, the RT-11 handlers have to use the concept of disk partitioning.

This session answers the following questions:

- o How is partitioning used?
 - o How are JREAD/JWRITE used?
 - o How is the BUP backup/recovery utility used with disks that are larger than 65K blocks?
-

RT027 **RT-11 FEEDBACK SESSION**
Time: Friday Room: Palos Verdes
 12:00 noon-1:00 p.m. Anaheim Hilton
Spkr: RT-11 Engineer Chair: Bradford Lubell
 Digital Equipment Corp. L.A. Heart Lab, UCLA

Abstract:

In this session RT-11 Engineering reviews the customer wishlist items which accrued during the week and since the previous DECUS. Wishlist items can be deposited in the wishlist box located in the booth area next to the RT-11 Demo System or can be given to an RT-11 SIG member or representative from RT-11 Engineering.

RT034 **REAL WORLD DISK COMPARISONS**
Time: Friday Room: Palos Verdes
 10:00-11:00 a.m. Anaheim Hilton
Spkr: Robert Peckham Chair: Selina Tourjee
 Computer Programming Services Multiware, Inc.

Abstract:

Many computer users are interested in the actual data transfer rates achieved when real controllers and disks operate with a real

operating system, doing real data transfers, as compared with the data transfer rates claimed in manufacturers' literature. This session presents the results of running a series of test programs to exercise the various operational parameters of a disk under real-world situations. The programs were run under RT-11 and TSX-PLUS on a wide variety of disks by about twenty DEC end-user sites. Disk devices tested ranged from RX01 thru Winchester and memory disks, and even included an Ethernet virtual disk. The results are presented in tabular form so that direct comparison is possible.

RT035 RT-11 SIG BUSINESS MEETING

Time: Monday 9:00-9:30 a.m. **Room:** Santa Monica Anaheim Hilton

Spkr: John Rasted
JTR Associates

Abstract:

This session begins with an overview of the RT-11 Special Interest Group (SIG), followed by SIG activity at the symposium and those areas of SIG activity which are not related to the symposium. These areas include:

- o Minitasker (the SIG Newsletter);
- o SIG tape copy;
- o SIG DECUS Library activity;
- o Local User Groups (LUGs); and
- o VAX/RT.

In this session, the SIG also begins the planning for the next DECUS symposium.

RT036 RT-11 SIG ROADMAP

Time: Monday 9:30-10:00 a.m. **Room:** Santa Monica Anaheim Hilton

Spkr: John Rasted
JTR Associates

Abstract:

This session is designed to help the attendee obtain the most benefit from the symposium. Veteran attendees discuss the tried and true techniques that new attendees can use to make the most of the week and still survive the experience. There is a brief description of those sessions which are relevant to RT-11 users. Schedule changes and possible session repeats are also discussed. Plan to attend so you can avoid the disappointment of missing an important session.

RT037 RT-11 SIG SYMPOSIUM WRAP-UP

Time: Friday 1:00-1:30 p.m. **Room:** Palos Verdes Anaheim Hilton

Spkr: John Rasted
JTR Associates

Abstract:

This is your chance to respond to the SIG and Digital presentations at the symposium and to influence future plans. The SIG is looking for input from the attendees to aid in selecting desirable sessions for the next symposium. At this session you have the opportunity to have questions answered that may have arisen during the symposium. Representatives from Digital are also present.

Topics Include:

- o SIG activities;
- o RT-11 and layered products;
- o Pre-symposia Seminars; and
- o Future DECUS symposia.

RT038 USING COBOL-PLUS

Time: Wednesday 2:00-3:00 p.m. **Room:** El Capitan Anaheim Hilton

Spkr: Laura DeChellis-Barry
MDB Systems Inc. **Chair:** Bill Leroy
The Software House, Inc.

Abstract:

This session will discuss the use of S&H Computer Systems' COBOL-PLUS for applications development. Discussion includes features, limitations, hints and a wishlist. Other users are invited to add their input in an open discussion.

Topics include, but are not limited to:

- o Code generation
- o Debugging tools
- o RTSORT (high speed sort facility)

RT039 **RUNNING IN RT-11**

Time: Monday Room: Malibu
 4:00-5:00 p.m. Anaheim Hilton

Spkr: RT-11 Engineer Chair: Nick Bourgeois
 Digital Equipment Corp. NAB Software Services

Abstract:

This presentation covers the different types of jobs that RT-11 allows and some implications of using each. File formats, job loading, memory utilization (including overlays), job priorities, context switching, and I/O restrictions are discussed. The use of VBGEEXE to allow XM jobs to avoid memory constraints along with constraints attendant to using VBGEEXE are included in this presentation.

RT040 **RT-11 MAGNETIC TAPE USAGE**

Time: Friday Room: Palos Verdes
 11:00-12:00 noon Anaheim Hilton

Spkr: RT-11 Engineer Chair: Nick Bourgeois
 Digital Equipment Corp. NAB Software Services

Abstract:

This talk includes a discussion of how the RT-11 operating system interfaces to various magnetic tape devices. An overview of tape handlers, tape-related SPFUNs, as well as PIP and BUP tape formats are presented.

RT041 **RT-11 RUNNING ON THE KXJ**

Time: Tuesday Room: Palos Verdes
 12:30-1:30 p.m. Anaheim Hilton

Spkr: RT-11 Engineer Chair: Robert Walraven
 Digital Equipment Corp. Multiware, Inc.

Abstract:

This session discusses, in general terms, the support of KXJ11-CAs by RT-11. RT-11 will run on one or more KXJs on a single QBUS system. The RT-11 systems can be run effectively "stand-alone" on the KXJs, after "booting" from the host processor, or can be run in an "RTEM-like" (closely coupled) relationship with the host system. The types of services available among KXJs and between a KXJ and the host processor are discussed. These include file, mailbox and event flag services. Some support for peripherals local to the KXJ will be discussed. These include the parallel port, the serial ports and the DMA engine.

RT042 **USING KERMIT WITH TSX-PLUS**

Time: Wednesday Room: El Capitan
 4:30-5:00 p.m. Anaheim Hilton

Spkr: Tim Clarke Chair: Jim Crapuchettes
 Omnex Corporation Omnex Corp.

Abstract:

This session is a tutorial on how to use Kermit with TSX-PLUS. While the issues of modem, modem control, and serial line interfaces are treated the main emphasis is on how to set up the TSX-PLUS environment. Subjects included are:

- o 'CL' lines
- o 'DTR' control
- o Use of command files to set up the TSX environment

There is a question and answer session at the end. Users that currently have questions or problems that may require testing are invited to submit them to the speaker in advance.

RT043 **USING THE RT-11 ETHERNET HANDLERS**

Time: Monday Room: Malibu
 5:00-5:30 p.m. Anaheim Hilton

Spkr: Jim Crapuchettes Chair: Dennis Jensen
 Omnex Corporation AMES Labs. ISU/USDOE

Abstract:

This session discusses use of the RT-11 Ethernet handlers NQ, NC and NU, with primary emphasis on the NQ handler. The session begins with a brief Ethernet overview. Then the available handlers and their associated hardware are identified. The session continues with a discussion of each of the handler special functions, some examples of calling sequences for these functions and some possible pitfalls in usage.

RT044 **PASCAL/MACRO-11 TECHNIQUES USED IN IMPLEMENTING AN RT-11/
TSX-PLUS ETHERNET DISK SERVER**

Time: Monday Room: Malibu
 5:30-6:00 p.m. Anaheim Hilton

Spkr: John R. Rose Chair: Dennis Jensen
 Omnex Corporation AMES Labs. ISU/USDOE

Abstract:

EtherDisk is a software system that provides RT-11/TSX-PLUS virtual disks across an Ethernet network, using DEC DEQNA interfaces and

emulators. EtherDisk was coded mainly in Pascal-2 (TM Oregon Software, Inc.), with hardware-control routines in Macro-11.

This session describes some of the design and implementation features of EtherDisk. Topics include modular coding techniques, Pascal linked lists, Pascal calls to Macro-11, Macro-11 calls to Pascal and Pascal completion routines. Some familiarity with Pascal and Macro-11 are helpful but not required.

RT045 **NUGGETS AND GEMS FROM THE WORLD OF RT-11**

Time: Thursday Room: Avila
 8:00-9:00 p.m. Anaheim Hilton

Chair: Ralston Barnard
 Sandia National Laboratories

Abstract:

This is a chance for RT-11 "experts" to describe some of their tricks and time-saving techniques. It is also a chance for anyone to ask the experts for suggestions on problems they may have with their systems. Some of the topics which may be covered are:

- o How to set up a print spooler
- o How to use Kermit
- o Getting the most from UCL
- o Hints of how to connect peripherals
- o Why use XM
- o Making a useful XM job run

RT046 **ANALYSIS OF RT-11/TSX-PLUS SYSTEM OVERHEAD**

Time: Thursday Room: Palisades
 6:00-7:00 p.m. Anaheim Hilton

Spkr: Jim Crapuchettes Chair: Jim Lindesmith
 Omnex Corporation Monsanto Research Corp.

Abstract:

Using a programmable clock, the system execution time was measured for I/O and some other system service calls for the various RT-11 monitors and for TSX-PLUS. In each case, the measured times for 10,000 system calls were recorded and histogrammed. Because the purpose of these tests was to measure system overhead, the I/O calls were primarily to the NL:, VM: and QM: handlers so that the effects of device latency would not be included in the timings. For the same reason, the other system calls timed were primarily those which did no I/O.

This session presents the results of these timing tests and discusses the implications of the results on system throughput.

John M. Crowell
RT-11 Newsletter Editor
c/o Multiware, Inc.
2121-B Second St., Suite 107
Davis, California 95616

September 4, 1987

Dear John:

I have always had difficulty obtaining RT-11 Sig swap tapes because I have not had access to a tape drive. Certainly other RT-11 users have had the same problem. There has been great resistance to the thought of distributing rx02/rx01 floppies because of the number of disks needed. The online distribution was great while it lasted (albeit slow and expensive). However, the advent of the rx33 drive has made distribution on floppies feasible.

I have received from Ralston Barnard most of the Spring 1977 RT-11 Sig swap tape on four rx33 disks (8897 blocks). In return I have promised to copy these for anyone who complies with the following instructions: send me five (5) rx33 floppies which have been formatted and initialized with an addressed (preferably stamped) mailer. I will copy the files and return the floppies in the same mailer. Hopefully this will be a useful service and can be continued with future swap tapes.

Anyone who doesn't have an rx33 drive, should consider getting one. If one has a ba23 or bal23 and an rdx3 already, the expense (by DEC standards) is almost trivial. These drives are small, quiet, cheap, several times faster than an rx02, hold 2 1/2 times as much as an rx02, and are hardware compatible with IBM PC (AT) type floppies. Mine even works.

Sincerely,

Harold Z. Bencowitz
810 Hospital Drive, 240
Beaumont, Texas 77701



EYE & EAR HOSPITAL OF PITTSBURGH
230 Lothrop Street Pittsburgh, Pennsylvania 15213-2592

September 2, 1987

John M. Crowell
RT-11 Newsletter Editor
Multiware, Inc.
2121-B Second St., Suite 107
Davis, CA 95616

Dear Mr. Crowell,


I notice in the September Minitasker that DEC still hadn't gotten back to you on the FPJ-11 bug I reported. Thanks to your mention of an ECO in the August issue, I got the DEC offices in both New York and Pittsburgh to work on the problem.

According to Pittsburgh, a Field Change Order, KDJ-11-AC-R001, has been issued to correct both the FPJ-11 problem and an associated FPJ-11/DMA problem with the dual-width 11/73 processor board. This is a "required" FCO, so the fix is free (though DEC may charge to actually "install" the board if you do not have a DEC service contract).

The part number for this kit, which contains a new KDJ-11 board (at rev level D) and a new FPJ-11 chip (version 5), plus new manuals, is EQ-01447-02. The New York office sent such a kit to me -- since I do not have a DEC service contract, I had to do the board swap myself to avoid having to pay for a service call.

Thanks for helping to notify the DEC user community of this problem. Also, thanks to DEC for finally providing a fix. I wonder how many users out there know about either the problem, or the solution? How does one find out about "designed-in bugs"? (Detroit sends out recall letters to all those who bought seriously-flawed products).

Sincerely,


Robert H. Schor

 OSTERGAARD ASSOCIATES
CONSULTANTS IN ACOUSTICS

24 July 1987

Mr. John M. Crowell
RT-11 Newsletter
Multiware, Inc.
2121-B Second Street, Suite 107
Davis, CA 95616

Dear Mr. Crowell:

For the last five months I have been, by default, handling our RT-11/TSX system. Our experts left or were "let go" so I'm trying to learn more of RT-11 and restrain efforts to go to all on a PC system without carefully examining what we have.

I don't get the RT-11 MINITASKER but I did get a copy in the DECUS SIG NEWSLETTER. In it I saw the RT-11 WISH LIST.

Under 3.3, DIR in the WISH LIST item g. wants some way to search for files in logical devices. Maybe we can go you one better with our HUNT. For example, HUNT DR*:*.FOR will search for all .FOR files on DRO, DR1, DR2, DR2 and in all LD's on all four disks. HUNT DR1:*.FOR will find all .FOR files on DR1 and in all LD's on DR1. The program returns to the screen, the device where the file is located, the logical disk, if any, the file name, number of blocks and date of creation.

Interested? If so, what would you like from me?

The rest of the WISH LIST, odds and ends, are interesting but I'm not sophisticated enough with the system to have too many needs.

Cordially yours,

OSTERGAARD ASSOCIATES


Paul B. Ostergaard, P.E.

PBO:kk

Editor's Note: I've asked Dr. Ostergaard to submit his program to the DECUS library. If you're desperate, you may be able to get an advanced copy from him.

SCURT
Southern California Users of RT-11
Digital Equipment Users Society

A SIMPLE TIME CARD SYSTEM

Mike LeRoy
4909 Sea Wolf Drive
Santa Rosa, CA 95405

***** SPECIAL SYMPOSIUM MEETING *****

To: All RT-11 Users and Other Interested Parties (TSX+)
Subject: Special Meeting of SCURT at the Anaheim Symposium
Date & Time: December 8, 1987 Tuesday at 3 to 5 p.m.
Location: RT-11 SIG Suite, Anaheim Hilton Hotel

***** RT-11 SOFTWARE WORKSHOP *****

Attendees may ask for help, or discuss problems they solved. Regular SCURT members will be on hand to answer questions, and give examples. A DEC RT-11 representative will be there. You do not need to be registered at the Symposium to attend. You may come to see DEXPO, sample the Symposium atmosphere, and attend an important SCURT LUG meeting, all in one day. Ask for directions to the RT SIG Suite at the General Information Booth.

Come to meet the people of your local RT-11 support group.

***** REGULAR SCURT MEETING DATES AT CALTEC PASADENA *****

***** FOURTH TUESDAYS AT 10:00 AM *****

December 15, 1987: RT-11 Languages, a view from DEC

January 26, 1988: PDP-11 to IBM PC communication networks

February 23, 1988:

***** FOR INFORMATION CALL *****

Chairman: Shal Farley (818) 351-5493 Pasadena
Program Chairman: Gary Sallee (714) 970-2864 Yorba Linda
Membership Chairman: Greg Adams (805) 296-0170 Sagus
Tape Librarians: Lisa Schultz (818) 358-1871 Monrovia
Daniel Fishman (818) 358-1871 Monrovia

Is there anyone out there that needs a simple time card system? Oh, well, I decided that I needed one and my system could handle it a lot nicer than I did. I've included three versions: Fortran-77 V5.0a, DISC's DBL version 2.2 and Oregon Software's Pascal V1.2 (Sorry not in TECO). It runs on TSX-Plus version 6.2.

There are three commands to this system: TIMIN, TIMEOUT, and CALTIM.

TIMIN job<cr> - starts a new job, also logs out previous job.
job could be up to 6 word phrase or connected with dashes like Big-Payroll-for-Bert

TIMOUT<cr> - logs out current job

CALTIM<cr> - calculates totals for this period

This system uses a file called TIME.DAT. On my machine it lives in the subdevice MEM:. If a group of people use the same machine, have each user would set up their own MEM: subdevice. Since, TIME.DAT is just a standard text file - it can be changed using any editor. The format is 'XX-AAA-XX XX:XX:XX AAAAAAAAAAAAAAAAAAAAAAAAAAAAA'. Which is the system date, system time and job description. This allows you to add transactions for time away from the machine.

There are two programs necessary that can be found in the DECUS C tools or in the book 'The Software Tools in Pascal' by Kernighan & Plauger. They are TR - transliterate and SORT - simple file sort.

I use this to keep track of time on development jobs as well as customer support calls. Start the day working on Jones Payroll, I would type 'TIMIN JONES Payroll<cr>', at 8:35 a customer calls so I type 'TIMIN CPA-support<cr>', then I take lunch 'TIMOUT<cr>'. The TIME.DAT file would look like:

01-Sep-87 06:15:37 JONES Payroll
01-Sep-87 08:35:12 CPA-support
01-Sep-87 11:47:13 off

CALTIM will produce the following:

Hours thru 01-SEP-87

Job	time	hours
cpa-support	3:12:01	3.20
jones payroll	2:19:35	2.32
Total	5:31:36	5.52

Since I am lazy the description will be converted to lower case for calculating end of period hours.

I did some timing comparisons between the three version with a 105 line TIME.DAT file, it came up in a dead heat at about 17 seconds.

TIMIN.COM

```
-----
r ind
SY:timin.ind ^1 ^2 ^3 ^4 ^5 ^6
^c
```

TIMIN.IND

```
-----
.ENABLE SUBSTITUTION
.OPENA MEM:TIME.DAT
.DATA '<DATE>' '<TIME>' 'P1' 'P2' 'P3' 'P4' 'P5' 'P6'
.CLOSE
```

TIMOUT.COM

```
-----
r ind
SY:timout.ind
^c
```

TIMOUT.IND

```
-----
.ENABLE SUBSTITUTION
.OPENA MEM:TIME.DAT
.DATA '<DATE>' '<TIME>' off
.CLOSE
```

CALTIM.COM

```
-----
ass dk old
ass vm dk
copy/nolog mem:time.dat dk:
tr <time.dat >time.x 'A-Z' 'a-z'
sort <time.x >time.dat
cal
! for f77
! or calctm for dbl
! or calp for pascal

sort <hours.dat >hours.dat
tot
! for f77
! or tottim for dbl
! or totp for pascal

type thours.dat
ass old dk
```

CAL.F77

```
-----
*
* cal.f77
*
* date: 1-sep-87
*
* author: Mike LeRoy
*
* this program takes time.dat and produces hours.dat
*
* input: time.dat
* format: 'XX-AAA-XX XX:XX:XX AAAAAAAAAAAAAAAAAAAAAAAAAAAAAA'
*
* output: hours.dat
* format: 'AAAAAAAAAAAAAAAAAAAAAAAAAAAAA XX-AAA-XX XX:XX:XX'
*
* this program is run after tr and the sort of time.dat
*
* build with f77 version 5.0a
*
integer*4 inday
character*30 indesc
integer*4 inhour
integer*4 inmin
character*7 inmmyy
integer*4 insec

integer*4 outday
character*30 outdesc
integer*4 outhour
integer*4 outmin
character*7 outmmyy
integer*4 outsec

integer*4 curtime
integer*4 lasttime
integer*4 time

open (unit=1, file='time.dat', status='old', access='sequential')
open (unit=2, file='hours.dat', status='new', access='sequential')

lasttime = 0

10 continue
read (1,100,end=900) inday, inmmyy, inhour, inmin, insec, indesc
100 format (i2, a7, x, i2, x, i2, x, i2, x, a30)

if (outdesc .ne. indesc) then
if (lasttime .ne. 0) then
curtime = insec + (inmin * 60) + (inhour * 3600)
time = curtime - lasttime
call cnvt (time, outhour, outmin, outsec)
write (2,300) outdesc, outday, outmmyy, outhour, outmin, outsec
300 format (a31, i2, a7, ' ', i2, ':', i2, ':', i2)
endif
endif
```

```

if (indesc .ne. 'off') then
  outday = inday
  outmmy = inmmy
  outdesc = indesc
  lasttime = (inhour * 3600) + (inmin * 60) + insec
else
  lasttime = 0
endif
goto 10

900 continue
close (unit=1)
if (lasttime .ne. 0) then
  type 950, indesc
950 format (' You are still working on job : ', a30)
  type 960
960 format (' ')
endif
close (unit=2, dispose='keep')
end

```

```

subroutine cnvt (itime,ihour,imin,isec)
integer*4 itime,ihour,imin,isec
ihour = itime / 3600
imin = (itime - (ihour * 3600)) / 60
isec = itime - (ihour * 3600) - (imin * 60)
end

```

TOT.F77

```

*
* tot.f77
*
* date: 1-sep-87
*
* author: Mike. LeRoy
*
* this program takes hours.dat and produces thours.dat
*
* input: hours.dat
* format: 'AAAAAAAAAAAAAAAAAAAAAAAAAAAAA XX-AAA-XX XX:XX:XX'
*
* output: thours.dat
* format: 'AAAAAAAAAAAAAAAAAAAAAAAAAAAAA XXX:XX:XX XXX.XX'
* plus heading and total lines
*
* run after cal and the sort of hours.dat
*
* build with f77 version 5.0a
*

```

```

character*31 indesc
character*10 indate
integer*4 inhour
integer*4 inmin
integer*4 insec

character*31 lastdesc
integer*4 lasttime
integer*4 hour
integer*4 min
real realtime
integer*4 sec
integer*4 time
character*9 today
real totrealtime
integer*4 tottime

open (unit=1, file='hours.dat', status='old', access='sequential')
open (unit=2, file='thours.dat', status='new', access='sequential')

time = 0
tottime = 0
totrealtime = 0.0

* write heading

call date (today)
write (2,5), today
5 format ('Time thru ', a9)
write (2,7)
7 format (' ')
write (2,9)
9 format ('Job time hours')
write (2,7)

10 continue
read (1,100,end=900) indesc, indate, inmmy, inhour, inmin, insec
100 format (a31, i2, a7, x, i2, x, i2, x, i2, x)

if (lastdesc .ne. indesc) then
  if (time .ne. 0) then
    call cnvt (time, hour, min, sec)
    realtime = time
    tottime = tottime + time
    totrealtime = totrealtime + realtime
    write (2,300) lastdesc, hour, min, sec, realtime / 3600.0
300 format (a31, i3, ':', i2.2, ':', i2.2, ' ', f5.2)
    time = 0
  endif
  lastdesc = indesc
endif
time = time + (inhour * 3600) + (inmin * 60) + insec
goto 10

```

```

900      continue
      close (unit=1)
      if (time .ne. 0) then
        call cnvt (time, hour, min, sec)
        realtime = time
        write (2,300) lastdesc, hour, min, sec, realtime / 3600
        tottime = tottime + time
        totrealtime = totrealtime + realtime
      endif
      call cnvt (tottime, hour, min, sec)
      write (2,950)
950      format ('          -----')
      lastdesc = ' total hours'
      write (2,300) lastdesc, hour, min, sec, totrealtime / 3600
      close (unit=2, dispose='keep')
      end

```

```

subroutine cnvt (itime,ihour,imin,isec)

integer*4      itime,ihour,imin,isec

ihour = itime / 3600
imin = (itime - (ihour * 3600)) / 60
isec = itime - (ihour * 3600) - (imin * 60)
end

```

CALCTM.DBL

```

;
; calctm.dbl
;
; date: 15-oct-86
;
; author: Mike LeRoy
;
; calculate time from time.dat and store in hours.dat
;
; time.dat format = 'XX-AAA-XX XX:XX:XX AAAAAAAAAAAAAAAAAAAAAAAAAAAAA'
;
; hours.dat format = 'AAAAAAAAAAAAAAAAAAAAAAAAAAAAA XX-AAA-XX XX:XX:XX'
;
; the description 'off' will restart calculation
;
; build using dbl version 2.2
;

```

```

RECORD IN
  INDATE ,A9
        ,A1
  INTIME ,A8
        ,A1
  INDESC ,A30

```

```

RECORD,X
  INDAY ,D2
        ,A8
  INHR ,D2
        ,A1
  INMIN ,D2
        ,A1
  INSEC ,D2

```

```

RECORD OUT
  OUTDES ,A30
        ,A1
  OUTDAT ,A9
        ,A1
  OUTTIM ,A8

```

```

RECORD,X
  OUTHR ,A41
        ,A2
        ,A1
  OUTMIN ,A2
        ,A1
  OUTSEC ,A2

```

```

RECORD
  CURTIM ,D8
  DIFF ,D8
  HOURS ,A9 ,'HOURS.DAT'
  LASTIM ,D8 ,0
  OFF ,A30 ,'off'
  TIMIN ,A8 ,'TIME.DAT'

```

```

PROC
  OPEN (15, I, 'TT:')
  OPEN (1, I, TIMIN)
  OPEN (2, O, HOURS)
  DO FOREVER
    BEGIN
      READS (1, IN, DONE)
      IF (LASTIM .NE. 0) THEN
        BEGIN
          CURTIM = INSEC + (INMIN * 60) + (INHR * 3600) + (INDAY * 86400)
          DIFF = CURTIM - LASTIM
          OUTTIM = ' : : '
          OUTHR = DIFF / 3600
          OUTMIN = (DIFF - ((DIFF / 3600) * 3600)) / 60
          OUTSEC = (DIFF - ((DIFF / 60) * 60))
          WRITES (2, OUT)
        END
      END
      IF (INDESC .NE. OFF) THEN
        BEGIN
          OUTDAT = INDATE
          OUTTIM =
          OUTDES = INDESC
          LASTIM = INSEC + (INMIN * 60) + (INHR * 3600) + (INDAY * 86400)
        END
      END
    END
  END

```



```

ELSE
  BEGIN
    LASTIM =
    OUT =
  END
END
DONE,
IF (LASTIM .NE. 0) THEN
  BEGIN
    DISPLAY (15, 'You are still working on job : ', INDESC, 13, 10, 7)
  END
CLOSE 1
CLOSE 2
XCALL FLAGS (1000000)
STOP

```

TOTTIM.DBL

```

;
; tottim.dbl
;
; date: 15-oct-86
;
; author: Mike LeRoy
;
; calculate total hours from hours.dat put in thours.dat
;
; hours.dat format = 'AAAAAAAAAAAAAAAAAAAAAAAAAAAAA XX-AAA-XX XX:XX:XX'
;
; thours.dat format = "AAAAAAAAAAAAAAAAAAAAAAAAAAAAA XXX:XX:XX  XXX.XX"
;
; build using dbl version 2.2
;
RECORD IN
  INDESC ,A30
        ,A1
  INDATE ,A9
        ,A1
  INTIME ,A8
RECORD,X
  INHR ,A41
        ,D2
  INMIN ,D2
        ,A1
  INSEC ,D2
        ,D2
RECORD HDR
  ,A30 , 'Job'
  ,A2
  ,A8 , ' time'
  ,A2
  ,A6 , ' hours'

```

```

RECORD UNLINE
  ,A32
  ,A8 , '-----'
  ,A2
  ,A6 , '-----'
RECORD OUT
  OUTDES ,A30
        ,A1
  OUTTIM ,A9 , ' : : '
        ,A2
  OUTHRS ,A6
RECORD,X
  ,A31
  OUTHR ,A3
        ,A1
  OUTMIN ,D2
        ,A1
  OUTSEC ,D2
RECORD
  BLANKS ,A30
  HOURS ,A9 , 'HOURS.DAT'
  THOURS ,A10 , 'THOURS.DAT'
  TMPHR ,D3
  TODAY ,A9
  TOTTIM ,D9
  XHOURS ,D9
  XHR2 ,D9
PROC
  OPEN (15, I, 'TT:')
  OPEN (1, I, HOURS)
  OPEN (2, O, THOURS)
  XCALL DATE (TODAY)
  DISPLAY (2, 'Hours thru ', TODAY, 13, 10, 13, 10)
  WRITES (2, HDR)
  DISPLAY (2, 13, 10)
  DO FOREVER
    BEGIN
      READS (1, IN, DONE)
      IF (OUTDES .NE. INDESC) THEN
        BEGIN
          IF (OUTDES .NE. BLANKS) CALL WRITE
          OUTDES = INDESC
          TOTTIM =
        END
        TOTTIM = TOTTIM + INSEC + (INMIN * 60) + (INHR * 3600)
      END
    END
  DONE,

```

```

CALL WRITE
WRITES (2, UNLINE)
OUTDES = ' Total '
TMPHR = XHOURS / 3600
OUTHR = TMPHR, 'ZZX'
OUTMIN = (XHOURS - (TMPHR * 3600)) / 60
OUTSEC = XHOURS - ((XHOURS / 60) * 60)
OUTHRS = XHR2, 'ZZZ.XX'
WRITES (2, OUT)
CLOSE 1
CLOSE 2
XCALL FLAGS (1000000)
STOP

```

```

WRITE,
TMPHR = TOTTIM / 3600
OUTHR = TMPHR, 'ZZX'
OUTMIN = (TOTTIM - (TMPHR * 3600)) / 60
OUTSEC = TOTTIM - ((TOTTIM / 60) * 60)
OUTHRS = TOTTIM / 36, 'ZZZ.XX'
WRITES (2, OUT)
XHOURS = XHOURS + TOTTIM
XHR2 = XHR2 + (TOTTIM / 36)
RETURN

```

CALP.PAS

```

(* calp.pas *)
(* ----- *)
(* calp.pas *)
(* *)
(* date: 2-sep-87 *)
(* *)
(* author: Mike LeRoy *)
(* *)
(* calculate time from time.dat and *)
(* store in hours.dat *)
(* *)
(* build with Oregon Software Pascal *)
(* version 1.2 *)
(* *)
(* ----- *)

```

```

Program calp;
var
  curtime : real;
  inday : integer;
  indesc : array [1..30] of char;
  inhour : integer;

```

```

inmin : integer;
inmmyy : array [1..7] of char;
insec : integer;
lasttime : real;
off : array [1..30] of char;
outday : integer;
outdesc : array [1..30] of char;
outhour : integer;
outmin : integer;
outmmyy : array [1..7] of char;
outsec : integer;
size : integer;
x : char;
xhour : real;
xmin : real;
xsec : real;
xtime : real;

```

Procedure cnvtime (xtime : real; var xhour, xmin, xsec : integer);

```

var
  offset : real;
  off2 : real;
  rhour : real;
  rmin : real;
  rsec : real;
begin
  rhour := xtime / 3600.0;
  xhour := trunc (rhour);
  offset := xhour * 3600.0;
  rmin := (xtime - offset) / 60.0;
  xmin := trunc (rmin);
  off2 := xmin * 60.0;
  rsec := xtime - (offset + off2);
  xsec := trunc (rsec);
end;

```

```

begin
  reset (input, 'time.dat', , size);
  if (size > 0) then
    begin
      off [1] := 'o';
      off [2] := 'f';
      off [3] := 'f';
      for inhour := 4 to 30 do off [inhour] := ' ';
      rewrite (output, 'hours.dat');
      lasttime := 0.0;
      while (not eof) do
        begin
          read (input, inday, inmmyy, x, inhour, x, inmin, x, insec);
          readln (input, x, indesc);

```

```

if (outdesc <> indesc) then
  begin
    if (lasttime <> 0.0) then
      begin
        xhour := inhour;
        xmin := inmin;
        xsec := insec;
        curtime := (xhour * 3600.0) + (xmin * 60.0) + xsec;
        xtime := curtime - lasttime;
        cnvtime (xtime, outhour, outmin, outsec);
        write (output, outdesc, ' ', outday:2, outmmyy);
        write (output, ' ', outhour:2, ':', outmin:2, ':');
        writeln (output, outsec:2);
      end;
    end;
    if (indesc <> off) then
      begin
        outday := inday;
        outmmyy := inmmyy;
        outdesc := indesc;
        xhour := inhour;
        xmin := inmin;
        xsec := insec;
        lasttime := (xhour * 3600.0) + (xmin * 60.0) + xsec;
      end
    else
      lasttime := 0.0;
    end;
    if (lasttime <> 0) then
      writeln ('You are still working on job : ', indesc);
      close (output);
    end
  else
    writeln ('time.dat file not found.');
```

TOTP.PAS

```

-----
(* totp.pas *)
(* ----- *)
(* totp.pas *)
(* ----- *)
(* date: 2-sep-87 *)
(* ----- *)
(* author: Mike LeRoy *)
(* ----- *)
(* calculate time from hours.dat and *)
(* store in thours.dat *)
(* ----- *)
(* build with Oregon Software Pascal *)
(* version 1.2 *)
(* ----- *)
(* ----- *)
```

Program calp;

```

type
  a9 = array [1..9] of char;
  a30 = array [1..30] of char;
var
  curtime : real;
  indate : a9;
  indesc : a30;
  inhour : integer;
  inmin : integer;
  insec : integer;
  lasttime : real;
  outdesc : a30;
  outhour : integer;
  outmin : integer;
  outsec : integer;
  size : integer;
  today : a9;
  toftime : real;
  x : char;
  xhour : real;
  xmin : real;
  xsec : real;
```

Procedure date (var day : a9);
FORTRAN;

Procedure cnvtime (xtime : real; var xhour, xmin, xsec : integer);

```

var
  offset : real;
  off2 : real;
  rhour : real;
  rmin : real;
  rsec : real;
```

```

begin
  rhour := xtime / 3600.0;
  xhour := trunc (rhour);
  offset := xhour * 3600.0;
  rmin := (xtime - offset) / 60.0;
  xmin := trunc (rmin);
  off2 := xmin * 60.0;
  rsec := xtime - (offset + off2);
  xsec := trunc (rsec);
end;
```

procedure print (xtime : real; xdesc : a30);

```

var
  hour : integer;
  min : integer;
  sec : integer;
```

The Flit Gum

```

begin
  cnvtime (xtime, hour, min, sec);
  write (output, xdesc, ' ', hour:3, ':');
  write (output, chr((min div 10) + ord('0')), (min mod 10):1, ':');
  write (output, chr((sec div 10) + ord('0')), (sec mod 10):1, '');
  writeln (output, (xtime / 3600.0) : 8:2);
end;

begin
  reset (input, 'hours.dat', , size);
  if (size > 0) then
    begin
      rewrite (output, 'thours.dat');
      lasttime := 0.0;
      tottime := 0.0;
      date (today);
      writeln (output, 'Time thru ', today);
      writeln (output);
      writeln (output, 'Job           time      hours');
      writeln (output);
      while (not eof) do
        begin
          readln (input, indesc, x, indate, x, inhour, x, inmin, x, insec);
          if (outdesc <> indesc) then
            begin
              if (lasttime <> 0.0) then
                begin
                  print (lasttime, outdesc);
                  tottime := tottime + lasttime;
                  lasttime := 0.0;
                end;
              outdesc := indesc;
            end;
          xhour := inhour;
          xmin := inmin;
          xsec := insec;
          curtime := xhour * 3600.0 + xmin * 60.0 + xsec;
          lasttime := lasttime + curtime;
        end;

          if (lasttime <> 0) then
            begin
              print (lasttime, outdesc);
              tottime := tottime + lasttime;
            end;
          writeln (output, '
          print (tottime, '      total hours
          close (output);
        end
      else
        writeln ('hours.dat file not found.');
```

The problem with the FPJ-11 floating point chip has been fixed. A couple of months ago I reported that the sign bit was dropped whenever the STCDF command changed the exponent. There is a new version of the chip which corrects this problem; and if you purchased your 11/83 or KDF11-A new after the first of the year, there should be no problem. If yours is an older processor, you are entitled to a "free" upgrade. See Bob Schor's letter on Page RT-18 for the FCO numbers for the dual-wide board. I foolishly let the field service engineer leave after fixing our 11/83 without getting the FCO number for it. I'll have it for you next month one way or another. By "free" upgrade I mean that there is no charge for the upgraded parts (even including the board swap for the 11/73). But if your system is not under contract, they may try to stick you for a service call. Whether you can argue them out of those charges depends upon your rapport with your local field service office.

Now for the NEW bugs. Quick Henry, the Flit!

I received the following piece of electronic mail from Nick Bourgeois which I pass on to you.

From: TOPAZ:BOURGEOIS
 To: CROWELL
 Subj: TSHLOB Patch

DECUS - INTEROFFICE MEMORANDUM

Date: 8-Sep-1987 01:56pm EST
 From: N. A. (Nick) Bourgeois
 BOURGEOIS

Subject: TSHLIB Patch

Rally Barnard has reported a bug in the TSHLIB U6.2 routine, ISTD. The following patch when installed in the source module, MNTDEV.MAC, in the DOIT2 subroutine will correct the error.

MOVB @2(R5), R0

^
 |___ Insert the address mode character.

Nick

P.S. John, you might put this information in the Minitasker.

And now, with all humility, I give the "Obscure Bug of the Month" award to myself for unearthing the following quirk:

If the unit number is omitted from the device name in a SET command (e.g. SET LD CLEAN), the SET code in the handler is supposed to be entered with the value 100000 in R1 (where the unit number is usually to be found). This works only for handlers with names AA[x] through US[x]. For devices with names UT[x] through ZZ[x], you get R1 = 0. This is not good if the set code is supposed to do something particularly different whenever the unit number is specified. The problem is in KMOVLY.MAC where you'll find this piece of code:

```

OVCMD SET
4$: CLR (PC)+
UNUM: .WORD 0 ; unit number goes here
.
.
.
CMP @R2,#<^RUSR> ; is it 'SET USR' ?
BEQ 13$ ; go to USR stuff
.BR NXTINS


NXTINS: ROR UNUM ; Rotate C-bit into unit
etc.
```

Needless to say, the C-bit will be set only if the handler name is alphabetically before 'USR' but not for devices UT through ZZ.

It is not my intention (heh-heh) to ridicule the programmer who let this one slip by, but to remind you that it happens to all of us. So for all you neophyte programmers out there, don't get too frustrated when you find those "dumb" mistakes. It happens to the best of 'em. (And I count the RT-11 developers among the very best.) And for all you hot shots who think you've arrived just because you've written a working device handler and found something useful to do with the MARK instruction, don't get too cocky. These little monsters will bite you at the most inopportune times.



DECUS



Letter from the Editor

Welcome to the November issue of the Site Management and Training SIG Newsletter. After a minor sabbatical, we are back to publishing.

This month's issue features an article by Phil Beetley of Purdue University on the subject of student employees. Thanks to Phil for sharing his experience and expertise with others through this medium.

Much has been happening in the DEC world as of late. This has served to increase the opportunities and enhanced the position of DEC system, site, and training managers as a whole. This fact has been driven home to the Site SIG Steering Committee recently. All of us have been overloaded with work and increased responsibilities. Hope you are benefiting from DEC's success as well.

Since each of us tend to be overloaded with work, it comes as no surprise that most of never find the time to share our knowledge and experience with others through published articles. However, the rewards of publishing an article are many. Increased esteem with and for your employer, good references for your resume, increased awareness and respect from your peers, and self satisfaction in sharing your experience with others.

Not to be left out is the hands on experience you receive by organizing your thought and recording them in a logical, coherent manner, on paper. It's good training. And who doesn't need more training in more effective communication.

Please think about it seriously. The SIG membership will benefit by your participation. Here are some general topics that might interest the SIG membership:

- Managing people
- Managing resources
- Software project management
- Managing software maintenance
- Managing hardware maintenance
- System Management methodology
- Network management methodology
- Disaster recovery planning
- Site preparation and design
- Capacity planning

These are just a few possibilities to get you started.

I am at a new address and have different capabilities for receiving electronic submissions. I can accept VMS Backup tapes (6250 BPI or 1600 BPI), PRO RX50 floppies, MS-DOS floppies (ASCII files, Symphony, Word Star, Word Perfect, Mass 11, or Micro Soft Word). VAX word processing files supported are: WPS Plus, Mass11, ASCII, and RUNOFF. Hardcopy is

also acceptable.

Send articles and comments to:

Gregory N. Brooks
Washington University
School of Medicine
400 South Kingshighway
St. Louis, MO. 63110
(314) 454 2237

Student Employees at a University Computing Facility:
Recommendations from One Site's Experience

Phil A. Beetley
Agricultural Data Network
Purdue University

BACKGROUND

Agricultural Data Network (ADN) is a computing, data communications, and technical consulting facility in the Purdue University School of Agriculture. ADN supports faculty and staff who use computers in research, regulatory, teaching, and extension activities. ADN is part of the Department of Agriculture Communications Services and has a staff of fifteen full-time employees.

ADN grew out of a Project MIRACLE, a National Science Foundation project to provide automatic data collection that began in 1971. From the beginning, full-time technical and professional staff have been augmented by student employees enrolled at Purdue. At present, student employees include two electronic technicians, two training materials developers, and one data base programmer.

ADN's experience with student employees has consistently been very positive. This article states the policies and procedures that ADN has followed in selecting and managing student employees. While ADN's situation may be in some ways unique, it is likely that some of our policies and procedures are relevant to other university-based facilities.

POLICIES AND PROCEDURES

1) Set priority of education over work.

Student employees are just that: students first and employees second. The work experience is to enhance and reinforce what is taught in the classroom and laboratory, not detract from it. Keep work schedules flexible to allow for class assignments, tests, and projects. At ADN, students are never assigned to critical path tasks. These tasks are often on tight schedules and would put undue pressure on a student employee.

2) Know the relevant university curricula.

When seeking student employees, it helps to know the relevant curricula that the university offers. At Purdue, for example, difference aspects of computing are taught in at least seven different schools, with numerous options and plans of study. Knowing the basic differences among the curricula helps when announcing openings and matching students to specific positions.

3) Establish faculty contacts.

Whenever possible, make contact with at least one faculty member in the relevant curriculum areas. Make this person aware of positions available, requirements, and type of work the student is likely to do. Once such a contact is established, the faculty may take the initiative to inform you of promising students who are looking for work. Check back with this contact to see what students are saying about their work experiences.

4) Formally interview student applicants.

Student applicants should be interviewed just like applicants for full-time positions. A student should bring the following to the interview: transcript of coursework, plan of study, and at least one sample of current work (e.g., program, design project, technical document, etc). This kind of interview helps identify the better applicants and is good practice for the student as well.

5) Look for long-term employment potential.

When selecting student employees, there is a tendency to hire the more advanced student, to choose the junior or senior over the sophomore. (ADN rarely hires freshmen unless there is prior military or work experience.) But any student will have to learn on the job. Sometimes it is better to hire a younger, but promising, student who could be a productive employee for two or three years.

6) Match area of study to work responsibilities.

Students who like their respective areas of study generally do better in school than those who do not. The same is true when matching a student employee to a position. Use the interview to discover interests and abilities that may not be clear from transcripts and sample projects. Don't assume that you know what kind of work a student can or wants to do - ask.

7) Students need mentors, not just managers.

Students need to be taught and guided more than they need to be managed in the way of full-time employees. With full-time staff, it's very important that they know why something needs to be done. But much of the "how" is left to the employee's training and experience. Student employees often need both "why" and "how" to be explained.

In some cases, managers are not the best supervisors for students. This is especially true in the case of working managers who have specific technical, as well as managerial, responsibilities. At ADN, we try to match a student employee with a staff member in his or her area of interest who is both a productive worker and a good communicator. Being the mentor for a student employee is both an important and rewarding assignment. Take care in choosing mentors.

8) Formally evaluate student employee performance.

Don't let a student employee disappear inside the organization. Student employees need to be formally evaluated at least twice per semester (16 weeks) or at least once per quarter (12 weeks). Both the mentor and the student employee need to be involved - separately if there may be sensitive or interpersonal problems to be resolved.

The purpose of the formal review is fivefold. First, to ensure that work is not interfering with education. Second, to see how the student is interacting with his or her staff mentor. Third, to confirm that the current job assignment is a good match for the student. Fourth, to assess job performance. Fifth, to identify areas of student progress and growth, both on the job and with respect to coursework.

Once an evaluation is complete, act on it. If rewards are in order, reward both the student and mentor as resources allow. If adjustments are needed, make them as soon as possible. If a student cannot balance work and education, education must take precedence.

9) Conduct an exit interview.

When a student employee leaves for whatever reason - graduation, difference student employment, educational demands, change in educational goals, etc. - be sure to hold an exit interview. Ask the student to evaluate his or her work experience. Encourage suggestions about how to make student employment better. If the student has done a good job, be sure that you communicate your satisfaction and appreciation. Indicate that you are available as a reference for subsequent employment.

10) Keep contact with students after graduation.

Whenever possible, maintain contact with student employees after they graduate. Try to know their respective job titles, areas of responsibility, and company or institution. If you publish a site newsletter, make sure they receive a copy and include a form for them to indicate any changes in employment or address. Keep in touch and be aware of their achievements. Be alert for former student employees who may want to return to a university setting.

11) Develop full-time employees.

Colleges and universities can rarely compete with business and industry with respect to salaries and other financial incentives. But

the academic environment has a lot to offer, too. Working in an academic setting often provides a degree of freedom and intellectual excitement not found in other jobs. If a student employee is interested in full-time employment after graduation, be sure he or she understands the trade-offs. In ADN's experience, student employees make very good full-time staff. A positive student employment experience can provide the basis for a good career.

CONCLUSION

Student employees, if properly selected and guided, can be a major asset to university computing facilities. By paying close attention to student needs and abilities, a site can benefit from students' skills and desire to learn, while being an integral, supportive part of the learning experience. A site that treats its student employees well will find a valuable source of talent that is renewed with each incoming class.

AUTHOR INFORMATION

Phil A. Beetley
Agricultural Data Network
105 Smith Hall
Purdue University
West Lafayette, IN 47907
(317) 494-8333

NEWSLETTER OF THE VAX SYSTEMS SIG



Our Mascot

Pageswapper




Table of Contents

Undocumented System Space Watchpoint Utility . .	VAX-3
Editor's Workfile	VAX-7
Migration to MicroVMS from DEC Operating Systems	VAX-8
Online Security Monitoring System	VAX-25
VAX/VMS System Performance Management	VAX-35
Converting Code for Parallel Processing	VAX-48
INPUT/OUTPUT	VAX-62

Forms at the End

INPUT/OUTPUT Submission Form	QU-25
System Improvement Request Submission Form . .	QU-27

To register for on-line submission to the Pageswapper dial:

(617) 262-6830

(in the United States) using a 1200 baud modem and log in with the username PAGESWAPPER.

Articles for publication in the Pageswapper can be sent (US mail only -- no "express" services please) to:

Larry Kilgallen, PAGESWAPPER Editor
Box 81, MIT Station
Cambridge, MA 02139-0901
USA

Preference is given to material submitted as machine-readable text (best is Runoff source). Line length should not exceed 64 characters and the number of text lines per page should not exceed 48 (these limits are particularly important for sample commands, etc. where simple text justification will not produce a meaningful result).

Please do not submit program source, as that is better distributed on the VAX SIG tape.

Please do not submit "slides" from DECUS Symposia presentations (or other meetings) as they are generally a very incomplete treatment for those readers of the Pageswapper who are not so fortunate as to be able to travel to Symposia. Please DO write articles based on such slides to get the content across to a wider audience than is able to attend.

Change of address, reports of non-receipt, and other circulation correspondence should be sent to:

DECUS U.S. Chapter
Attention: Publications Department
249 Northboro Road (BPO2)
Marlborough, MA 01752
USA

Only if discrepancies of the mailing system are reported can they be analyzed and corrected.

Undocumented System Space Watchpoint Utility

Lee K. Gleason
Principal, Control-G Consultants
2416 Branard D
Houston TX 77098

I received my VMS 4.6 kit yesterday, and immediately copied it to a disk, to look for new features. I noticed a new driver in SYS\$SYSTEM - WPDRIVER.EXE, and along with it, WP.EXE. My first guess was that it was related to Word Processing. A little study showed that this guess was all wrong.

It turned out to be a driver for a "software" device, like NOA0 or SWA0 - a device driver that has no physical device associated with it, but exists to provide a convenient QIO oriented interface to some complicated bit of internal trickery. The WP driver implements a system space watchpoint utility. QIOs directed to it (by its command interface utility, WP.EXE) cause it to monitor and log changes to specified locations in S0 space.

I did some more looking, and found the details of its operation in SYS\$HELP:WP.HLB and WP.HLP.

To enable use of WP, first load the driver and create a unit.

```
$ MCR SYSGEN
SYSGEN>CONNECT WPA0:/NOADAPTER
SYSGEN>EXIT
$
```

To enter WP commands, run the command interpreter...

```
$ RUN SYS$SYSTEM:WP.EXE
```

Once in WP, the command syntax is as follows.

The WATCH command sets up a watch point

```
WATCH watch addr      address to watch
                   /BYTE      watch a byte sized field
                   /WORD      word
                   /LONG      long (the default)
                   /QUAD      quad
```

```

/SILENT      take no extreme action on watch (default)
/XDELTA      enter XDELTA on watch
/FATAL       do a BUGCHECK on watch
  
```

The SHOW command shows information about a watch point

```

SHOW watch_addr      address being watched
  /CONTROL_BLOCK     show watchpoint control block (default)
  /TRACE_TABLE       show watchpoint trace table
  /FULL              show everything
  
```

The SET command controls logging of output to a file, WP.LOG

```

SET
  /LOG           write output to log file and terminal
  /NOLOG        close log file
  
```

The IGNORE command cancels a watch point

```

IGNORE watch_addr      addr of watch point to cancel
  
```

The EXIT command does what you would guess - returns you to DCL.
 Control-Z will do the same thing.

To get a feel for this, I needed a location in system space that I knew would change often enough to be interesting. The answer was no further away than my copy of "VAX/VMS Internals and Data Structures" (and let me assure you, it is never very far away). A look through it suggested something that would change on a pretty regular basis - the entries in the Timer Queue. The Timer Queue has a listhead that would change as Timer Queue Elements are added and deleted at the head of the list. I used ANALYZE/SYSTEM to get the address of the listhead.

```

$ ANALYZE/SYS
VAX/VMS System analyzer
SDA>EVALUATE EXE$G_L_TQFL
Hex = 80002B58   Decimal = -2147472552           EXE$G_L_TQFL
  
```

And then used that value to set a watch point with WP.

```

$ run sys$system:wp.exe
Watch Point Utility Version X-1N2
  
```

```

WP> WATCH 80002B58
  
```

I waited a few minutes, to let a few TQEs pass into history, and then did a SHOW on that watchpoint.

```

WP> SHOW/FULL 80002B58
  
```

```

Base Address      = 80002B58           Length      = 04
Address Touched   = 80002B58           Type        = SILENT
Time Touched      = 03:33:19.60        Touched Count = 00000509
  
```

Watch Point Contents

```

Initial = 00000000802569E0  Previous = 00000000802569E0
Post    = 000000008014B720
  
```

Register Contents

```

R0 = 00000000  R1 = 7FFB09B6  R2 = 00000000  R3 = 7FFB08B8
R4 = 8010A850  R5 = 802569E0  R6 = 7FFB0580  R7 = 7FFB0646
R8 = 7FF41D38  R9 = 7FFB0208  R10 = 7FFB0400  R11 = 7FFE0270
AP = 80045052  FP = 7FFE9DE4  SP = 802947D4  PC = 8000A160
PSL = 04180004
  
```

Instruction Stream

```

0050AF500BA50200EF1208DA55650F .PP.....Ue. 8000A160
      8000A160:      REMQUE (R5),R5
      8000A163:      MTPR #08,#12
      8000A166:      EXTZV #00,#02,0B(R5),R0
%WP-W-INSKIPPED, unreasonable instruction stream-220 bytes skipped
Base Address      = 80002B58
  
```

Ref. Byte	Ref. Count	Ref. Time	Op. Code	PC	PSL	Previous Contents
00	00000509	02:33:19.60	0F	8000A160	04180004	802569E0
00	00000508	02:33:19.39	0F	8000A160	04180004	8014B720
00	00000507	02:33:19.38	0F	8000A160	04180004	80002B58

The information displayed is pretty much self explanatory. I was impressed to see that it decoded the instructions that did the modification that triggered the watchpoint. I recognized the REMQUE as part of the software timer interrupt service routine.

This utility will be useful for figuring out what's happening in the Executive. Using it to help debug a device driver comes immediately to mind, and it would be useful for debugging user written system services as well.

A few caveats are probably in order, for the incautious. The help file for this utility has even more disclaimers from DEC than usual, and a look at the commands show that crashing a system with this tool would be no effort at all. Proceed, at your own risk, and with caution. There may be a reason for this being undocumented and unsupported.

This utility can also cause a drastic system slowdown if you watch any location that is modified often at high IPL. I found this to be the case when I set a watchpoint at the listhead of my terminal's UCB IRP queue. The system slowed down to a crawl, and output to my terminal became very leisurely. The system recovered when I used the IGNORE command to cancel the watchpoint.

WP implements watchpoints by changing the memory protection on the page that contains the location being watched, and analyzes each access violation that occurs to see if it would result in the modification of the watched variable. In order to do this, the WPDRIVER "steals" the interrupt vector for Access Violation - that is, it replaces the address at 20+@EXE\$GLSCB with the address of its own routine, so that it can examine every access violation that occurs, and deal appropriately with the ones caused by watchpoints.

As far as I could tell, there is no way to list the watchpoints you have set up, so if you don't record the addresses where you set them, you can't find them to SHOW or IGNORE them.

Since I haven't installed VMS 4.6 yet, these tests were done by copying the components of WP to a 4.5 system. It seems to work just fine there. I presume, that since WP came with the 4.6 kit, it would work on that version as well.

Editor's Workfile

Computer Security

The (DoD) National Computer Security Center has finally acknowledged the importance of computer security concerns other than secrecy. Even after receiving some duties regarding civilian computer security in the controversial presidential order NSDD 145, the military-oriented agency has kept its "Orange Book" concentrated on preventing secrets from leaking, rather than dealing also with questions of integrity (avoiding improper MODIFICATION of data) or denial of service.

A new publication called the "Trusted Network Interpretation", however, discusses plans for accrediting network security, and in the process gives plenty of attention to integrity and denial of service issues. It seems as though they didn't want to actually admit the Orange Book was deficient but they realized they would not make it anywhere with a commercial audience talking only about secrecy.

Symposium Transcripts this Issue

Many thanks to Ron Frederick for providing the Pageswapper with transcriptions of several Symposium sessions for this issue. Let me know how you feel about transcripts in the Pageswapper. I get very little feedback about what people do or do not want to see, other than one recurring comment about I/O entries being difficult to understand if the antecedent entries were from a previous month. I'm working on that one.

Larry Kilgallen
Pageswapper Editor

Migration to MicroVMS from DEC Operating Systems

Spring 1987 US DECUS (Nashville) Symposium Session V172

presented by Bill Taber, W. I. Taber Inc.
and Henry Schneiker, Pivotal

transcribed by Ron Frederick
N.I.P.E.R.
P.O. Box 2128
Bartlesville, OK 74005

We are going to just take questions. We need to do the standard DECUS procedure of coming up to the microphone, one question per iteration through the line. Please identify yourself and the company you work for. Come up and ask your questions about various migration things.

Q: Laura Berry, MDB Systems, currently we have a PDP-11/73 running RT11 and TSX. We would like to convert to MicroVAX-II, but have somewhere between 700 and 1000 COBOL programs that need to be converted. Does anybody know how difficult that's going to be?

A: You will have to recompile them all, plus you will have to account for any differences between the compilers from your current system to the MicroVAX. There is no getting around that you are going to be doing a bit of work, having both machines sitting side by side and setting someone to the task of converting one at a time. You will have to test them also.

There is a filter program, but I'm not sure whether it exists for the VAX. We might be able to twist DEC's arm to get it out. This program filters from COBOL11 to COBOL81.

Q: Is COBOL81 what the MicroVAX-II uses?

A: It's COBOL8X or whatever they are up to. That program does exist. I'll go back to the PDP-11 group and see if that does exist on the VAX as well.

Q: Russ Hanson, Mayo Clinic. I've programmed an RSX for a while in MACRO and I have to convert to a MicroVAX, but

I get to convert to a MicroVAX-II. I haven't looked at all at MACRO in VMS. Is it much different?

A: Totally different. You are going to get very confused, very fast. Moving R0 to R1, for example, works in reverse. It will take R1 and put it in R0, not R0 into R1.

Q: Are there any good introductory books?

A: There happens to be a very good one downstairs in the bookstore. It has to do with the VAX architecture.

Q: I'm Ed McKay from Galveston College in Galveston, Texas. We are a small 2000-student, two-year junior college, annual budget of about \$6 million or so. We do everything on a PDP-11/70. That means student records, budgeting, accounting, our own payroll, as well as teaching FORTRAN, BASIC, PASCAL, COBOL81 on the PDP-11/70. We have two DH emulators and 8 DZs. That's something like 104 terminals and we typically run 50 jobs on the 70. We are converting to VAX hardware primarily for political reasons. The other schools in our area have all gone to VAX and we are kind of in the Dark Ages if we don't have VAX hardware. My first MicroVAX is sitting in a box down there to be installed next week when I get back and I don't know whether it's going to take two, three, or perhaps, four MicroVAX-II's to replace the 70. I'm confused at this point about what way I ought to be going. Should I plan just to connect my MicroVAX-II's with DECNET over the Ethernet, or I should go to a local area cluster over the Ethernet? Perhaps I should be looking at Ultrix because of the much, much lower cost and some of the basic products that are available on Ultrix? I would like to have the panel vote which of the three ways you would go -- whether you think it's going to take two, three, or four MicroVAXes to replace my 70. I have to prepare the budget for the next year's addition as soon as I get back.

A: One MicroVAX should be able to effectively replace the CPU power of your 11/70. That shouldn't be that much of a problem. If you want to add additional CPUs, however, clustering is a very convenient way to go because all of your resources are available to all of the MicroVAXes. Just remember that you have to manage a cluster as if it were a single processor. It is basically a distributed

operating system. You have to manage the entire thing as if it were one machine because when you cluster, it actually is. You can incrementally increase the amount of power you want, including peripherals and CPU. VAX clusters are a very nice way to go in that regard.

I want to add something. You said you were coming off RSTS. Most RSTS programmers make use of the operating system itself to do things that are very, very nice, like little programs, lots of changing. That works on an RSTS system; that hits the VAX in the worst possible place -- image activation. You will need to make your running applications all one large program in a hurry so you can activate them all at once. Do a lot of calls to subroutines. Another thing, the compilers on the MicroVAX are much, much faster than the 11 compilers. You will pick up a lot of throughput there.

Q: Ron Frederick, NIPER in Oklahoma. We have a PDP-11/70 running under RSX- 11M-PLUS right now. We write a lot of command files. Is there a difference in the DCL from one DCL to the other. I understand there are some differences. What's the situation there.

A: Are you referring to DCL or indirect on RSX?

Q: I am primarily asking about DCL at this point. Is DCL on the RSX different from . . .

A: Yes. There are differences. The DCL on the VAX is much more powerful. You will run into syntax problems. It shouldn't take very much to convert a DCL command file from one operating system to the other.

Q: Ron Frederick, NIPER. You asked about indirect. Indirect versus the indirect command process? I am confused at this point by your bringing up indirect.

A: Okay. There is a DCL command parser and there is also the indirect, which is the utility, the parser under RSX with the dot commands and so forth. They are totally different. You are going to do a lot of work if you convert an indirect command file as opposed to a DCL command file.

Q: I wonder if you could address the pros and cons of MicroVAX-II clusters that would allow us to go to the

VAX in pieces, as opposed to just jumping in the water with the big one. Somebody told us that for our needs, we are going to need an 8350. We currently have two 11/44s, one of them has 3 meg the other, four meg. We are making use of the virtual disk. We've got Eagle disks on both machines, one also has an RA80. One machine has about 25 users, with 10 or 12 hard copy printers attached as terminal lines. On that system it's BASIC, BASIC+2, primarily inquiry, maintenance, report generation. We also have a fairly large DIBOL package that could have as many as four or five users. The second machine is our development machine, primarily where we write and compile, taskbuild our programs. It's got the 3 meg, an Eagle disk. It's also an 1144.

A: The kind of applications you run determines whether to go with one big CPU or a lot of smaller CPUs and cluster them together. If you are going to write a lot of small, independent jobs that don't need lots of CPU individually, it's often much more efficient to run them parallel on separate processors of a cluster. If you've got one big job that sits there and runs all day long, then you want one big processor to run it on. Having several processors won't do you any good because it's only going to run on one and the others will just sit there idle. If your job mix allows you to spread those jobs over multiple processors, such as you have lots of students and you can log five onto one processor or something like that, quite often it is advantageous to go to a cluster system and buy the first MicroVAX. When you run out of CPU, buy a second one and stick it on a cluster, a third one, a fourth one, and so on. The advantage of doing clustering versus networking is that you can access all the resources if they are a part of that CPU. You have to manage that as one big system.

The other thing to watch out for when you go to a cluster, I am going to be very conservative at this point, you probably are going to lose about 20 percent of the horsepower of the machine. Two MicroVAX are equivalent to 1.6 MicroVAX.

Also, you need to consider the I/O bandwidth because you have very little disk space. The MicroVAX uses the Q-bus, which tends to be a bit slower than what other systems use, HSC-50's for instance. That is another consideration on total system throughput.

Q: Bob Van Curan, Userware. One comment on that. If they are separate applications that don't have shared data, you might even consider having separate MicroVAXes and not putting them in a cluster.

Q. Jim Cronin, Carolina Power and Light. I've got an application that runs currently on a 780 and I need to move it to a MicroVAX-II. I've been told, and I read, that a VAX is a VAX is a VAX. Should I expect any problems moving this application from a 780 to a MicroVAX?

A: While there are some distinctions between VAX, in general, they are correct -- a VAX is a VAX is a VAX. The problem comes in certain areas, such as how they actually implement it and which bus they are using. DEC rates MicroVAX as it comes configured about .9 versus a main 780, partly because it uses slower peripherals. We maintain that if you put fast peripherals on a MicroVAX, it will go about as fast as a 780. The second major distinction is the number of instructions and which instructions they have actually implemented in hardware. The VAX has the capability of having the software implement instructions. This saves on the hardware. That's what they did with the MicroVAX and a good number of instructions just aren't implemented. I think 40 percent of the instructions are missing and are implemented in hardware, but that shouldn't bother you. They did a study of what applications use which instructions and implemented the popular ones. One instruction missed was the "string move" instruction or "string compare." I don't remember which. At the time they did the study, none of the compilers generated the instruction so it wasn't a problem. About the time MicroVAX was released, the optimizing FORTRAN compiler came out. The optimizing FORTRAN compiler uses that instruction. If you write FORTRAN programs that do a lot of string manipulation, for instance, you will find that it runs about half as fast as a 780. There are certain considerations like that. In general, a VAX is a VAX is a VAX. It's the same operating system and everything else. You should also remember none of the PDP-11 emulation instructions are on the MicroVAX. If you are thinking of running compatibility mode, don't bother to think about it. Go out and buy a real 11.

Q: Harold Addleman, National Computer Systems. I wonder if any

of you would care to comment on the recent, very large increase in fees for license of MicroVMS end user and what the Digital strategy is in doing that.

A: I don't like high fees any more than you do. DEC tried to equalize all of their licensing and the MicroVAX got stuck in there with some higher prices, especially for the licenses. Beyond that, I am not sure I want to comment much further. As I mentioned, I don't like higher prices any more than you do.

Yes, it is because you can't put more terminals on it.

Q: Russ Hanson, Mayo Clinic. I've had my MicroVAX-II for about six months and I have had one TK50 problem. It didn't let the cartridge back and I had to have a service guy come out to release it. How many should I expect? I haven't been using the tape very much, but I am going to be. How often should I expect a failure.

A: The TK50s are a problematic drive at best and are notorious for having problems. There are certain things that you should know if you have a TK50. There actually is a procedure for causing errors and there are two types of errors you will typically run into. In one, you stick the tape halfway into the drive, change your mind, and pull it out. The system thinks there is a tape in there and won't let you insert the tape again. The second problem -- it's got your tape and won't give it back. Quite often half of the tape is on the internal reel so you can't pull it out anyway. There would be tape everywhere if you did. These two conditions are cleared by pushing the red, or dismount button. Push it in, push it out, push it in, push it out, and let the thing whir-buzz-click until it finishes. At that point the error should be cleared. If the tape was inside the drive, it should be possible to pull the tape. If the tape could not be inserted into the drive, the drive should have reset itself so that you can.

There is also a difference physically on the cartridges. DEC has an old-style cartridge and a new-style cartridge. If you look on the edge there are little indentations. Originally they thought "Oh, well, we want people to be able to pull the cartridges back out." If you turn it sideways so you can see into the depression, there is a ramp on one side and a flat side

on the other. On the old style, the ramp was on the out-facing side so you could pull it back out. On the new style, it's on the inward side so you can push it in. They decided it was more important to let you get it all the way in and not have the situation where you couldn't stick it in any more than try to get it back. I have had experience when it wouldn't give back your tape and nothing else would clear it. You actually had to go in there, physically take the machine apart, reach in, and grab the solenoids without powering the machine down. Powering down the machine while the TK50 is actually doing something is not good for the poor tape. You risk mangling your tape in the process. In general, they are problematic, but if you are nice and kind to them -- sit there and don't try to rush them or anything, put your tape in, wait until it settles, push the button, wait until it does its thing, then run your program. When you come back, make sure your program has stopped accessing the tape, push the button in, let it do all of its rewinding. Wait until you hear the final click before you pull up the latch and pull the tape out. If you are very smooth about it and don't rush the tape at all, the chances are you will have fewer problems. The human interface problem with that thing is pretty bad.

Q: Bob Van Curan, Userware International. A couple of comments. I have heard throughout the week about the pricing issue. From what I have heard the change in price is prompted by realignment of all the VMS prices for the entire series of computers because MicroVMS is going away in the sense it is just going to be part of VMS itself. There won't be any formal distinction between it and the other parts of VMS. I've heard a lot of comments from people that this price change will slow their migration to VMS because the pricing balance between staying on a PDP-11 and going to a VAX changes. I also attended a BOF session where some people were considering moving to a UNIX on the VAX instead of VMS specifically because of the increased prices. There were some rumors that a third party vendor has a version of BASIC on UNIX that is very similar to BASIC+. I think this is going to slow down our customers in their migration to VMS. They are always looking at "how much bang for my bucks am I going to get if I move over to VMS" and this changes the whole equation. Some interesting things as a result of this price change.

A: That is indeed the fact. You should see, I believe, around 4.6 disappearance of MicroVMS altogether and there will be one common command procedure. I imagine that tailoring for the 730 is going to go away at the same time and you will have one command procedure that installs everything. That is going to be nice.

Q: That frightens me. I only have 8000 blocks left on the system disk.

Q: Bob Perry, Tektronics. Just a few comments on some things talked about here so far. We have a large number of MicroVAXes at Tektronics and realistically we have only seen something less than a 10 percent failure rate on the TK50s. Primarily those are from the grabber mechanism that comes out and tries to grab the tape initially. A lot of times it won't let loose. We have had cases where it has ripped the leader. I don't think I have ever seen any of the wrap problems on our systems, but you have to realize that a TK50 is a delicate mechanism. You don't want to slam the tape in, you don't want to pull it out fast. If you do, you are going to jam it, guaranteed. You've got to push it carefully -- just shove it with one finger until it snaps in. Then push the button in. When you are taking them out, don't jiggle them from side to side because they will jam. You have to be very careful with them. If you do that, you won't have as high a failure rate. We are careful in the way we handle our TK50s in our TK50 drives and they are still a pain.

We have a number of migrations from different operating systems to MicroVMS, the toughest being from VMS to MicroVMS. In most of those cases the system parameters and quotas were different. When you are going from a larger (larger being physically larger) VAX to a MicroVAX, watch what has been set on your system. It can set sizes and some of the dynamic system parameters may be different, thus causing your task or your process (I work in RSX too!), to react differently under the MicroVAX system. Watch your quotas or you may run out of quotas. Make sure you have just about the same quotas that you had on the VMS system. We have had some migrations from both RSTS and RSX to VMS.

Some more comments on the TK50. I understand you can buy leaders from DEC Direct for \$2.00 each in packages of 10. DEC sells them for some ungodly price, I don't know what. If you are interested in replacing your leaders, buy a package of 10 for \$20.00 bucks or so, and replace them yourself. If you have a TK50 and the leader is pulled back so the internal mechanism cannot grab it, you will have to release two cams inside the TK50 box. There are little bitty holes with latches that you can stick a pen into and release both. It takes three hands to do this operation. One hand to hold each pen and another hand to turn the tape with the door open. Rotate the tape so that the leader is sticking out again, then release the door and all the cams. You can fix things that way also. When you own a TK50 you learn a lot of tricks to get it up and running again.

Q: I am willing to trade anybody my load device for a TK50. I have an RX50 load device.

Q: Bob Henbeck, Winrock International. I have been curious about the MicroVAX 2000. I haven't heard a lot about them this week and I wondered if you had any comments, especially where they might fit in small networks and clusters.

A: You mean a 2000.

The new, small MicroVAX 2000. The little, recently announced one.

The diskless version? Either that or whatever configuration. Those are actually very nice on a cluster. I think that is what DEC's intention was originally -- diskless cluster members. You put enough memory on them so they don't cause your Ethernet great consternation and they run just fine. Ethernet is an approximation of hard disk speeds for transferring a lot of things. You tend to not see degradation. By the way, some interesting timings. If you have a 13-node Ethernet cluster, you can quite typically run around 20 percent bandwidth of your Ethernet. When you look at bands of Ethernet, never try to get more than about 35 to 40 percent because the technology doesn't work that way. You will have increasing collisions, increasing delays, and everything else beyond that point. Most

people consider a utilization of 35 to 40 percent for such a technology to be the maximum. If your cluster node does not have sufficient memory or local paging and swapping, you will need to increase utilization of the bandwidth substantially. It will have to page and swap. This is why DEC recommends that whenever you have a cluster node, you have a minimum of 10 to 16 megabytes of memory or a local hard disk to do your paging and swapping.

Q: Bob Perry, Tektronics again. Another note for those of you making your migration from RSX to VMS. DEC has not formally announced the product, but there is a coprocessor board to fit the Q-bus of the MicroVAX-II and run RSX-11M-PLUS. It's a J11 board and talks to VMS. You can give your commands, log onto the VMS system, and talk to the M+ system. They haven't announced the product, but they have announced the study program and they have a working model on the floor. We are interested in it from a real time standpoint. We do a lot of data acquisition with RSX. It would be nice to be able to sit a MicroVAX out and be able to do processing on a MicroVAX. Look for the product announcement along about the end of this year. There is a second version of this system. They were talking about not even having to put it in the backplane, just having an Ethernet connection between the MicroVAX and the RSX system. It is like an AME under VMS, but a little more involved than that. You might want to talk to the folks down on the exhibit floor.

Q: Barbara Lawrence, Starsburg Electric System again. I know my questions are very naive, but I really know very little about what we are fixing to do. Number 1: how much of our existing hardware would be practical, as well as possible, to take over to VAX. We've got LP25 printers and LP26 MS11 tape drive, the Eagle disk. I know the RA81 will plug right into the VAX. In terms of other peripheral equipment, of course, we've got multiplexers and such things. How much of that will be usable on a VAX?

A: None of your multiplexers because they are Unibus not Q-bus. None of the controllers for your disks or your tape drives. You can, however, get controllers from others vendors for your Eagle disks that work very nicely.

- Q: Oh, good. I wondered if we couldn't do that.
- A: There is a controller for it. It's a J11 anyway.
- Q: Well, I know it's not everybody's favorite. At least temporarily.
- A: The printers. Again, you would have to get proper controllers for them, either through DEC or some other third party people.
- Q: Is that the practical thing to do rather than replace them?
- A: Peripherals are rather expensive to replace, so it is a matter of can you afford to replace them or is it cheaper to just buy new controllers. You might also check the used market. You can quite often get either used or refurbished equipment for very nice prices.
- Q: That's probably where we are going to go for everything, certainly everything we can.
- A: The Eagles do run. I've got a MicroVAX-II with one Supereagle and one Eagle on it.
- Q: Fantastic. I don't see anybody else waiting; I wonder if I could ask something else. You may not be able to answer it, not being fully aware of what we are doing and what have you, but I will describe it to some degree. We are primarily I/O-bound as opposed to compute-bound. How many users could I expect to support on a MicroVAX-II?
- A: It really depends on the application. I have an application right now that supports probably 45 active users on a MicroVAX.
- Q: Okay, what kind of application is that?
- A: A single application. It's a lot of data inquiry, credit collection bureau.
- Q: Okay, that's probably comparable. We are doing utility billing and all that goes with it.
- A: Two things you must consider as far as how many users to specify for a MicroVAX, or any other piece of computer gear for that matter are: what hardware you actually

have and how much memory you have. These things have a direct effect on how much swapping and paging you do for many applications, which directly affect how much of your disk bandwidth is chewed up for other things than doing useful work. There are things you can trade off. You will probably have to put the users on, look at any bottlenecks you might have, release those bottlenecks, and go from there. Your typical performance thing. Watch the system, make some changes, watch the system, make some changes, and grow the system as necessary.

The other thing on the system I currently have running -- it was a ground zero start-up. It was designed with the MicroVAX in mind. Everybody sits in one image, everybody is running one shared image. I have 16-meg on it. Normally 4 meg is free most of the time with 35 users on it.

- Q: What is the maximum number that you can put on the MicroVAX-II.
- A: 16 meg.
- Q: Is that maximum?
- A: That is the address bus width. There is a physical hardware limitation that won't allow you to address any more than 16 physical megabytes.
- Q: Does DEC now sell 8 meg boards?
- A: Yes, DEC now sells 8 meg boards along with all of DEC's competitors.
- Q: When we put it together, we couldn't get DEC memory for it.
- Q: ** with Caterpillar Incorporated. I recently got a MicroVAX Workstation2. What would be a reasonable page faulting number to shoot for with graphics running and DECNET.
- A: If you look in the performance manual in the big manual set -- all of you who have MicroVAX need access to a full manual set. I highly recommend that you go out and buy one. You should have a minimum of one big manual set per VAX site. It is very important because a lot of things are not in those two user manuals that you got.

Q: It's four now.

A: Yes, but the other two don't count. I mean they are FORTRAN manuals, but no FORTRAN compiler. Going back to performance, the manual says that a maximum number to consider for paging or swapping on any CPU is enough soft page faults to constitute 5 percent of the CPU time in processing those page faults. A soft page fault is generated through the page cache, the free list or the modified list. I'm not sure if you are familiar with any of these terms.

Q: Not really. No.

A: Okay. Since I don't have 20 minutes to explain this, just believe that a page taken out of your working set is placed in this cache area and is available when the system needs memory. It can pull these out and use them. Finding a page in that cache is far faster than going all the way out to disk to get one. Once you have figured what 5 percent is, and on a VAX780-MicroVAX-type machine 5 percent represents around 100 soft page faults per minute. Some around 100 or 120, plus or minus. You are allowed 10 percent of the number of soft page faults as hard page faults. This means somewhere between 10 and 12 hard page faults. Those are the ones you get by going all the way out to disk.

Q: Per minute?

A: No, per second.

Q: Per second. Okay.

A: Per second. Do a SHOW SYSTEM or SHOW PROCESS/CONTINUOUS. This also affects how fast your disk device is and whether to raise or lower these numbers. The guidelines indicate this is a starting point and you as a system manager must decide how much paging you are willing to have on your system. Also understand that you can never get rid of paging to activate an image. The system sets up all the paging tables and then faults that particular image into memory. Whenever you activate an image you see a large increase in pages that slowly drops to a base level. You will never be rid of paging. Once a process has been activated and is up and running, you should set things so it does not do more than these

guidelines for paging.

Q: Thank you.

Q: Fritz Merkle, Temple Union High School District. Could you elaborate more on the image activation on RSTS and include information about where you have written programs, heavily overlaid techniques, resident libraries, things like that. From a very naive, talk about it naively because I...

A: Throw your overlays away. They don't really exist on a VAX. Two things that RSTS did very well were single character I/O from a terminal and starting a program. A run command or a chain command to activate a program was very, very fast on RSTS; it is extremely slow under the VAX. There is an awful lot of set up to get a program running internally in a VAX. Are you programming in BASIC+2 or something like that? A perfect example of what occurs as far as the I/O goes is if you take your BASIC+2 program and go -- 10 PRINT"."; 20 GO TO 10 -- run it, it will take about 3 percent of the CPU on an 11. You have just gone 95 percent CPU-bound on a VAX because of the QIO processor. These are things you have to be aware of and you have to look at them when you start. Quite often on a RSTS system you ran out of memory, then you broke your programs into model streams and ran one program right after the other, chaining back and forth. All these programs you were chaining can be called as subroutines on the VAX. You will run a lot faster.

Almost everything you have in the VAX is a resident library. There are things you need to do when you get running...My favorite command for tuning a system is to say -- SHOW DEVICE <system-disk-name>/FILE/NOSYSTEM -- and find all the VAX VMSRTL program files that multiple people have copies of. Immediately find those, reinstall them -- SHARED/OPEN -- then only one copy resides in memory. This is more effective use of your CPU memory. Oh, yes, /HEADER.

Q: Mel Potter, Atomic Energy of Canada. Are there any limitations on a MicroVAX-II with equal drives? Can the drive be the system drive?

A: Yes, the drive can be a system drive.

Q: Do you have any comments on using the CDC9772 disk system on the MicroVAX-II?

A: Have not run across it.

I haven't run across it either. If the hardware is smart enough to be told to go out and access the boot files, you must not be required to load microcode before using the disk. If that is the case (all of DEC's drives are that way), then you can boot off it and there is no problem. I have a reference here for DIGITAL REVIEW, January 26, 1987. There is apparently an article about that.

Q: Thank you.

A: Talks about the drives and controllers. Anybody looked that article up?

Q: Russ Hanson, Mayo Clinic. I've been comparing transfer rates over the network between an 11/73 using a RD53, and a MicroVAX-II using a RD53; I get something like two or three times faster with the MicroVAX. Is that to be expected? I get about seven times faster than an 1123 with an RL02.

A: Which controllers are you using?

Q: The second version.

A: RQDX2?

Q: Right.

A: I believe the other controller is the better controller, the -3, and goes faster.

Q: The -3 is, you say.

A: Yes. RQDX3 is about three to four times faster than the RQDX1. I don't know the difference between the RQDX1 and the RQDX2 -- someplace in the middle.

Q: Do you know if version 3 on an 11/73 versus a MicroVAX with the same hardware going the same place are different? Using identical hardware, my MicroVAX is at least twice as fast when moving big files across.

A: With the same hardware on an 11 and on a MicroVAX, DMA transfer should be equally fast as long as you are transferring to onboard memory because the hardware is basically identical. If you are transferring into Q-bus memory, it may be a tad slower. Not too much slower, otherwise the disk would overrun the memory. That should never be the case. The Q-bus is plenty fast to do those transfers. The only problem would be the operating system setting up and making actual use of that information. The actual transfer time should be the same because the hardware is identical.

Q: Mel Potter, Atomic Energy of Canada. I'm going to have a lot of people using the MicroVAX that don't necessarily want to learn all the command language. What about the A-Z product, is that a good solution?

A: Anybody here work with A-Z? I don't know anything about A-Z. As another possibility, write command procedures, sometimes even CAPTIVE command procedures, to do what you need the user to do. Give them a menu of what you want for an interphase, give them some prompt, make it look like some favorite computer these guys know and love. It will slow things down a little bit, but sometimes the convenience is well worth the effort.

A: Most of my systems are written with CAPTIVE processes. People are getting CAPTIVE and can't do anything other than what they are given the capability of doing. A lot of procedures are then passworded at the next level down. System accounts for my nontechnical managers -- I have given them menus in the system so they can add users, delete users, run backup, do most of the system management from command processes.

A: Anybody else have a question?

Q: Laura Berry, MDB Systems. You mentioned that you can't put TS-11 on a MicroVAX-II. Is that correct?

A: The reason is the controller DEC supplies is a Unibus controller, not a Q-bus controller. If you can find somebody who supplies a Q-bus controller, you can put it on.

Be aware that there is no support for any tape drive whatever for MicroVMS on Digital. DEC does sell a TSV05, which means when you buy the TSV05 from DEC you have to buy an additional software package, the TSV05 tape driver. There is no tape support other than the TK50 that comes with MicroVMS.

Also, they don't include device drivers with MicroVMS because you haven't got the hardware and you can't possibly put it on. They eliminated that so it wouldn't take up additional disk space.

Q: Ron Frederick, NIPER. We are in the process of looking into a new tape drive and we are looking at the TU81+ to go on a MicroVAX and also onto our PDP. Does what you just said about tape drives mean we are going to have to get some software in addition to the hardware and the controller?

A: Yes. It's only like \$250, I think.

Q: I'm John Santos, from EGH. The TU81+ uses the same driver as the TK50 and that's already there. It uses the TMSCP tape driver and that isn't good in a MicroVMS so I don't know what they would give you, that would be additional.

A: Okay, I wasn't aware of that. Thank you.

Q: What happens if you have both, what's it going to do?

A: It will see two different units.

Q: Mel Potter, Atomic Energy of Canada. I was looking at some benchmarking MicroVAX-II programs and occasionally I came across the words "coded BLISS". What is "BLISS", a compiler or just a switch on the FORTRAN compiler?

A: It's a compiler and you can buy it from Digital.

Online Security Monitoring System

Spring 1987 US DECUS (Nashville) Symposium Session V004

presented by Linda Lankowicz
Computer Science Department of Springhill College
Mobile, Alabama
and Marino Nikolai
Computer Science Department
University of South Alabama
Mobile, Alabama

transcribed by Ron Frederick
N.I.P.E.R.
P.O. Box 2128
Bartlesville, OK 74005

This presentation is the result of on-going research efforts to develop an algorithm for identifying security violations. While VMS provides a means of auditing, our research is directed more to developing an on-line security system that will examine the activity of authorized users. Assistant managers' responsibilities include a number of areas that require knowledge of the parameters of resource utilization. From creating access control lists for directories to setting . . . limits for authorizing users, to examining the CPU time for cost accounting, and finally, tuning the system for performance. Security is just one aspect of the responsibilities of the system manager. A system manager might establish procedures for the physical security and rely on the mechanisms of the operating system for the internal system.

A system manager develops a certain sense of how his system behaves and how certain activities use his system. An experienced manager develops an intuition about what activity is taking place based upon resource utilization of a specific system. When I worked as system manager, it seemed to me that a process's use of system resources could be used to identify the activity of that user. Resources such as the disk, with page faults and . . . or the terminal usage with buffered . . . , the memory . . . , or the CPU with processor time. I thought other parameters that were not readily available, such as the page faults per image or the processor time per image, would be also good to look at and evaluate. A process might generate a lot of page faults, you have to have a very low ratio of page

faults per image. In some instances while monitoring the users, I was able to detect unusual behavior patterns based on abnormal patterns of resource utilization. It seemed that if some of these parameters could be used to identify the user activities, they could also be used to determine whether a user posed a threat to the security of the system. Our research was an effort to quantify this and included two areas.

First, we performed a statistical analysis of some of our user behavior to show that a combination of these parameters was sufficient to identify user activities. Second, we developed a model for security threatening measures and developed an on-line security monitoring system. The trusted computer systems evaluation criteria, as it is called, are in the orange book. The Department of Defense states that for a secure system, the trusted computing base should contain the mechanism that is able to monitor events in the system, events that might indicate an imminent security violation and notify the system manager. It is really not sufficient just to audit past activities of the system. It is necessary to be able to detect that a possible security violation occurred and then notify the manager.

If specific activities are identifiable, then unusual activities, activities that would pose a threat to the system, ought to be recognizable also. We selected an unusual activity, browsing, and sought to quantify this into a model applicable to other systems. Browsing consists of the attempts of a knowledgeable user to compromise the security of the system by performing activities that look much like activities of the system manager including copying files, searching for files, open files, files that are not protected, moving from directory to directory.

Linda is describing the concept we started to work with about two years ago concerning our ability to monitor the activities of the authorized user. We felt that was a problem and we heard that mentioned several times this morning. We were not really sure what level of background information the audience would have concerning some of the details of security policies established by the Department of Defense. I want to digress just two or three minutes to look at terms and concepts associated with the security violations and threats that we are going to discuss and track. When looking at this description, we realize the key is to be able to monitor the occurrence of imminent security policy violations. We would like to provide a tool that extends the use of the monitor utility to allow the system managers to set a threshold to actively observe the

performance of various users or processes on the system.

The Department of Defense clearly states that in order to achieve any level of certifiable system, main areas of security must be investigated. The first one deals with external security, the second, internal security. The two previous presenters talked a great deal about the external security measures concerning physical access in communications, the problems associated with networks, the ability to protect back-up problems. Our concern is not directed toward external security, although system managers must be aware of that as a primary obligation. As far as internal security is concerned, the Department of Defense feels there are several ways in which internal security needs to be monitored. These include access controls which most of you are familiar with; we will talk about them very briefly. An idea which may be new to some of you or may be not quite as familiar as access controls involves flow controls, then the idea of inference control, which is associated with the ability of a user to go into a database, make queries, infer from those queries, and gain access he otherwise would not be authorized to use. Of course, the need for . . . in addition to those protections and all of those internal and external securities are associated with certain . . . of the system.

Since we are dealing primarily with internal security, I thought a very, very brief history of the development of access control strategies would be appropriate. The earliest known access control strategy dealt with a ringed protection hierarchy. It was implemented on MULTICS operating system and at that time the idea was to separate, or segment, users from other users and users from operating systems to allow very limited communication between processes. MULTICS chose eight rings and their implementation VAX has four rings: a user ring, a supervisor ring, an executive, and a kernel. No one ring can talk to a ring that is two removed unless it goes through a user gateway in the intermediate ring. In addition to the access control ring protection, there are several other capability lists which are very popular these days. The idea of a capability list is to associate with each process a list of objects and access types allocated for those objects. As a process enters, it has a description of which objects it has access to. Of course, most of you are familiar with ACLS, Access Control Lists, which is just the reverse of that in the sense that for each object in the system there is a list that describes which processes have access to it.

The concept of flow control may be new to you. The DoD specifies that all of these must be in place for a certifiable system. Access lists or access controls usually keep users from interfering with other users or from interfering with the operating system. They do not discourage or disallow a user from giving away information. That's an important concept in flow control because we are interested in maintaining the need-to-know status. A typical example is in a student or academic environment where one student will make or run a program then pass it to friends through the use of a mail utility. When the program goes through mail, all the security and protection tags are stripped and it becomes an unprotected piece of information. In a sense, that is a security violation.

Some strategies for dealing with flow control. Develop authority levels for every file and then for each of those files make sure the sender and receiver have the same authority level, whether it be a classification or category. The problem with that approach is that it generally leads to a higher classification or lower classification of the entire process because everything is moved up to the highest level. In addition to maintaining flow control, one would like to be able to develop a trace of all program flow to make sure that terminations occur normally, that all loops terminate and that there is no ability to escape from the isolated domain. The fourth area that is mentioned is one of domain isolation. We heard this morning of the use of the Trojan horse, which is the writing or presenting to a user with a privilege to . . . a very innocuous looking program. I ask that system manager to run the program for me. Imbedded in that code is a piece of software that writes some information out of user space so I've gained information through using the manager's account that I would not have normally been allowed to access. In addition to those four control strategies, DoD also emphasized the need for a kernel, a secure kernel, in which all secure operations are centralized. Essentially, that kernel must be completely isolated and correct. There is a great deal of work being done at the theoretical level to prove the correctness of various types of kernels. All security operation would be performed at the kernel level. The kernel would run directly on hardware and would be completely separated from applications by an operating system emulator. In this case the application would never see the operating system. In addition to the software kernel, some hardware considerations are important: the efficient method of contact switching, the need for argument validation and the need for access types.

Just a very brief list of some work that has been done beginning in 1973 to show you there are some systems around trying to satisfy these requirements. The SRI provably secure operating system . . . the SRIPSOS system deals with capabilities we just discussed: access to control strategies and a tag architecture. That specifically addresses the need for flow control.

Honeywell . . . was a communications processor developed for the US Air Force based on the MULTICS ring structure. It has a secure kernel and currently has the highest DoD classification of any system available. That was developed in 1975. In 1976, MITRE and UCLA jointly developed a secure Unix operating environment that incorporates a kernel, a policy management system and an authentication process. System Development Corporation in 1976 developed KVM which is an extension of VM370. It is a VM370 with security modifications and uses the idea of security relevant and non security relevant domains, as well as a secure kernel. In 1978 MITRE developed a kernelized, secure operating system which included kernel check access interprocesses, information flow, and extensive audit capabilities. The Department of Defense developed in 1972, and modified in 1978, Directive 5200.28 specifying the security requirements needed to attain various levels of classification. These have been incorporated into the orange book which was developed in 1983. They specify particularly that physical and internal security requirements must be satisfied in four environmental areas: dedicated environment, system high, controlled, and multilevel. Dedicated environment essentially is where all users have the same security clearance and the same need to know. It's a very, very restricted environment, it's ideal, primarily concerned with physical security. You have already limited the users to the information that is there. The other, you have a multilevel environment where you mix all security types, all need-to-know classifications with all manner of objects. There you are concerned primarily with internal security measures. In 1984 President Reagan signed Directive 145, thereby establishing the National Security Agency as the agency responsible for communications and computer security. This directive also established a set of vendor criteria which must be met in order to sell to the government.

The certification process is managed through the Department of Defense Computer Security Center at Fort Meade. They have a two-phase process. The first phase involves direct interaction with the operating system designers. Problems associated with the operating system are resolved at this point, they go back

and forth. There is no published result at this stage. When the operating system is finally submitted for certification second phase, the security classification is submitted and a penetration risk analysis is performed. Our particular study was really concerned with penetration risk analysis.

A couple of very brief definitions. The definition of a secure system is a system which will control through the use of specific security features access of information such that only properly authorized individuals or processes operating on their behalf will have access to read, write, create or delete information. There were six stated requirements to satisfy those goals. Two in the area of access control and four in the area of accountability and assurance. First is that an explicit well-defined security policy is required which will enforce rules governing access of objects by subjects, by processes. Access control lists must be associated with each object under accountabilities. The requirement for individual identification bi-level, requirement for augmenting available and protected files, need for evaluation of control mechanisms and the need to protect security mechanisms from alteration.

In the area of background the Department of Defense has defined four classes of internal security mechanisms, A through D. You have heard this several times this morning. It is the minimal protection, you automatically get D if you do not achieve any of the other levels. Division C is discretionary protection; it is based on security levels and need to know. There are two subclassifications, C1 and C2. I won't go through these in detail. I believe we have heard this morning that the VAX 4.3 is at a C2 level now. Division B, which is the next higher classification, has the following requirements and must enforce integrity against acts of users. We were concerned with that area -- monitor of real time activities, protection of logical data structures, access control lists, hardware detection of malfunction, proven design specs, unannounced system checks.

Three areas: B1, 2, and 3, culminating the security domains. Division A, which is the highest, requires a verify design; I'm not sure that anybody really knows what that means. Right now, . . . is the only processor that satisfies the criterion and requires a mathematical proof that the model is consistent with the axioms and sufficient to support the security policy. This requires a set of verifiable, formal, top-level specifications and a trusted computing base which satisfies the top-level specifications. Linda will continue with a detailed description of our research project.

First, I examined the accounting output from student activities for over a semester to attempt to quantify them into a profile. You are familiar with the output from accounting. We get a large amount of information about each user process, the record is written at the termination of each process. We are looking at peak values since it is not recorded until the process terminates. When we are looking at accounting data, we are looking at the peak working set. It would be nice to look at the working set over the life of a process. In looking at some of the monitoring data, the system manager is able to see all of the processes on the system, as well as the shared and the working set pages, the direct I/O count, the faults and the CPU time for the processes. It would be nice also to be able to see the number of images generated by these processes and, perhaps, the buffered I/O for each of these processes because that would be interesting from a security standpoint. The information provided by monitoring is hard to digest. It appears on the screen interactively or you can capture it to a file, but you accumulate a large amount of misinformation. The most effective use is usually to follow one or two user processes that you would like to investigate further.

Our on-line detection routine would alert the system manager as to existing problems. It is really difficult for the system manager just to watch monitors and get that kind of information. We performed a discriminate analysis for three different groups of users to determine which of the parameters for resource utilization could be used to identify their activities. Identification was possible because we knew what their activities were. We chose a senior English class which had received instruction in word processing and did not know how to use the other features on the system, even as far as using the mail. Another group of students was involved in editing and programming with PASCAL. This was a beginning group of students and they did not know how to use the other resources. The third group was a business statistics class who were given accounts to use . . . , an interactive statistical software package. These three groups of users were quite different and we wanted to evaluate those three groups to see if their resource utilization parameters could be used to determine their activities. By looking at the means for 11,006 different log ins over a period of a semester, you can see the groups did have quite different means. Even the peak working set from accounting ranged from 895, 516, 378. As I said before, sometimes it is more useful to calculate a parameter such as the buffered I/O per image. A student may have been logged on for a long time and have a large amount of buffered I/O and, yet, the amount per image might be

low or the process of time per image. We ran a discriminant analysis of this and saw that a combination of these parameters could be used to identify the activities of the user. We were very pleased with the outcome. We were able to identify a user word processing 91 percent of the time. We have this information available in thesis if you would like to examine that. Once it was established that the resource utilization could be used to identify the activity of the user, I attempted to determine how an activity might appear over the lifetime of that activity rather than just look at the termination of the process. I also attempted to look at browsing and how it affected system resources. I looked at the system manager's account, as well as some students who were involved in browsing. In order to do this, I took output captured from the monitor, recorded it in a file, and because that data was not in a usable form, I imported that into Lotus 1-2-3, parsed it into columns, and then sorted so I could examine one process at a time and look at how that process changed over its life.

We looked at a number of parameters. This is just one. There is another example in the DECUS notes with another graph. We generated many of these graphs from many users for a number of parameters we determined to be critical. This particular graph shows the working set of six different users. I put these graphs together so you could see the difference between these users. The working set default was 200, the working set quota was 500, and the working set extent was 1,000. Those are the figures on the left side of the graph. The top three graphs are for users that were word processing. On our particular system, users that were word processing with this type software would quickly fault in the number of pages needed. They would always get a very large working set even if a one-word document was created in the word processor, always far above the working set quota. Then their working set would remain steady until they exited the word processor and, you can see in one of them, they would trim the working set fault and expand again if they re-entered the word processor. If a student edited the very same file, the student generated a different profile editor. The two straight lines across the middle are students that are editing. They would always be in a range below working set quota and allowed to grow, but generally would maintain a steady image in the range from 350 to 550.

The bottom line shows a browser. A browser's working set looks quite different because of the trimming he experiences as he generates images. Because he generates so many images, changing his activity, changing directories, trying to type files, looking at file protections, copying files, he is constantly being trimmed. Actually, the browsing profile generally is below working set default, usually no more than working set default. Even though a browser would begin at working set default, he would be trimmed with each image generation. The maximum they ever displayed was below 300. In an academic environment using a stepwise discriminate analyses

(inaudible or cut from tape when turned)

implement a system for detecting browsing.

We begin by looking at what kind of goals we would have. There was always the goal of little overhead. Any new feature for security is going to increase system overhead and we would like to minimize that as much as possible while realizing there would be a cost involved. The other two goals were simplicity and consistency, treating all processes alike. Two different approaches were considered. In one approach, a process creation (a shadowing process) would be created. You would have twin processes and one would constantly monitor the other process. This offers the advantage of individualized treatment because that one process could check more often for someone who was demonstrating browsing. This also means there would be an inconsistent treatment of the different processes. This would be a disadvantage because the process doing the monitoring could be lulled into a false sense of security as the process got into a word processing mode and then changed to browsing and, also, there would be overhead from the large number of processes on the system. You would be doubling the number of processes on your system.

We did not implement this one, but instead implemented a single process taking snapshots of all the processes on the system. No process, therefore, is bypassed. There is consistency because every account, even the system manager's, is treated the same. The disadvantage is that this one process has to maintain information about all the other processes. We tried to address this problem. The monitoring process using the four parameters we specified hibernates, wakes up, and takes snapshots of all the processes on the system. It determines whether any of the processes might be browsing and sets a bit in a bit map if it is determined any of those processes meet the four parameters. The

monitoring process schedules its next wake-up and hibernates again.

The monitoring process uses GET JOB PROCESS information. It obtains information about the four items that are needed: the working set size, the number of images, the fault rate, as well as the name of the image being executed by that process. If the monitoring process does determine a security problem, it will write it to OPA0, to the console. Each bit map is associated with a port, so that when the monitoring process is taking place you have a window of the activity on that port. We rotate the bits each time a snapshot is taken. If you are taking a snapshot at one minute intervals (for 32 bits), it would mean that you would have a window, or picture, of the past 32 minutes of that particular port. If any process was generated and that port was determined to be browsing, a bit would be set.

There also needs to be some criteria for deciding when to notify the system manager. We decided that for our implementation every user would occasionally demonstrate behavior that might be classified as browsing: moving, changing from directory to directory, copying files, etc. We were concerned about sustained activity in this area and gave some consideration as to what would constitute enough browsing activity on that particular system to notify the system manager. Our system was set at four consecutive bits, meaning the browsing activity was detected for four minutes. The implementation did have a very low overhead since there was no interprocess communication -- that was one of the aims. A bit map only using 32 bits per port maintained all of the information. The code was written in MACRO, we used GET JOB PROCESS INFORMATION, wrote the information using QIO to the console, and all processes were treated equally for consistency.

Note that the system manager forgot to log off at some remote site and some student was in that account browsing, that activity would be detected and recorded. We were able to develop a methodology from modeling different user types. Our discriminate analysis showed it is possible to identify users based on system resource utilization parameters and we implemented this with an on-line monitoring program. We need to further refine and extend our model of browsing and perhaps other security threatening behaviors. We would like to develop other profiles of other user groups, extend the work possibly to other systems and also to networks (we did nothing as far as networks) and to be able to select the optimum parameters for early detection of browsing.

VAX/VMS System Performance Management

Spring 1987 US DECUS (Nashville) Symposium Session V041
presented by Steve Neubauer
VMS Performance Group

transcribed by Ron Frederick
N.I.P.E.R.
P.O. Box 2128
Bartlesville, OK 74005

Good morning. The title of today's talk is "A Practical Approach to VAX Performance Management" or "How to Identify a System-Limiting Resource and Why It Is Limited." As the sophisticated system manager knows, this is the first and most important step in performance management. The resources we are going to look at are CPU, memory, and I/O subsystems. Identifying the limiting resource is important when looking at specific performance problems and a general performance strategy. We want to anticipate problems before they happen. The main objective of this talk is to present an overall methodology for identifying the limiting resource.

A limiting resource is just that, a resource that limits the performance of your system. We will see later in the talk exactly what is meant by "a limiting resource." This presentation will include only the tools on the standard VMS kit. We will not talk about the layered products, VPA, or SPM. We may make passing reference to them, but they will not be the focus of today's talk. We will talk primarily about the VMS monitor utility and other DCL SHOW commands, however, this is not intended to be a monitor tutorial. If you have more detailed questions on monitoring, please bring them to the VAX campground. I'll be there today from 1 to 3 P.M. Since this is only a one hour talk, we will cover only the major points. We are not necessarily going to talk about solutions to your problems, but will focus on problem identification. This talk is really meant to get you started, to whet your appetites. The topics of the talk include preparation, steps to take before you begin a performance investigation, problem identification, and identifying the limiting resource on your system. We will then look at an investigation of the three major resources.

Let's look at the things you should do before you begin a performance investigation. Collecting historical monitor data is one of the most important things you can do. This will enable you to understand your workload and know what is typical on your system. When you have a particular problem, you can look for deviations from those typical values. If you are looking at kernel mode time or compute queue, you should get a feel for what is typical on your system when performance is good. Monitor your system at system start-up. There are command files in the directory, SYS\$EXAMPLES, to help you do this. These command files start MONITOR with a large sampling interval, with a 10-minute sampling interval. MONITOR is running as a background process on your system, collecting performance data.

There are really two aspects to using MONITOR: collecting or recording data, and summarizing data. You should do that and look at the information in those reports on a daily basis. If you are collecting this monitor data in background, you can look at the recording files on a live basis. These files allow for multiple readers, so you can actually summarize the data from a live recording file. This command, MONITOR/INPUT equals (FILE NAME)/BEGIN equals some incremental time, will let you look at recent system usage. Sometimes historical data is not sufficient to help resolve a particular performance problems. Things slow down for a short period of time, for example, for a minute. Sometimes the 10-minute recording intervals will not enable you to focus on the particular problems and dilute the information over the entire interval. It's nice to have some command files around to invoke during problem periods. If you anticipate problems, then you can have these files ready. Hack what's in the SYS\$EXAMPLES directory and replace the interval period of 10 minutes with a shorter period of time. These files will provide the finer granularity data that's often needed when evaluating problems. The examples we are going to show in this talk will be from some of these time monitor files.

You should also run AUTOGEN of course. AUTOGEN generates system parameters based on your hardware configuration for typical workloads. AUTOGEN does not know about your own workload, therefore, it is important that you know these characteristics and make any necessary changes in the parameter file MODPARAMS.DAT. If you have a high level of sharing and locking on your system, for example, you may want to increase the SRP count and your lock ID table, may be your resource hash table. It's important to know your workload to know where your workload may deviate from the typical. Wherever possible, use the

ED symbol in AUTOGEN. These enable you to specify incremental amounts of change based on what AUTOGEN calculates. AUTOGEN will make a calculation and by using these symbols you will be able to specify incremental amounts. It is good to use these so you won't override the values that AUTOGEN calculates. In some particular cases, like PAGEDYN, one particular example where a hardcoded value of PAGEDYN went in MODPARAMS.DAT may override some of the work that AUTOGEN does.

Hardware problems may cause performance problems. You can use the DCL command SHOW ERRORS to see if a device is producing errors. You can use ANALYZE/ERROR to check the error logs for more information. SHOW MEMORY will show improperly configured memory on your system and any bad pages. Also, the console logs and the operator logs are useful places to check for hardware problems. I wanted to mention this now. Typically, you start looking at a performance problem, then it leads you to investigate a hardware problem.

Let's move on to the major part of the talk -- identifying the limiting resource on your system. We see this particular system manager busily typing away, however he has taken the steps we just talked about. He is prepared for an investigation. You should take a top-down approach to identifying the limiting resource on your system for each major resource -- CPU, memory, and I/O.

There are really two questions to ask. One is a question of capacity, the other is a question of efficiency. Determine if a particular resource has reached or is reaching capacity. You can do this with the two monitor classes: MONITOR SYSTEM and MONITOR DISK. If you find that a particular resource has reached capacity, go on to step two and see if that resource is being used efficiently. Even if you find that no resource has approached capacity, take step two to ensure the resources are being used efficiently. You may have plenty of free memory on your system, but you may have set working set sizes too small. You may be incurring page faulting or you may have some idle time, but you are spending a lot of time in kernel mode doing system operations. We are going to look at these particular issues shortly and take a top-down approach for each of the major resources and some of the examples that follow, concentrating primarily on single system performance.

VAX clusters really demand a similar approach. To identify the limiting resource on your VAX cluster, focus on that resource for more detail. When summarizing monitor data on VAX cluster systems, use the multi-file summaries, especially if you are looking at disk statistics. The disk information is recorded on individual nodes. Roll that up and look at the total cluster-wide access in individual disks. You can also do live cluster monitoring starting in version 4.4 with a MONITOR CLUSTER command. This is actually a command you would use on a live basis for doing a performance investigation. You can look at problems as they happen, if they happen.

This is an example of the MONITOR CLUSTER screen. Here we see an overview of the major resources of the CPU, I/O, memory, and some locking statistics. We won't discuss this in any more detail, but some of the information to be presented later is also available in this screen. This would be the first place to look when evaluating a cluster performance.

Memory investigation. I chose to talk about memory first because in some sense it is the most involved resource. Even though I'm not going to get into the details of memory management, I still have a lot of things to say about memory. The capacity question, "Is there any free memory available," can be answered by looking at the free list in MONITOR SYSTEM or MONITOR PAGE. The efficiency question really boils down to: are you doing paging or swapping?

In these example screens, I will focus on what's illustrated in yellow. Here we have an example of MONITOR SYSTEM in its graphical form. We see there are a little over a thousand pages on the free list where there is memory available on this system. The right-hand side of the screen shows the number of available pages outside that allocated to VMS. Memory on a VMS system can be broken down into three major areas. One area covers what's allocated to VMS (includes nonpaged dynamic memory and the look aside list -- this is the number given on the bottom of your SHOW MEMORY command; the number of pages allocated to VMS; it is also the part allocated to PROCESS WORKING SETS). What's left over is either on the free list or the modified list. The number on the right-hand side of the monitor system on the free list part of the screen represents the area dedicated to the two latter regions. Anything other than that is allocated to VMS.

"Are you using memory efficiently, is the system paging." There are two types of paging to consider: system paging and process paging. First, let's look at system paging. The system working set is a fixed portion of the part of memory allocated to the VMS. This is where the page pool and page dynamic memory reside when a part of physical memory. When the system tries to access page in page pool that's not in memory, a system fault occurs. That's the SYSGEN parameter. SYSMWCNT controls the size of the system working set. It's important to run AUTOGEN in relation to this parameter because AUTOGEN adjusts the system working set based on the size of page pod. If you make any change in the system that increases page pod, these parameter rules will be increased proportionately. You should also watch the system fault rate on the monitor page screen.

This is an example of MONITOR PAGE. We see that the average system fault rate here is .39 faults per second. I want you to focus on the second column on the screen -- the average column -- on most of these screens. That's usually the number we will look at. Here we see .39 faults per second. Conventional wisdom says to keep the system fault rate low, below a few faults per second, it is not clear, however, that on faster processors higher fault rates may not be intolerable, especially if these faults happen to be soft. This is really an area requiring further investigation. It's possible that the system faults may be generating lots of modified pages and flushing the modified page list. We know that a low system fault rate is good and at this point we are not ready to go against conventional wisdom.

Let's look at process paging and process working sets. Setting working set parameters is one of the most important things a system manager can do so I am going to take a moment to consider them. On the lower right-hand part of the screen, we see the user authorization working set parameters, working set default (the size of the working set when you are not executing an image), working set quota (the size the working set is normally permitted to grow), and working set extent or WS extent (the size the working set can grow if there is sufficient free memory available on your system). The area between working set quota and working set extent can really be considered a low region -- you are allowed this many pages provided there is sufficient free memory available. The WSMAX SYSGEN parameter governs the maximum size any working set can grow to. The actual set size can vary at any given time between working set fault and working set extent. There is an F\$GETJPI parameter, WSSIZE, which tells you the current size of your working set. We are going to show

results from command procedure WORKSET.COM. This very handy command produces a good overview of memory status on your system. It's relatively easy to write your own command procedures using F\$GETJPI commands and this is a very good example of one. This is an example output from the WORKSET.COM command procedure. I would like you to focus on the headings and the columns for the time being. We are going to take a look at this particular example in the context of a specific problem later and look at some of the values. This sample gives the user name process, the state the process is in, the working set extent, the working set quota, and the working set default values, which are either the user authorization values or the values that may have been changed by the working set command.

The next is working set size, obtained from the WSSIZE GETJPI parameter. WSI tells the current size of the working set. The pages in the working set tell the number of pages you are actually using, the number of pages that are actually mapped and valid. It also gives the number of page faults and the image being executed. We are going to come back to this example.

There are two separate considerations: setting of individual process working set sizes and setting of all process working set sizes. Make sure your processes display reasonable faulting behavior with working WSQUOTA pages in individual process working sets. Don't rely on the working set extend mechanism -- when memory gets tight on your system, the processes will not be able to grow to this value. Be sure processes display good faulting behavior with working set quota pages. The sum of active working set quotas should be less than the available free memory for all processes. This is to prevent excessive paging or swapping. Here we are talking about active processes, not all processes. You can estimate active processes by looking in either MONITOR SYSTEM or MONITOR STATES and considering only the processes not in the LEF (i.e., local event flag) or HIB (i.e., hibernate) states. Those can be considered your active processes. With this general rule, we have a little bit of slop left over for global pages. If you are really ambitious, you can look at MONITOR PROCESSES for a breakdown of private pages versus global pages.

This is an example of a particular image -- it's the number of page faults versus the working set size. As we increase the working set size, eventually the page faults drop off. Theoretically, the working set quota would be set to the knee* in the curve for this particular image of about 700 pages. It is not reasonable to expect system managers to know the shape of

this curve for every image, but your user guides or application manuals will give guidelines for setting working set parameters.

Now we want to focus on a process that's doing a lot of page faulting. A similar procedure to identify a particular process on your system will also be used in the area of I/O and CPU. When you see high fault rates in MONITOR SYSTEM or MONITOR PAGE and want to identify the highest fault generators on the system, use the monitor processes /TOPFAULT command. Once you have identified the highest fault generators, then check to see what those processes are doing. You can do that with SHOW PROCESS/CONTINUOUS. This is an example from MONITOR SYSTEM and it's in tabular form. By using MONITOR SYSTEM/ALL you will see this form of monitor system rather than the graphical form. Here we see a system that's generating lots of page faults and very few of them are hard faults. That is very few of them are page read I/O faults. We have plenty of memory available on this system so we are doing lots of page faulting.

This is a MicroVAX-II system and all the examples we have are from a MicroVAX system. A page fault rate of over 100 faults per second is generally considered to be excessive on this system. When you see those high rates, you generally would like to see who is doing the faulting. This provides a little bit more information. This is from MONITOR PAGE. We see most of the faults are from the free page list and the modified list. In this particular screen the top five items in the middle part of the screen are all soft page faults. Faults that don't require disk I/O account for the majority of soft page faults. Then we can look at MONITOR PROCESSES/TOPFAULT to see which processes fault. We see the process OPUS is generating most of the faults. This slide was played back from the recorded information. At the time we recorded this information, we didn't know which process was faulting, but we were recording the processes class. We played back the top page faulters.

This is an example from the graphical form of MONITOR SYSTEM. I put this up to show most of this information is on this screen. If we had more faulters, we would have identified those on the MONITOR PROCESSES/TOPFAULT screen. On this screen, the fault bar was on a bar under the page fault rate. It's to the right of the margin bar (I don't know if this is that legible on this screen). This bar indicates the number of hard page faults on the system. Now we can look at SHOW PROCESS/CONTINUOUS and see how many pages are in the working set. When we look at this screen on a live system, we see the live page fault rate or the page fault count ticking away and we see some simulator program

is being run. By looking at the WORKSET.COM command procedure, we can see this process has grown. We see what the working set parameters are -- working set extent 512, working set quota 256, working set default 128, working set size, the current size is also 512. We see this process has grown to its working extent and is using all the pages in its working set. We would like to increase the working set parameters for this process. This is the output list of this command file.

Some of the things to look for when you see a process actually growing beyond the working set quota: you can determine swapping or trimming is occurring, if the number of pages in the working set are actually growing to the working set size, if the processes are actually using all the pages in the working sets. Command procedure gives you a good grasp of what's happening with respect to memory on your system. We increased OPUS working set parameters, we increased working set quota to 700. This particular system had lots of free memory and would have grown to its working set extent anyway. We increased the working set quota to 7000 and our faulting behavior -- our faulting leveled off and we see good faulting behavior. We should also check swapping in addition to paging to see if memory is being used efficiently. The level of swapping is given by the inswap rate of MONITOR IO. If you are swapping, one of the things to check for is to make sure you have enough balance set slots per processes in memory, that you are not artificially inducing swapping.*****

Take a look at some examples. This is from SHOW SYSTEM and often is your first clue that swapping is occurring. It shows the number of processes that are swapping. Having out swap processes is not necessarily a bad thing. In some systems it may be desirable to have out swap processes. If you have a lot of users idle lots of the time, it may be preferable to have out swap processes. The rate of swapping given in the MONITOR IO screen is more critical to system performance. A reasonable threshold for swapping is below one inswap per second. Here we are well below that limit and don't really have a swap rate problem on this system. We can go to SHOW MEMORY and see the reason we did any swapping -- we ran out of balance set slots. We have lots of memory available, we just don't have any more slots in memory for those processes. That concludes our memory discussion.

Now we will look at doing a disk I/O investigation. The main issue concerns keeping disk response times down. Response time is the time required to get data off the disks. The SPM does provide a measure of response time. If you are using that tool you could do an analysis based on response time with MONITOR. Because MONITOR does not provide a response time, we will look at I/O operation rates. This graph shows the average response times for a typical VMS I/O as we increase activity to the disk. By typical I/O's we mean those small I/O's common in most environments and with a fair amount of seeking. As we increase the operation rate, disk queues start to form, requests start waiting for other requests, and eventually response time will start to increase significantly. The knee in the curve of typical I/O's is around 25 I/O's per second. It is really workload dependent. Some workloads may be a little less, some a little more, depending on your I/O sizes and the distribution of requests across the disk. This graph compares RA60 response times to RA81 response times. The RA60 is not quite as fast seeking as the RA81 so its access capacity, the number of requests per second it is able to satisfy, is a little bit less. Typically the knee in the curve is closer to 20 I/O's per second. We begin our disk I/O investigation by looking at the MONITOR DISK command. Early in the talk I mentioned a few places to start a performance investigation were MONITOR SYSTEM, MONITOR DISK. MONITOR SYSTEM does give some indication of I/O activity, but it does not breakdown I/O rates to individual disks. We need to look at MONITOR DISK for that information.

Disk efficiency is really a question of whether some I/O can be avoided or better balanced among individual disks to prevent a single disk from being a bottleneck. In this example from MONITOR DISK, we see an average I/O rate of 22 I/O's per second and a current rate around 30. We want to investigate what is happening. This is an example of the MONITOR DISK/ITEM = QUEUELENGTH. This is another place you may look to see how big your queues are. This includes the current process. We see the queue length is two, indicating we have requests waiting for other requests. You can look at either the I/O rate or the queue length. If you wanted to get fancy, you could even compute the response times using queueing theory. I guess we don't expect most system managers to get out the calculator and start doing that. When you find a heavily used disk, you want to see what is going on on that disk. Is it system I/O or is it user I/O? System I/O consists of primarily two types: memory management I/O or file system I/O. These I/O's do not show up as direct I/O's in the MONITOR I/O or MONITOR SYSTEM screen. Total disk I/O is the sum of direct I/O's plus memory management

I/O's, plus file system I/O's. To see if you are in current memory-related management I/O's, find out if your page and swap files are located on your active disk. If you don't know where they are, SHOW MEMORY will show the location of those files. If they are located on a hot disk, check for paging or swapping by looking at the MONITOR I/O screen. We talked about most of those items. The disk read I/O rate or the page fault read I/O rate, the inswap rate and also the page write rate, which indicates the number of modified page writes occurring. A memory investigation can be done if you see lots of activity here, look at your working set sizes and do the things we talked about in the memory section. Keep in mind that faults incurred during image activation are page read I/O faults.

The other source of system I/O is file system I/O. The monitor FCP screen is the first place to check for file system activity (FCP stands for file control primitives). The efficiency of your file system caches are next to be checked, use the monitor file system cache command to do this. This is an example from MONITOR FCP, the disk activity is given by the disk read rate, the disk write rate, and the erase rate. The erase rate on the bottom includes operations necessary to support volume high watermarking. If you don't want this security feature and you are seeing lots of erases, turn off volume high marking. Excessive disk read rates may be an indication your file system caches are not configured properly. Check the MONITOR FILE screen. Here we have highlighted the actual hit rates on your file system cache. You also want to consider the attempt rates in order to estimate or calculate the I/O operations being caused by cache misses and increase the appropriate cache sizes. The buffer supplement on performance management was distributed at the VAX System Software area. That handout has an excellent discussion of the VMS file system and things to consider there. If you see high direct I/O rates, home in on the processes doing I/O. This investigation is similar to the one used to identify the highest fault generators. MONITOR PROCESSES/TOPDIO identifies the highest I/O generators. You can see what those particular processes are doing by using SHOW PROCESS/CONTINUOUS. In this output from MONITOR SYSTEM, we see a high direct I/O rate -- 28 I/O operations per second. This is from the same example as before when all the I/O's were being generated on one disk. We see one particular process is generating most of these direct I/O's. The same information is available in MONITOR PROCESSES/TOPDIO and we could also see other processes generating I/O's on this screen. From here it is useful to see which files are being used, which files are open on the disk. This information can be obtained from the SHOW DEVICE/FILE

command. See if it is possible to load balance these files, to move them around to different disks and offload the individual disk experiencing the high I/O rates. This is an example output from SHOW DEVICE/FILES. In this example we see the files being accessed by a process and we also see the files accessed by other processes. We don't see operation rates in this particular screen, but we get some ideas as to which files are being used. Often knowledge of your work load will let you know which file should be offloaded to other disks.

The last area is the area of CPU investigation. The CPU controls the other resources. When we begin a CPU investigation, we may end up doing a memory investigation or disk I/O investigation. The question of CPU capacity can be answered by looking at the MONITOR SYSTEM and MONITOR MODES screens. CPU BUSY TIME, and also IDLE TIME are on those screens. Is there contention for the CPU or is there high usage in areas generally reserved for the system? Here is an example from MONITOR SYSTEM. The CPU usage is at 100 percent and the top user is the process cutter John, who is using 42 percent of the CPU. We see there is a compute queue of four; this also includes the null process. Apart from the null process, there are three computable jobs. Is this good or bad? It depends on what the typical rates are on the system and you know what the processes are doing. Are there is any background batch processes, or is this all interactive? You should be familiar with this number on your system and watch for any deviations. Get to know a typical value for your own system. Just as with memory or I/O, we can see who the top CPU users are. We won't go through an example of this particular case. This system is being used to get a breakdown of CPU time by different processor modes -- look at MONITOR SYSTEM/ALL or MONITOR MODES. We want to look for high usage in this area. Because supervisor mode is used by the command language interpreter, a high rate there may be an indication you are using lots of DCL command procedures. Executive mode is used by RMS and some other database products -- DBMS, RDB, ACMS also use the executive mode. We expect this to be higher on systems using those products than on systems that aren't. Kernel mode is used for page faulting, locking, and most of the VMS system functions.

We are going to look at an example of excessive kernel mode usage. Interrupt stack time is really kernel mode time -- time that can't be charged to any process. It is time that uses device I/O's in a VAX cluster to respond to requests from other nodes. We see fairly high kernel mode usage in this example from MONITOR SYSTEM -- 44-45 percent of the CPU is being used in

kernel mode. We expect this to be below 25-30 percent on most systems. MONITOR does not provide a direct breakdown of kernel mode time, that means we have to do a little bit of detective work. If you had SPM and had anticipated these problems, some PC sampling would identify this kernel mode time. Some kernel mode usage is page faulting. Is the system page faulting -- we have looked at this area before. Are you doing lots of logical name translations; this is available on MONITOR I/O. If there is lots of file system activity, look at the CPU tick rate and MONITOR FCP. We have some examples of all these screens coming up.

What's contributing to this 45 percent kernel mode usage? I'm presenting some numbers based on the capacity of the MicroVAX to do these systems. They are based on the CPU time to do specific operations. You would scale these numbers appropriately for other systems.

This is an example from MONITOR PAGE. The MicroVAX can do about 2000 soft page faults per second and about 400 hard page faults per second based on the CPU capacity to process page faults. By looking at the number in these screens, we see there is only a fraction of CPU being used to process page faults. Page faulting is not a problem. We looked at memory in the memory investigation earlier by looking at this particular screen to check the number of page faults occurring, looking at the kernel mode time, and then moving on to doing a memory investigation. This is an example from the monitor I/O screen. An average logical translation rate is about 7 per second; a MicroVAX2 can do approximately 1000 logical name translations per second. We see less than one percent of the CPU is being spent processing logical names.

This is an example from the MONITOR FCP screen. The CPU tick rate is the number of 10 milliseconds of CPU ticks spent in the file system and on the single processor system. This can be used as the percentage of time spent in the file system. We are spending less than two percent of our time in the file system. The file system is not a problem here. Are some other areas locking? Look at MONITOR LOCK, look at DECnet activity, and also I/O activity. The time doing all these processes is spent in kernel mode.

Here is an example from the MONITOR LOCK SCREEN. MicroVAX can do approximately 1500 ENQUEUE operations per second, including the corresponding DEQUEUE operation associated with that. So 1500 \$ENQ's, plus \$DEQ's, and about 5000 converted ENQUEUE operations. Locking is not a problem for this particular system.

This is an example from MONITOR DECnet. We see average packet rates of less than one packet per second. A MicroVAX can handle over 300 packets per second, thus DECnet is not a problem. Direct I/O and buffer I/O are also processed in kernel mode.

This is an example from MONITOR I/O. MicroVAX can handle over 300 I/O operations per second. Less than a few percent of the system is being spent in kernel mode -- so none of the things we have looked at are a problem on this system.

Where is this kernel mode time being spent? The last thing we can do, or perhaps the first thing, is check the top CPU users. Kernel mode time is charged to a particular process, and we can see which processes are using the CPU. Sure enough. We have our culprit here. It is the process Oliver Wendell. Forty-five percent of the CPU was in kernel mode. He is using 80 percent of the CPU, so it is likely that process was generating most of the kernel mode usage. Not only is this process doing a lot of change mode to kernel, we see this is a privileged user, doing privilege operations. We may also find that he has raised his priorities and can use more of the CPU than other processes. Notice we can identify which process is using the CPU because kernel mode time is charged to a process. This concludes the CPU investigation part.

I want to mention one other thing -- image level accounting. This is very useful for identifying individual resource use on a particular image basis. You can see which images use lots of CPU, which images generate lots of direct I/O, which images do lots of faulting, and you can sort the output from image level accounting.

Converting Code for Parallel Processing

Spring 1987 US DECUS (Nashville) Symposium Session
transcribed by Ron Frederick
N.I.P.E.R.
P.O. Box 2128
Bartlesville, OK 74005

There is a very important question to ask yourself before you begin to parallel an application. Do you really want to parallel this application? Whenever you approach parallelism, be aware you are making a trade-off. You lower your overall system throughput slightly to benefit a single application and give it the quickest turnaround time possible. As an example of this, let's assume you have five VAX8650s on a cluster. If you run five different batch jobs using distributed queue through the batch jobs you will get approximately 25 times the 780 compute capacity as system throughput. Each of those individual jobs will run at the speed of one 8650 or five times the 780.

By employing parallel programming, we concentrate all five of those processes onto a single job with the goal of speeding the turnaround time for that job. This slide suggests 25 times the work of the 780 on your single job. That is, of course, the theoretical maximum you could do given this particular cluster configuration, however, there are various forms of inefficiencies. You won't really get 25 times 780, other than in very rare circumstances, more likely you will average somewhere around 18 to 20. Before you begin paralleling, ask yourself if your site is compute bound as is. Do you need to get 25 times 780 overall system throughput or is the single job important enough that you want to get 20 or so times 780 system throughput to concentrate on that one job. We expect most of the paralleled applications now coming into the market place to run single stream for most of the job. It's those last minute, very time critical jobs that will be run in parallel. You have to decide this before you begin paralleling. Once you make the decision to parallel, what do we have to help you.

First, our Aid Services Division offers courses on parallel programming. We have technical documentation available through our sales force. There are examples in the DECUS library on doing parallel programming in shared memory systems. We have library calls in the same package in the DECUS library to provide routines to help you call these VAX system services.

Rather than calling the system services directly, these routines package them in the way they are most commonly needed and used. What does this do to the issue of VAX compatibility? Any application paralleled for, say the VAX8800, will run on every other VAX processor all the way down to the VAX Station 2000 and the MicroVAXs. Parallel programming uses only a few features of the operating system. Normally we use the ability of setting up global sections, creating subprocesses, and coordinating the activities of the subprocesses through a event flags. Those three features are found on every member of the VAX family. That same job will run on the MicroVAX as well as on the 8800 in parallel.

Next, we will look at the structure of a parallel application. We are going to look at basic questions of paralleling. How much of it has to be changed, how can you estimate the efficiency, and, once you know how to estimate, how do you improve efficiency? Finally, when is it reasonable to parallel an application?

How much of a program has to be changed? The best news I can give you today is, thankfully, very little has to be changed. Most of the third party applications coming to market, like Meta Software's H FIST and the like, say one to two percent of their code actually needs to be changed. We are really adding a structure to the program to control how the existing algorithm will be run in parallel in different parts of the data set. Very rarely do you have to change the fundamental algorithms of the application. You also may have to change the data structure slightly to better facilitate parallelism.

I have some examples of that and a few slides. The execution flow in a single stream code is very simple as shown on the slide. The job starts on the left-hand side of the slide, continues running for some time and, eventually, completes with the right answer. In parallelism, the execution flow gets a bit more complex. The job still starts out in single stream code, but at some point it fans out into parallel sections -- those fan-out sections are the five horizontal lines -- comes back into more single stream code, fans out again in parallelism, and continues this until finally the job completes.

I do have five horizontal lines on there. Only two would imply that these techniques only work for dual processors, which I do not wish to imply. These techniques work for any number of processors. If I had any other even power of two, you may infer to the next processors coming from Digital. Five is a nice safe

number to get the point across without implying too much. Obviously, going from an execution flow in single stream code, like this, to the more complex parallelism execution flow will require use of controlling constructs.

There are two basic types of constructs you may wish to add. One is task scheduling, which we will get to in a couple of minutes. It's a way of improving your efficiency. The second is barrier synchronization. As you can see where the lines are coming out of the work barrier synchronization, we control the transitions from single stream code into parallel sections and out of the parallel section into the single stream code. In more formal terms, barrier synchronization blocks further execution until all of the processes involved in this application have reached a certain location. Let's go back one slide, assume we are dealing with that left parallel section and that we are ready to enter into the single stream code. Let's say that single stream code in the middle is going to write a temporary file giving our results at that point. It would be very impolite for the main process to write out that answer before all the answers are calculated. Barrier synchronization ensures all the processes involved in doing the calculations in that left parallel section have completed their work before we use that data for something else. Barrier synchronizations normally are implemented by using the event flags. Event flags are provided by the operating system and are a very easy means of putting processes to sleep and waking them up later when we are ready to begin the next parallel section.

There is another other decision to be made when implementing parallelism. What will the last process do when it finishes the parallel section? Let's say that middle section -- the single stream code in the very middle of the slide -- was only re-initializing an array for some function that did not have anything to do with the I/O. Any processor or subprocess participating in our application could implement this function. The most efficient way to parallel this case would be to have last subprocess to complete this parallel section, then do the single stream code, and then awaken all the other processes and subprocesses at the start of the next parallel section. If, however, we were writing on a temporary file, the I/O channels would be run by the main process. In this case, the last subprocess to complete would wake up the main process.

The issue of which process continues after the last process has finished the section really is a simple one. The question is what happens in that single stream code. Can any process do it? If so, for efficiency let the last process continue. If this must be done by the main process, again I/O is the usual case, you have to waken the main process. Those are really the only two issues on implementing barrier synchronization. What does the last process do when it completes and how do you implement what it does? Usually these are implemented by the event flags. There might be other, better ways for implementation in your particular case depending upon your application.

Task scheduling is generally the next controlling construct added. A very inefficient way to parallel something, as we will see, (say we are on a dual processor system) is to chop the work in half and assign one-half to the main process and one-half to a subprocess. That can induce a lot of inefficiencies. The opposite of that approach is called task scheduling. Each subprocess, or the main process, dynamically determines which one will do what piece of work simply by which one is finishing up pre-existing work. Putting that again, every process allocates a piece of work to be done and whenever it finishes its own piece, comes back to allocate the next piece available for work. This sounds complex to implement. The five lines of FORTRAN on the screen usually are all that is required to implement this. Use the interlocked instruction to create what is called a critical section -- a piece of code that guarantees only one process can be in at a time. We don't want multiple processes here because we don't want to have two processes all doing task number 17 for example. You want to guarantee that one does 17 and one does 18, therefore, we use a critical section to guarantee only one process can come in at a time. The next line simply takes the shared variable that will contain the next task number, grab a local copy of it (say it is number 17), increment it for the next task, and free up the locks so the next process can come in. Then you want to make sure the number you received is still a valid task number. If there are only 100 tasks defined in this parallel section, you don't want to execute number 105. We will see that, task scheduling can return a great efficiency for the penalty of a very simple code being added to the program.

There are some potential data structure changes you may wish to make and these data structures usually take two separate forms. The first one involves eliminating subprocesses that require an answer from another subprocess and the second type involves multiple subprocesses trying to write into the same shared

variable at the same time. Look at the first of these, the subprocess requiring an answer from another subprocess. Let's pretend for a moment we are going to parallel the following do-loop -- just going from one to 10,000, take a location in the A array, A_i equals the next location at the same array multiplied by B_i in the end of the loop. I hope you will see later on this loop is too small to efficiently parallel directly. I'm merely going to use this as an example of a subprocess requiring an answer from another subprocess. The problem here is the " A_i plus one term" being used as input to the operation of the particular iteration. What happens if the particular iteration has already completed work. Let's look at the fifteenth and sixteenth iteration. If the sixteenth iteration is assigned to some process and has finished and modified A_{16} , whenever the fifteenth iteration tries to access A_{i+1} for that same sixteenth location, it will receive the wrong information. Subprocesses requiring answers from another process have two forms, one of which is an actual value that gets passed to the other process, but the other form is more tricky. The requirement is the knowledge that the other process has continued beyond a certain point. The fifteenth iteration has for example already read the sixteenth value of A before the sixteenth iteration can continue. Paralleling directly and putting in locks to force the fifteenth to finish before the sixteenth can continue would be very tricky here.

As in a lot of cases of data dependency, There are other very simple ways around data dependency. One way relies on the realization that in the single stream code, i always goes from 1 to 10,000. Incrementing that, the A_{i+1} term that is referenced always references a value in A that was there before the do-loop was even started. There is an easier way to parallel this. Rather than trying to put controlling constructs into this do-loop to prevent them stepping on each other, simply copy the A array into a temporary array and modify the original do-loop to reference the temporary values rather than the original A values. Now we have two loops that can each be run in parallel very efficiently with no data dependency. You can now run that second loop in reverse from 10,000 to 1 and it will make no difference. Every iteration will get the proper value it expects -- the original, proper value in the A array. The use of temporary arrays is a very common technique for getting around data dependency problems in paralleled code. What looks complex is the locking mechanism. You can very often get good efficiency by using temporary arrays and eliminating data dependency.

The other type of data dependency is multiple subprocesses writing to a shared variable at the same time. Here is a classic case in computer science of two processes trying to do " i equals $i+1$ " at the same time. The results are really predictable if two or more processes are trying that at the same time. It simply becomes a race condition of who has read memory, incremented it in its microcode, and written it back out. Chances are very good that race condition eventually will turn around and hurt you by giving the incorrect answers to the program. Finding areas where multiple processes could be writing to the same shared variables and stopping that from occurring are key. Pretend we are going to calculate pi to some absurd number of decimal places, like a million or so. The classic equation for this being that pi is 4 times the series $1-1/3+1/5-1/7$ -- on ad infinitum until we finally complete the four millionth decimal place. We are trying to stop two or more processes from being able to write their value into the same array location, in this case array location 1,000, at the same time. To calculate this to a million decimal places we store the numbers in a million-entry array for extended precision calculation.

There are three different techniques you can apply in parallel programming to eliminate this kind of problem. Not all three work in all cases. The first case involves segmenting the work to eliminate the data share. This is by far the most common way of approaching this problem and works in a vast majority of cases. In this particular case, sad to say, it won't work. What do I mean by segmenting the work? Let's say we are calculating the "1 divided by 21st" term. We could have all of the processes participating calculate "one divided by 21" out to a million decimal places and then participate in adding or subtracting from the main pi array. If this is broken into groups of 1,000 digits, the first 1,000 have to be completed before we can start to calculate the 1,001 digit. Therefore, we are really back to single stream code. Segmenting the work is not useful for us in this case because part A has to be completed before part B can begin working. Look for places where both A and B can be done at the same time if you segment the work.

The second common way is to have a separate lock for output location. In this case, we will assign each subprocess a particular term in the series. One process will do the entire "1 divided by 21" out to a million decimal places. Before it updates this $pi_{1,000}$, the process will access a lock, gain exclusive ownership of that lock, write the value (adding or

subtracting as the case maybe), release the lock, and continue on calculating the 1,001 digit. This technique will work for an even number of variables, though if we are dealing with a million lock variables the continual locking, updating, unlocking, and moving on to the next digit will be prohibitively expensive. This method works for a small number of variables but is not good when you are dealing with arrays the size of a million.

The third common method, and the one that will work quite well in this case, is for each process to have its own temporary work array with a single permission to modify the master array lock being in place. One subprocess calculates "1 divided by 21" out to a million decimal places into its own temporary array, allocates a single lock that grants it the right to modify the master pi array, does its addition or subtraction, and then frees up that lock. In this case, there is more memory to choose from. Each process must have an additional one-million element array. Paralleling this way will give you good efficiency with very little overhead. These are the three most common techniques of approaching the problem. Which of these is best suited for an application depends totally upon the application.

By looking at this slide, we begin to see the degree of the original program to be modified. On the left, we show what you could consider to be the virtual address range of the original single stream code. In the center and to the right we modify it by adding the parallelism additions in the yellow column. There are various controlling constructs in the code that have to be added as the yellow shows in the vertical bar. Some of these techniques have been mentioned using temporary arrays as ways around parallelism problems. These are in the yellow box off to the right. Generally, there are more temporary arrays added to the program. If you look at the percentage of yellow versus orange, the percentage of yellow on this slide is highly exaggerated. As I suggested before in the earlier application, one to two percent needed to be added. McNeil-Schnedler's is in the same category -- about one to two percent. Abacus is the same way. We felt we used enough applications to begin to see a pattern. Very rarely, if ever, will you need to go as high as five percent addition or modification to the original single stream code. The best news is that time investment to the original single stream code for the most part is preserved. You are not throwing out the original code and rewriting it. You are merely adding parallels and controlling constructs around the algorithms that already exist.

The next section is how to estimate efficiency and, as a side corollary, how to prove the efficiency. There are really four important metrics to consider. How much work is in each task we are doing the allocation on, how expensive is this task scheduling, how many tasks are there compared to the number of processors in the machine, and, finally, how expensive is the barrier synchronization concept. We will look at most of these in a few slides. How expensive is the task scheduling? Recall the five lines of FORTRAN code on the screen. The compiler will turn those into 16 to 20 VAX instructions to implement the task scheduling. By implementing this in MACRO, you can call the interlocked instructions directly and bring the number down to about six or seven. We will average this, for convenience, to 10 instructions required to perform the task scheduling overhead. The absolute number is not important. What is important is the ratio between the size of the task scheduling and the amount of work in each task. How much work is there going to be in each task? In general, the larger the task, the better off you are going to be. You can go overboard though and start to run yourself into inefficiencies by having tasks that are too large.

The ratio between the task scheduling overhead and the amount of real work being allocated in that task is one of the key factors in the efficiency you receive from parallel programming. This is by far the most important thing to keep in mind. What can happen if you don't have a good ratio? The top of the slide shows the very bad situation of a high scheduling overhead. Each block refers to one VAX instruction (for convenience sake I've made them all the same size) and the two lines are the two processors in today's dual processor systems. The task scheduling for purposes of this talk will average 10 instructions. The 10 instructions of overhead are shown in the 10 yellow blocks that occur before the real work begins. The real work in this case shows as the three instructions in green. We are going to go through 10 instructions of overhead to execute three instructions of work. That results in seven instructions equivalent worth of time blatantly wasted. The process cannot execute because there is no more work already assigned to it, nor can it begin the task scheduling -- that is something that only one process at a time can do. If you mentally overlap those two lines we are running this original program, we are running the parallel program at 30 percent the speed of the original program and burning up two processing elements getting there. This is not a situation to get into when you are parallel programming. It is very important to make sure that your task size and the amount of real work compares

favorably to the size of the overhead.

We see a situation at the bottom of the slide where the task scheduling overhead is significantly better. Before anyone accuses me of cheating and dropping task scheduling to a single instruction, please note the scale at the bottom of the slide has changed. This new scale is one block equals 10 instructions. This shows the relative efficiency of having that 10 instructions of overhead, but now instead of only doing three instructions of real work we are going to allocate 100 instructions of real work. This shows quite graphically that just making that slight change in the program dramatically increases our efficiency from running at only 30 percent of the speed of the single processor version to now running about 1.8, 1.9 times efficiency, if not better.

How do you change an algorithm to get significantly larger tasks in only three instructions? We are going to look at a matrix multiply operation, envision different ways to parallel it, and see what the results would be. This is going to be a 100 by 100 matrix multiply. The inner loop, the calculation to do a single output location, is shown there on the screen. As quick reminder, in a matrix multiply operation we multiply each of 100 elements in the highlighted row in matrix A by the appropriate 100 elements in the highlighted column B and add it into the single highlighted location in the C array. We could parallel each one of those do-loops with each iteration of the do-loop as a separate task. If you look at the number of instructions required to do the multiplication and addition, you see we are down to around three instructions. Given the average of 10 instructions of overhead, this leads us directly into the overhead problem on the top half of the slide. Clearly that is not an efficient way of paralleling. We could assign the tasks as "do all the calculations required for a particular location in C." In this case, one task would do that entire inner loop for the highlighted location in C and the next task would do all the calculation for an adjacent location in C, or somewhere else in the array. Now instead of having three instructions with the 10 instructions of overhead, we are up to about 300 instructions. Now we have one-third the scheduling of overhead shown at the bottom of the slide. That slight change in approach to paralleling leads us from the horrible inefficiency at the top of the slide to only one-third the overhead that's at the bottom of the slide. It is very important to keep the task size in mind when you are doing parallelism.

There are other ways we could parallel this. We could parallel a task being one output row or column and end up with 30,000 instructions in a task. At that time you almost wouldn't see yellow any more. We could parallel it where a task is one-tenth the output matrix and be up to 300,000 or we could go overboard as shown in the next slide and simply divide the output matrix in half (if we are on an 8800 today) and have 1.5 million instructions per task. There are inefficiencies involved in this level of parallelism, the high end, but I will be showing this on the next slide.

How many instructions per task do we see on actual applications? We see multiple thousand instructions. In HSPICE, the calculation for each instruction is on the order of 2,500 to 3,000 VAX instructions. Paralleling it was a natural way that the program was paralleled. We didn't have to do anything special to get it to that level. We simply looked at it and said what is a reasonable way to parallel this and that task size naturally came out of it. Simply looking at any level other than the very tight three instruction do-loop I showed you earlier often will result in task sizes large enough to drive the parallelism inefficiency down to a level that is hardly noticeable. It is very easy to get tasks of appropriate size. What is a good ratio between the number of tasks and the number of processors you can have participating in this? A one-to-one ratio as in the last slide. Simply saying "We are on a dual processor today -- let's cut the array in half and assign one-half to each side" is very inefficient. The problem here is twofold. First off, you the programmer would have to decide where that halfway point is. In something as straightforward as a matrix multiply, this is not difficult to decide, however, it is nearly impossible at programming time to know what the customer's circuit simulation is going to look like and be able to divide it in half. It is a very difficult thing to do in most circumstances.

You hope computing progress will be identical for every process and subprocess involved in your application. A goal of parallel programming is to have every piece of your job that is participating in a parallel section complete that parallel section in as short a time as possible, or in as similar time as possible.

If you divide the matrix in half and one half hits a half dozen page faults, that particular process (half of the array) is not going to complete at the same time as the other. While those page faults are being serviced and that half of the array is being calculated, the other processor is not doing anything. This is very inefficient for parallelism. A one-to-one ratio is clearly not good. You have to make too many assumptions, both in the program and in the computing progress of your processes. A 10 to 1 ratio is a nice figure to try for. With a 10 to 1 ratio, if you do hit some page faults or if one of your processors does go out on the compute queue for a little while, there are enough other tasks left around the processes that are running on the processors still have more work to do and you can keep your efficiency very high. It doesn't mean you have to shoehorn your program to get exactly a 10 to 1 ratio. If I were paralleling the 100 by 100 matrix multiply, I'd probably make each task an entire output column. This would give us 100 tasks or, on today's 8800, a ratio of 50 to 1. Ten to one is a nice number to shoot for; there is nothing magical that says a 9 to 1 or 11 to 1 is going to run you into inefficiency.

Finally, the expense of barrier synchronization. Barrier synchronizations are implemented using the event flags as system services much of the time. As with any system service, there is little overhead in calling them. In this example there are about 500+ instructions hit whenever you do an event flag system service. You therefore should make sure your parallel section is going to be large enough to justify this additional overhead and still return good efficiency. In a 100 by 100 matrix multiply with about 3 million instructions, clearly 500 instructions of overhead will not factor into anything substantial. If you were trying to parallel a three by three matrix multiply, or a four by four, the event flag system service would utilize most of the time and would be very inefficient. Make sure the parallel section is big enough to justify the overhead of starting a parallel section.

There will be times when you will not want to parallel something. That sounds like a strange way of ending my talk. I've been saying it's easy, you only have a half dozen things to watch out for and parallel programming really is easy. There are only a half dozen things to watch out for, but there are some situations which you should not or do not want to parallel. The first one simply involves the overhead consideration. If paralleling results in tasks too small relative to the overhead of scheduling the next tasks, you would not wish to parallel in that way. You might be saying, it's easy to get 100

instructions versus 10 instructions of overhead in the shared memory systems. In particular I am thinking about paralleling across the VAX cluster or, worse yet, paralleling across DECnet. Now the amount of time to allocate the next task is not 10 instructions; it can be measured in the thousands. If the amount of real work to be allocated does not scale up to the appropriate ratio and overhead is doing the task allocation, this is going to become too great. There are some cases where overhead considerations will be so high you really won't want to parallel the application in the style that you are approaching.

The second consideration is the overall system throughput. There is a trade-off in parallel programming as I showed you in the example of the five 8650 cluster. You are saying I'm willing to accept slightly lower overall system throughput in exchange for a significantly better single job turnaround. In most cases I would suggest that running multiple batch jobs on the 8800 is the most efficient way of using it. Parallel programming really pays off for those critical last minute jobs at the end of the design cycle, when you are willing to trade away a few percent in throughput for this overnight job being able to run during the daytime.

Finally, the third consideration for not paralleling something is called Amdahl's Law of Parallelism Paybacks. In short, this says the potential payback from parallelism is directly related to the amount of the program that can be run in parallel. Let's say we have a hypothetical program with 16 minutes of fully parallelable material and zero inefficiency calculations. There are 10 minutes of initialization code and, at the end, 10 minutes of results storing, writing out a report file, or whatever. We could parallel that parallelable computation section and not get down to 30 minutes of real time. Instead of this job dropping from 80 minutes of execution time to 40 with 100 percent parallelism over two processors, we see the execution time really drops to only 50 minutes. We will end up with a 1.6 times faster application. We have perfectly paralleled the parallelable areas, our tasks are large enough, there is zero scheduling of overhead, and we still did not achieve a 2.0 times throughput. We are still stuck with the 10 minutes of initialization code and 10 minutes of results storing. A 1.6 times speed up on an 80 minute job now taking 50 minutes is a considerable speed up and is probably worthwhile for a critical application. If this only had six minutes of calculation that we could drop down to three, our overall job time is only going from 26 minutes to 23. That is hardly worth the hassle of parallel programming. Before you begin to

parallel something, be sure the process is going to spend enough time in the particular sections you want to parallel for parallelism to be worthwhile.

In summary, parallel programming is doable, is doable today under version 4.x whatever at VMS, as is shown by some of these commercial applications being marketed in parallel form. It does not mandate a total rewrite of these applications, most of these programs are saying definitely under five percent. Usually one or two percent of the code needs to be modified and can return a very high efficiency. Before you ask about the performance of those third party applications, I can't answer. You have to be sure any new programming technique is applicable to the program. Be sure you are willing to make the few percent trade-off on system throughput and that you are spending enough time in the paralleled sections for truly good payback. With that I will open the floor for any questions you may have.

Q: Dave Marlin, American Electric Power. There are two aspects to parallelism. One of them is what you have been talking about, performance increase; the other is fault-proof systems. You are running the same applications in parallel and comparing output. Is DEC's approach to that type of parallelism going to be enhanced in future cluster releases?

A: We can't comment on future directions.

Q: Is there going to be any direction on that or is DEC going to depend on the very low mean time failure of the processors to ensure that this is not really necessary.

A: Answering your question would be answering in a way I can't answer. We can't comment on future directions, though basically high availability is what you are asking for. I cannot answer that.

Q: In the other aspect, the performance aspect, you have two processors and, am I to utilize an 8350 or an 8000, is it going to be necessary for me to go back and rewrite a lot of my user mode applications. As I understand kernel mode, a lot of my own intensive processes cannot be used in that second processor. Will I be required to rewrite a lot of my applications or

A: Rewrite when?

Q: When I add the second processor or is VMS going to throw user mode processes out there.

A: If you wish to take advantage of parallel programming to speed up single job turnaround, you have to implement parallel programming within it yourself. If you wish to use it as a multiple batch job compute engine, you do not have to do any reprogramming. VMS will schedule your jobs automatically on the separate processors. Parallel programming you must do yourself.

Q: Right now, out of the VAX environment, in an RST environment we run 14 KXTs on a single backplane. KXTs -- an important difference. That's doing very well for us. We are anxious to get into the VAX at that level.

A: I'm sure you will find that parallelism on the VAX is very easy and gives you good efficiency.

Q: Nowlin Olsen, Scientific Computer Systems. How strongly do you emphasize checkpoint and restart and how is that significantly changed when going to parallelism?

A: Checkpoint restart is more like a VMS question than a parallel programming question.

Q: If you have a very large compute intensive job, I guess I'm asking how much emphasis needs to be applied to building in checkpoint and restart to do that job?

A: You mean adding your own checkpoint and restart to a particular application? If I have a very large compute-bound application, just internal number crunching very similar to the pi application, I do my own checkpoint restarting within it. If a job is going to run for multiple days, multiple weeks, it would make sense to protect your investment. That's my personal opinion. That says nothing about where DEC's direction is going in checkpoint and restart.

Q: In that level of checkpointing restart, it seems to me the complexity increases when you go into the parallel programming.

A: Really, I don't believe the complexity would increase that much. Parallel applications fan out into the parallel section to compute the cycle of circuit simulation, for

example, then come back into single stream code to do some maintenance and clean-up. I would see checkpoint restarting as something to be done in the maintenance portion in single stream code where everything is at a known state. It really wouldn't be hard to add.

INPUT/OUTPUT

A SIG Information Interchange

A form for INPUT/OUTPUT submissions is available at the back of the issue.

To register for on-line submission to the Pageswapper dial:

(617) 262-6830

(in the United States) using a 1200 baud modem and log in with the username PAGESWAPPER.

```
=====
Note 421.1      ReGIS to GKS/0b Conversion      1 of 5
"Mack W. Haynes, Jr."      10 lines  22-SEP-1987 09:17
      -< GKS/0b - A suggestion >-
-----
```

I have submitted SPR's to DEC requesting that all software supporting Regis output also support GKS/0b since Regis is now a dead product and DEC has thrown it's support behind the GKS standard. Additionally, SPR's have been submitted requesting that the DEC GKS/0b software support standard output devices such as Electrostatic plotters and Laser Printers. Many of us in the industry use these devices as standard output.

I would like to request that others with similar situations make like requests to DEC. So far, DEC's response has been "We're investigating the possibility".

Mack W. Haynes, Jr.
Tenneco Oil Co.
1100 Louisiana

MS: IFP2954
P.O. Box 2511
Houston, TX 77001
(713) 951-1118

```
=====
Note 421.2      ReGIS to GKS/0b Conversion      2 of 5
"Bill Mayhew"      21 lines  22-SEP-1987 10:06
      -< ReGIS: Dead or Alive? >-
-----
```

Is there more detail about this "ReGIS is a dead product" notion? I have seen the press notices that the "ReGIS Graphics Library" is a dead product, but not ReGIS itself. The distinction is that the RGL is a distinct layered software product providing an interface to ReGIS for the applications programmer, while ReGIS itself is a graphics communications protocol embedded in the firmware of the VT125/VT24x/VT3[34]x et al., which causes a stream of bytes to be converted to a graphic image.

My perception of GKS is that it, too, is an application programmer's interface, not a communications protocol. Thus, GKS could be implemented to talk to any of several device-dependent or -independent communications protocols, and that, therefore, there would need to be something like a GKS-to-Sixel software module to talk to LN03s, for instance.

I am certainly not a graphics software wizard, however, and what graphics I do tends to be either directly at the protocol level (i.e. ReGIS) or depends on intervening higher-level products.

Is my perception accurate? What sort of communications protocol is being advanced as a replacement for ReGIS? PostScript?

Bill Mayhew
Village Systems Workshop Inc
PO Box 642
Natick MA 01760
617-237-0238

=====
Note 421.3 ReGIS to GKS/0b Conversion 3 of 5
"Bob Hassinger" 32 lines 22-SEP-1987 11:42

-< Some thoughts about GKS... >-

1) The DEC product has not been "GKS/0b" for quite some time (that was V1 I think). Version 3.0 is in the field ("VAX GKS"). It now implements level 2c, the highest level defined in the GKS standard.

2) VAX GKS V3 supports user written output drivers so you can support what ever you like if you want to do it badly enough. It is not likely we are going to see DEC supporting too many non-DEC printer protocols I would think. V3 does support the LN03-Plus and the LN03R as well as Postscript printers in general. They do not seem to support the LN03 however.

3) There is work being done on standardized device level graphics interfaces that would more or less do the same job Regis does. The future there is not real clear to me yet.

4) As to software products that speak Regis getting a GKS interface: In general the software has to be linked to the GKS library so, in the VAX GKS case, users must be licensed for the run time library at least if not the full package. Any product that is built this way could be used by any user with the RTL to work with any supported device. The RTL license is a serious issue here however, as it is for FMS applications for example. It is also very tricky to write code that uses GKS and which produces GOOD results with all supported devices - their characteristics vary FAR too much. Regis is a MUCH simpler case because the number and type of devices and the variation in their characteristics is MUCH smaller! GKS is a big step forward but true device independent graphics is still a long way off.

Bob Hassinger
Liberty Mutual Research Center
71 Frankland Road
Hopkinton, MA 01748
617-435-9061

=====
Note 421.4 ReGIS to GKS/0b Conversion 4 of 5
"Chris Erskine" 17 lines 25-SEP-1987 08:42

-< Differences between REGIS and GKS >-

Looking at the VT330 and VT340 terminals, I do not think you can say that REGIS is dead. DEC has just learned how to make REGIS run faster.

As stated, REGIS is the device protocol and has nothing to do with GKS. This is like the Tektronix 4010 graphics protocol or the HPGL protocol that is used by other display devices.

GKS is the graphics interface to the application programmer and is used to replace the different graphics that everyone had for their machine. This standard is designed that a person who writes code on a VAX with DEC's GKS package can take the same program and use the code on another machine with a GKS package. It is the device drivers which is part of the GKS system which talks REGIS.

Chris Erskine
6001 Adams Rd.
Bloomfield Hills, MI 48013
(313) 258-4049

=====
Note 421.5 ReGIS to GKS/0b Conversion 5 of 5
"Jack Patteeuw" 15 lines 29-SEP-1987 13:23

-< More features from the VT330 >-

re: .3

"There is work being done on standardize device level graphics interfaces ..."

In fact the GKS subcommittee (is it from ANSI or ISO or both ?) has proposed (and I think include in level 2c) a "metafile" format for the passing of graphics information between system (software and/or hardware).

I have heard a rumor that the VT330/340 supports this metafile format directly (or will in the near future). This would mean a significant improvement in drawing speed because the CPU would (theoretically) have less layers of translation to go thru to get to the final device format.

Jack Patteeuw
Ford Motor Co.
Electrical and Electronics Division
31630 Wyoming
Livonia, MI 48150
313-323-8643

```
=====
Note 601.6          high-speed dialup modems?          6 of 6
"MICHAEL GRATTAN"          20 lines    3-SEP-1987 08:55
          -< Check the Mileage >-
-----
```

We have been using Anderson Jacobson 9631-S between our central office here in New Bedford, Mass. and our warehouse in L.A. on the dial-up network. (The 9631 is a 9600 baud sync. modem for leased or dial lines.) We are going dial-up while we wait for the leased line to be installed.

Dial-up has been VERY disappointing. We have had to use the fallback speed (4800 baud) and the line quality has been so bad that the line has dropped in the middle of a session. It's a huge headache.

I think that the modems are just fine. However, I would strongly caution anyone considering a high speed modem for dial up to look at the distance involved. Local connections were fine and fast, but long distance is a killer.

Murphy's nth Law of Data Communications: Line quality is inversely proportional to distance. ;-)

MICHAEL GRATTAN
FAIRHAVEN CORP.
358 BELLEVILLE AVE.
NEW BEDFORD, MA. 02742
617-993-9981 EXT 106

```
=====
Note 631.2          RA81 FCT File Format          2 of 2
"Offline Submission"          13 lines    29-SEP-1987 12:35
          -< RA81 FCT File Format >-
-----
```

The factory volume id occupies the first 64 bits of the RCT - block 0. This can be read, a block at a time, using a physical function from the area immediately following user data. (for RA81, P3 = 891072).

Bill Noble
Shell Canada
Box 100
Calgary, Alberta T2P 2H5

Telephone: 403-232-4370

September 11, 1987

```
=====
Note 693.18        Opening a file with NIL sharing      18 of 19
"Jack Patteeuw"          20 lines    22-SEP-1987 17:58
          -< Problem solved ! >-
-----
```

VMS V4.6 solves the problem !

To quote the release notes, page 2-17:

"The Version 4.6 VAX C Run-Time Library supports file sharing when you use record mode to access files; you must use the ctx=rec file attribute with all file open functions. Specify the shr=xxx (ie. shr=nil in this case; jp) file attributes as appropriate."

In addition ...

"Version 4.6 improves stream I/O facilities in the VAX C Run-Time Library. You can now specify the mbc=xxx file attribute when opening stream files. The value for this attribute specifies the number of blocks to allocate for I/O buffer. Reads and writes are performed doing this block size."

It is implied that these functional improvements are only available if VAX C V2.3 and VMS V4.6 are both installed.

Jack Patteeuw
Ford Motor Co.
Electrical and Electronics Division
31630 Wyoming
Livonia, MI 48150
313-323-8643

```
=====
Note 702.6           Planning an Ethernet           6 of 8
"R. Michael Dupont"      18 lines   1-SEP-1987 17:59
      -< More features for less spectrum >-
-----
```

We are currently using our multi-plant broadband system for video and data, and since we will be adding even more channels our spectrum space is considered scarce. We initially looked at both the Digital and Chipcom modems, but felt the 18 MHz could be better utilized.

Even before Chipcom announced their new line which uses only 12 MHz, we found and tested a product that only requires 6 MHz of broadband and still gives us the full 10 MHz of baseband Ethernet.

The unit also acts as a bridge, filtering local baseband segment traffic from the broadband "backbone" segment. Several months of beating on the units (BETA) have shown us the units induce no performance degradations to the Ethernet (802.3 as well) traffic.

Non-disclosure prohibits my revealing any further details, but I will say the product is not from DEC or Chipcom.

R. Michael Dupont
EDS
2925 West Minnesota
Indianapolis, IN 46241

```
=====
Note 702.7           Planning an Ethernet           7 of 8
"Chris Erskine"         23 lines   2-SEP-1987 08:53
      -< A poor way to reach broadband >-
-----
```

Even before Chipcom announced their new line which uses only 12 MHz, we found and tested a product that only requires 6 MHz of broadband and still gives us the full 10 MHz of baseband Ethernet.

Ungerma Bass has a device called a buffered repeater h works in a 6 MHz channel. We found that it has a couple of problems. They do not do any collision detection on the broadband section of the media. This breaks part of the ethernet standard which states that the message will be transmitted since the trashed packet from the buffered repeater is dropped. UB states that it is the responsibility of the upper protocol to handle the lost packet. This is fine with ethernet not always receiving the packet but will cause some problems with DECnet on RSX machines. The minimum retransmit time for RSX is about 4 sec. I was told by DEC that this time would be fixed but so far I have not heard of the fix.

I do not know if there are any other protocols which have such a long timeout value which will cause the same type of delay but I do not like to add devices which help add such a delay to the network.

Chris Erskine
6001 Adams Rd.
Bloomfield Hills, MI 48013
(313) 258-4049

```
=====
Note 702.8           Planning an Ethernet           8 of 8
"R. Michael Dupont"      25 lines   4-SEP-1987 17:59
  -< Better (much better) than a Buffered Repeater >-
-----
```

Ungerma Bass has a device called a buffered repeater which works in a 6 MHz channel. We found that it has a couple of problems. They do not do any collision detection on the broadband section of the media. This breaks part of the ethernet standard which states that the message will be transmitted since the trashed packet from the buffered repeater is dropped.

We also have tried the Buffered Repeater, with the drawback of being on the same broadband channel as all of our terminal traffic (400+ sync and async combined), and the buffered repeater pulls all of those packets onto the baseband segment, along with the desired traffic from the any remote buffered repeater.

The unit I alluded to before is an announced product from Applitek. It uses vector phase technology (much like the 9600 baud modems over phone lines) to squeeze 10 MHz into the 6 MHz broadband channel. They went "buffered repeat" one better by adding "Bridge" functionality to the unit. This keeps local traffic local and doesn't clutter up the broadband "backbone" connecting the local segments.

The bridge can be programmed to filter if desired, but we have no wish to do so at our site.

R. Michael Dupont
EDS
2925 West Minnesota
Indianapolis, IN 46241

```
=====
Note 719.3           BACKUP status codes           3 of 3
"Marc Lippmann"        9 lines   23-SEP-1987 10:02
  -< Improve F$MESSAGE >-
-----
```

One (I think) extremely valid possibility would be to allow the specification of a message file as an additional parameter to the f\$message lexical function, i.e.:

```
$ MESS = F$MESSAGE($STATUS,"SYS$MESSAGE:SYSMGTMSG")
```

or even better, allow a list, such as

```
$ MESS = F$MESSAGE($STATUS,-
  "SYS$MESSAGE:SYSMGTMSG,SYS$MESSAGE:CLIUTLMSG")
```

Marc Lippmann
PO Box 355
210 Grove Street
Franklin, MA 02038
(617)528-6200 X330

```
=====
Note 722.1           Long Distance Data Communications   1 of 1
"JIM PALMER"        23 lines  11-SEP-1987 13:05
  -< Satellites and DECnet >-
-----
```

We have been using a satellite service to connect DECnet nodes between here (Irvine, Ca.) and Miami, Florida for a number of years. We are running a single 9600 baud SYNC line via 2 DMF's and associated modems.

The service is fairly reliable. (about 97%) uptime. However, we use to have a dedicated AT&T land line between the exact same nodes. The land line was absolutely 100% reliable, with no light speed delay.

The reason that we switched was to save money. Since the DECnet traffic on that particular wire is fairly light, the satellite link is workable.

Our carrier is a company called American Satellite.

JIM PALMER
3 BROOKDALE
IRVINE, CA. 92714-3338
(714) 458-3028

=====
Note 740.0 Quasi-dynamic Priority Adjustment 3 replies
"Dale E. Coy (505) 667-3270" 23 lines 4-SEP-1987 14:15

We are experimenting with using a dynamic priority adjustment routine. The basic idea is to lower the priority of jobs which are compute-bound, so that they interfere less with people doing edits, etc. Of course, there's more complexity than that. In particular, lowered priorities sometimes need to be raised again.

We're using a MACRO code written a while ago, updated for VMS 4.x, but we're also looking at SMAUG from the VAX86C tape. Both our code and SMAUG use SYS\$SETPRI.

Very occasionally (like once a month) we see some process which should have a maximum base priority of 4 get raised to 15. Of course, nobody does anything until the process completes. Possibilities include: 1) A bug in our code which we haven't found yet, and 2) A tiny problem window somewhere in SYS\$SETPRI.

Before we chase too many phantoms, has anybody else seen a similar problem? Particularly with SMAUG?

Comments on whether we're totally crazy to consider this will also be considered. :-}

DALE E. COY
LOS ALAMOS NATIONAL LAB
PO BOX 1663, MS J957
LOS ALAMOS, NM 87545
505-667-3270

=====
Note 740.1 Quasi-dynamic Priority Adjustment 1 of 3
"Bruce Bowler" 20 lines 8-SEP-1987 09:15
 -< the culprit is \$GETJPI >-

Dale, We too experimented with a 'dynamic priority adjuster' from one of the sig tapes. I don't remember which tape it came from but the program is called ZEUS. We saw a similar problem with jobs occasionally getting a boost to 14/15 and after numerous calls to DEC discovered that the problem is not with \$SETPRI as I had expected, but rather with \$GETJPI and the scheduler. It turns out that in order to get some of the information that can be retrieved from \$GETJPI a (kernel-mode?) AST must be delivered to the target process and, due to the combined bug with \$GETJPI and the scheduler, this raises the priority of the target process to that of the requesting process. We were running ZEUS (and several other process monitoring processes) at priority 14/15 in order to get (what we thought were) the most accurate readings possible. When we lowered the priority on these guys down to 6 the problem still existed but to a much less damaging extent.

Interestingly enough, it appears that this is not a problem with the requesting process is run at one of the 'real-time' priorities since SPM runs at 24 and we see no inordinate priorities on our timesharing jobs.

Bruce Bowler
General Electric
1 River Road
Bldg 2 Room 609
Schenectady, NY 12345

```
=====
Note 740.2      Quasi-dynamic Priority Adjustment      2 of 3
"Jamie Hanrahan"      20 lines      8-SEP-1987 15:54
  -< It's not a bug; it's correct behavior >-
-----
```

< Note 740.1 by NODE::US140424 "Bruce Bowler" >...

..is correct; however, the reported behavior of \$GETJPI (the scheduler has nothing to do with it) is not a bug, but quite intentional. The reasoning goes like so: Process A does a \$GETJPI on process B, which has a lower base priority. To get the info, \$GETJPI sends a special kernel AST to process B. Thus, process A can't continue until process B does the AST. Process B's priority is boosted to equal process A's to avoid a potential deadlock which might be caused by having a high-priority process dependent upon action by a lower-priority one.

Note that B's boosted priority will fall back to its normal range as soon as B has been rescheduled a sufficient number of times (15 times will always be sufficient).

Running process A at a real-time priority (16 or above) will not cause B to get priority boosts for the \$GETJPI calls. This is because a priority boost that would take a process into the real-time range is never applied.

Jamie Hanrahan
Simpact Associates
9210 Sky Park Court
San Diego, CA 92123
619-565-1865

```
=====
Note 741.0      750's and LAVC's      7 replies
"JIM PALMER"      57 lines      7-SEP-1987 23:23
-----
```

We have just finished retrofitting our mixed network of 750's and MicroVAXes to connect via the NI/ETHERNET medium with encouraging results.

Naturally, the next logical step is to configure a LAVC.

According to all LAVC literature that I have, and all DEC sales reps that I have talked to, it is NOT possible for a 750 to participate as a satellite (as opposed to a boot member) in a LAVC.

From an engineering perspective, I can see no technical reason why such a constraint should exist. As far as I can determine, all that would be required to boot a 750 over an ethernet would be the addition of a boot ROM and/or VMB that knows how to place a boot request message on the net. (In fact, the DEUNA has a DIP switch provision for this, indicating that some one was thinking ahead...).

Albeit, some talk has surfaced regarding potential performance problems. But, then again, what NI/CI based network hasn't any? [1]

This is a shame, because we have numbers of loyal, reliable 750's that could be reincarnated as gateways to Brand-X and Plain Wrap networks, batch servers, CMS library servers, etc., etc., etc.

Thus I perceive this omission as purely DEC marketing oriented, concocted for reasons unknown.

Now, our company, like most, does not have an infinite amount of money to spend to change out all our present VAX-11 with the latest generation of trendy hardware that comes along. And, even if we did, we would find ourselves locking horns with the logistics of such.

Does anyone have any information/opinion to the contrary? Or, better yet, has someone booted a 750 over NI wire? Or, better still, is this feature actually supported by DEC but not widely publicized?

[1] To be sure, it probably wouldn't be a grand idea to place your share database on the slowest disk on the slowest node in the cluster....

JIM PALMER
3 BROOKDALE
IRVINE, CA. 92714-3338
(714) 458-3028

=====
Note 741.1 750's and LAVC's 1 of 7
"John Osudar" 8 lines 8-SEP-1987 15:21
 -< 750's in LAVC >-

DEC's official position is that the only CPU's that have the capability to boot over the Ethernet are MicroVAXes. As you point out, it would require a suitable boot ROM to get this to work from a 750 -- and if DEC has such a part, they aren't selling it. A related question, that has been asked before & elsewhere, is whether you can boot a non-MicroVAX from a local disk and then have it join an LAVc as a "satellite" that has its own local disk. That appears to be feasible (though officially unsupported).

John Osudar
Argonne National Laboratory
9700 S. Cass Ave.
Bldg. 205 A-051
Argonne, IL 60439-4837
(312) 972-7505

=====
Note 741.4 750's and LAVC's 4 of 7
"Dale E. Coy (505) 667-3270" 35 lines 10-SEP-1987 00:21
 -< A modest proposal, and more fuel >-

We have just finished retrofitting our mixed network of 750's and MicroVAXes to connect via the NI/ETHERNET medium with encouraging results.

This may not be much help, but you didn't say what your mix of hardware is...

A bit better than one 750 boot-node would be two 750s. LAVC 1.2 (VMS 4.6) allows two per cluster. Then, of course, you can have multiple clusters on the same Ethernet by giving each cluster a different group number. Thus, if you have at least 1 MicroVAX per each pair of 750s..... :-)

A couple of other notes on LAVC experience:

1. I recently ordered LAVC licenses, plus "Documentation and Media" (mag tape). The mag tape arrived, but no documentation. DEC's initial response was that "Documentation and Media" does not include documentation, which must be ordered separately.
2. I received this tape last week, and it's Version 1.1 - when I asked why I didn't get 1.2 (announced in June), I got a "return authorization" number, and was told to send it back and they would ship V1.2
3. Meanwhile, the local office told me that V1.0, 1.1, and 1.2 are the same thing (like a DECnet key). Unpacking the save set seemed to bear this out. Then the local office "kindly" furnished me with a line-printer copy of the documentation for V1.2 (less illustrations, of course) which says:

\$ @SYS\$UPDATE:VMSINSTAL LAV010 ddcu:

If there's a moral in this, I haven't found it yet.

DALE E. COY
LOS ALAMOS NATIONAL LAB
PO BOX 1663, MS J957
LOS ALAMOS, NM 87545
505-667-3270

=====
Note 741.6 750's and LAVC's 6 of 7
"John Osudar" 16 lines 10-SEP-1987 18:35

 -< added thoughts on LAVC's >-

Incidentally: before we got our LAVc, I had wondered about what exactly happens to the satellite nodes when the boot member goes down (either via SHUTDOWN or crash). The answer is... nothing! Although the satellites cannot do anything while the boot member is down, once the boot member comes back up (to the point where it loads the MSCP server and begins serving the required disks), the satellites all mount-verify the served disks, and resume operation as if nothing happened. (It's a little strange to read the console messages, indicating that the boot node is sending a membership request to a satellite node -- but it does seem to work!) We had been a little concerned, since we wanted to add an existing MicroVAX-II to the LAVc; that node is used to run some multi-CPU-day batch jobs, and we were afraid that our 785's typical reboot-every-two-weeks behavior would impact those jobs. It doesn't; the batch jobs will wait until the 785 boot node comes back up, and then will continue running.

John Osudar
Argonne National Laboratory
9700 S. Cass Ave.
Bldg. 205 A-051
Argonne, IL 60439-4837
(312) 972-7505

=====
Note 742.0 LSE Support for Non-DEC Compilers 4 replies
"Larry Kilgallen" 14 lines 8-SEP-1987 20:13

A DECworld speaker giving an Ada talk today said DEC is now supporting the generation of LSE diagnostics files (.DIA) by non-DEC compilers. Presumably this means the next version of LSE. The example given in 35mm slides was for a Jovial compiler.

Off-line a different DEC employee told me that the support mechanism will be a new ASCII input format for the diagnostics file because the existing binary format was so complicated. LSE will be able to read either the binary file produced by DEC compilers or the Ascii file produced by customer compilers. This will have the disadvantage of not allowing the concatenation of those Ascii files with the binary files.

Larry Kilgallen
Box 81, MIT Station
Cambridge, MA 02139-0901

=====
Note 742.2 LSE Support for Non-DEC Compilers 2 of 4
"Walter R. Crosby" 15 lines 14-SEP-1987 20:16
 -< Walter Crosby, How about other foreign languages? >-

Has anyone out there considered using LSE with an ADABAS/Natural type programming environment on the VAX? We have a very large Data Center that is centered around IBM 3090 Compatible MVS/XA monsters all running that environment.

MIS is investigating the use of departmental processors, but is concerned about the viability of using VAXen in conjunction with the big machines, if the VAXen will give only marginal contribution to programmer efficiency.

Walter R. Crosby
OMIS/BSPP
1 Ashburton Place, Room 1601
Boston, MA 02108

```
=====
Note 742.3      LSE Support for Non-DEC Compilers      3 of 4
"Dale E. Coy (505) 667-3270"      9 lines  15-SEP-1987 00:20
-< Should be straightforward >-
-----
```

We're not in that particular game, but have looked at integrating other stuff into LSE. Since LSE was designed to have other things added to it, it is really straightforward to do so (but don't confuse straightforward with "trivial").

For commercial products trying to sell into VAX space, the vendor SHOULD be interested in developing environment files for LSE. It's probably a 3 to 6 programmer-month job to do it elegantly.

DALE E. COY
LOS ALAMOS NATIONAL LAB
PO BOX 1663, MS J957
LOS ALAMOS, NM 87545
505-667-3270

```
=====
Note 742.4      LSE Support for Non-DEC Compilers      4 of 4
"Larry Kilgallen"      7 lines  15-SEP-1987 07:05
-< It takes too long for DEC to do it right >-
-----
```

I have seen some samples of LSE in action where it calls for (as an example) a gzorn. When one asks for expansion of what "gzorn" means, it response with the very unhelpful phrase "Fill in the appropriate value for a gzorn". I realize that this is tied up in the continuing myth of VMS inter-language compatibility, but "on-line" information should be "on-line".

Larry Kilgallen
Box 81, MIT Station
Cambridge, MA 02139-0901

```
=====
Note 743.0      WPS-Plus and Laser Jets      2 replies
"Lee K. Gleason"      13 lines  8-SEP-1987 23:44
-----
```

I'm looking for a set of WPS-Plus printer tables for the Hewlett Packard Laser Jet+. I'd sure like to avoid the tedious process of figuring out the escape sequences and entering them via TPU. I tried to tell the manager it would be easier (for me, at least) to just buy LN03s, but he wouldn't go for it.

Surely someone out there has already gone through this...

Lee K. Gleason
Control-G Consultants
2416 Branard #D
Houston TX 77098
713/528-1859

```
=====
Note 743.2      WPS-Plus and Laser Jets      2 of 2
"Jack Patteeuw"      16 lines  9-SEP-1987 11:48
-< They're Right Here !!! >-
-----
```

The Printer Attributes (.PRA) and Printer Characteristics (.PRC) table for an HP Laser Jet (don't know what the differences are with the + model) have been uploaded to [US176598]HP.PRA and [US176598]HP.PRC.

We have been using these tables for over a year and I know that several copies have been downloaded to other sites so I assume they must be correct. This tables are also on either the Spring86 or Fall86 Symposium tape under [MIVAXLUG.FORD].

Also in my directory are the tables for a XEROX 2700 (should work with 3700 from what I have been told) and for QMS 800 (should work with any QMS printer that has QUIC).

If anyone has problems with any of these tables please let me know !

Jack Patteeuw
Ford Motor Co.
Electrical and Electronics Division

31630 Wyoming
Livonia, MI 48150
313-323-8643

=====
Note 745.0 Ethernet Encryption No replies
"Larry Kilgallen" 13 lines 9-SEP-1987 15:42

In the DECworld session "Security Status and Trends" (10 am weekday mornings, today through Thursday September 17), DEC is giving a "program announcement" of DESNC Ethernet encryption. The hardware will be a box which connects up to 4 nodes to an Ethernet (in between the controllers and a single transceiver). The software will be key distribution software running on one or more (not clear if there is a limit) key distribution VAXen. In addition to the DES encryption you would expect, the eventual offering is planned to enforce that node n is plugged into port X of DESNC Y (by virtue of the fact that DECnet changes the physical Ethernet address on the controller to have the two low order bytes contain the node number).

Larry Kilgallen
Box 81, MIT Station
Cambridge, MA 02139-0901

=====
Note 747.0 Tape Mgmt Sys - (TAPESYS) 8 replies
"Jan C Ostendarp" 11 lines 11-SEP-1987 02:49

Does anyone have experience with, or a recommendation for, VAX tape management software? (i.e. TAPESYS from Software Partners, or DECUS Library software, etc.) Are there rumors circulating about DEC's intended functionality improvements in BACKUP (v5 and beyond)?

The documentation I've read on TAPESYS leads me to believe they've developed a full function system well integrated with VMS backup. If you've used it or other similar products available on VAX, tells us about it please.

Jan C Ostendarp
Massachusetts Financial Services
200 Berkeley Street

Boston, MA 02116

=====
Note 747.1 Tape Mgmt Sys - (TAPESYS) 1 of 8
"Larry Kilgallen" 6 lines 11-SEP-1987 20:34
-< Does TAPESYS make patches? >-

I have heard a rumor that TAPESYS patches VMS images in order to make it all work. This poses a support issue which might deter some people. Hopefully the Backup and Mount improvements outlined in Nashville for some future version of VMS would alleviate some of the need for third party vendors to go to such extremes.

Larry Kilgallen
Box 81, MIT Station
Cambridge, MA 02139-0901

=====
Note 747.2 Tape Mgmt Sys - (TAPESYS) 2 of 8
"John D. Ferriby" 29 lines 16-SEP-1987 03:32
-< Mixed feelings about TAPESYS and VMS BACKUP >-

Concerning the TAPESYS product and use at our site...

We have had some awful experiences with TAPESYS. Mostly with just its being a poor product with poor support and poor documentation. This experience was our initial, when we had to pull their teeth to get them to airfreight the *CORRECT* images that were not linked with debug or just plain did not work. It has been distributed to five of 130+ plants we support, and we have heard many initial horror stories. To my knowledge, it does not make any patches to VMS. If I find out they did, ohh.. well... *&^%*#!

On the bright side, in the last year and since the last release, life has been peaceful. It runs quietly as far as I can tell. We have one person devoted to administer/support it, but we have a large development community to support and this justifies the dedication. It does on occasion conflict with the users if, for some reason, a backup job is running during prime time.

On a related VMS BACKUP note, there is a problem with backup for those people with wide spread image cycles, and that is that BACKUP does not delete files from the previous pass during succeeding restores that were deleted by the user at that time. For some time our image cycle was one month, and if a problem arose we would have to go back up to 30 days. Once in fact, that last 30 day image was bad and we had to go back 60 days and reconstruct. (*NEVER* again.) The crux of the problem is that BACKUP/INCREMENTAL does not keep track of what the entire director's contents status is, only that particular files have been modified or that they have never been backed up.

John D. Ferriby~r
{
2871 Troy Centry
{#2010
Troy, MI 48084
(313) 362-2595

=====
Note 747.3 Tape Mgmt Sys - (TAPESYS) 3 of 8
"Larry Kilgallen" 10 lines 16-SEP-1987 06:39
-< VMS Backup used to do deletions correctly >-

R.E. 747.2 - VMS Backup

VMS Backup is SUPPOSED to do it right, and to my recollection at least at one point it did. The one thing it does not handle is deletions in the case where you RENAME a DIRECTORY, but in other cases restoring the full followed by restoring the incrementals (using /INCREMENTAL) in reverse order is supposed to delete files which were deleted. If you can come up with a solid counter-example, please send and SPR for the benefit of us all.

Larry Kilgallen
Box 81, MIT Station
Cambridge, MA 02139-0901

=====
Note 747.4 Tape Mgmt Sys - (TAPESYS) 4 of 8
"Brian Tillman, Smiths Industries." 7 lines 16-SEP-1987 08:27
-< BACKUP will delete. >-

We use image plus incremental and our restores DO properly delete files that were deleted between backups. You must have some other problem. Perhaps you're not following the proper procedures for the restore (i.e., as Larry says, restoring the full backup and then the incrementals). I believe you must also generate a journal file of the backups in order for files to be deleted properly upon restoration.

Brian Tillman
Lear Siegler, Inc.
4141 Eastern Ave. MS121
Grand Rapids, MI 49518-8727
(616)241-8425

=====
Note 747.5 Tape Mgmt Sys - (TAPESYS) 5 of 8
"Bob Hassinger" 49 lines 16-SEP-1987 13:13
-< A "counter-example"... >-

VMS Backup is SUPPOSED to do it right,restoring the full followed by restoring the incrementals (using /INCREMENTAL) in reverse order is supposed to delete files which were deleted. If you can come up with a solid counter-example, please send and SPR for the benefit of us all.

I think this discussion is about a problem that I did find, document, work with TSC, confirm with them, SPR, followup three times, ask about at my LUG meeting, post (here I think), and raise at the Spring Symposium in the Advanced Q&A! It has been more than half a year now and I STILL can not get the curtesy of any response at all, even one telling me I am just making a mistake or that it can not be reproduced!

I did get a call from someone in a DEC location I never heard of (a local office?) who was trying to come up with a work around, NOT a FIX, for someone (he was evasive about the details). He wanted to know what I knew and how I solved it (I wonder who should be paying whom!).

My case was: After an image backup a directory tree with files is deleted, then an incremental backup is done. I found that restoring the image, followed by the incremental, all done exactly as specified in the manuals, failed to delete the directory tree or the files in it. TSC confirmed this but it is still possible there is some simple problem we both missed, I don't know. There is a suspicion something real DID break in BACKUP some place along the way. I never did this kind of restore before early 87 so I am not sure if this is a new or an old problem. It turned out to be a real problem for me when the deleted trees involved hundreds of thousands of blocks on full disks!

Note re: .4 - I find that individual files deleted from directories that remain on the disk ARE deleted as they should be. My problem involves deleted directory trees and the files in them. It seems to be a problem in BACKUP involving the walking of directory trees to delete the files they contain and/or a problem in the deleting of the .DIR file along with what seems to be a failure to report any problem doing these operations. I think this distinction has been the source of more than a little confusion on the subject.

(BTW, in my case it is directory trees but I think the problem even shows up in the simple case of an individual directory with files in it if they are both deleted from the disk between the image and the incremental.)

Bob Hassinger
Liberty Mutual Research Center
71 Frankland Road
Hopkinton, MA 01748
617-435-9061

```
=====
Note 747.6          Tape Mgmt Sys - (TAPESYS)          6 of 8
"Larry Kilgallen"  7 lines `16-SEP-1987 21:04
                  -< BACKUP directory deletion >-
=====
```

RE: 747.5

I do believe that directory deletion is another instance of the same problem as directory renaming. I think VMS developers have stated that it has never worked. I think I have also heard VMS developers describe it as at least a known limitation.

Larry Kilgallen
Box 81, MIT Station
Cambridge, MA 02139-0901

```
=====
Note 747.7          Tape Mgmt Sys - (TAPESYS)          7 of 8
"John D. Ferriby"  5 lines 21-SEP-1987 15:56
                  -< Correction to .2 >-
=====
```

Our staff has closely documented this error and they are working with some folks from DECl and to resolve the situation. Our situation is more accurately described by Bob in .6 than what I said in .2

John D. Ferriby~r
{
2871 Troy Centry
{#2010
Troy, MI 48084
(313) 362-2595

=====
Note 747.8 Tape Mgmt Sys - (TAPESYS) 8 of 8
"Mack W. Haynes, Jr." 10 lines 22-SEP-1987 09:10
 -< Tapesys - A reply >-

I have been using Tapesys for about two years without any major problems. There have been a few typo's in the command files and a few minor bugs in some of the programs. However, I can attribute that to the fact that I receive pre-release versions of the software. The package uses a modified version of VMS backup and requires FMS be installed as a minimum.

Generally, I have been very pleased with the package. Phil Jamieson, President of the company has been very receptive to suggestions and very supportive of any minor problems I have experienced. I can recommend this package very highly.

Mack W. Haynes, Jr.
Tenneco Oil Co.
1100 Louisiana
MS: IFP2954
P.O. Box 2511
Houston, TX 77001
(713) 951-1118

=====
Note 749.6 VMS 4.6 - where are you? 6 of 7
"Stuart Renes" 24 lines 24-SEP-1987 13:48
 -< VMS version V4.6 experience >-

We got VMS version 4.6 the week of September 14 and have installed it successfully on three machines. Two problems that arose are as follows:

1. If you have a VAXcluster and DO NOT do a rolling upgrade AND have the logical name SHUTDOWN\$INFORM_NODES defined, the upgrade will abort during the first shutdown/reboot (due to a FATAL error from OPCOM not recognizing the CLUSTER node name from the logical).

This causes VMSINSTAL to abort and delete its efforts in [SYSUPD...] and you get to start over!

(Admittedly this is a rare problem, but it bit me in 4.4 and I had forgotten!)... so GET RID OF THE LOGICAL PRIOR TO THE UPGRADE!

2. If you have SI Eagle disks that use a patched DRDRIVER, you cant patch the new one with the SI kit... but the V4.5/4.5 DRDRIVER seems to work fine until SI comes up with the patch.

Otherwise, the upgrade works as expected.

Stuart Renes
AT&T Technologies, Inc.
Mail Stop: 2793
3000 Skyline Dr.
Mesquite, TX 75149
(214) 288-2286

=====
Note 752.0 Q-bus versus BI-bus 15 replies
"MICHAEL GRATTAN" 25 lines 15-SEP-1987 09:42

We are looking at getting a mid-range VAX, and as such we saw the new VAX 3600 at DECworld. It is very nicely priced and has the peripherals that we would like. As we were originally thinking of getting an 8250, we are now discussing the merits of the Q-bus vs. the BI-bus.

We see several points in each direction:

The Q-bus is:

1. older and supposedly slower
2. open, lots of third party peripherals

The BI-bus is:

1. new and fast
2. closed, only DEC peripherals or licensed vendors

Since none of us here are real technical, can anyone point out anything that I've missed? (I'm sure there's something.)

Does anyone have any thoughts on this? I would greatly appreciate any comments. Thanks.

MICHAEL GRATTAN
FAIRHAVEN CORP.
358 BELLEVILLE AVE.
NEW BEDFORD, MA. 02742
617-993-9981 EXT 106

=====
Note 752.1 Q-bus versus BI-bus 1 of 15
"Brian Tillman, Smiths Industries." 9 lines 15-SEP-1987 12:27
-< One reason for BI >-

Of course, with BI, there is the added advantage of multiprocessing that the QBUS will never support. The 8250 is field-upgradable to the 8350, which is a multiprocessor and can use the SMP (symmetric multiprocessing) features that VMS 5.0 will have. If your processing loads are such that you would be able to take advantage of SMP (generally compute-intensive), would would do well to purchase the entry-level BI machine (8250), provided you didn't have an already substantial investment in other bus technology devices.

Brian Tillman
Lear Siegler, Inc.
4141 Eastern Ave. MS121
Grand Rapids, MI 49518-8727
(616)241-8425

=====
Note 752.2 Q-bus versus BI-bus 2 of 15
"Bill Mayhew" 26 lines 15-SEP-1987 16:24
-< MicroVAX 3500/3600 vs. VAX 8xxx >-

If your processing loads are such that you would be able to take advantage of SMP, you would do well to purchase the entry-level BI machine (8250), provided you didn't have an already substantial investment in other bus technology devices.

True enough, except that for around or less than the same cost as the 8250, you get roughly 3x the CPU power today in the 3000 series. Interesting tradeoff. I have also read, in the past week, that Digital has more-or-less written-off the 8250 as an entry-level machine, and now positions the 8350 as "the" entry-level BI system. I interpret what I read as a marketing posture, rather than a marketing or technical decision.

Would anybody like to comment on the typical kinds of applications (or mix thereof) that would do well under SMP, as opposed to under a cluster, which of course can be done today (ummm... when Digital delivers 'em) with 3000-series machines?

Back to the original question, some of the other considerations are:

Clustering performance (Ethernet vs. CI, roughly 1:7 raw bit speed)

Disk capacities with Digital disks

Memory size limits (32 vs. up to 256Mb)

Operating environment (office vs. "computer room")

Bill Mayhew
Village Systems Workshop Inc
PO Box 642
Natick MA 01760
617-237-0238

=====
Note 752.3 Q-bus versus BI-bus 3 of 15
"Larry Kilgallen" 12 lines 16-SEP-1987 06:43
 -< Consider HSC-only Features >-

At the present time DEC does not offer CI interconnect for the MicroVAX. In the past the reason has been given that the CI adapter would cost more than the MicroVAX, but it seems to me the MicroVAX 3 pricing has solved that problem.

So if you want either of the HSC-based features:

1. HSC-backups (boo)
2. Volume-shadowing (yay)

now or in the future, that should make you choose a non-Qbus machine.

Larry Kilgallen
Box 81, MIT Station
Cambridge, MA 02139-0901

=====
Note 752.4 Q-bus versus BI-bus 4 of 15
"M. Erik Husby" 22 lines 16-SEP-1987 09:24
 -< Other features to consider >-

Some features of MicroVAX III to consider.

When the memory boards are available, you will be able to support 64megs of memory in one.

You will be able to do SMP with the MicroVAX III because it is not strictly a Q-BUS machine like the MicroVAX II is. DEC has some products in the works which will make use of that feature.

To me the decision on whether to go 8000 class or MicroVAX III class will depend on the type of processing you want to do. If you envision a shop where you have a lot of people doing word processing, electronic mail, moderate size business type computing, and moderate amounts of disk space then a MicroVAX

III should be sufficient. It can grow in to a LAVC.

If you need a lot of disk space and a lot of cpu crunching power, look closely at the 8000 class machines.

M. Erik Husby
Project Software & Development
14 Story St.
Cambridge, MA. 02138
(617)-661-1666

=====
Note 752.5 Q-bus versus BI-bus 5 of 15
"Bob Hassinger" 9 lines 16-SEP-1987 10:03
 -< How about an HSC with NI interface? >-

At the present time DEC does not offer CI interconnect for the MicroVAX. ...

Question: Rather than looking for a CI interface for the MicroVAX, why not look for an Ethernet (NI) interface for the HSC? In smaller clusters that might be a viable option that would be much cheaper. Are there technical problems with this?

Bob Hassinger
Liberty Mutual Research Center
71 Frankland Road
Hopkinton, MA 01748
617-435-9061

=====
Note 752.6 Q-bus versus BI-bus 6 of 15
"Walter R. Crosby" 10 lines 16-SEP-1987 18:10
 -< 3600 vs. 8XXXxxx >-

In some benchmarks, that I have seen, the indication is that the additional 8350 processor does absolutely no good for typical business data processing applications.

Unless you're willing to go with the big bucks for 8500 or larger, it would appear that you would be best off with the 3600 for business data processing.

The lesson of DECWorld appears to be that if you want to do number crunching, you're better off with VAXstations galore!

Walter R. Crosby
OMIS/BSP
1 Ashburton Place, Room 1601
Boston, MA 02108

=====
Note 752.7 Q-bus versus BI-bus 7 of 15
"Jamie Hanrahan" 15 lines 16-SEP-1987 19:25
-< multiprocessing will be better under V5 >-

The benchmarks referred to have all been done under VMS's current asymmetrical multiprocessing scheme. Things will work much better under 5.0. In general, if you have two or more processes (other than NULL) in the COM state a lot of the time, the second processor in an 8350 WILL help, and will add about 0.8 of a processor to your machine. (And there's the potential for adding yet more processors; the developer in charge of SMP said that in their tests they found that the 8200 chassis can support at least 4, and usually 5 8200 cpus before it runs out of bandwidth for I/O! Think of a 5 mips VAX in an 8200 footprint... of course, whether DEC will sell us those processors is another question; they'd prefer we buy 8500s and up...)

Jamie Hanrahan
Simpact Associates
9210 Sky Park Court
San Diego, CA 92123
619-565-1865

=====
Note 752.8 Q-bus versus BI-bus 8 of 15
"Jack Patteeuw" 24 lines 18-SEP-1987 17:34
-< My 2 cents >-

It is obvious that the 3000 series has "temporarily" surpassed the 800 series in the VAX line. However, please note the "hole" left in the 8000 series number between the 83xx and the 85xx ! If the the CVAX chip set is placed on the BI bus it would make a wonderful 8400 !!

I agree with 752.3. The big differences between the the 3000 and the 8000 series are:

	3000	8000
Disk	4	255 (via CI)
RAM	32M-128M(?)	128M(BI)-256M(Nautilus)
SMP	Don't think so	Just around the corner

So if you feel that you will need to grow into lots of disk or RAM stay with the 8000 series but wait for future announcements for a bit if you can.

I really don't expect DEC to support SMP on the QBus. After all it (multiprocessing) is the biggest reason they invented the BI bus. Rumor also has it that their 3D high end graphics machine will be BI based, not QBus as all previous graphics workstations.

Jack Patteeuw
Ford Motor Co.
Electrical and Electronics Division
31630 Wyoming
Livonia, MI 48150
313-323-8643

RT-11 DUCM/DYC GRAF11 DEL DIR PLOT-10 IMAGE LIBED FSTATS MS/DOS TIC-TAC-TOE QIX VAX-LIB-DATMAN/VAX EDTPlus SPICE2 TREEDUPL LISPEX MCLS TYPE PLUS MINC DISK USE FRAG EDTEX PORT LOCATOR TECO CHPLOT NANNY DIR11-W WATCHDOG INACTIVE ACCOUNTS IMGSPICE ICE TEX EDITOR VAX-LIB-4 GRAPHIC UTILITIES SETAUX.ARC STATPK FIGure KERMIT Distribution TENBACKU JUICER VTEDIT 2022 VAX-LIB-3 VISTA EDITOR MTU TDE RSTSOPEN DRAWTREE WATCHDOG PRM-1 SMARTMAILER TEN SPELL DECPoint of Sale JUICER PARALLEL Library V2 RTMULTI and Addo SMARTMAILER for RSTS/E CU FILTRA Spring 86 RT-11 SIG CP/M KERMIT S Invasion for PRO Bonner La SPLICE RUNOFF VAX-LIB-3 VAX-LIB-2 IMAGE SPELL TURBOCOM FDNFIL PC-8088 Collection #10 VT20 TOOLKIT PLATOOLS SMARTMAILER DEPROC LaTeX KERMIT-11 FANCY FONTS XMIT CU ReGis to HPG CED International RUNITOFF JP5-JP6 FODT PASCAL-OS/8 ANISMT WPSIM PARALLEL LIBRARY DECSYSTEM-20 SIG Spring 85 CAMERA DELPHIN HACK BIBENTRY APFELN DIGITIZING Acid Docume Generator VAX-LIB-2 AMAR-10 AMAR-20 DATMAN/VAX IMAGE RT-11 DUCM/DYC GRAF11 DEL DIR PLOT-1 IMAGE LIBED FSTATS MS/DOS TIC-TAC-TOE QIX VAX-LIB-5 DATMAN/VAX EDTPlus SPICE2 TREEDUP LISPEX MCLS TYPE PLUS AMAR-20 DISK USE FRAG EDTEX PORT LOCATOR TECO CHPLOT NANN DIR11-W WATCHDOG INACTIVE ACCOUNTS IMGSPICE ICE TEXT EDITOR VAX-LIB-4 GRAPHIC UTILITIE SETAUX.ARC STATPK FIGure KERMIT Distribution TENBACKUP JUICER VTEDIT 2022 VAX-LIB-3 VIST EDITOR MTU TDE RSTSOPEN DRAWTREE WATCHDOG PRM-11 SMARTMAILER TEN SPELL DECPoint of Sa JUICER PARALLEL Library V2 RTMULTI and Addons SMARTMAILER for RSTS/E CU FILTRA Spring 86 RT-1 SIG CP/M KERMIT S Invasion for PRO Bonner Labs APFELN RUNOFF VAX-LIB-3 VAX-LIB-2 IMAGE SPEL TURBOCOM FDNFIL PC-8088 Collection #10 VT200 TOOLKIT PLATOOLS SMARTMAILER DEPROC LaTe KERMIT-11 FANCY FONTS XMIT CU ReGis to HPG CED International RUNITOFF JP5-JP6 FODT PASCAL-OS/ ANISMT TECO WPSIM DECSYSTEM-20 SIG Spring 85 CAMERA DELPHIN HACK BIBENTRY APFELN KERMI S DIGITIZING Acid Document Generator VAX-LIB-2 AMAR-10 AMAR-20 DATMAN/VAX IMAGE VT200 TOOLKI COMPRO EVENTS PC8088 Collection #9 TECO Cher Tree Workstation Bookings System EXPORT Data Inputt Generator CMSBROWSE PERSONNEL INVENTORY MS/DOS COMMS Selection Electronic Grade Book CP/ KERMIT LaTeX JUICER SPELL PORTACALC DPRINT DUNGEON MINC BUDGET BUG CALC C Langua System DPROC "DEP" DECENC DECmate II OS/278 DIAL DTC GAMMA-11 GDADL LISP for RSX-11 MEM KERMIT S VAX-LIB-6 SPICE 3A6 VT200 TOOLKIT RUNOFF SPLICE SPY:RSX TCOPY SPELL VT-200 COMPR EVENTS CMSBROWSE UNDELETE DIAL BLOCKER SCAN CODER BITMAP DTC/PC ADDRESS BOO LaserWriter PORTACALC SPICE 3A6 PRO/Smart Mailer CBASIC2 Accts JP5-JP6 Payable/Receivable McGraw-Hi Payroll SEDT: EDT/WPS Screen CLNDRS:A Calendar Program INDEX AKCOUNT CORPHONE E-Systems Grab Ba RGT RDG PLTXSMB ICON DEVICS DATATRIEVE Library Collection CMSBROWSE EXPERT FPaint IMAG DBMS/Spreadsheet for MS/DOS AMAR-1 AMAR-20 RDIR/SQMAP PC-8088 Collection #11 UP TIME REPORTE RT-11 DUCM/DYC GRAF11 DEL DIR PLOT-10 IMAGE LIBED FSTATS MS/DOS TIC-TAC-TOE QIX VAX-LIB-DATMAN/VAX EDTPlus SPICE2 TREEDUPL LISPEX MCLS TYPE PLUS EXPORT DISK USE FRAG EDTEX PORT LOCATOR TECO CHPLOT NANNY DIR11-W WATCHDOG INACTIVE ACCOUNTS IMGSPICE ICE TEX EDITOR VAX-LIB-4 GRAPHIC UTILITIES SETAUX.ARC STATPK FIGure KERMIT Distribution TENBACKU JUICER VTEDIT 2022 VAX-LIB-3 VISTA EDITOR MTU TDE RSTSOPEN DRAWTREE WATCHDOG PRM-1 SMARTMAILER TEN SPELL DECPoint of Sale JUICER PARALLEL Library V2 RTMULTI and Addo SMARTMAILER for RSTS/E CU GRAPHKIT FILTRA Spring 86 RT-11 SIG CP/M KERMIT S Invasion for PR Bonner Labs RUNOFF VAX-LIB-3 VAX-LIB-2 IMAGE SPELL TURBOCOM FDNFIL PC-8088 Collection #10 VT2 TOOLKIT PLATOOLS SMARTMAILER DEPROC LaTeX KERMIT-11 FANCY FONTS XMIT MEMO ReGis to HPG CED International RUNITOFF JP5-JP6 FODT PASCAL-OS/8 ANISMT CODER WPSIM DECSYSTEM-20 SIG Sprin 85 CAMERA DELPHIN HACK BIBENTRY APFELN REPORTER DIGITIZING Acid Document Generator VAX-LIB-AMAR-10 AMAR-20 DATMAN/VAX IMAGE VT200 TOOLKIT COMPRO EVENTS PC8088 Collection #9 TECO Cher Tree Workstation Bookings System EXPORT Data Inputter Generator CMSBROWSE PERSONNEL INVENTOR MS/DOS COMMS Selection Electronic Grade Book CP/M KERMIT LaTeX JUICER SPELL PORTACALC DPRINT DUNGEON MINC BUDGET BUG CALC C Language System DPROC "DEP" DECENC DECmate II OS/278 DIA DTC GAMMA-11 GDADL LISP for RSX-11 MEMO PORTACALC VAX-LIB-6 SPICE 3A6 VT200 TOOLKI RUNOFF SPLICE SPY:RSX TCOPY SPELL VT-200 COMPRO EVENTS CMSBROWSE UNDELETE DIA BLOCKER SCAN CODER BITMAP DTC/PC ADDRESS BOOK LaserWriter PORTACALC SPICE 3A6 PRO/Sma Mailer CBASIC2 Accts TREEDUPL Payable/Receivable McGraw-Hi Payroll SEDT: EDT/WPS Screen CLNDRS: Calendar Program INDEX AKCOUNT CORPHONE E-Systems Grab Ba RGT RDG PLTXSMB ICON DEVIC DATATRIEVE Library Collection CMSBROWSE EXPERT FPaint IMAGE DBMS/Spreadsheet for MS/DOS AMAR-1 AMAR-20 RDIR/SQMAP PC-8088 Collection #11 UP TIME REPORTER RT-11 DUCM/DYC GRAF11 DEL DIR PLO 10 IMAGE LIBED FSTATS MS/DOS TIC-TAC-TOE QIX VAX-LIB-5 DATMAN/VAX SPICE2 RT-11 DUCM/DYC G



DECUS

Program Library

SOFTWARE ABSTRACTS

"A Solution with a Future"

LIB

DECUS PROGRAM LIBRARY

NEW LIBRARY PROGRAMS AVAILABLE FOR THE VAX/VMS FAMILY OF COMPUTERS

DECUS NO: V-SP-65 TITLE: Symposium Collection from the RSX SIG, Spring 1987, Nashville Version: 1, August 1987

Author: Various

Submitted by: Glenn C. Everhart, Ph.D.

Operating System: IAS, RSX-11M, RSX-11M-PLUS, VAX/VMS Source Language: BASIC-11, C, FOCAL, FORTRAN 77, FORTRAN IV, FORTRAN IV-PLUS, MACRO-11, VAX FORTRAN Keywords: Symposia Tapes - RSX-11

Abstract: This is the RSX SIG tape from the Spring 1987 DECUS Symposium in Nashville. The tape consists of two parts. The first is the files submitted to teacopy in Spring 1987. These consisted of about 22,000 blocks. Since there was room on the tape, the second part was added. These are files which appeared on the RSX SIG tapes in the period from Fall 1977 to Spring 1979 (plus maybe a couple of later items). The files in this group are selected as those which appear still useful (frequently in HOLs). The 1977-1979 tapes were never available via the DECUS Library, so this material has generally not been available via regular DECUS channels. To order the BRU version, order DECUS No. 11-SP-98.

Area I: New Items for Spring 1987

- [5,4] DECUS C updates for I/D space.
- [5,15] DECUS C updates for I/D space.
- [5,16] DECUS C updates for I/D space.
- [5,24] DECUS C updates for I/D space.
- [307,20] Gary Maxwell's upgraded virtual disk package for M+ VF.
- [312,315] Virtual disk driver for VMS, RSX FOCAL, old TECO Doctor, a MAKE program src., program to read VMS Backup tapes under un*x, UUCP lookalike PD program archives, DISOWN, and an RSX task disassembler, submitted by Glenn Everhart.
- [312,361] Public domain UUCP clone sources. Not specifically for RSX but may be possible to get working.
- [312,371] Fix to RECALC files for AnalytiCalc - minor bugfix.
- [321,5] Structured Macro library. Routines to set time on Qbus clock, submitted by William Kyle.
- [337,50] Jim McGlinchey's Hitchiker's Guide to RSX.
- [344,*] RSX KMSKIT - lots of stuff, submitted by Jim Downard, KMS Fusion.
- [350,340] Pipe Driver vx: for task to task comm. update to previous driver, (by Dave Healey, Utah Power + Light), submitted by Eddy Fey.
- [350,124] & [350,125] ODS-2 ACP for RSX. (.SLP files only), submitted by Dan Eisner.
- [351,73] AUX (keypad cmd language) and ECR (enhanced MCR) for IAS; Skeleton IAS handler, submitted by F. Borger.
- [351,144] TEM terminal emulator for RSX, submitted by Tom Wyant.

- [351,145] Session notes & examples for sessions RX001, RX002 on indirect command processor, submitted by Tom Wyant.
- [356,40] RSX KERMIT, submitted by Brian Nelson.
- [356,41] VMSTPC tape <-> disk <-> tape utility for VMS, submitted by Brian Nelson.
- [356,42] Bitnet servers sources, submitted by Brian Nelson.
- [356,45] IAS KERMIT-11, submitted by Frank Borger.
- [370,352] CLE, MYMACS.MLB. Cmd line editor, submitted by Steven Jobes.
- [370,365] FORTRAN aids and tools, submitted by Richard Neitzel, Golden, Colorado. SST handlers, DL driver fix, undeletion, SEARCH, binary file compare, more.

AREA II: Files collected from older RSX SIG tapes and related sources (highlights only, not all listed here)

- [264,2] 3D plot package from DECUS Europe (Amsterdam) tape, 1981.
- [300,17] FLECS (FORTRAN Language with Extended Control Structures) FORTRAN preprocessor. Source, doc.
- [300,47] Code to intercept illegal instructions plus document.
- [300,51] Design spec generators, document maintenance system, source code configurators (for several languages), source code managers, the above in DATATRIEVE, plus some TECO macros of use, submitted by Dan Curtis.
- [301,16] SSP - Scientific Subroutine Package sources for Digital Equipment Corporation FORTRAN (but no comments), submitted by Charles South.
- [301,27] Set of FORTRAN callable matrix subroutines.
- [303,1] Document of how to run giant (100K lines of FORTRAN) programs under RSX-11M.
- [303,40] RSX mailbox handler.
- [307,22] Disk disaster recovery tools for ODS-1 disk disasters.
- [307,26] SKED project scheduler and resource/milestone tracker.
- [312,356] Infinite precision calculator in FORTRAN.
- [312,366] Virtual disk for RSX11D and IAS, submitted by Shack Toms.
- [321,2] RATFOR (RATional FORtran) preprocessor for RSX.
- [321,3] SUPERMAC structured MACRO-11 assembly macros and doc.
- [323,2] CSMP - Continuous Systems Modeling Program, models systems of continuously varying parameters.
- [330,11] FORTRAN resequencer RESEQ. for F4P programs.
- [332,100] Directory and selective restore from DSC tapes.
- [334,2] OBR - Reads .OBJ files, reporting globals and global defs.
- [334,6] LIBSEE - Query a library for a module or global symbol.

- [340,1] ARC MAIL mail utility (DECnet not needed).
- [341,307] ELIZA (or DOCTOR) program in PL/I with objects. The computerized psychoanalyst.
- [342,2] TECO V36. The full TECO V36 distribution including machine readable manual file.
- [344,51] How to do transient libraries under RSX-11M, submitted by Jim Downard.
- [346,100] Ralph Stamerjohn's collection. ACP manuals, virtual disks, loadable XDT, SIG tape index of early RSX tapes, CDA workbook, and more.
- [355,2] File structure editing/fixing tools BM, Fiddle, VMS like DUMP, execution profiler, disk usage summary.
- [360,214] FORTRAN conditional compilation preprocessors for multiple level conditionals.
- [364,20] Binary semaphore directives for RSX-11M plus docs.
- [370,70] Description of FORTRAN OTS.
- [370,130] INDEX - FORTRAN cross reference program. Handles lots of analysis, staticcode checking, call trees, and much more for PDP-11 FORTRAN, for FORTRAN IV and F4P.
- [372,4] SAMSTAT source for statistics package (a STAT-11 variant).
- [373,3] FORCE, forces commands to a terminal.
- [373,5] RTR, program to read RSTS/E disks from RSX, and program to convert files so read to RSX form for input to BP2.
- [373,7] File recoverer - undeletes a freshly deleted file.
- [373,10] SND - command interface to send/receive directives for software debug.
- [373,17] Show what pool is being used for. Can also follow FCB pointers through F11ACP to find file control structures.
- [373,21] Block by block comparison of binary files, or whole directories full of files.
- [373,101] Macro library covering data conversion, string manipulation, sorting. Help file for your help system documents it. First appearance of help libraries as does for utility libraries.

Notes: Most submissions include source: a few do not. Source code is present where it is supplied. ODS-2 ACP is only difference files to Digital Equipment Corporation source code.

Complete sources not included.

Media (Service Charge Code): 2400' Magnetic Tape (PS) Format: VMS/BACKUP, TK50 Tape Cartridge (TC) Format: VMS/BACKUP

DECUS NO: VAX-274 TITLE: POPUP: A DCL Popup Menu Utility Version: July 1987

Submitted by: John Reece, Intel, Santa Clara, CA

Operating System: VAX/VMS V4.5 Source Language: C Keywords: DCL, Menu Control, Utilities - VMS

Abstract: POPUP is a menu utility that can be installed as a foreign DCL command and used to create elegant pop-up menus in DCL procedures. User options, a menu title, and the screen coordinates are specified as DCL command line parameters and

the resulting user selection is returned in a global symbol. The user selects an option from the resulting menu by either moving a lightbar with the cursor keys to a choice and pressing return, or by typing the first letter of the desired choice. Broadcast messages are trapped and displayed in a box at the bottom of the screen.

POPUP uses no graphics packages other than the SMG functions in the VMS Run-Time Library. It has been tested on VT100 and VT200 series terminals, and on the PC terminal emulators PROcomm, SmarTerm 100 and SmarTerm 240. It works in 132 column mode.

Release notes are distributed with each order.

Notes: Operating system VAX/VMS V4.4 or higher is required.

Sources not included.

Media (Service Charge Code): 600' Magnetic Tape (MA) Format: VMS/BACKUP

DECUS NO: VAX-275 TITLE: DTR Version: X1.0-0, March 1987

Submitted by: Digital Equipment Corporation.

Operating System: MicroVMS, VAX/VMS V4 Source Language: MACRO-32 Keywords: Data Communications

Abstract: DTR is a privileged program which asserts the DTR modem control line for specified terminal communications options. DTR supports the following communication options: DZ-11, DZQ-11, DHU-11, and DHV-11. The user must have CMKRNL privileges to run DTR. This program will turn on the DTR control line (similar to SET TERMINAL/MODEM), except the DTR line will not drop when a login timeout occurs. This is used in conjunction with the RF-FOAFB-AA fiber optic adaptor only.

Release notes are distributed with each order.

Notes: Operating system VAX/VMS V4.0 or higher is required.

Documentation available in hardcopy only.

Media (Service Charge Code): Source Listing (BA), One RX50 Diskette (JA) Format: VAX/ANSI, 600' Magnetic Tape (MA) Format: VAX/ANSI

DECUS NO: VAX-277 TITLE: GameParse Version: 1.0, August 1987

Submitted by: Michael Levin, Swampscott, MA

Operating System: MicroVMS V4.6 Source Language: C Software Required: C Compiler Keywords: Games

Abstract: GameParse is a parser designed to work with text adventure games, such as Dungeon and Adventure. It allows the user to write an adventure game in the C language, by providing a parser and an easy way of teaching it words appropriate to that adventure and the relationships between them.

It consists of an .H file, and an .OBJ file. The user writes a program in C, and uses the "

include" statement to include START.H at the beginning of his program. Then, he compiles and links his program with PARSE.OBJ using the VMS linker. His program can then use function calls to PARSE(), to get commands from the user. The

parser can also be used for other applications which require language parsing.

The parser is taught new words by editing START.H. The parser understands verbs, nouns, adjectives, prepositions, determiners, and can resolve pronoun usage. Methods are provided of specifying which verbs are useful with which nouns, and which are valid but useless. It can also use intransitive verbs, prepositional phrases, and ask intelligent questions. Complete instructions for its use, as well as a sample program which uses the parser, and a dialog which shows the parser's features are included.

Notes: The parser itself is an .OBJ file, source module is not included. The sources needed to call it from any program are included.

Restrictions: Can only be called by C programs.

Documentation available in hardcopy only. Complete sources not included.

Media (Service Charge Code): User's Manual (EA), One RX50 Diskette (JA) **Format:** VAX/ANSI, 600' Magnetic Tape (MA) **Format:** VAX/ANSI

DECUS NO: VAX-278 **TITLE:** VMAP - SCREEN MAPPING DEVELOPMENT TOOL FOR VT100 **Version:** 1.0, August 1987

Submitted by: Jesus Lu, California State University, Los Angeles, CA

Operating System: VAX/VMS V4.0 **Source Language:** MACRO-32, VAX, COBOL **Hardware Required:** VT100 or compatible terminals **Keywords:** Tools - Applications Development, VT100 Routines

Abstract: VMAP is an application development tool for creating on-line screens for VT100 terminals. It facilitates the development of COBOL programs for on-line displays and data entries. Version 1.0 supports field protection, video attributes, line drawings, function key supports (numeric or application mode), map tables, 80 or 132 display columns, graphic symbols, and others.

Included on the distribution media are the VMAP documentation, the VMAP translator program (in COBOL), SEND and RECV utilities (in MACRO-32), and a demo map and program.

The procedure for building and installing VMAP is explained on the last chapter of the VMAP documentation.

The VMAP translator program translates VMAP source statements and creates three output files: the screen map file, the symbolic input (data) file, and the symbolic map control file. These files are used in the application COBOL program by use of the COPY statements.

The SEND utility displays screen maps to the terminal, sets terminal keypad modes, and displays COBOL-type descriptor strings. The RECV utility accepts characters from the terminal, deposits them into the respective fields, marks them as 'entered', and returns a function code or terminator code when a keypad key was pressed.

Notes: Operating system VAX/VMS V4.0 or higher is required.

Media (Service Charge Code): 600' Magnetic Tape (MA) **Format:** VAX/VMS BACKUP

DECUS NO: VAX-279 **TITLE:** WEVE - WONDERFUL EVE **EDITOR Version:** 1.0, June 1987

Submitted by: Messrs. K. Swystun & A. Baillie, Saskatoon Cancer Clinic, Saskatoon, Saskatchewan, Canada S7N 0X0

Operating System: VAX/VMS V4.4 **Source Language:** VAX-TPU **Hardware Required:** VT100 or VT200 compatible terminals **Keywords:** Editors

Abstract: WEVE (Wonderful EVE Editor) is an editor interface that has been designed to emulate and extend the EDT editor. It is based on the EVE editor which has been enhanced with several user written VAXTPU procedures. This software is intended to give current EDT users an interface emulating EDT, but also incorporating the more powerful features of VAXTPU, such as windowing; multiple buffers intimately related to specific files; spawn; and the ability to run DCL commands from within the editor. Functions have also been written to do things such as: automatic indenting; jump to previous buffer; delete buffer; clear buffer; automatic jump to file that cursor points to; show current line number; join line; begin of line only find; alternate cursor behavior option; show all buffer names; and automatic documentation template insertion. In addition to giving the EDT user immediate added functionality, it also gives him the ability to enhance or customize the editor by writing further procedures.

Notes: Operating system VAX/VMS V4.2 or higher is required.

Media (Service Charge Code): 600' Magnetic Tape (MA) **Format:** VMS/BACKUP

DECUS NO: VAX-281 **TITLE:** WCC: A C-Subset Compiler **Version:** 1.0, AUGUST 1987

Submitted by: Lutz Hamel, CSPI, Billerica, MA

Operating System: ULTRIX V.1.2A, VAX/VMS V.4.5 **Source Language:** C, LEX, YACC **Memory Required:** 2MB **Keywords:** Compilers

Abstract: WCC is a small, experimental compiler for a functional subset of the C programming language. The current implementation of the compiler generates code for the VAX-11 computer running either the VMS or the ULTRIX operating system. The WCC compiler itself is written in C (maybe one day it will be able to compile itself).

Language Summary:

Program Control:

- if (expression) statement
- if (expression) statement else statement
- while (expression) statement
- break
- continue
- return

Data types:

- char
- short
- int
- long
- float

One dimensional arrays of these primitive types are allowed, pointers to these types are allowed. No complex types are implemented.

All arithmetic operators are implemented except bit manipulation and address arithmetic. Function calls are supported.

Notes: This tape is in VMS/BACKUP format for use on a machine running VMS.

Media (Service Charge Code): 600' Magnetic Tape (MA) **Format:** VMS/BACKUP

NEW LIBRARY PROGRAMS AVAILABLE FOR THE PDP-11 COMPUTER FAMILY

DECUS NO: 11-SP-97 **TITLE:** Symposium Collection from the RT-11 SIG, Spring 1987, Nashville **Version:** Spring 1987

Submitted by: R.W. Barnard, Sandia National Laboratories, Albuquerque, NM

Operating System: RT-11 V5 **Source Language:** C, FORTRAN 77, FORTRAN IV, MACRO-11 **Memory Required:** Various, specified in submission **Software Required:** Will be specified, if required. **Hardware Required:** Special requirements will be specified in the submissions. **Keywords:** FORTRAN, Plotting, Symposia Tapes - RT-11

Abstract: The symposium swap tape from the RT-11 SIG contains twenty-five packages in subdevice format. The tape includes an annotated directory TAPDIR.TXT, and instructions for RT-11 and RSTS users on recovering files from subdevices. The file TAPDIR.TXT includes a summary, cross-reference and index section. The tape contains the following submissions:

VIRTUL	This program allows RSTS/E users to break down the subdevice files from this tape after they have been copied to disk.
DIRTWO	Contains annotated directories of the DECUS Symposia RT-11 tapes from the Fall of 1981 through the Fall of 1986.
WSHLST	RT-11 wish list survey.
FONT	Downloadable VT-200 character font.
SPELL	Spelling-checker with dictionary.
CALEND	Calendar display program.
DFIND	Subdevice directory program.
RDMF77	Directory and other utilities.
MAIL	On-line message facility for TSX-PLUS.
TAPE	Tape utilities to backup specific disk devices to magtape. Also includes ANSIR and ANSIW, or reading and writing unlabelled ANSI magnetic tapes, and T10IBM, for reading EBCDIC IBM tapes.
ACODES	On-line telephone area codes retriever.
TIMING	RT-11/TSX-PLUS System Timing Studies.
TSXLIB	FORTRAN-Callable TSX-PLUS EMT's.
DROIDS	A game which pits your (or your kid's) skills against a planetfull of droids bent on your destruction.
UCLPLS	User command language (UCL) program.
PM	RT-11 monitor prompt handleroid.
PLT	File oriented plotting utility for RT.
FLXIND	IND control files for FLECS processing.
F77IND	IND control files for FORTRAN-77 compilations.
BAKAL	IND control file to automate backups.
THESIS	RUNOFF macros for formatting a thesis.

GKS RT-11 implementation of GKS plotting standard.
INDFIL IND control files for manipulating subdevices.
DIAL Terminal emulator front end.
KERMIT File transfer protocol for PDP-11's. This is release 2.44 of KERMIT-11.

Restrictions: Will be specified in submissions, if any.

Media (Service Charge Code): Write-Up (AA), 2400' Magnetic Tape (PS) **Format:** RT-11, TK50 Tape Cartridge (TC) **Format:** RT-11

DECUS NO: 11-SP-98 **TITLE:** Symposium Collection from the RSX SIG, Spring 1987, Nashville **Version:** 1, August 1987

Author: Various

Submitted by: Glenn C. Everhart, Ph.D.

Operating System: IAS, RSX-11M, RSX-11M-PLUS, VAX/VMS **Source Language:** BASIC-11, C, FOCAL, FORTRAN 77, FORTRAN IV, FORTRAN IV-PLUS, MACRO-11, VAX FORTRAN **Keywords:** Symposia Tapes - RSX-11

Abstract: This is the RSX SIG tape from the Spring 1987 DECUS Symposium in Nashville. The tape consists of two parts. The first is the files submitted to tapecopy in Spring 1987. These consisted of about 22,000 blocks. Since there was room on the tape, the second part was added. These are files which appeared on the RSX SIG tapes in the period from Fall 1977 to Spring 1979 (plus maybe a couple of later items). The files in this group are selected as those which appear still useful (frequently in HOLs). The 1977-1979 tapes were never available via the DECUS Library, so this material has generally not been available via regular DECUS channels. To order the VMS/BACKUP version, order DECUS No. V-SP-65.

Area I: New Items for Spring 1987

[5,4]	DECUS C updates for I/D space
[5,15]	DECUS C updates for I/D space
[5,16]	DECUS C updates for I/D space
[5,24]	DECUS C updates for I/D space
[307,20]	Gary Maxwell's upgraded virtual disk package for M+ VF:
[312,315]	Virtual disk driver for VMS, RSX FOCAL, old TECO Doctor, a MAKE program src., program to read VMS Backup tapes under un*x, UUCP lookalike PD program archives, DISOWN, and an RSX task disassembler, submitted by Glenn Everhart.
[312,361]	Public domain UUCP clone sources. Not specifically for RSX but may be possible to get working.
[312,371]	Fix to RECALC files for AnalytiCalc - minor bugfix.
[321,5]	Structured Macro library. Routines to set time on Qbus clock, submitted by William Kyle.
[337,50]	Jim McGlinchey's Hitchiker's Guide to RSX.
[344,*]	RSX KMSKIT - lots of stuff, submitted by Jim Downard, KMS Fusion.
[350,340]	Pipe Driver vx: for task to task comm. update to previous driver, (by Dave Healey, Utah Power + Light), submitted by Eddy Fey.
[350,124] & [350,125]	ODS-2 ACP for RSX, (.SLP files only), submitted by Dan Eisner.

[351.73] AUX (keypad cmd language) and ECR (enhanced MCR) for IAS; Skeleton IAS handler, submitted by F. Borger.

[351.144] TEM terminal emulator for RSX, submitted by Tom Wyant.

[351.145] Session notes & examples for sessions RX001, RX002 on indirect command processor, submitted by Tom Wyant.

[356.40] RSX KERMIT, submitted by Brian Nelson.

[356.41] VMSTPC tape <-> disk <-> tape utility for VMS, submitted by Brian Nelson.

[356.42] Bitnet servers sources, submitted by Brian Nelson.

[356.45] IAS KERMIT-11, submitted by Frank Borger.

[370.352] CLE, MYMACS,MLB. Cmd line editor, submitted by Steven Jobs.

[370.365] FORTRAN aids and tools, submitted by Richard Neitzel, Golden, Colorado. SST handlers, DL driver fix, undeletion, SEARCH, binary file compare, more.

AREA II: Files collected from older RSX SIG tapes and related sources (highlights only, not all listed here)

[264.2] 3D plot package from DECUS Europe (Amsterdam) tape, 1981.

[300.17] FLECS (FORTRAN Language with Extended Control Structures) FORTRAN preprocessor. Source, doc.

[300.47] Code to intercept illegal instructions plus document.

[300.51] Design spec generators, document maintenance system, source code configurators (for several languages), source code managers, the above in DATATRIEVE, plus some TECO macros of use, submitted by Dan Curtis.

[301.16] SSP - Scientific Subroutine Package sources for Digital Equipment Corporation FORTRAN (but no comments), submitted by Charles South.

[301.27] Set of FORTRAN callable matrix subroutines.

[303.1] Document of how to run giant (100K lines of FORTRAN) programs under RSX-11M.

[303.40] RSX mailbox handler.

[307.22] Disk disaster recovery tools for ODS-1 disk disasters.

[307.26] SKED project scheduler and resource/mile-stone tracker.

[312.356] Infinite precision calculator in FORTRAN.

[312.366] Virtual disk for RSX11D and IAS, submitted by Shaek Toms.

[321.2] RATFOR (RATional FORtran) preprocessor for RSX.

[321.3] SUPERMAC structured MACRO-11 assembly macros and doc.

[323.2] CSMP - Continuous Systems Modeling Program, models systems of continuously varying parameters.

[330.11] FORTRAN resequencer RESEQ. for F4P programs.

[332.100] Directory and selective restore from DSC tapes.

[334.2] OBR - Reads .OBJ files, reporting globals and global defs.

[334.6] LIBSEE - Query a library for a module or global symbol.

[340.1] ARC MAIL mail utility (DECnet not needed).

[341.307] ELIZA (or DOCTOR) program in PL/1 with objects. The computerized psychoanalyst.

[342.2] TECO V36. The full TECO V36 distribution including machine readable manual file.

[344.51] How to do transient libraries under RSX-11M, submitted by Jim Downward.

[346.100] Ralph Stamerjohn's collection. ACP manuals, virtual disks, loadable XDT, SIG tape index of early RSX tapes, CDA workbook, and more.

[355.2] File structure editing/fixing tools BM, Fiddle, VMS like DUMP, execution profiler, disk usage summary.

[360.214] FORTRAN conditional compilation preprocessors for multiple level conditionals.

[364.20] Binary semaphore directives for RSX-11M plus docs.

[370.70] Description of FORTRAN OTS.

[370.130] INDEX - FORTRAN cross reference program. Handles lots of analysis, static code checking, call trees, and much more for PDP-11 FORTRAN, for FORTRAN IV and F4P.

[372.4] SAMSTAT source for statistics package (a STAT-11 variant).

[373.3] FORCE, forces commands to a terminal.

[373.5] RTR, program to read RSTS/E disks from RSX, and program to convert files so read to RSX form for input to BP2.

[373.7] File recoverer - undeletes a freshly deleted file.

[373.10] SND - command interface to send/receive directives for software debug.

[373.17] Show what pool is being used for. Can also follow FCB pointers through F11ACP to find file control structures.

[373.21] Block by block comparison of binary files, or whole directories full of files.

[373.101] Macro library covering data conversion, string manipulation, sorting. Help file for your help system documents it. First appearance of help libraries as does for utility libraries.

Notes: Most submissions include source: a few do not. Source code is present where it is supplied. ODS-2 ACP is only difference files to Digital Equipment Corporation source code.

Complete sources not included.

Media (Service Charge Code): 2400' Magnetic Tape (PS) Format: BRU Version 3.2, TK50 Tape Cartridge (TC) Format: BRU Version 3.2

DECUS NO: 11-892 **TITLE:** LOST: An Adventure Game **Version:** 2, August 1987

Submitted by: P.A. Edwards, Weardrive Ltd., Hints, Staffordshire, England 878 3DW

Operating System: RSX-11M V4.1, RSX-11M-PLUS V3.0 **Source Language:** CORAL **Memory Required:** 32KW **Keywords:** Games

Abstract: The game of "LOST" is a database driven Adventure style game which takes its parameters from files written by the

user with a suitable text editor such as EDT, EDI or TECO. Two sample databases are supplied as an introduction to the facilities of the game, and as a guide to the preparation of new databases.

Release notes are distributed with each order.

Media (Service Charge Code): 600' Magnetic Tape (MA) Format: BRU

REVISIONS TO LIBRARY PROGRAMS

DECUS NO: RB-117 **TITLE:** Vehicle Records **Version:** 11, August 1987

Submitted by: Bruce W. Roeckel, Florida Power Corp., St. Petersburg, FL

Operating System: MS/DOS V2.11 **Source Language:** FORTRAN 77 **Memory Required:** 192KB **Software Required:** Microsoft FORTRAN is required to recompile and relink. **Keywords:** Business Applications

Abstract: The Vehicle Records program is designed to allow a user to store mileage and maintenance information for up to twenty-five vehicles. A full-screen editor is utilized for the addition, editing and selling of vehicle entries in the master file. Prompts are used for data to be entered when updating mileage or maintenance records for each individual vehicle.

Mileage information is broken up into two categories; city and trip mileage. When reports and/or graphs are generated, these two categories are always kept separate. Also, when entering trip mileage, the user is prompted for a description of the trip.

Maintenance information is also broken up into two categories; recurring items and special repairs. For the recurring items, the only data stored is that which pertains to the last time you performed the task. Typical recurring maintenance items are oil changes, lube jobs, tire rotations, etc. Each time you update the recurring items, you're prompted for any notes that you may want to store, (i.e. the type of oil used), as well as the date of repair, cost and odometer reading. For the special repair items, you are asked for a description of the repair in addition to other data, i.e. date, cost, etc.

Summary reports can be obtained for any vehicle and include:

- A maintenance records report
- A city or trip mileage report
- A city or trip mileage graph

The graphs can be displayed directly on the screen without the need for the Rainbow Graphics Option Card. All of the data on any of the summary reports is sorted by odometer reading.

Also included with this disk is a complete library of FORTRAN 77 subroutines developed by this author. These routines range from simple screen attribute calls (bolding, blinking, double-height double width characters, etc.) to complicated routines such as on-screen graphs.

Notes: Documentation is available by either typing the file VEHICLE.HLP or requesting HELP from within the program.

Changes and Improvements: Improved MMI. Includes source code to all library routines.1

Media (Service Charge Code): One RX50 Diskette (JA) Format: MS/DOS

DECUS NO: V-SP-40 **TITLE:** PRAXIS: An Alternative to Ada **Version:** 7.7, July 1987

Submitted by: Frederick Holloway, Lawrence Livermore National Laboratory, Livermore, CA

Operating System: VAX/VMS V4.5 **Source Language:** PRAXIS **Keywords:** Programming Languages, Tools - Applications Development

Abstract: PRAXIS is a modern block structured controls-oriented language similar to Ada (registered DoD) for distributed control system applications on VAX/VMS, PDP-11/RXS, LSI-11/RXS, and LSI-11 stand-alone computers. It is also useful as a training aid and stepping-stone to Ada.

PRAXIS was developed for and used extensively on the Nova High Energy Laser Project at the Lawrence Livermore National Laboratory by Bolt, Beranek and Newman, Inc. It has been extensively improved recently at LLNL in collaboration with other users. Features include: separate compilation of modules, strong type-checking, user defined types, encapsulation, guard and exception blocks for error control, segment control, clean interface to other languages, and ROM-able output code. In addition to the compilers, the release contains test suites, run-time support, text I/O routines (terminal and file), documentation sources (RUNOFF), and other support utilities. The compiler runs under VMS and can generate code for any of the above combinations.

Version 7.7 adds support for the VMS symbolic debugger, run-time library, extensive enhancements to the compiler including optimized rangechecking, and a test suite of over 400 example programs.

Direct contact with the submitter is encouraged for further information and assistance.

Notes: Sources of example programs and run-time support are included.

Changes and Improvements: Support VMS symbol debugger, VMS RTL, rangechecking, over 400 example programs, extensive enhancements.

Complete sources not included. **Media (Service Charge Code):** 2400' Magnetic Tape (PA) Format: VMS/BACKUP

DECUS PROGRAM LIBRARY CHANGES:

DECUS NO. VAX-235, "CAYENNE", the media code is listed as being available on 600' Magtape (MS) in the Catalog. Change the media code as being available on 600' Magtape (MC).

DECUS NO. V-SP-64, "Symposium Collection for the VAX SIG, Spring 1987, Nashville", the media code is listed as being available on 2400' Magtape (PS). Change the media code as being available on 2400' Magtape (PB).

DECUS NO. VAX-6, "SPICE 3A7 Version: 3A7", the media code for the User's Manual is listed as being available on (EB) and (EC). The Media code (EB) is no longer available.

SUBMITTING ARTICLES TO HARD NEWS

The purpose of HARD NEWS, the HMS SIG newsletter, is to serve as a forum to share information related to DEC hardware with the members of the SIG. As such, the existence of the newsletter is entirely dependent on your contributions. If you have an HHK item, a better or safer way to do something, product news, a tutorial article of general interest, etc., we would like to publish it in the newsletter. We hope that HARD NEWS will be published at least six times a year.

You can submit material to the editor, Carmen Wiseman, or to the HMS SIG chair, Bill Walker. We can accept submissions in a wide variety of formats:

- o Items can be sent to the editor on VMS-format RX50s, TK50 cartridges, or IBM PC format 5 1/4" floppies. The SIG chair prefers RT-11 floppies but can handle any reasonable media.
- o Hard copy, like cash, is always acceptable. Camera-ready copy will save us a lot of typing, but we don't insist on it. You can also use the Hardware Submission Form in the "Questionnaires" section of the combined SIGs Newsletters.
- o Those of you with access to DCS can send things to WALKER or WISEMAN. DCS is usually checked on a daily basis.
- o You can reach the SIG chair on CompuServe as "Bill Walker 71066,24" or via EasyLink mailbox 62752448 or MCI Mail account 333-1675. You can reach the editor via EasyLink mailbox 62960090 (be sure to say ATTN: or TO: Carmen Wiseman somewhere in the body of the message).

If you have anything to submit, send it! If it is a mess, but we can read it, we will get it into the newsletter somehow. Finally, if you have any questions about submitting material, call one of us. The telephone numbers are listed below.

Contributions can be sent to:

William K. Walker		Carmen D. Wiseman
Monsanto Research Corp.	OR	Digital Review
P.O. Box 32 A-152	==	Prudential Tower, Suite 1390
Miamisburg, OH 45342		800 Boylston Street
(513) 865-3557 (work)		Boston, MA 02199
(513) 426-7094/0344 (home)		(617) 375-4361 (work)

How to Submit articles to the RSTS Newsletter

The RSTS SIG newsletter solicits contributions of items of interest including, but not limited to, articles, DCL magic, copies of SPR's, and war stories.

You may electronically submit material by calling the SIG newsletter system at (201) 435-2546 at either 300 or 1200 baud. Press a few RETURN's until you get the RSTS banner, then sign on with account 2,1. No password is required. KERMIT is available for uploading material. Then you can use MAIL to compose a cover letter for your material and send it to NEWS.

You may also reach the editor as user KENNEDY on both DCS and DECUServe, if you have access to either of those systems.

You may also submit material on RX50's (in RSTS or RT11 format), on 800, 1600, 3200, or 6250 BPI 9-track tape (in DOS, ANSI, BRU, RSTS BACKUP or VMS BACKUP format), or on PC-DOS floppies (5¼ or 3½ inch format). If you are really desperate, I can also accept RSTS or RT11 format RL02 and RK07 disks. You may also submit hardcopy documents, but these will take longer to get into print.

If you are sending media you want returned, please insure it.

If you want your submission returned, please include a completed airbill billed to you, or include reasonable funds for insurance and return.

The address for sending material via US Mail is:

Terence M. Kennedy
St. Peter's College
Department of Computer Science
2641 Kennedy Blvd.
Jersey City, N.J. 07306
(201) 435-1890

The address for sending material via UPS, FedEx, etc. is:

Terence M. Kennedy
St. Peter's College
Department of Computer Science
121 Glenwood Avenue
Jersey City, N.J. 07306
(201) 435-1890



DECUS U.S. CHAPTER SUBSCRIPTION SERVICE SIGS NEWSLETTERS ORDER FORM (U.S. Members Only)

As a member of DECUS U.S. Chapter, you are entitled to contribute and subscribe to the DECUS monthly publication, **SIGs Newsletters**. You also have the opportunity to subscribe to the Symposia Proceedings which are a compilation of the reports from various speakers at the U.S. National DECUS Symposia.

- **No Purchase Orders will be accepted.**
- The order form below must be used as an invoice.
- All checks must be made payable to DECUS.
- All orders **MUST** be paid in full.
- Minimum of \$25.00 for orders placed via a credit card.
- No refunds will be made.
- The address provided below will be used for all DECUS mailings, i.e. Membership, Subscription Service and Symposia.
- SIGs Newsletters Price is for a one-year subscription beginning the month following receipt of payment.

Name _____ DECUS Member # _____
 Company _____
 Address _____

 City _____ State _____ Zip _____
 Telephone # _____

Subscription Service Offering	Unit Price	Quantity	Total
SIGs Newsletters	\$35.00	_____	_____
Fall '86 Proceedings (FA6)	15.00	_____	_____
Spring '87 Proceedings (SP7)	15.00	_____	_____
Fall '87 Proceedings (FA7)	15.00	_____	_____
Spring '88 Proceedings (SP8)	15.00	_____	_____
Total Amt.			\$ _____

MASTERCARD VISA DINERS CLUB/CARTE BLANCHE®
 Credit Card # _____ Expiration Date _____

I understand that there will be no refunds even if I decide to cancel my subscription.

Signature _____

FOR DIGITAL EMPLOYEES ONLY

Badge # _____ Cost Center _____
 Cost Center Mgr. Name _____ Cost Center Mgr. Signature _____

MAIL TO: Subscription Service, DECUS (BP02), 219 Boston Post Road, Marlboro, MA 01752-1850, (617) 480-3659.

FOR DECUS OFFICE ONLY

Check Number _____ Bank Number _____
 Amount \$ _____



DECUS U.S. CHAPTER APPLICATION FOR MEMBERSHIP

New Membership Update to current membership profile Current DECUS Member.# _____

Please provide a complete mailing address, include zip code in accordance with postal regulations for your locality.

Are you an employee of Digital Equipment Corporation? YES NO

NOTE: Please print clearly or type!

Name: _____
(First) (Middle Initial) (Last/Family Name)

Company: _____

Address: _____

City/Town/State/Zip: _____

Telephone: Home () _____ Work () _____

How Did You Learn About DECUS? Please Check Applicable Item.

- | | | |
|---|---|---|
| <input type="checkbox"/> 1 ANOTHER DECUS MEMBER | <input type="checkbox"/> 4 DIGITAL SALES | <input type="checkbox"/> 13 LOCAL USERS GROUP |
| <input type="checkbox"/> 2 SYMPOSIA | <input type="checkbox"/> 5 HARDWARE PACKAGE | <input type="checkbox"/> 14 SPECIAL INTEREST GROUP |
| <input type="checkbox"/> 8 DECUS CHAPTER OFFICE | <input type="checkbox"/> 6 SOFTWARE PACKAGE | <input type="checkbox"/> 7 SOFTWARE DISPATCH (Digital Newsletter) |
| <input type="checkbox"/> 10 DIGITAL STORE | <input type="checkbox"/> 12 ADVERTISING | |

Do you wish to be included in mailings conducted by Digital (for marketing purposes etc.?) Permission Refusal

Type Of Digital Hardware Used: Please Check Those Applicable To You.

- | | | | |
|--|---|--|------------------------------------|
| <input type="checkbox"/> 20 DECMATE | <input type="checkbox"/> 52 LSI-11 | <input type="checkbox"/> 21 PROFESSIONAL | <input type="checkbox"/> 5 WPS-8 |
| <input type="checkbox"/> 82 DECSYSTEM-10 | <input type="checkbox"/> 3 PDP-8 FAMILY | <input type="checkbox"/> 22 RAINBOW | <input type="checkbox"/> 51 WPS-11 |
| <input type="checkbox"/> 83 DECSYSTEM-20 | <input type="checkbox"/> 50 PDP-11 FAMILY | <input type="checkbox"/> 54 VAX FAMILY | |

Major Operating Systems? Languages Used: Please Check Those Applicable To You.

- | | | | | |
|-----------------------------------|--|-------------------------------------|------------------------------------|--|
| <input type="checkbox"/> 1 ADA | <input type="checkbox"/> 26 CORAL-66 | <input type="checkbox"/> 47 FOCAL | <input type="checkbox"/> 67 OS/8 | <input type="checkbox"/> 109 RT-11 |
| <input type="checkbox"/> 2 ALGOL | <input type="checkbox"/> 28 COS | <input type="checkbox"/> 48 FORTRAN | <input type="checkbox"/> 68 PASCAL | <input type="checkbox"/> 97 TECO |
| <input type="checkbox"/> 5 APL | <input type="checkbox"/> 34 DATATRIEVE | <input type="checkbox"/> 51 GAMMA | <input type="checkbox"/> 72 PL-11 | <input type="checkbox"/> 70 TOPS-10 |
| <input type="checkbox"/> 7 BASIC | <input type="checkbox"/> 35 DBMS | <input type="checkbox"/> 110 IAS | <input type="checkbox"/> 92 RPG | <input type="checkbox"/> 71 TOPS-20 |
| <input type="checkbox"/> 17 BLISS | <input type="checkbox"/> 38 DECNET | <input type="checkbox"/> 53 IQL | <input type="checkbox"/> 81 RSTS/E | <input type="checkbox"/> 111 ULTRIX/UNIX |
| <input type="checkbox"/> 19 C | <input type="checkbox"/> 43 DIBOL | <input type="checkbox"/> 58 MACRO | <input type="checkbox"/> 83 RSX | <input type="checkbox"/> 104 VMS |
| <input type="checkbox"/> 22 COBOL | <input type="checkbox"/> 45 DOS-11 | <input type="checkbox"/> 65 MUMPS | <input type="checkbox"/> 91 RMS | <input type="checkbox"/> 107 WPS-8 |

Type Of Business (Environment)/Computer Applications
Please Check That Which Best Describes Your Business/ Application.

- | | | |
|--|---|---|
| 21 <input type="checkbox"/> ACCOUNTANCY | 1 <input type="checkbox"/> EDUCATION/PRIMARY | 23 <input type="checkbox"/> NUMERICAL CONTROL |
| 7 <input type="checkbox"/> BANK | 2 <input type="checkbox"/> EDUCATION/SECONDARY | 68 <input type="checkbox"/> OEM-COMMERCIAL |
| 64 <input type="checkbox"/> BUSINESS/COMMERCIAL | 61 <input type="checkbox"/> EDUCATION-TECHNOLOGY | 78 <input type="checkbox"/> OEM-TECHNICAL |
| 74 <input type="checkbox"/> BUSINESS/INFORMATION SYSTEMS | 3 <input type="checkbox"/> EDUCATION/UNIVERSITY | 56 <input type="checkbox"/> PHYSICAL SCIENCES |
| 57 <input type="checkbox"/> CHEMISTRY | 67 <input type="checkbox"/> ENGINEERING | 20 <input type="checkbox"/> RESEARCH/DEVELOPMENT |
| 54 <input type="checkbox"/> CLINICAL LABORATORY | 65 <input type="checkbox"/> FINANCE/ACCOUNTING | 10 <input type="checkbox"/> RETAIL |
| 63 <input type="checkbox"/> COMPUTATION | 77 <input type="checkbox"/> GOVERNMENT | 73 <input type="checkbox"/> SOFTWARE DEVELOPMENT |
| 11 <input type="checkbox"/> CONSUMER ELECTRONICS | 75 <input type="checkbox"/> GRAPHICS | 53 <input type="checkbox"/> TELECOMMUNICATIONS |
| 18 <input type="checkbox"/> CONSULTANT | 4 <input type="checkbox"/> HOSPITAL | 19 <input type="checkbox"/> TELEPHONE/UTILITIES |
| 72 <input type="checkbox"/> DATA ACQUISITION | 62 <input type="checkbox"/> INDUSTRIAL | 51 <input type="checkbox"/> TIMESHARING |
| 52 <input type="checkbox"/> DATA COMMUNICATIONS | 55 <input type="checkbox"/> LABORATORY/SCIENTIFIC | 80 <input type="checkbox"/> TRAINING/INSTRUCTION |
| 13 <input type="checkbox"/> DATA PROCESSING SERVICES | 14 <input type="checkbox"/> LIBRARY | 66 <input type="checkbox"/> TYPESETTING/PUBLICATION |
| 71 <input type="checkbox"/> DATA REDUCTION | 58 <input type="checkbox"/> LIFE SCIENCES | |
| 17 <input type="checkbox"/> DIGITAL EMPLOYEE-ENGINEERING | 70 <input type="checkbox"/> MANUFACTURING | |
| 15 <input type="checkbox"/> DIGITAL EMPLOYEE-MARKETING | 79 <input type="checkbox"/> MARKETING | |
| 16 <input type="checkbox"/> DIGITAL EMPLOYEE-SERVICE GROUP | 59 <input type="checkbox"/> MEDICAL RESEARCH | |
| 60 <input type="checkbox"/> EDUCATIONAL ADMINISTRATION | 6 <input type="checkbox"/> MILITARY INSTALLATION | |

Special Interest Groups (SIGs) Enrollment

I Wish To Participate In The Following DECUS U. S. Chapter Special Interest Groups.

- | | | |
|--|--|---|
| 3 <input type="checkbox"/> ARTIFICIAL INTELLIGENCE | 11 <input type="checkbox"/> HARDWARE AND MICRO | 36 <input type="checkbox"/> PERSONAL COMPUTER |
| 7 <input type="checkbox"/> BUSINESS APPLICATIONS | 35 <input type="checkbox"/> IAS | 18 <input type="checkbox"/> RSTS/E |
| 2 <input type="checkbox"/> COMMERCIAL LANGUAGES | 27 <input type="checkbox"/> LARGE SYSTEMS | 17 <input type="checkbox"/> RSX |
| 6 <input type="checkbox"/> DATA MGMT. SYSTEMS | 16 <input type="checkbox"/> L & T | 19 <input type="checkbox"/> RT-11 |
| 31 <input type="checkbox"/> DAARC (LABS) | 14 <input type="checkbox"/> MUMPS | 32 <input type="checkbox"/> SITE MGMT. & TRNG |
| 5 <input type="checkbox"/> DATATRIEVE/4GL | 15 <input type="checkbox"/> NETWORKS | 21 <input type="checkbox"/> UNISIG |
| 8 <input type="checkbox"/> EDUSIG | 34 <input type="checkbox"/> OFFICE AUTOMATION | 26 <input type="checkbox"/> VAX |
| 10 <input type="checkbox"/> GRAPHICS APPLICATIONS | | |

Job Title/Position - Please Check:


- | | |
|---|---|
| 1 <input type="checkbox"/> CORPORATE STAFF | 101 <input type="checkbox"/> CORPORATE DIRECTOR OF DP/MIS |
| 2 <input type="checkbox"/> DIVISION OR DEPARTMENT STAFF | 102 <input type="checkbox"/> ADMINISTRATIVE ASSISTANT |
| 3 <input type="checkbox"/> SYSTEMS ANALYSIS | 103 <input type="checkbox"/> TECHNICAL ASSISTANT |
| 4 <input type="checkbox"/> APPLICATIONS PROGRAMMING | 104 <input type="checkbox"/> SERVICES COORDINATOR |
| 5 <input type="checkbox"/> SYSTEMS ANALYSIS/PROGRAMMING | 105 <input type="checkbox"/> MANAGER |
| 6 <input type="checkbox"/> OPERATING SYSTEM PROGRAMMING | 106 <input type="checkbox"/> ANALYST |
| 7 <input type="checkbox"/> DATABASE ADMINISTRATION | 107 <input type="checkbox"/> PROGRAMMER |
| 8 <input type="checkbox"/> DATA COMMUNICATIONS/TELECOMMUNICATIONS | 108 <input type="checkbox"/> DATABASE MANAGER |
| 9 <input type="checkbox"/> COMPUTER OPERATIONS | 109 <input type="checkbox"/> DATABASE ADMINISTRATOR |
| 10 <input type="checkbox"/> PRODUCTION CONTROL | 110 <input type="checkbox"/> MANAGER OF DP OPERATIONS |

Citizen of The United States of America? YES NO Country: _____

Signature: _____ Date: _____

Forward To: DECUS U. S. Chapter
 Digital Equipment Computer Users Society
 Membership Processing Group
 219 Boston Post Road, BP02
 Marlboro, MA 01752-1850
 Phone: (617)480-3418
 DTN: 8-296-3418

STEERING COMMITTEE LISTS



ARTIFICIAL INTELLIGENCE SIG

CHAIR

Cheryl Jalbert
JCC
128 West Broadway
Granville, OH 43023
(614) 587-0157

VICE-CHAIR

OPS5 WORKING GROUP CHAIR
Don Rosenthal
Space Telescope Science Inst.
Homewood Campus
Baltimore, MD 21218
(801) 338-4844

NEWSLETTER TASK FORCE CHAIR

ADMINISTRATIVE ASSISTANCE
Becky Wise
Amdalh CSD
2200 North Greenville Ave.
Richardson, TX 75081
(214) 699-9500 x 272

NEWSLETTER EDITOR

Terry Shannon
Digital Review
Prudential Tower
800 Boylston St. Suite 1390
Boston, MA 02199
(617) 375-4321

SYMPOSIUM COORDINATOR

Pam Vavra
Hughes Aircraft EDSG
P.O. Box 902 E52/D220
El Segundo, CA 90245-0902
(213) 616-7071

MEMBERSHIP COORDINATOR

SUITE COORDINATOR
Chris Goddard
Simpact Associates
9210 Skypark Court
San Diego, CA 92123
(619) 565-1865

SESSION NOTE EDITOR

George Humfeld
Naval Sea Systems Command
PMS 350 ED Dept of the Navy
Washington, DC 20362-5101
(202) 692-0137

ASS'T SESSION NOTES EDITOR

David Frydenlund

STORE REPRESENTATIVE

Sally Townsend
Inst. Defense Analysis
1801 N. Beauregard St.
Alexandria, VA 22311
(703) 845-2122

PUBLIC DOMAIN SOFTWARE TF CHAIR

LIBRARY REPRESENTATIVE

Jim Sims
Space Telescope Science Ins.
3700 San Martin Drive
Baltimore, MD 21218
(801) 338-4949

AI LUG COORDINATOR

ASSISTANT STORE REP.

Dennis Clark
RT2 Box 264
Kingston, TN 37763
(615) 576-7384

REPORTER TO THE UPDATE.DAILY

Bill Lennon

SEMINAR UNIT REP. CAMPGROUND COORDINATOR

Leona Fluck
Educational Testing Service
Rosedale Road
Princeton, NJ 08540
(609) 734-1243

DEC COUNTERPART

Art Beane
Hudson, MA

MEMBERS-AT-LARGE

David Slater
George Winkler
Jeff Fox
John Williamson
Wayne Graves
Matt Mathews
Dave Campbell
Shirley Bockstahler-Brandt
Barry Breen
Tom Viana



BUSINESS APPLICATIONS SIG

CHAIRMAN

George Dyer
Gallaudet University
800 Florida Ave, NE-EMG Bldg
Washington, DC 20002
(202) 651-5300

COMMUNICATIONS COORDINATOR

Steve Lacativa
Price Waterhouse
153 East 53rd Street
New York, NY 10022
(212) 371-2000 x 3107

SYMPOSIUM COORDINATOR

Mark Hulst
USSA Administrative Systems
USSA Bldg. B01E
San Antonio, TX 78288
(512) 498-8725

LUG COORDINATOR

Patrick LeSene
U.S. Coast Guard
Room 1416E 2100 2nd St. SW
Washington, DC 20593
(202) 267-0354

MARKETING COORDINATOR

Tom Byrne
L. Karp & Sons
1301 Estes
Elk Grove Village, IL 60007
(312) 598-5706

PROGRAM PLANNING COORDINATOR

Stuart Lewis
Douglas Furniture Corp.
P.O. Box 97
Bedford Park, IL 60499
(312) 458-1505

SEMINARS COORDINATOR

Daniel Esbensen
Touch Technologies, Inc.
9990 Mesa Rm., Rd. #220
San Diego, CA 92121
(619) 455-7404

LRP COORDINATOR

Arnold I. Epstein
D-M Computer Consultants
Rolling Meadows, IL 60008
(312) 394-8889

LIBRARY REPRESENTATIVE

David Hittner
Projects Unlimited
3680 Wyse Road
Dayton, OH 45414
(513) 890-1800

CL SIG LIAISON

Becky Burkes-Ham

DMS SIG LIAISON Joe Sciuto MEMBERS-AT-LARGE

Robert D. Lazenby
Dixie Beer Dist., Inc.
Louisville, KY

Robert Kayne
Gallaudet College
Washington, DC

Ray Evanson
Paragon Data Systems
Winona, MN

DEC COUNTERPARTS

Sue Yarger
Digital Equipment Corporation
Merrimack, NH 03054-0430

Paula Daley
Digital Equipment Corporation
Merrimack, NH 03054-0430

Pam Kukla
Digital Equipment Corporation
Maynard, MA 01754

Wombat Magic



DATATRIEVE/4GL SIG

CHAIRMAN

Joe H. Gallagher
Research Medical Center
2316 East Meyer Blvd
Kansas City, MO 64132
(816) 276-4235

PAST SIG CHAIRMAN

Larry Jasmann
U.S. Coast Guard
10067 Marshall Pond Rd.
Burke, VA 22015
(202)267-2626

SYMPOSIUM COORDINATOR

Lisa M. Pratt
Vitro Corporation
Nuwes Code 3144
Keyport, WA 98345
(206) 396-2501

ASS'T SYMPOSIUM REPRESENTATIVE

T.C. Wool
E.L. duPont DeNemours & Co.
Engineering Dept.
P.O. BOX 6090.
Newark, DE 19714-6090

COMMUNICATIONS REPRESENTATIVE

NEWSLETTER EDITOR

LRRP

Donald E. Stern, Jr.
Warner Lambert Company
10 Webster Road
Milford, CT 06460
(203) 783-0238

ASSOCIATE EDITOR

Steve Cordiviola
Kentucky Geological Survey
311 Breckinridge Hall
Lexington, KY 40506
(606) 257-5863

Pasquale (Pat) F. Scopelliti
Corning Glass Works
Mail Stop MP-RO-01-1
Corning, New York 14831
(607) 974-4496

Lorey B. Kimmel
Thomson-CGR Medical Corp.
10150 Old Columbia Road
Columbia, Maryland 21046
(301) 290-8757

ASST. VOLUNTEER COORDINATOR

Susan Krentz
NKF Engineering
12200 Sunrise Valley Dr.
Reston, VA 22091
(703) 620-0900

SEMINARS

Dana Schwartz
15719 Millbrook Lane
Laurel, MD 20707
(301) 859-6277

SESSION NOTES EDITOR

Wanda Anderson
SRI International MS: pN841
333 Ravenswood Avenue
Menlo Park, CA 94025
(415) 859-2577

CAMPGROUND

Bert Roseberry
Commandant (G-APA-1)
2100 2nd Street, S.W.
Washington, DC 20593-0001
(202) 267-2629

WW EDITOR

**PIR COORDINATOR
LRRP**

Philip A. Naecker
Consulting Engineer
3011 N. Mount Curve Ave.
Altadena, CA 91001
(818) 791-0945

DEC COUNTERPARTS

DATATRIEVE

Mary Ann Fitzhugh
110 Spit Brook Road
ZK2-2/M28
Nashua, NH 03060
(603) 881-2329

ARTIST & LIBRARY REP.

Bart Z. Lederman
I.T.T. World Communications
67 Broad Street (28th Floor)
New York, NY 10004
(212) 607-2657

POWERHOUSE W/G CHAIR

Fred Vietri
TSO Financial Corp.
Five TSO Financial Center
Three Hundred Welsh Road
Horsham, PA 19044-2009
(215) 784-0520

DMS & CL SIG LIAISON


William Tabor
Computer Products
Pompano Beach, FL

SMARTSTAR WORKING GROUP CHAIR

Leslie Byars
Corning Electronics
3900 Electronics Drive
Raleigh, NC 27604
(919) 878-6301

ACCENT-R USER GROUP LIAISON

Winston Tellis
Fairfield University
North Benson Road
Fairfield, CT 06430
(203) 255-5411



DAARC SIG

CHAIRMAN

James Deck
Inland Steel Research Lab.
3001 East Columbus Drive
East Chicago, IL 46312
(219) 392-5613

SYMPOSIA COORDINATOR

Mack Overton
FDA
Chicago, IL

COMMUNICATIONS REPRESENTATIVE

NEWSLETTER EDITOR

Dale Hutchison
Cummins Engine Co.
4720 Baker St., Ext.
Lakewood, NY 14750
(716) 466-2191

DEC COUNTERPART

Bill Forbes
Marlboro, MA

HARDWARE & INTERFACING

Peter Clout
Los Alamos National Lab
Los Alamos, NM

MATH STATISTICS & ANALYSIS

Herbert J. Gould
C.C.F.A. Univ. of Ill. Medical Ctr.
Chicago, IL

PROCESS CONTROL-INDUSTRIAL AUTOMATION

Bill Tippie
Kinetic Systems Corp.
Lockport, IL

RS-1

George Winkler
CPC International
Argo, IL



DATA MANAGEMENT SYSTEMS SIG

CHAIRMAN

Joseph F. Sciuto
Army Research Institute
5001 Eisenhower Ave
Alexandria, VA 22333
(202) 274-8221

COMPTROLLER

Alan Schultz
Land Bank National DP Center
7300 Woolworth Ave
Omaha, NE 68124
(402) 397-5040

SYMPOSIA COORDINATOR

Keith Hare
JCC
P.O. Box 463
Granville, OH 43023
(614) 587-0157

SYMPOSIA COORDINATOR

Barbara Mann
TRW
Redondo Beach, CA
(213) 532-2211

COMMUNICATIONS REPRESENTATIVE

NEWSLETTER EDITOR

Mark S. Crego
Mantech Int'l Corp.
2121 Eisenhower Ave.
Alexandria, VA 22314
(703) 838-5677

SESSION NOTE EDITOR

Mark Morgan
Farm Credit Banks
P.O. Box 141
Springfield, MA 01102
(413) 732-9721

MEMBERSHIP COORDINATOR

Vacant

PRODUCT DIRECTION COMMITTEE

PAST SIG CHAIRMAN

Steve Pacheco
Ship Analytics
North Stonington, CT 06359
(203) 535-3092

**WORKING GROUP COORDINATOR/
DATABASE WORKING GROUP**

Jim Perkins
PSC, Inc.
20 Kimball Ave., Suite 305
Shelburne, VT 05401
(802) 863-8825

FORMS WORKING GROUP

ANSI STANDARDS COORDINATOR

Paul W. Plum, Jr
Lukens Steel Company
Coatesville, PA
(215) 383-2024

NON-DIGITAL WORKING GROUP

Doug Dickey
GTE Government Systems
1700 Research Blvd
Rockville, MD 20850
(301) 294-8400

RMS WORKING GROUP COORDINATOR

Allen Jay Bennett
Lear Siegler Rapistan
555 Plymouth N.E.
Grand Rapids, MI 49505
(616) 451-6429

PRE-SYMPOSIUM SEMINAR COORDINATOR

Rocky Hayden
Userware International
Escondido, CA
(619) 745-6006

AI SIG LIAISON

David Slater
Institute for Defense Analysis
Alexandria, VA
(703) 845-2200

DATATRIEVE SIG LIAISON

William I. Tabor
W.I. Tabor, Inc.
Coral Springs, FL
(305) 755-7895

DEC COUNTERPART

Wendy Herman
Nashua, NH



EDUSIG

CHAIRMAN

Robert A. Shive, Jr.
Millsaps College
Jackson, MS 39210
(601) 354-5201

SYMPOSIA COORDINATOR

Mary Jac Reed
Off Comp Based Instruction
University of Delaware
305 Willard Hall
Newark, DE 19716
(302) 451-8161

COMMUNICATIONS REPRESENTATIVE

Robert W. McCarley
Millsaps College
Jackson, MS 39210
(601) 354-5201

NEWSLETTER EDITOR

Fred Bell
Taft College
29 Emmons Park Drive
P.O. Box 1437
Taft, CA 93267
(805) 763-4282

PSS COORDINATOR**VAX SYSTEMS SIG LIAISON**

Donald C. Fuhr
Tuskegee Institute
Tuskegee Institute, AL 36088
(205) 727-8242

ADMINISTRATIVE APPLICATIONS COORD.

Dave Cothrun
Taft College
29 Emmons Pk Drive
P.O. Box 1437
Taft, CA 93268
(805) 763-4282

SESSION NOTE EDITOR

Paula Barnes
Guilford College
5800 West Friendly Avenue
Greensboro, NC 17410
(919) 292-5511

DEC COUNTERPART

Gary Finerty
Marlboro, MA



**Graphics
Applications
SIG**

GRAPHICS APPLICATIONS SIG**CHAIRMAN**

William Kramer
NAS Systems Engineering Branch
NASA Ames Research Center
Moffett Field, CA 94035
(415) 694-5189

SYMPOSIA COORDINATOR

Bijoy Misra
Smithsonian Institution
60 Gordon St., MS39
Cambridge, MA 02138
(617) 495-7392

COMMUNICATIONS REPRESENTATIVE**NEWSLETTER EDITOR**

Carol Schwob
Florida Atlantic Univ.
Bldg. 22, Room 106
Boca Raton, FL 33431
(305) 393-2640

NEWSLETTER EDITOR (acting)

Jim Flatten
Digital Equipment Corp.
110 Spithook Road
ZK02-3/R56
Nashua, NH
(603) 881-0049

ASSOCIATE NEWSLETTER EDITOR

Charles D. Carter
Huntington Alloys, Inc.
Technology Dept.
P.O. Box 1958
Huntington, WV 25720
(304) 526-5721

WORKSTATION WORKING GROUP COORD.

Bob McCormick
Video Communications, Inc.
1325 Springfield Street
Feeding Hills, MA 01030
(413) 786-7955

ENGINEERING GRAPHICS WORKING GROUP COORD.

Eric Rehm
Gonzaga University
SPOCAD
E 502 Boone
Spokane, WA 99258
(509) 484-6814

SESSION NOTE EDITOR

Carol Schwob
Florida Atlantic University
Academic Computing
500 N.W. 20th Street
Boca Raton, FL 33431
(305) 393-2640

LIBRARY COORDINATOR

Mike McPherson
Michigan University
269 Engineering Bldg.
East Lansing, MI 48824
(517) 353-9769

STANDARDS COORDINATOR

Jim Flatten
Ames Lab
2581 Metals Dev.
Ames, IA 50011
(515) 294-7908

VOLUNTEER COORDINATOR

Dick McCurdy
University of Florida
Room 216, Larsen Hall
Gainesville, FL 32611
(904) 392-4915

LIBRARY COMMITTEE

James M. Turner
Saber Technology
San Jose, CA

DEC COUNTERPART

Rick Berzle
Spit Brook, NH

INFORMATION OFFICER

Mike York
Boeing Computer Services
P.O. Box 24346 M/S 03-73
Seattle, WA 98124
(206) 342-1442

HUMAN INTERFACE WORKING GROUP COORD.

Dottie Elliott
Northrop Services Inc.
P.O. Box 12313
Research Triangle PK, NC 27709
(919) 541-1300

DATA DISPLAY WORKING GROUP COORD.

Joy Williams
Eaton Corp.
P.O. Box 766
Southfield, MI 48037

HARD NEWS

HARDWARE MICRO SIG**CHAIRMAN**

William K. Walker
Monsanto Research Corp.
Miamisburg, OH

PRODUCT PLANNING COORDINATOR

George Hamma
Synergistic Technology
Cupertino, CA

PRE-SYMPOSIUM SEMINAR COORDINATOR

James R. Lindsmith
Monsanto Research Corp.
Miamisburg, OH

COMMUNICATIONS COORDINATOR

John G. Hayes
Information Systems
South Central Bell
Birmingham, AL

NEWSLETTER EDITOR

Carmen D. Wiseman
Digital Review
Boston, MA

SESSION NOTE EDITOR**DAARC SIG LIAISON**

Bill Tippie
Kinetic Systems Corp.
Lockport, IL

STANDARDS COORDINATOR**CAMAC WORKING GROUP COORDINATOR**

Peter Clout
Los Alamos National Lab
Los Alamos, NM

LUG COORDINATOR

Gregg Giesler
Los Alamos Science Lab
Los Alamos, NM

TOEM (CHIPS & BOARDS)

Jack J. Peterson
Horizon Data Systems
Richmond, VA

HHK (HARDWARE HINTS & KINKS)

Wayne Kesling
Monsanto Research Cor.
Miamisburg, OH

UNIBUS HARDWARE

Ron Bogue
LJV Aerospace & Defense Co.
Dallas, TX

PERFORMANCE MEASUREMENT COORD.

William Wallace
600 W. Washington Street
Peoria, IL

CSS COORDINATOR

Pratap Gohel
E.I. duPont
Ingleside, TX

NETWORKS SIG LIAISON

Sandra Traylor
Target Systems
Yorba Linda, CA

VAX SIG LIAISON

Dave Schmidt
5100 Centre Avenue
Pittsburgh, PA

UNISIG LIAISON

Jim Livingston
1 Results Way
Cupertino, CA

SITE SIG LIAISON

Emily Kitchen
A.H. Robins Co.
Richmond, VA

RT-11 SIG LIAISON

Gary Sallee
Sallee Software Consulting
Yorba Linda, CA

RSX SIG LIAISON

Hans Jung
Associated Press
New York, NY

MEMBERS-AT-LARGE

Mike Rembis
Americán Dade
Costa Mesa, CA

Hans Dahlke
Richland, WA

Jim Cutler
EDS Tower
16533 Evergreen
Southfield, MI

DEC COUNTERPARTS TERMINALS

Nina Abramson
Maynard, MA

TOEM (Chips & Boards)

Art Bigler
Marlboro, MA

DIAGNOSTIC

George D. Cooke
Maynard, MA

STORAGE

Marilyn Fedele
Maynard, MA

MSD (Micro Systems Develop.)
Roy Rodgers
Maynard, MA

PRINTER PRODUCTS
Frank Orlando
Maynard, MA

DECUS EUROPE LIAISON
Hans Zoller



IAS SIG

CHAIRMAN

Mike Robitaille
Grumman - CTEC, Inc.
6862 Elm Street
McLean, VA 22101
(703) 556-7400 x 431

NEWSLETTER EDITOR

Frank R. Borger
Radiation Therapy
Michael Reese Medical Center
Lake Shore Drive @ 31st St.
Chicago, IL 60616
(312) 791-2515

WHIMS COORDINATOR

Kathleen Anderson
Eaton Information Management
System Division
Hampton, VA
(804) 826-1941

RSX LIAISON

Ray French
Boeing Computer Services
Seattle, WA
(206) 655-6228

MEMBER-AT-LARGE

Doug Reno
Abbot Laboratories
North Chicago, IL
(312) 937-7504

CHAIRMAN EMERITUS

Bob Curley
Division of Medical Physics
University of Pennsylvania
Philadelphia, PA
(215) 662-3083

SYMPOSIA COORDINATOR

Lynda L. Roenicke
Special Systems Branch
Naval Ocean Systems Center
271 Cataline Blvd Code 742
San Diego, CA
(619) 225-7569

LIBRARY COORDINATOR

Ted Smith
The University of PA
Philadelphia, PA 19101
(215) 662-3500

MEMBER-AT-LARGE

Kerry Wyckoff
Salt Lake City, UT

DEC COUNTERPART

Nancyfaye Autenzio
Stow, MA
(617) 496-9606



LANGUAGES AND TOOLS SIG

CHAIRMAN

Sam Whidden
American Mathematical Society
201 Charles St.
P.O. Box 6248
Providence, RI 02940
(401) 272-9500

VICE CHAIR

SYMPOSIA COORDINATOR

Earl Cory
Eaton Corporation
31717 Latienda Dr.
Westlake Village, CA 91359
(818) 706-5385

STORE REPRESENTATIVE

Chair, TECH. PROD OF DOC. W/G

Howard Holcombe
RCA
Front & Cooper Sts.
Camden, NJ 08055
(609) 338-4946

NEWSLETTER EDITOR

Al Folsom, Jr.
Fischer & Porter Co.
E. County Line Rd.
Warminster, PA 18974
(215) 674-7154

SESSION NOTES EDITOR

Mark Katz
GTE Government Systems
100 First Avenue
Waltham, MA 02154
(617) 466-3437

AUSTRALIAN L&T INTERFACE

Gordon Brimble
Bldg. 180 Labs Area
Defence Research Centre
Box 2151 GPO
Adelaide, S.A. Australia 5001
(61)(8)259-6119

INTERSIG COORDINATOR

Dorothy Geiger
Wollongong Logistics Group
49 Showers Drive # 451
Mountain View, CA 94040
(415) 948-1008

EUROPEAN METHODS

L&T INTERFACE

Bernd Gliss
Max-Planck-Institute
Heisenbergstra Be 1
7000 Stuttgart 80, W. Germany
(711) 686-0251

PAST CHAIR

PRODUCTIVITY TOOLS COORDINATOR

Kathy Hornbach
Digital Equipment Corporation
110 Spit Brook Rd., ZK02-2/R55
Nashua, NH 03062
(603) 881-2505

CHAIR TECO WORKING GROUP

Mark J. Hyde
Advanced Computing Services
209 Ardsley Drive
DeWitt, NY 13214
(315) 446-7223

CHAIR, BASIC WORKING GROUP

MEMBER, ANSI BASIC X3J2 STDS, COMM.

Stephen C. Jackson
SCJ Consulting, Inc.
7260 University Avenue N.E.
Suite 105
Minneapolis, MN 55432
(612) 571-8430

CHAIR, CONFIG. MGMT. WORKING GROUP

Mark Alan Kidwell
Texas Instruments Inc.
P.O. Box 869305 M/S 8435
Plano, TX 75086
(214) 575-3547

DEVEL. COUNTERPART, PDP-11 SOFTWARE

Joe Mulvey
Digital Equipment Corp., ZK01-3/J10
110 Spit Brook Road
Nashua, NH 03062-2642
(603) 881-1218

LISP/AI COORDINATOR

Don Rosenthal
Space Telescope Science Institute
Homewood Campus
Baltimore, MD 21218
(301) 338-4844

LIBRARY REPRESENTATIVE

SIG TAPE LIBRARIAN

PUBLIC DOMAN SOFTWARE W/G CHAIR

Tony Scandora
Argonne National Laboratory
CMT 205
Argonne, IL 60439
(312) 972-7541

CHAIR, C WORKING GROUP

James Maves
Eaton Corp, IMSD
31717 Latienda Drive
Box 5009
Westlake Village, CA 91359
(818) 706-5395

COMMCOMM REP.

STANDARDS COORDINATOR

TOOLS INTEGRATION W/G CHAIR

Jay Wiley
Bechtel Power Corp
12400 East Imperial Highway
Norwalk, CA 90650
(213) 807-4016

ASST TO CHAIR

JR Westmoreland
UTAH Power & Light
1407 West North Temple
Annex 6/208
Salt Lake City, UT 84116
(801) 535-4784

RECORDING SECRETARY

Melodee Westmoreland
Custom Software Products
1456 E. Hilda Drive
Fruit Heights, UT 84037
(801) 533-2350

SESSION CHAIRS COORDINATOR

Gary C. Lelvis
IMSL
2500 Park West Tower One
2500 City West Blvd.
Houston, TX 77042-3020
(713) 782-6060

CHAIR, FORTRAN WORKING GROUP

Scott Krusemark
8473 Daisywood Ave N.W.
North Canton, OH 44720
(216) 499-6251

VOLUNTEERS COORDINATOR

Louise Wholey
Measurix Corp
1 Results Way
Cupertino, CA 95014
(408) 255-1500 x 5571

CHAIR, LOW LEVEL LANGUAGES W/G

Gerald Lester
Computerized Processes Unlim.
2901 Houma Blvd. Suite 5
Metairie, LA 70006
(504) 889-2784

DEVEL. COUNTERPART, COMM. LANG.

Jim Totton
Digital Equipment Corp.
ZK02-3/K06
110 Spit Brook Road
Nashua, NH 03062

HUMAN INTERFACE COORDINATOR

Jim Wilson
Pfizer
QC Division
P.O. Box 88
Terre Haute, IN 47808
(812) 299-2121 x 271

VICE CHAIR**SYMPOSIUM LOGISTICS COORD.**

Terry Medlin
Survey Sampling, Inc.
1 Post Road
Westport, CT 06482
(203) 255-4200

CHAIR, DIBOL WORKING GROUP

Bruce Mebust
Vicom
9713 Valley View Road
Eden Prairie, MN 55344
(612) 944-7135

MASTERS COORDINATOR**CHAIR, PROJECT MGMT W/G**

Dena Shelton
Cullinet Software Inc.
2860 Zanker Rd, Suite 206
San Jose, CA 95134
(408) 434-6636

APL WORKING GROUP CHAIR

Bob Van Keuren
UserWare International, Inc.
2235 Meyers Ave.
Escondido, CA 92025
(619) 745-6006

WISHLIST COORDINATOR

Shava Nerad
MIT
77 Mass Avenue 10-256
Cambridge, MA 02139
(617) 253-7438

WORKING GROUPS COORD.**ASST. SYMPOSIUM LOGISTICS COORD.**

Joseph Pollizzi, III
Space Telescope Science Institute
3700 San Martin Drive
Homewood Campus
Baltimore, MD 21218
(301) 338-4901

CHAIR, MODULA II WORKING GROUP

Jack Davis
NAP Consumer Electronics Corp.
Video Business Unit
111 North Shore Drive
Knoxville, TN 37919
(615) 558-5206

CHAIR, SCAN WORKING GROUP**CHAIR, PL/I WORKING GROUP**

David Ream
Lexi-Comp
26173 Tallwood Drive
N. Olmsted, OH 44070
(216) 777-0095

CHAIR, PASCAL WORKING GROUP

E. Wayne Sewell
E-Systems, Garland Div.
Box 660023, MS 53730
Dallas, TX 75266-0023
(214) 272-0515 x3553

CHAIR, PDP-11 LAYERED PROD. W/G

William I. Tabor
W.I. Tabor, Inc.
11652 N.W. 30th St.
Coral Springs, FL 33065
(305) 755-7895

CHAIR, SOFTWARE METRICS W/G

Michael S. Terrazas
LDS Church
50 E. North Temple, 27th Floor
Salt Lake City, UT 84150
(801) 531-3246

MEMBER, ANSI COBOL X3J4, STDS, COMM.

Bruce Gaader
Donahue Enterprises, Inc.
2441 26th Ave., S.
Minneapolis, MN 55406
(612) 721-2418

CHAIR, RPG WORKING GROUP

Charles Williamson
Hargray Telephone Co.
P.O. Box 5519
Hilton Head Is., SC 29938
(803) 686-1204

PRE-SYMPOSIUM SEMINAR COORD.

Barry C. Breen
Sundstrand Data Control, Inc.
15001 N.E. 36th Street
P.O. Box 97001
Redmond, WA 98073-9701
(206) 885-8436

CHAIR COBOL WORKING GROUP

Edward W. Woodward
SAIC
10210 Campus Point Drive
M/S # 24
San Diego, CA 92121
(619) 546-3758

CHAIR, LARGE SYS. MAIN. W/G**CLINIC DIRECTOR**

George Scott
Computer Sciences Corp.
304 West Route #38, P.O. Box N
Moorestown, NJ 08057
(609) 234-1100

ASSISTANT CAMPGROUND COORD.

Keith Batzel
Crowe, Chizek & Co.
330 E. Jefferson
Box 7
South Bend, IN 46624
(219) 232-3992

STEERING COMMITTEE MEMBERS

Ted Bear
Ramtek
2211 Lawson Lane
Santa Clara, CA 95050
(408) 988-2211

DEVEL. COUNTERPART TECH. LANG.

Leslie J. Klein
Digital Equipment Corp.
ZK02-3/N30
110 Spit Brook Road
Nashua, NH 03062

DIGITAL COUNTERPART

Celeste La Rock
Digital Equipment Corp.
ZK02-3/008
110 Spit Brook Road
Nashua, NH 03062

**LARGE SYSTEMS****CHAIR**

Leslie Maltz
Stevens Institute of Technology
Computer Center
Hoboken, NJ 07030
(201) 420-5478
BITNET:MALTZ@JUNCC.CSC.ORG
ARPANET:SIT.MALTZ@CU20B.COLUMBIA.EDU

ASST. SEMINARS REPRESENTATIVE

Robert C. McQueen
Stevens Institute of Technology
Computer Center
Hoboken, NJ 07030
(201) 420-5454
BITNET:RMCQUEEN@SITVXB:
ARPANET:SIT.MCQUEEN@CU20B.COLUMBIA.ED

COMMUNICATIONS REPRESENTATIVE**NEWSLETTER EDITOR**

Clyde T. Poole
The University of Texas at Austin
Department of Computer Science
Taylor Hall 2.124
Austin, TX 78712-1188
(512) 471-9551
ARPANET/CSNET:ctp@sally.utexas.edu

MENU COORDINATOR

OPEN

36 BIT SYSTEMS

Clive Dawson
Microelectronics & Computer Technology Corp.
9430 Research Blvd.
Echelon Bldg. #1, Suite 200
Austin, TX 78759
(512) 343-0860
ARPANET/CSNET:CLIVE@MCC.COM

LANGUAGES COORDINATOR

OPEN

SYMPOSIUM REPRESENTATIVE

Betsy Ramsey
American Mathematical Society
P.O. Box 6248
Providence, RI 02940
(401) 272-9500 x 295
ARPANET:EWR@XX.LCS.MIT.EDU

SYSTEM SOFTWARE COORDINATOR

OPEN

SUPERCOMPUTING

E.F. Berkley Shands
Washington University
Department of Computer Science
P.O. Box 1045
St. Louis, MO 63136
(314) 889-6636
UUCP:BERKLEY@WUCS.UUCP
BITNET:Berkley@Wunet

NETWORKS COORDINATOR

Don Kassebaum
Computation Center
University of Texas at Austin
Austin, TX 78712
(512) 471-3241
ARPANET:CC.KASSEBAUM@A20.CC.UTEXAS.F

SITE SIG LIAISON

Gary C. Bremer
Emerson Electric Co.
Electronics and Space Division
8100 W. Florissant
St. Louis, MO 63136
(314) 553-4448/4447

SIG VICE-CHAIRMAN**CORPORATE ISSUES**

Ralph M. Bradshaw
Johnson & Johnson
Research & Scientific Services
Management Information Center
Raritan, NJ 08869-1489
(201) 685-3434

LIBRARY REPRESENTATIVE**SIR/MENU BALLOT**

Jack Stevens
The Gillette Company
Technical Services, 4U-3
1 Gillette Park
Boston, MA 02106-2131
(617) 463-2089

SIG MARKETING

Steve Attaya
Weiner Enterprises
P.O. Box 23607
Harahan, LA 70183
(504) 733-7055
ARPANET:g.attaya@R20.UTEXAS.EDU

DEC COUNTERPARTS

Dave Braithwaite
Digital Equipment Corporation
Marlboro, MA
Rich Whitman
Digital Equipment Corporation
Marlboro, MA
Reed Powell
Digital Equipment Corporation
Marlboro, MA



MUMPS SIG

CHAIRMAN

Mark Berryman
Plessey Peripheral Systems
Irvine, CA

SYMPOSIA COORDINATOR

Chris Richardson
Computer Sciences Corp.
Ridgecrest, CA

NEWSLETTER EDITOR

Mark J. Hyde
Advanced Computing Services
209 Ardsley Drive
DeWitt, NY 13214
(315)46-7223

VAX LIAISON

Coyett A.J. Dese
VA DM&S Verification & Dev. Ctr.
San Francisco, CA

DEC COUNTERPARTS

Beatrice Walther
Marlboro, MA
Diane Brown
Marlboro, MA



NETWORKS SIG

CHAIRMAN

Stuart Lewis
Douglas Furn. Corp.
(312) 458-1505

COMMUNICATIONS COMMITTEE REP.

Bob Gustafson
Northeast Utilities
(203) 665-5082

NEWSLETTER EDITOR

Judi Mandl
UConn Health Center
263 Farmington Ave. Bldg. 19
Farmington, CT 06032

SEMINAR UNIT REP & VICE (BACKUP) SIG CHAIR

Sandy Traylor
Target Systems, Inc.
(714) 921-0112

SYMPOSIA COORDINATOR

Bill Hancock
(817) 261-2283

STANDARDS COORDINATOR

Jim Ebright
Software Results Corp.
(614) 267-2203

ASSISTANT NEWSLETTER EDITOR

Judi Mandl
UConn Health Center
(203) 674-3912

SESSION NOTES EDITOR

Mary Marvel-Nelson
General Motors Research Lab.
(313) 986-1382

DEC COUNTERPART

Monica Bradlee
(617) 486-7341



OFFICE AUTOMATION SIG

CHAIR

Katherine "Kit" Trimm
Pivotal, Inc
Tucson, AZ
(602) 886-5563

VICE CHAIRMAN

Ralph Bradshaw
Johnson and Johnson
Raritan, NJ
(201) 685-3434

COMMUNICATIONS REPRESENTATIVE

E. Catherine Ditamore
ARA Services
Philadelphia, PA
(215) 238-3638

SYMPOSIA COORDINATOR

Mitch Brown
GenRad, Ind.
Waltham, MA
(617) 369-4400 x3052

NEW MEMBER COORDINATOR

Tricia Cross
American Mathematical Society
P.O. Box 6248
Providence, RI 02940
(401) 272-9500

BOF COORDINATOR

Ray Kaplan
PIVOTAL, Inc.
Tucson, AZ
(602) 886-5563

NEWSLETTER EDITOR

Therese LeBlanc
T.M. LeBlanc & Assoc.
Wheeling, IL
(312) 459-1784

LIBRARY

Bob Hassinger
Liberty Mutual Research Center
Hopkinton, MA
(617) 435-9061

OA TAPE COORDINATOR

Mary Jane Boiling
Foreign Mission Board
3806 Monument Avenue
Richmond, VA 23230
(804) 353-0151

SYMPOSIA ASSISTANT

Sal Gianni
Northeast Utilities
Hartford, CT
(203) 665-5652

STORE COORDINATOR

Mike Jackson
Air Force Operational
Test and Evaluation Center
Kirtland AFB, NM
(505) 846-5641

PERSONAL COMPUTER SIG LIAISON

Cheryl Johnson
Grinnell College
Grinnell, IA
(515) 236-2570

OA LUG COORDINATOR

Tom Orlovski
American Council on Education
1 DuPont Circle (Suite 110)
Washington, DC
(202)939-9371

OA SIG COORDINATOR

Joe Whatley
Neilson Media Research
375 Patricia Avenue
Dunedin, FL 33528
(813)734-5473



PERSONAL COMPUTER SIG

CHAIR

Barbara Maaskant
UT Health Science Center—CR
7703 Floyd Curl Drive
San Antonio, TX 78284
(512) 567-2200

PRO WORKING GROUP CHAIRMAN

Thomas R. Hintz
Univ. of Florida
IFAS Computer Network,
Bldg. 120
Gainesville, FL 32611
(904) 392-5180

DECMATE WORKING GROUP CHAIRMAN PRE-SYMPOSIUM SEMINAR COORD.

Vince Perriello
Crosfield CSI
570 Taxter Rd.
Elmsford, NY 10523
(914) 592-3600

VICE, CHAIR RAINBOW W/G CHAIRMAN

Lynn Jarrett
Union Tribune Publishing Co.
P.O. Box 191
San Diego, CA 92108
(619) 299-3131 x1130

VAXMATE WORKING GROUP CHAIRMAN

Frederick G. Howard
Eastman Kodak Company
901 Elm Grove Road, D845-LP
Rochester, NY 14650
(716) 253-2363

VOLUNTEER COORDINATOR

Pierre M. Hahn
SUNY HSC-T10-028-8101
Stony Brook, NY 11794

LIBRARIAN/Committee Rep.

Ron S. Hafner
Hafner and Associates
P.O. Box 2924
Livermore, CA 94550
(415) 449-4178

COMMUNICATIONS REPRESENTATIVE STORE REPRESENTATIVE

Kenneth LeFebvre
Sytek, Inc.
19 Church St.
P.O. Box 128
Berea, OH 44017
(216) 243-1613

NEWSLETTER EDITOR

Gary Rice
McDonnell Douglas
5555 Garden Grove Blvd.
MS: K20 77/200
Westminster, CA 92683
(714) 952-6582

DECMATE W.G. CHAIR PRE-SYMPOSIA SEMINAR COORDINATOR

Vince Perriello
Crosfield Composition Systems
570 Taxter Road
Elmsford, NY 10523
(914) 592-3600

SYMPOSIA COORDINATOR

Rick Eliopoulos
5258 Vickie Drive
San Diego, CA 92109
(619) 225-7867

SESSION NOTE EDITOR

Dr. Tomas L. Warren
Oklahoma State Univ.
Dept. of English
Div. Tech. Writing Program
Stillwater, OK 74078
(405) 624-6138

SYMPOSIA COORDINATOR 12/87

Jim Wilson
Ntl Tech Inst for the Deaf
Rochester Inst. of Tech
P.O. Box 9887
Rochester, NY 14623
(716) 475-6241

MEMBERS-AT-LARGE

Michael Bowers
Univ. of California
Animal Science Department
Davis, CA 95616
(916) 752-6136

Theodore Needleman
Hardcopy Magazine
Seldin Publishing, Inc.
1061 S. Melrose, Suite D
Placentia, CA 92607
(714) 356-6331

DEC COUNTERPARTS**PRO**

To Be Identified
Digital Equipment Corp.
ML021-2/U12
146 Main Street
Maynard, MA 01750

PERSONAL COMPUTING SYSTEMS GRP.

Anita Uhler
Digital Equipment Corp.
LJ02/13
30 Porter Road
Littleton, MA 01460
(617) 486-2451

CAMPGROUND COORDINATOR

Jim Hobbs
Technical Support Group
Adolf Coors Co. (BC380)
Golden, CO 80401-1295
(303) 277-2855

**RSTS SIG****CHAIRMAN**

Charles Mustain
Stark County School system
Louisville, OH

SYMPOSIA COORDINATOR

Glenn Dollar
Digital Computer Consultants Inc.
21363 Lassen St., Suite 205
Chatsworth, CA 91311
(818) 341-9171

ASS'T SYMPOSIA COORDINATOR

Wef Fleischman
Software Techniques
Cypress, CA

NEWSLETTER EDITOR

Terence M. Kennedy
St. Peter's College
Department of Computer Science
2641 Kennedy Blvd.
Jersey City, NJ 07306
(201) 435-1890

LIBRARY REPRESENTATIVE

Susan Abercrombie
Ventrex Laboratories Inc.
Portland, ME

PRE-SYMPOSIA SEMINAR COORDINATOR

Scott Castleberry
1750 North Collins
Suite 108
Richardson, TX 75080
(214) 437-3477

WISH LISTS COORDINATOR

Neal E. Goldsmith
Software Techniques, Inc.
Cypress, CA

VICE CHAIRMAN**WISH LISTS COORDINATOR**

Lynnell Koehler
Campus America
POISE Prod. Ctr.
201 North Nevada Avenue
Roswell, NM 88201
(505) 625-5500

EDUSIG LIAISON

George Wyncott
Purdue University Computer Center
W. Lafayette, IN

RSTS PRODUCT PLANNING COORDINATOR

Errol E. Ethier
Information Design and Management, Inc.
23 Hunting Avenue
(617) 842-4220
Shrewsbury, MA 01545

DEC COUNTERPART

Kathy Waldron
Digital Equipment Corporation
Continental Blvd.
Merrimack, NH 03054

MEMBERS-AT-LARGE

Edward F. Beadel, Jr.
Manager
Instructional Computing Center
S.U.N.Y. College at Oswego
Oswego, NY 13126
(315) 341-3055
Mark Hartman
Jadtec Computer Group
546 W. Katella Avenue
Orange, CA 92667
(714) 997-8928
Jeff J. Killeen
Information Design & Management, Inc.
31 Hopedale Street
Hopedale, MA 01747

**RSX SIG****CHAIRMAN**

Dan Eisner
Perkin-Elmer Corp.
Garden Grove, CA

SYMPOSIA COORDINATOR

Rick Sharpe
Toledo Edison
Toledo, OH

PRE-SYMPOSIUM SEMINAR COORDINATOR

Hans Jung
Associated Press
New York, NY

COMMUNICATIONS REPRESENTATIVE

Jay Allen Bennett
Lear Siegler Rapistan
Grand Rapids, MI

NEWSLETTER EDITOR**MULTI-PROCESSORS WORKING GROUP COORDINATOR**

Bruce Mitchell
Machine Intelligence & Industry Magin
Byron, MIN

STORE COORDINATOR

Jim Hopp
Carlton Financial Computation
South Bend, IN

SESSION NOTE EDITOR

Burt Janz
BHJ Associates
Nashua, NH

LIBRARIAN

Glenn Everhart
Mt. Holly, NJ

CAMPGROUND COORDINATOR

Jerry Ethington
Prolifix Inc.
Frankfort, KY

DEC COUNTERPARTS

Lin Olsen
Nashua, NH
Dick Day
Nashua, NH

WORKING GROUP COORDINATOR

Sharon Johnson
Epidemiology
Minneapolis, MN

WORKING GROUP CHAIR

Evan Kudlajev
Philadelphia Electric Co.
Philadelphia, PA

RSX GROUP CHAIR SOFTWARE CLINIC COORD.

Roy S. Maull
U.S. Air Force
Offutt AFB, NE

SOFTWARE CLINIC COORDINATOR

Bruce Zielinski
RCS
Moorestown, NJ

VOLUNTEER COORDINATOR

Gary Maxwell
U.S. Geological Survey
Menlo Park, CA

SRD WORKING GROUP COORDINATOR

Bob Turkelson
Goddard Space Flight Center
Greenbelt, MD

ACCOUNTING & PERFORMANCE WORKING GROUP COORDINATOR

Denny Walthers
American McGraw
Irvine, CA

MENU COORDINATOR

Ed Cetron
Center for Biomedical Design
Salt Lake City, UT

MEMBERS-AT-LARGE

Jim McGlinchey
Warrenton, PA
Jim Neeland
Hughes Research Labs.
Malibu, CA
Anthony E. Scandora, Jr.
Argonne National Laboratory
Argonne, IL
Ralph Stamerjohn
Creve Coeur, MO

**RT-11 SIG****CHAIRMAN**

John T. Rasted
JTR Associates
58 Rasted Lane
Meriden, CT 06450
(203) 684-1632

COM. COM VOTING REP.**COBOL CONTACT**

Bill Leroy
The Software House, Inc.
P.O. Box 52661
Atlanta, GA 30355-0661
(404) 231-1484

STANDARDS COORDINATOR

Robert Roddy
Naval Ship Research Ctr.
Bethesda, MD 20084
(301) 227-1724

MACRO CONTACT

Nick Bourgeois
NAB Software Services Inc.
P.O. Box 20009
Albuquerque, NM 87154
(505) 296-2346

NEWSLETTER EDITOR
TECO CONTACT
PRODUCT PLANNING CONTACT
John M. Crowell
Multiware, Inc.
2121-B Second St. Suite 107
Davis, CA 95616
(916) 756-3291

NETWORKING CONTACT
Jim Crapuchettes
Omnex Corp.
2483 Old Middlefield Way
Mountain View, CA 94043
(415) 966-8400

WISH LIST CONTACT
UNIX/ULTRIX CONTACT
Bradford Lubell
L.A. Heart Lab, UCLA
10833 Le Conte Avenue
Los Angeles, CA 90024-1760
(213) 206-6119

TSX & C CONTACT
Jack Peterson
Horizon Data Systems
P.O. Box 29028
Richmond, VA 23229
(804) 740-9244

RUNOFF CONTACT
John Davis
Naval Ship Research Center
Code 2950
Bethesda, MD 20084
(301) 227-1592

LUG CONTACT
Ned Rhodes
Software Systems Group
2001 North Kenilworth St.
Arlington, VA 22205
(703) 534-2297

PERSONAL COMPUTERS
Dennis V. Jensen
AMES Labs. ISU/USDOE
310 Metallurgy
Ames, Iowa 50011
(515) 294-4823

SYMPOSIA COORDINATOR
Milton Campbell
Talisman Systems
Drawer CP-255
Manhattan Beach, CA 90266
(213) 318-2206

TAPE COPY GENERATION
TAPE COPY DISTRIBUTION
RT DECUS LIBRARY CONTACT
Tom Shinal
Syntropic Technology
P.O. Box 198
Waterford, VA 22190
(703) 882-3000

PRE-SYMPOSIUM SEMINAR
RT-11 SUITE MANAGER
Bruce Sidlinger
Sidlinger Computer Corp.
4335 N.W. Loop 410, #209
San Antonio, TX 78229
(512) DIG-ITAL

BASIC CONTACT
Ralston Barnard
Div 7523
Sandia Labs
Albuquerque, NM 87185
(505) 844-5115

PRO RT-11 & HARDWARE
Bill Walker
Monsanto Research Corp.
P.O. Box 32, A-152
Miamisburg, OH 45342
(513) 865-3557

FORTRAN CONTACT
Robert Walraven
Multiware, Inc.
2121-B 2nd St. Suite 107
Davis, CA 95616
(916) 756-3291

OTHER LANGUAGES
Gary Sallee
19912 Fernglan Drive
Yorba Linda, CA 92686
(714) 970-2864



SITE SIG
CHAIRMAN
DMS SIG Liaison
Larry W. Hicks
Relational Database services
P.O. Box 644
121 S. Main St.
Kernersville, NC. 27285-0644
(919) 996-4882

SYMPOSIA COORDINATOR
Sue Abercrombie
48 Malilly Rd.
Portland, ME 04103
(207) 772-2837

SESSION NOTE EDITOR
LARGE SYSTEMS SIG LIAISON
Gary Bremer
Emerson Electric Co.
8100 W. Flerisant
St. Louis, MO. 63136
(314) 553-4448

NEWSLETTER EDITOR
NETWORKS SIG LIAISON
OA SIG LIAISON
Gregory N. Brooks
Washington University
Behavior Research Labs
1420 Grattan St.
St. Louis, MO. 63104
(314) 241-7600 ext. 257

LIBRARY COORDINATOR
RSTS SIG LIAISON
Timothy Frazer
Specialized Bicycle Components
15130 Concord Circle #77
Morgan Hill, CA. 95037
(408) 779-6229

HARDWARE COORDINATOR
HMS SIG Liaison
Emily Kitchin
A.H. Robins Co.
1211 Sherwood Ave. RT-2
Richmond, VA. 23220
(804) 257-2925

COMMUNICATIONS COMMITTEE REPRESENTATIVE
AI SIG Liaison
Terry C. Shannon
Digital Review
160 State St.
6th Floor
Boston, MA. 02109
(617) 367-7190

PRE-SYMPOSIUM SEMINAR COORDINATOR
Phillip Ventura
STAFF MANAGEMENT
Adam Zavitski
Simmonds Precision ICD
3100 Highland Blvd.
Raleigh, NC. 27625
(919) 872-9500

MEMBERS-AT-LARGE
Ann Goergen
Texas Instruments
13510 N. Central
M/S 437
Dallas, TX. 75266
(214) 995-4629

HMS SIG Liaison
RT SIG Liaison
David Hunt
Lawrence Livermore National Lab
MS L-230
P.O. Box 808
Livermore CA. 94550
(802) 656-3190
Gary Siftar
Digital Equipment Corporation
Tulsa, OK.

DEC COUNTERPARTS
Joe Allen
Stow MA.
Lil Holloway
Bedford MA.
Susan Porada
Marlboro, MA.



UNISIG
CHAIRMAN
Kurt Reisler
Hadron Incorporated
9990 Lee Highway
Fairfax, VA 22030
(703) 359-6100
decvax!seismo!hadron!klr

SYMPOSIA COORDINATOR
Stephen M. Lazarus
Ford Aerospace, MS X-20
3839 Fabian Way
Paulo Alto, CA 94804
(415) 852-4203
ihnp4!fortune!wdll!sml

SESSION NOTE EDITOR
Bill Cheswick
Director of Academic Computing
Computer Services Dept.
323 Martin Luther King Blvd.
Neward, NJ 07960
(201) 596-2904

COMMUNICATIONS REPRESENTATIVE
NEWSLETTER EDITOR
James W. Livingston
Measurix Corporation
1 Results Way
Cupertino, CA 95014
(408) 255-1500 x 5556
ihnp4!decwrljwl

ADMINISTRATIVE DAEMON
Dorothy Geiger
The Wollongong Group
49 Showers Drive, 451
Mountain View, CA 94040
(415) 948-1003
ihnp4!decwrl!dgeiger

TAPE LIBRARIAN
Carl Lowenstein
Marine Physical Laboratory
Scripps Institute of Oceanography, P-004
LaJolla, CA 92093
(619) 294-2678
(ihnp4!decvax!akgul!dedwest!ucbvax)
!sdevax!lmpvax!cdl

USENET LIAISON
Joe Kelsey
FlexComm Corporation
711 Powell Avenue, SW
Renton, WA 98055
allegra!fluke!joe

STANDARDS REPRESENTATIVE
Ed Gould
Mt. Xinu
2560 Ninth Street
Suite 312
Berkeley, CA 94710
(415) 644-0146

MINISTER WITHOUT PORTFOLIO

Norman Wilson
 Bell Laboratories, 2C-529
 600 Mountain Avenue
 Murray Hill, NJ 07974
 (201) 582-2842
 (decvax!ihnp4)!research!norman

SEMINARS REPRESENTATIVE

Steven Stephanik

DEC COUNTERPART

Roseann Maclean
 Merrimack, NH
 (603) 884-5702
 decvax!maclean

**VAX SYSTEMS SIG****SYMPOSIUM COORD., ASSISTANT**

David Cossey
 Computer Center
 Union College
 Schenectady, NY 12308

SESSION NOTES EDITOR

Ken Johnson
 Meridian Technology Corp.
 P.O. Box 2006
 St. Louis, MO 63011

NEWSLETTER EDITOR

Lawrence J. Kilgallen
 Box 81, MIT Station
 Cambridge, MA 02139-0901

LIBRARY WORKING GROUP

Glen Everhart
 25 Sleigh Ride Road
 Glen Mills, PA 19342

VAXcluster WORKING GROUP

Thomas Linscomb
 Computation Center
 University of Texas
 Austin, TX 78712

NETWORK WORKING GROUP

Bill Hancock
 Dimension Data Systems, Inc.
 P.O. Box 13557
 Arlington, TX 76094-0557

MicroVAX WORKING GROUP

Ray Kaplan
 Pivotal, Inc.
 6892 East Dorado Court
 Tucson, AZ 85715-3264
 (602) 886-5563

SYSTEM IMPROVEMENT REQUEST (CORE)

Mark D. Oakley
 Battelle Memorial Institute
 505 King Avenue
 Columbus, OH 43201-2669

MULTIPROCESSOR WORKING GROUP

Eugene Pal
 U.S. Army
 CAORA (ATORCATC)
 Fort Leavenworth, KA

PRE-SYMPOSIUM SEMINAR COORD. HISTORIAN

Jeff Jalbert
 JCC
 P.O. Box 381
 Granville, OH 43023

PRE-SYMPOSIUM SEMINAR COORD. (ACTING)

June Baker
 Computer Sciences Corp.
 6565 Arlington Blvd.
 Falls Church, VA 22046

FIELD SERVICE WORKING GROUP

Dave Slater
 Computer Sciences Corp.
 6565 Arlington Blvd
 Falls Church, VA 22046

LARGE SYSTEMS INTEGRATION WORKING GP

Leslie Maltz
 Stevens Institute of Tech.
 Computer Center
 Hoboken, NJ 07030

VOLUNTEER COORDINATOR

Elizabeth Bailey
 222 CEB
 Tennessee Valley Authority
 Muscle Shoals, AL 35661

COMMERCIAL WORKING GROUP

Bob Boyd
 GE Microelectronics Center
 P.O. Box 13409, MS7T3-01
 Research Triangle Park, NC 27709-3049

SECURITY

C. Douglas Brown
 Sandia National Labs
 Division 2644
 P.O. Box 5800
 Albuquerque, NM 87185-5800

MIGRATION AND HOST DEVELOPMENT**VAXintosh WORKING GROUP**

Jim Downward
 KMS Fusion Incorporated
 P.O. Box 156D
 Ann Arbor, MI 48106

REAL TIME/PROCESS CONTROL

Dennis Frayne
 McDonnell Douglas
 5301 Bolsa Avenue
 Huntington Beach, CA 92646

Larry Robertson
 Bear Computer Systems
 56512 Case Avenue
 North Hollywood, CA

INTERNALS WORKING GROUP

Carl E. Friedberg
 Seaport Systems, Inc.
 165 William Street, 9th Floor
 New York, NY 10038-2605

COMMUNICATIONS ASSISTANT

David L. Wyse
 Professional Business Software
 3680 Wyse Road
 Dayton, OH 45414-2539

CAMPGROUND COORDINATOR

Kirk Kendrick
 Shell Oil Co.
 333 Highway G, MS D-2146
 Houston, TX 77082-8892

PAST CHAIR

Marge Knox
 Computation Center
 University of Texas
 Austin, TX 78712

SYSTEM MANAGEMENT

Steve Tihor
 251 Mercer Street
 New York, NY 10012

ADVISORS

Joseph Angelico
 U.S. Coast Guard Detachment
 National Data Buoy Center
 NSTL Station, MS 39529-6000

Art McClinton
 Mitre
 1820 Dolley Madison Blvd.
 McLean, VA 22102

Al Siegel
 Battelle Memorial Institute
 505 King Avenue
 Columbus, OH 43201-2693

CHAIR (CORE)

Susan T. Rehse
 Lockheed Missiles & Space Co.
 0/19-50, B/101, P.O. Box 3504
 Sunnyvale, CA 94088-3504

VICE-CHAIR (CORE) WORKING GROUP COORD.

Ross Miller
 Online Data Processing Inc.
 N 637 Hamilton
 Spokane, WA 99202

SYMPOSIA COORD. (CORE)

Jack Cundiff
 Horry-Georgetown Tech. College
 P.O. Box 1966
 Conway, SC 29526

COMMUNICATION COORD. (CORE)

Don Golden
 Shell Oil Company
 Westhollow Research Center
 P.O. Box 1380, Room D2158
 Houston, TX 77251-1380

LIBRARIAN

Joseph L. Bingham
 Mantech International
 2320 Mill Road
 Alexandria, VA 22314

LUG COORDINATOR (CORE)

Dave Schmidt
 Management Sciences Associates
 5100 Centre Avenue
 Pittsburgh, PA 15232

**Ask the WOMBAT WIZARD
Submission Form**

To submit a problem to the WIZARD, please fill out the form below and send it to:

WW Editor, Philip A. Naecker
Consulting Software Engineer
3011 North Mount Curve Avenue
Altadena, CA 91001
USA

Name: _____ DECUS Membership No. _____

Affiliation: _____

Address: _____

Telephone Number: _____

Statement of Problem: _____

Please following the following guidelines when submitting support material:

1. If you are trying to demonstrate a method or a concept, please simplify the procedures, records, and other information to the shortest form possible.
2. Annotate your attachments. Simple comments or hand-written notes ("Everything worked until I added this statement.") go a long way toward identifying the problem.
3. Keep an exact copy of what you send. And number the pages on both copies. But send everything that is related to your question, even remotely.
4. If you would like a direct response or would like your materials returned, please don't forget to include a stamped, self-addressed envelope large enough to hold the materials you send.

DATATRIEVE/4GL SIG
Product Improvement Request Submission Form

Submittor:
Address:

DECUS Membership Number:
Firm:

Phone:

Product or Products:

How to write a PIR

A PIR should be directed at a specific product or group of products. Be sure to give the full name of the product(s) and version numbers if applicable. Describe the functionality you would like to see in as complete terms as possible. Don't assume that the PIR editors or software developers know how it is done in some other software product - state specifically how you want the software to function. Provide justification of your request and give an example of its use. If you can, suggest a possible implementation of your request.

Abstract: (Please limit to one or two short sentences.)

Description and Examples: (Use additional pages as necessary.)

[Put my name and address on reverse side, thus:]

PIR Editor, Philip A. Naecker
Consulting Software Engineer
3011 North Mount Curve Avenue
Altadena, CA 91001
USA

H M S S I G

HARDWARE SUBMISSION FORM -- A SIG INFORMATION INTERCHANGE

Message _____

Contact Name _____

Address _____

Telephone _____

Type of equipment _____

SUBMIT ANY TYPE OF HARDWARE PROBLEMS AND/OR FIXES.

SEND TO:

William K. Walker
Monsanto Research Corp.
P.O. Box 32 A-152
Miamisburg, OH 45342

OR
==

Carmen D. Wiseman
Digital Review
Prudential Tower, Suite 1390
800 Boylston Street
Boston, MA 02199

IAS WHIMS

WHAT: (Describe your WHIM) (Please print or type)

WHY: (Describe the reason for the WHIM)

HOW: (Make any suggestions for a possible implementation)

Name: _____

Company: _____

Address: _____

Phone: _____

Please mail to:

Kathleen M. Anderson
EATON Information Management
Systems Division
2017 Cunningham Drive
Suite 208
Hampton, Virginia 23666

Phone: (804) 326-1941

Languages & Tools SIG

MASTERS APPLICATION

Name: _____ Title _____

Company: _____

Address: _____

_____ Phone: () _____

Network Address: _____ Date: _____

The Languages & Tools SIG has established the designation "LANGUAGES AND TOOLS MASTER", to be applied to selected, qualified people willing to share their expertise in various subjects with others. Masters are people who are knowledgeable enough in one or more languages or tools to be comfortable answering questions about them. The qualifications of an L&T Master are: expertise in a specific area, a willingness to have his/her name published as a Master, and a willingness to volunteer services in different ways. Each product may have several Masters, and there is an overall Masters Coordinator who is a member of the L&T Steering Committee.

Masters are asked to serve other users (and, under some circumstances, DEC), as a resource on products within their competence. In addition to being listed in the L&T Masters Directory (published in the newsletter) as available for occasional telephone consultation, Masters may act as 'Doctors' at Symposium Clinics, present Symposium sessions on the products of interest to them, field test products, interact with DEC product managers when appropriate, or act as a reference for a product for Digital salespeople. Especially on mature products, the SIG is anxious for knowledgeable users to offer product tutorial sessions at Symposia, and Masters can be of great help here. At Symposia, Masters will wear an identifying button bearing the legend "Ask Me About....." and the name of the language or tool in which he/she specializes.

If you'd like to serve as an L&T Master, please mark the products on which you are willing to answer questions with an "M" (for Master). Please mark any other products running at your site with an "A" (for "also running") to provide users with a broader picture of your facilities. (Although not an L&T product, Mumps is included here at the request of the Mumps SIG as a service to Mumps users). You may request removal of your name from the Masters Directory at any time, although you may continue to be listed for a month or two, because of publication lead times.

I am qualified to act as an L&T Master for the following products: Mumps

- | | | | | | |
|------------------------------------|--------------------------------|------------------------------|-------------------------------|---|---|
| <input type="checkbox"/> Debug | <input type="checkbox"/> Bliss | <input type="checkbox"/> CMS | <input type="checkbox"/> TPU | <input type="checkbox"/> C | <input type="checkbox"/> Test Manager |
| <input type="checkbox"/> Pascal | <input type="checkbox"/> Basic | <input type="checkbox"/> MMS | <input type="checkbox"/> EVE | <input type="checkbox"/> Ada ¹ | <input type="checkbox"/> Runoff & DSR |
| <input type="checkbox"/> Fortran | <input type="checkbox"/> Cobol | <input type="checkbox"/> LSE | <input type="checkbox"/> EDT | <input type="checkbox"/> APL | <input type="checkbox"/> T _E X & L ^A T _E X |
| <input type="checkbox"/> Document | <input type="checkbox"/> Dibol | <input type="checkbox"/> SCA | <input type="checkbox"/> TECO | <input type="checkbox"/> RPG | <input type="checkbox"/> Cobol Generator |
| <input type="checkbox"/> VAX Notes | <input type="checkbox"/> Emacs | <input type="checkbox"/> PCA | <input type="checkbox"/> PL/I | <input type="checkbox"/> Scan | <input type="checkbox"/> Software Project Mgr |

Briefly describe your experience with those you checked. _____

How long have you held your present position? _____

Are you able to attend at least one symposium each year? _____

Users are encouraged to seek assistance with products by calling appropriate Masters listed in the Directory. As a Master, your name and telephone number will be published in the Masters Directory, and users will call on you for limited help from time to time. Please check, below, any additional activities you might do:

- Field-test new versions of your product at your work site.
- Provide feedback on the product when needed by its DEC product manager.
- Act as a reference for the product at the request of Digital Sales or Marketing people.

Mail to: Dena Shelton, L&T SIG Masters Coordinator, Cullinet Software, Inc., 2860 Zanker Road, Suite 206, San Jose, CA 95134.

¹Ada is a trademark of the DoD

Languages & Tools SIG

WISHLIST QUESTIONNAIRE

Name: _____ Title _____

Company: _____

Address: _____

_____ Phone: () _____

Network Address: _____ Date: _____

The Languages & Tools SIG is principally concerned with the DEC and public domain software products listed below. If your request directly involves one of these products, please check which one (if you have more than one request, please use a separate form for each):

- | | | | | | |
|------------------------------------|--------------------------------|------------------------------|-------------------------------|---|---|
| <input type="checkbox"/> Debug | <input type="checkbox"/> Bliss | <input type="checkbox"/> CMS | <input type="checkbox"/> TPU | <input type="checkbox"/> C | <input type="checkbox"/> Test Manager |
| <input type="checkbox"/> Pascal | <input type="checkbox"/> Basic | <input type="checkbox"/> MMS | <input type="checkbox"/> EVE | <input type="checkbox"/> Ada ¹ | <input type="checkbox"/> Runoff & DSR |
| <input type="checkbox"/> Fortran | <input type="checkbox"/> Cobol | <input type="checkbox"/> LSE | <input type="checkbox"/> EDT | <input type="checkbox"/> APL | <input type="checkbox"/> T _E X & L _A T _E X |
| <input type="checkbox"/> Document | <input type="checkbox"/> Dibol | <input type="checkbox"/> SCA | <input type="checkbox"/> TECO | <input type="checkbox"/> RPG | <input type="checkbox"/> Cobol Generator |
| <input type="checkbox"/> VAX Notes | <input type="checkbox"/> Emacs | <input type="checkbox"/> PCA | <input type="checkbox"/> PL/I | <input type="checkbox"/> Scan | <input type="checkbox"/> Software Project Mgr |

If your request or suggestion doesn't relate to one of the products listed above, check which one of the following Language & Tools SIG topics it concerns:

- | | | |
|---|---|---|
| <input type="checkbox"/> Newsletter | <input type="checkbox"/> Symposium Sessions | <input type="checkbox"/> Pre-Symposium Seminars |
| <input type="checkbox"/> Masters Program | <input type="checkbox"/> Working Group Activities | <input type="checkbox"/> Session Notes |
| <input type="checkbox"/> Information Folder | <input type="checkbox"/> SIG Tape | <input type="checkbox"/> DECUS Store Item |
| <input type="checkbox"/> Other L&T SIG topic: _____ | | |

Wish List Request—brief description: _____

Complete description—please explain your request thoroughly; don't assume we know details of other products or services; give examples. _____

Mail to: Shava Nerad, L&T Wishlist Coordinator, MIT, 77 Mass Ave. W91-219A, Cambridge, MA 02139; (617)253-7438

¹Ada is a trademark of the DoD

DATAGRAM

DATAGRAMs are short messages, comments, requests, or answers that are published in NETwords. Please fill in the sections below and send the DATAGRAM to:

Vickie Hess
NETWords Editor
2510 Limestone Ln.
Garland, Tx. 75040

Title: _____

Message: _____

Your Name: _____

Address: _____

Telephone: _____

If this is a reply to a previous DATAGRAM, what #? ___

Signature: _____ **Date:** _____

**Place
Stamp
Here**

**Vickie Hess
NETWords Editor
2510 Limestone Ln.
Garland, Tx. 75040**

Fold Here

OFFICE AUTOMATION SIG SYSTEM IMPROVEMENT REQUEST BALLOT

DECUS Membership Number _____

INSTRUCTIONS: System Improvement Request (SIR) Ballots allow you, the user, to assist in the prioritization of the submitted SIR's before they are forwarded to Digital. The total number of points which you may allocate on this ballot may not exceed 100 points (absolute value). No more than 10 points may be given to any single SIR. Your ballot must be received by MARCH 1 to be counted.

SIR NUMBER	POINTS
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____

TOTAL

100 POINTS

E. Catherine Ditamore
ARA Services
Corp MIS
The ARA Tower
1101 Market Street
Philadelphia, Pa. 19107

PC POSTSCRIPT

PC Postscripts are short requests, comments and responses to be published in the *Postscript Section* of the PC SIG Newsletter. Please respond to the following:

Y/N This is a reply to a previous Postscript. Issue No. No.

Title: _____

Message: _____

Name: _____

Address: _____

Phone: (_____) _____

Signature: _____ Date _____

RSTS Newsletter Reader Survey

In order to serve you better, the newsletter editor solicits the following information:

I would to see a newsletter article on _____

I am interested in a Symposium session on _____

I am willing to write an article on _____

I/my company have _____ machines running RSTS version _____

I attend the DECUS Symposia: __ Always __ Sometimes __ Never

Do you/your company use DECmail-11? _____

What other operating systems do you use MAIL with? _____

Do you/your company use DECNET/E? _____

What other operating systems do you use DECNET with? _____

What other layered products do you use? _____

Would you be willing to serve as an 'expert' on one of the above products? _____ If so, which one(s)? _____

If so, please give contact information:

Name: _____

Address: _____

Phone: (____) _____ - _____ Ext. _____

Other comments: _____

fold here

Terence M. Kennedy
St. Peter's College
Dep't. of Computer Science
2641 Kennedy Blvd.
Jersey City, N.J. 07306

fold here

RSTS/E Wishlist Form

<u>Cusps</u> Editors _____ Error Reporting _____ Spooling _____ Backup Package _____ Other _____ <i>Please check only one box per form.</i>	<u>Monitor</u> Init _____ Installation _____ Executive _____ Terminal Service _____ DECnet/E _____ <i>Do Not submit language or layered product wishes.</i>	<u>Others</u> DCL _____ RMS-11 _____ Sort/Merge _____ RSX Emulator _____ RT Emulator _____ <u>Services</u> Documentation _____ Distribution _____ Update Svcs. _____
---	---	---

Wish:

Reason:

Digital/Decus Use Only

Technical _____ Product _____ Fun _____

Fall 87 DECUS - Anaheim

fold here

Terence M. Kennedy
St. Peter's College
Dep't. of Computer Science
2641 Kennedy Blvd.
Jersey City, N.J. 07306

fold here

RT-11 WISH LIST SURVEY

Name (optional) _____

Address (optional) _____

DECUS Number (optional) -----

1.1	___	3.1	___	3.7u	___	3.13a	___	5.1b	___
1.2	___	3.2a	___	3.7v	___	3.13b	___	5.2a	___
1.3	___	3.2b	___	3.7w	___	3.13c	___	5.2b	___
1.4	___	3.2c	___	3.7x	___	3.13d	___	6.1	___
1.5	___	3.2d	___	3.7y	___	3.14	___	6.2a	___
1.6	___	3.2e	___	3.7z	___	3.15	___	6.2b	___
1.7a	___	3.3a	___	3.7aa	___	3.16	___	6.2c	___
1.7b	___	3.3b	___	3.7bb	___	3.17a	___	6.2d	___
1.8	___	3.3c	___	3.7cc	___	3.17b	___	6.3	___
1.9a	___	3.3d	___	3.7dd	___	3.17c	___	6.4a	___
1.9b	___	3.4a	___	3.7ee	___	3.17d	___	6.4b	___
1.9c	___	3.4b	___	3.8a	___	3.17e	___	6.4c	___
1.9d	___	3.4c	___	3.8b	___	3.17f	___	6.4d	___
1.10	___	3.5a	___	3.8c	___	3.18	___	6.5	___
1.11	___	3.5b	___	3.9a	___	3.19a	___	6.6a	___
1.12	___	3.6a	___	3.9b	___	3.19b	___	6.6b	___
1.13	___	3.6b	___	3.9c	___	3.19c	___	6.6c	___
1.14	___	3.6c	___	3.9d	___	4.1	___	6.6d	___
2.1	___	3.6d	___	3.9e	___	4.2a	___	6.7	___
2.2	___	3.6e	___	3.9f	___	4.2b	___	6.8a	___
2.3	___	3.6f	___	3.9g	___	4.3	___	6.8b	___
2.4	___	3.6g	___	3.9h	___	4.4a	___	6.8c	___
2.5	___	3.7a	___	3.9i	___	4.4b	___	6.8d	___
2.6	___	3.7b	___	3.9j	___	4.5a	___	6.8e	___
2.7	___	3.7c	___	3.9k	___	4.5b	___	7.	___
2.8	___	3.7d	___	3.10a	___	4.6	___	8.	___
2.9	___	3.7e	___	3.10b	___	4.7a	___	9.1	___
2.10	___	3.7f	___	3.10c	___	4.7b	___	9.2a	___
2.11	___	3.7g	___	3.10d	___	4.7c	___	9.2b	___
2.12	___	3.7h	___	3.10e	___	4.7d	___	9.3a	___
2.13	___	3.7i	___	3.10f	___	4.7e	___	9.3b	___
2.14	___	3.7j	___	3.10g	___	4.7f	___	10.1	___
2.15	___	3.7k	___	3.10h	___	4.7g	___	10.2	___
2.16	___	3.7l	___	3.10i	___	4.7h	___	10.3	___
2.17	___	3.7m	___	3.10j	___	4.7i	___		___
2.18	___	3.7n	___	3.10k	___	4.7j	___		___
2.19	___	3.7o	___	3.10l	___	4.7k	___		___
2.20	___	3.7p	___	3.10m	___	4.7l	___		___
2.21	___	3.7q	___	3.10n	___	4.7m	___		___
2.22	___	3.7r	___	3.11a	___	4.7n	___		___
2.23	___	3.7s	___	3.11b	___	4.7o	___		___
2.24	___	3.7t	___	3.12	___	5.1a	___		___

Send Responses to: RT-11 Wish List Survey
 Multiware, Inc.
 2121-B Second St. Suite 107
 Davis, CA 95616

INPUT/OUTPUT Submission Form

A SIG Information Interchange

Please reprint in the next issue of the Pageswapper

If this is a reply to a previous I/O, which number? _____

Caption: _____

Message: _____

Contact:

Name _____

Address _____

Telephone _____

Signature _____ Date _____

Mail this form to: Larry Kilgallen, PAGESWAPPER Editor
Box 81, MIT Station, Cambridge, MA 02139-0901, USA

To register for on-line submission, dial (in the United States):
(617) 262-6830 and log in with the username PAGESWAPPER.

PAGESWAPPER - November 1987 - Volume 9 Number 4
INPUT/OUTPUT Submission Form

Tear out or photocopy reverse to submit an I/O item

Larry Kilgallen, PAGESWAPPER Editor
Box 81, MIT Station
Cambridge, MA 02139-0901
USA

System Improvement Request Submission Form

Page 1 of _____

Submittor: _____ Firm: _____

Address: _____ Phone: _____

How to write an SIR:

Describe the capability you would like to see available on VAX systems. Be as specific as possible. Please don't assume we know how it's done on the XYZ system. Justify why the capability would be useful and give an example of its use. If you wish, suggest a possible implementation of your request.

Abstract (Please limit to four lines):

Description and examples (use additional pages if required)

PAGESWAPPER - November 1987 - Volume 9 Number 4
System Improvement Request Submission Form

Tear out or photocopy reverse to submit an SIR

Mark D. Oakley
Battelle Columbus Division
Room 11-6-008
505 King Avenue
Columbus, Ohio 43201-2369
USA

Printed in the U.S.A.

“The Following are Trademarks of Digital Equipment Corporation”

AI	MicroVAX	TOPS-10
ALL-IN-1	MicroVAX-1 (et.al)	ULTRIX
DATATRIEVE	MicroVMS	UNIBUS
DDCMP	PDP	VAX
DEClander	P/OS	VAX-11/730 (et.al)
DECnet	Q-bus	VAX 8600 (et.al)
DECservice	Rainbow	VAXcluster
DECsystem-10	ReGIS	VAX DATATRIEVE
DECSYSTEM-20	RMS-11	VAXMATE
DECUS	RSTS	VAXstation
IAS (et.al)	RSX	VAX/VMS
LN03	RSX-11M-PLUS	VMS
MASSBUS	RT-11	WPS
MicroPDP-11	RX02 (et.al)	

**Copyright©DECUS and Digital Equipment Corporation 1987
All Rights Reserved**

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation or DECUS. Digital Equipment Corporation and DECUS assume no responsibility for any errors that may appear in this document.

It is assumed that all articles submitted to the editor of this newsletter are with the authors' permission to publish in any DECUS publication. The articles are the responsibility of the authors and, therefore, DECUS Digital Equipment Corporation, and the editor assume no responsibility of liability for articles or information appearing in the document. The views herein expressed are those of the authors and do not necessarily express the views of DECUS or Digital Equipment Corporation.

CP/M is a registered trademark of Digital Research, Inc.; IBM is a registered trademark of International Business Machine Corporation; Microsoft is a registered trademark of Microsoft Corporation; MUMPS is a registered trademark of Massachusetts General Hospital; MODULA-2 is a trademark of Interface Technologies Corporation; SMP is a trademark of Inference Corporation; TSX-PLUS is a trademark of S&H Computer Systems, Inc.; UNIX is a registered trademark of American Telephone & Telegraph Company.

Production Staff:

Beverly Welborne: Communications Committee Chair
R.E. Center

Don Golden: SIG Publications Chair
Shell Oil

Judy Arsenault: Publications Manager
DECUS

Judy Tessier: Phototypographer/Graphics Designer
DECUS

Circulation: 6443



DECUS SUBSCRIPTION SERVICE
DIGITAL EQUIPMENT COMPUTER SOCIETY
219 BOSTON POST ROAD, (BP02)
MARLBORO, MA 01752-1850

Bulk Rate
U.S. Postage
PAID
Permit No. 18
Leominster, MA
01453

STATUS CHANGE

Please notify us immediately to guarantee continuing receipt of DECUS literature. Allow up to six weeks for change to take effect.

- Change of Address
- Please Delete My Membership Record
(I Do Not Wish To Remain A Member)

DECUS Membership No: _____

Name: _____


Company: _____

Address: _____

State/Country: _____

Zip/Postal Code: _____

Mail to: DECUS - Attn: Subscription Service
219 Boston Post Road, BP02
Marlboro, Massachusetts 01752-1850
USA



label
not
old
ude
ion,
sity.