

# Central Support ON-LINE

Published for System Software Users

**McDonnell Douglas Field Service Company**

Volume 1

Number 3

**Irvine, Ca.**

3rd Quarter 1988

## LOGON:

### SOFTWARE SUPPORT SURVEY

In June, McDonnell Douglas Field Service Company mailed a Software Support Survey to 600 randomly chosen customers. Our desire was to obtain feedback concerning our Centralized Software Support Service. The return rate of the survey was 27%.

#### THE SUPPORT OPERATION

The questions were focused on the critical elements of the support operation. Customers were asked to rate these elements for both importance and satisfaction on a scale of 1 to 10. Weighted against importance scores of 10, satisfaction scores of 8, 9, and 10 are considered to be good to excellent. A score of 7 is acceptable, but should be considered an area to watch closely. A score of 6 or below is a danger signal.

The "gap" indicator is the difference between the importance and satisfaction rating. A gap of two points or greater for any element warrants concern. For instance, if the importance rating is two points higher than the satisfaction rating, that area should be targeted for improvement actions.

Following is a matrix of the average importance and satisfaction ratings for the seven questions asked in the survey.

SOFTWARE SERVICE	Importance	Satisfaction	Gap
Response Time	9.57	7.61	1.96
Telephone Manner	8.15	8.97	-.82
Technical Expertise	9.61	8.26	1.36
Resolution Time	9.42	7.46	1.96
<b>GENERAL CONSIDERATIONS</b>			
Overall System Performance	9.66	7.91	1.75
Availability of Management Personnel	8.53	7.24	1.29
Overall Software Service	9.31	7.73	1.58
<b>Average</b>	9.18	7.88	1.30

In the "Software Service" category, the area we passed with flying colors is Telephone Manner, with a negative gap of .82. The satisfaction rating is higher than the importance rating in your opinion. Since customer satisfaction is a strategic element of the Field Service Company's philosophy, our analysts have excellent customer satisfaction skills.

Technical expertise is also an element reflecting a good rating, with a gap of 1.3. We have worked hard over this past year to provide our analysts continuous training on the internals of the operating system.

The ratings for Response Time and Resolution Time both received borderline gaps of 1.96. With a gap of two representing a potential trouble spot, more in-depth evaluation is warranted here.

Response time is being improved by several methods. We have increased the staffing of this group by adding five new analysts since June. By instituting a "first-call screener" concept, we now have the flexibility to add more first-call screeners when the call load is heavy. Additionally we have switched more analysts to the early morning shift to accommodate the Eastern and Central time zones.

Resolution Time will continue to improve as our expertise is further enhanced. Also, as our staffing increases, we will be better able to follow through and expedite your problem toward a quicker resolution.

In the area of General Considerations, Overall System Performance relates to how your system and software perform for you. There are many complex factors which make this very unique for each customer. As the quality of our operating system and utility software packages improve along with your own application software, your system performance will improve.

Availability of management personnel had a rating of 7.24, with a gap of 1.29. Although they are somewhat invisible externally, our management staff is very active internally in customer support. Cathy Acret wrote a letter to all customers in August introducing herself as the new manager of the operating system Central Support group. Steve Gill manages the group supporting the application overlays such as WORDMATE, A\*L\*L, REALLINK, REALCALC and REALGRAPH. Mauro Mendoza manages the group that supports datacommunications and the 7000 product line. I manage the entire sustaining support effort for the MDCSC product line. Each of us can be reached by calling McDonnell Douglas Field

Service Company in Irvine, California at (714) 250-1000. As always, your local Field Service District Manager is there for your hardware problems and to collaborate with us in managing those problems involving both hardware and software.

Overall Software Service received a rating of 7.73, with a gap of 1.58. I expect that the improvements discussed above combined with our continuing campaign to achieve excellence in the support arena will improve this rating over time. I would like to thank our analysts for their dedication to the customer and a job well done.

#### ADDITIONAL STATISTICS

Eighteen percent of the respondents said they were interested in contracted After Hour's Software Support Service. Of the 18 percent that were interested, 55 percent said they would opt for weekend, 31 percent for evening, and 14 percent for graveyard coverages. We have recently added one person to the evening shift and one to the graveyard shift. Contracts and pricing for After Hour's Software Support are being written at this time.

Seventy-nine percent of the respondents said they were interested in an electronic bulletin board. We are currently developing a bulletin board which will be introduced first to customers of the 14/100 product who sign up for a software support contract (refer to the article on page 3). It will eventually be available to the entire user base.

Highlighted in the additional services/comments, the following list contains the most frequently stated ideas:

- \* Better documentation
- \* Automatic software and documentation updates
- \* More information on patches, workarounds, peripheral configurations, etc.
- \* Better communications with all areas of the company with action plans, escalation and followup.
- \* Better distribution of patches
- \* Bulletin Board System
- \* More and better classes and training
- \* Local on-site support
- \* Path to request enhancements
- \* Less expensive 24/7 software support
- \* Want specific analysts to talk to

Some of these requests are already being addressed, and the remainder will be evaluated and fine-tuned where applicable. You will be hearing more from us on these issues in upcoming newsletters.

**Your input is invaluable!** We thank you for taking time to respond. Of greatest importance is our constant striving to provide better support. We also know that there is room for improvement in specific areas. Armed with this information and these

measurements, we now have a more realistic target at which to aim.

- Christine Harding

## PRODUCT INTRODUCTION GROUP

Central Software Support has a sibling organization called the Product Introduction Department (PI). Created late in 1987, PI is the result of a restructuring of the Field Service Company's Product Support organization. PI's function is to focus on the importance of the early portion of a product's life cycle in the success of that product.

Fourteen staff, headed by manager Doreen Neck, are involved from the very earliest stages of development of new products. They provide product design input from a field service viewpoint. Issues such as hardware maintainability and support tools for software are introduced into the review process from the beginning.

As prototype products are produced, whether hardware or software, PI assists in quality assurance testing. As draft documentation is written, PI reviews it for accuracy and completeness. They also produce or coordinate the internal documentation needed by Field Service to support the product.

PI assists in the selection and qualification of alpha and beta test sites for new products. Internal to McDonnell Douglas, as an alpha site, the company's manufacturing plant was running on the 6.0 release of the Reality operating system on the new Series 18 hardware before either were installed on a customer site.

Beta sites are carefully selected customers willing to test a new product before it is production released. PI determines that beta candidates are well-suited to the rigors of beta testing. A desirable beta customer has good technical knowledge, will use new features of the product, and is willing to provide feedback to the development team.

Once a product is installed at a beta site, PI is the primary focus of support for the site. It coordinates other company resources, such as the engineering department, in support of beta customers. PI staff travel to the beta sites for installation and problem analysis when necessary.

In a product's life cycle, the final phase of PI's involvement is the transition of support to the Product Support team. PI staff coordinates training of Central Support and other staff. Both groups work closely together in supporting customer sites until Product Support is ready to assume full support and the product is officially released.

Creation of the Product Introduction Department represents the company's recognition that support of

customers begins long before a product is ever delivered. Its existence demonstrates commitment to producing the highest quality product possible by McDonnell Douglas.

• Doreen Neck

## MARKETING

### Welcome REALISM!

REALISM, (REALITY Integrated System Management) is an exciting new user interface for the REALITY operating system. Operating on the Series 6000, Series 9000, and Series 18 lines of hardware products, REALISM can be easily retrofitted to existing systems.

Designed to allow users full use of the features available in the REALITY Operating System, REALISM provides a layer of sophisticated and practical user functions across the entire operating system.

In addition, REALISM provides flexible and effective security against unauthorized access at both a hardware and software level. Hardware access can be restricted to individual ports and the time of day ports can be used. For software, access can be restricted down to the field level. Accounts, processes, files and TCL all run under the REALISM security system.

The effect of this new interface is that even if you are a new or non-technical user, you can quickly become confident in using all the facilities available. An experienced user can take advantage of an easy-to-use menu interface, particularly for infrequently performed tasks.

Because of this security, when logging on to a system with REALISM, the first screen displayed is the security screen:

- Only when a valid user ID and Password are entered will a user have access to REALISM menus.
- When troubleshooting a system with REALISM installed, the user must be given access to a valid user ID and Password.
- It is also important that the user ID being used has been given access to Verbs, PROCS and other processes which are required for trouble-shooting.

REALISM is available as an option on Series 6000 2.3, Series 9000 5.3, and Series 18 6.0 systems.

REALISM provides a number of tools and applications as standard:

**Office Automation:** This includes a sophisticated electronic mail facility to send or distribute mail. It also includes a Personal Diary, Personal Notepad, Address Book, Calculator, and Front Desk System. This includes a Paragraph Builder which enables you to build ENGLISH sentences by prompting each element in the sentence. This also includes Quick Data Entry, XTCL, and a Full Screen Editor.

**Toolkit:** This allows you to run any processes during non-peak periods of the day. There is a general batch queue and in addition, a user's personal queue.

**Night Moves:** Extensive online documentation, which can be viewed on the screen or printed. In either case, the user only receives the documentation pertaining to those facilities to which they have been given access. There are also a number of facilities that assist the user to write their own documentation for their own applications.

• Helen Scott

### SERIES 14/100

McDonnell Douglas' Series 14/100 system vastly expands the capability of a single user PC/AT by increasing its functionality to that of a multi-user, multi-functional data processing center. The 14/100 system is ideal for managing an entire department or a small business while providing an optimal growth path for expansion to larger McDonnell Douglas systems.

The Series 14/100 system is a low-cost, value-added enhancement to the IBM PC AT architecture that allows the user to take advantage of existing MS-DOS software and applications in addition to the benefits and performance of REALITY applications all on a single hardware system.

Consisting of a Motorola MC68020 based co-processor board with 1 MB of on-board memory, the 14/100 system turns a user's IBM PC AT (or 100% compatible) into a multi-function, multi-user data processing center capable of running the thousands of business software applications written for the REALITY Operating System.

Series 14/100 runs the REALITY 2.3 Operating System, the same software provided with the Series 6000 family, and supports up to eight terminal connections.

Enabling the PC user to switch between MS-DOS and REALITY applications with a single key stroke, the 14/100 bridge software provides users with the ability to use MS-DOS data with REALITY applications and REALITY data with MS-DOS applications.

With the 14/100 system, McDonnell Douglas provides the customer who needs a small computer system the following:

- A vertical solution targeted to a particular business. With over 3000 distinct applications, REALITY offers one of the most diverse environments from which to choose software solutions.
- A smooth upgrade path as the customer's business grows. McDonnell Douglas offers systems supporting from one to 400 simultaneous users with application and data portability from the smallest to the largest REALITY system. As growth to a larger system occurs, the user's investment in the PC is protected since we offer many different ways of connecting and integrating the PC equipment with larger REALITY systems.
- A smooth integration with existing PC applications. Only a 14/100 system from McDonnell Douglas offers the ability to execute any MS-DOS program from the PC console while up to eight users are simultaneously executing a multi-user REALITY application from remote terminals.

The Series 14/100 system is offered with the following configuration:

- A REALITY co-processor card with a 16 Mhz MC68020 microprocessor and one megabyte of on-board memory.
- One set of floppy disks containing the REALITY Operating system, MS-DOS bridge software and diagnostics.
- One set of documentation including: *A User's Guide* and an *Installation Guide*.

In addition, the Series 14/100 supports the following user supplied hardware:

- One 8-port terminal controller: Digiboard COM/8S or STARGATE +8
- One 60 MB 1/4" cartridge tape drive: Everex 1/4" Internal or external drive with D05 or D05D controller or Wangtec 1/4" Internal drive.
- One Uninterruptable Power Supply.

The Series 14/100 system is currently undergoing its final beta testing phase. You can find out more on the Series 14/100 system by contacting your local distribution channel.

- John Seaman and Jeff Brown

## ROS

### Dynamic vs. Dimensioned Arrays

DATA/BASIC allows you to manipulate variables in the form of dynamic arrays or as individual array elements. This topic discusses the tradeoffs involved in each format. (The authors assume you understand the REALITY file structure and have done some programming in DATA/BASIC.)

**Dynamic arrays** are stored in the following way:

**ABC^DEF]GHI]JKL^MNO\PQR]V2]SV1\SV2]V4^YZ**

**Dimensioned arrays** are stored as follows:

**1 ABC**

**2 DEF]GHI]JKL**

**3 MNO\PQR]V2]SV1\SV2]V4**

**4 YZ**

Dynamic arrays are primarily useful in interfacing to REALITY file items. By specifying a single variable name, you can read or write an entire item, or access individual items easily.

There are several functions available which scan dynamic arrays, looking for individual elements. This can sometimes be inefficient, depending on the application.

For instance, to extract attribute 10 from an item, DATA/BASIC first scans over attributes 1 through 9. When replacing data in an item, DATA/BASIC passes by all the data preceding the replacement, copies in the new value and then appends all the data that follows the attribute. This continual scanning becomes inefficient when you are accessing a large number of elements or processing a large item.

In such cases, it is better to read the item into a dimensioned array (using the MATREAD statement). MATREAD places each attribute of the item into separate variable locations that can be accessed individually, without scanning the item. This is particularly useful if several attributes need to be changed, because you can modify attributes without moving the rest of the item. The MATBUILD statement builds a string variable from a dimensioned array, while the MATPARSE statement assigns the elements of a string variable to the variables of a dimensioned array.

You can still use dynamic array references to extract, delete, insert and replace attributes, values and subvalues. You simply specify the array subscript.

For example, if the item is stored as a dynamic array, the expression ITEM<3,3,2> accesses the fourth subvalue in the second value in the third attribute. On the other hand, if the item is stored as a dimensioned array, the expression ITEM(3)<1,3,2> accesses the same element.

Reviewing both arrays shown above, this element has a value of SV2.

This method has its disadvantages too. Because there is a separate descriptor for each cell in the array, it uses more variable descriptor space. It also uses more freespace, because each attribute may have its own buffer in freespace.

In addition, items must be disassembled for a MATREAD and reassembled with a MATWRITE, which takes more time than a simple READ or WRITE.

The following rules should be used as guidelines in determining whether to use dynamic or dimensioned arrays:

1. Use dynamic arrays only when dealing with data read from or written to REALITY file items. Use standard dimensioned arrays when the need for an internal table arises.
2. If a value is to be used more than once, assign it to a variable instead of performing multiple extractions. Also, if the application is to reference many values or sub-values in an attribute, assign the attribute to a variable and perform the extractions and replacements on this smaller string (using an attribute value of one) rather than continually scanning the entire item.
3. Use dynamic arrays to extract several attributes from the beginning of an item or to replace four or five values in a large item. However, use a dimensioned array to construct new items or to access several different attributes.
4. Use dimensioned arrays only where necessary, because the object code required to reference them is greater than for single variables. However, avoid doubling the source code simply to avoid using subscripted variables.

The following table and programming example were prepared by Dick Vahlstrom, a Senior Scientist in our Software Engineering Department. Dick has been studying performance issues over the past few months and has gathered this data on a 9250.

FUNCTION	EXPRESSION	ARRAY ELEMENT NUMBER (L)				
		1	5	10	50	100
NON-ARRAY	B=D	1.0	1.0	1.0	1.0	1.0
DIMENSIONED	C(L)=B	3.2	3.2	3.2	4.2	4.2
ARRAY	B=C(L)	3.2	3.2	3.2	4.2	4.2
DYNAMIC	B=A<L>	4.6	5.2	5.4	11.2	18.4
ARRAY	A<L>=B	56.0	57.0	57.0	63.0	69.0

These are normalized, approximate times for simple operations using dimensioned and dynamic arrays. These arrays have 121 attributes.

The following is an example of a small programming change that can increase the performance of a particular piece of code by a factor of ten!

This code takes 156 milliseconds to run on a 9250:

```
B=""  
FOR X=1 TO 80  
B<X>=A<X>  
NEXT X
```

This code performs the same function and takes only 15 milliseconds:

```
DIM AA(80)  
MATPARSE AA FROM A  
B=""  
MATBUILD B FROM AA
```

Please keep these tips in mind when you are developing new applications, or when you are refining your existing applications. As you can see, different programming techniques can have a dramatic effect on your system performance.

• Cathy Acret and Dick Vahlstrom

## Current OS Releases And Patches

The following table contains the most current operating system (OS) revisions and patch levels for each supported system. If you do not have the current patch tape installed for your particular operating system contact the following:

For Independent Sales Organization (ISO) Customers, please contact your dealer.

For Branch Customers, please contact your field engineer through Central Dispatch to schedule a time for installation.

Series	Release	Patches (PP = Paper Patches)
4700	4.2RevD	RevB Tape, PP2
6000	2.3RevD 1.1RevD	RevB Tape (Includes PP1-68) RevB Tape
9000	5.3RevD 1.3RevC	RevB Tape (Includes PP1-65) RevA Tape
18	6.0RevF	RevA Tape (Includes PP1-47)

## PCI on Series 18

Correctly attaching a device to a system requires an understanding of baud rate and PCI. **Baud rate** is the number of pulses within a specified time period that can be transmitted over a line. In most systems, the transmission rate measured in baud is the same as that measured in bits per second (bps). The term **PCI** stands for **Programmable Communications Interface**. By using one of several data bit pattern lengths provided by PCI, a system becomes more flexible in its ability to communicate with other devices, even other systems.

After upgrading from a 9000 system to a Series 18, you may encounter a situation during which a serial port I/O device (e.g., a CRT or port printer) may not function properly. This may be traced back to the differences in how each type of system transfers data among devices.

For example, if a working device on the 9000 with a PCI setting of 74 (7-bit, no parity) fails to function when connected to the Series 18, it is most likely due to how the Series 18 interprets the bit pattern which it receives from the device.

The 9000 handles data in a rather unique manner when set up for 7-bit, no parity. This bit pattern represents a PCI of 74. (For additional information on PCI settings and their bit pattern representations, please see your Programmer's Reference Manual.) You may have used a PCI setting of 78 (8-bit, no parity) on some of your 9000 ports, especially, if you were using the 5.3 Operating System. Generally, however, it did not matter if your device was set for 8-bit, no parity while the port was set for 7-bit, no parity.

The reason for this is simple. When the 9000 transmits or receives data through a port set at 7-bit, no parity, it will set the 8th bit false. Although the system is transmitting and receiving 8-bit data words, it is essentially ignoring the high-order (8th) bit.

The Series 18, however, transfers data in a more conventional fashion: if you use 7-bit data, then set up for 7-bit data flow control (using PCI = 74); if you use 8-bit data, then set up for 8-bit data flow control (using PCI = 78). Consequently, both the device and the system must be set up with matching data bit lengths in order to function properly.

Consider, for example, a situation using XON/XOFF flow control. The hex equivalent\* used by the system for XON and XOFF is X'11' and X'13', respectively. However, the connected device may be set up using a hex X'91' and X'93' for XON and XOFF. Since the high-order bit is no longer ignored, flow control must be specified exactly (i.e., X'91' will not be interpreted as X'11' with the high-order bit (=1) being ignored).

In summary, you should match the data bit length among the Series 18 (PCI setting) and its attached devices in order to insure that correct data transmission will occur.

\* The XON decimal equivalent = 17, its hexadecimal equivalent = X'11'; represented as a Control-Q. Its binary equivalent = 0001 0001, while the X'91' = 1001 0001 (with high-order bit set to 1). Please refer to your Programmer's Reference Manual for more information on PCI.

• Gary Moote

## APPS

### WORDMATE 2.1C, RevA Patch Tape

The WORDMATE 2.1C Overload Rev. A patch tape has been production released. If you are running WORDMATE 2.1C, you should make arrangements to have the patch tape installed as follows:

If you are an Independent Sales Organization (ISO) customer, please contact your dealer.

If you are a Branch customer, please contact your field engineer through Central Dispatch to schedule installation during the next site visit.

### Identifying A\*L\*L Files

During the AF-RESTORE process, it is sometimes desirable to resize certain files that have grown unmanageable. Some of the files used by A\*L\*L can fall into this predicament. You will find listed below the files used by A\*L\*L, with a brief description, and whether they should be resized. If you are developing A\*L\*L functions, then you should compare weekly file statistics listings to determine the necessity of resizing any growing A\*L\*L file.

CBSPEC	- This file contains the control break information used in report and update functions.
DB.XREF	- This file contains a symbol table for each function that has been genned and is used by the A*L*L Debugger. This file frequently needs to be resized.
FDESC	- This file contains the description of each function defined through Function Definition.
FILEDEF	- This file contains the definitions for each field in a file defined through File Definition. This file sometimes needs to be resized.
FILEHDR	- This file contains the header information for each file defined through File Definition. This file sometimes needs to be resized.
FILESPEC	- This file contains the file access information for each file defined in a logical screen, report, or update.
FLDSPEC	- This file contains the definitions for each field defined in a screen and report function, as defined through Fields Definition.

FUNCTDEF	- This file contains the function definition information for each function. This file sometimes needs to be resized.
GENFILE	- This file contains the parameters necessary for the function to be executed. This file frequently needs to be resized.
HDR	- This file contains the information for menu, screen, and report formats. This file sometimes needs to be resized.
HDRIMAGE	- This file contains the report header layout for each report function.
LOGIC	- This file contains the logic statements for each logic routine in a function. This file frequently needs to be resized.
LOGONF	- This file contains the security information necessary to access A*L*L.
LSDEF	- This file contains the logical entity information for each logical screen, report, or update.
MENUDEF	- This file contains the names of the functions or programs to be executed from the menu and type.
MSGF	- This file contains the system help messages, system error messages, and user-defined messages. This file frequently needs to be resized.
SCRIMAGE	- This file contains the parameters necessary to print the screen and menu image in the function documentation for screen and menu functions.
SECURITY	- This file contains the definition of each category that an Operator ID can access.
TERMMAST	- This file contains the security information for terminal.
VARDEF	- This file is the Global Dictionary and contains file variables, company-dependent value variables, system-dependent value variables, and user-defined global display codes. This file frequently needs to be resized.

• Bryan Glassick

### Mandatory Upgrade

REALITY Integrated Office 2.3 Rev A has been production released and is available for customer shipping. The previous release, 2.1, will no longer be supported. Contact your local sales organization to order this upgrade at no cost.

### REALLINK and PS/2 50Z

REALLINK 2.0 will not install on the IBM PS/2 50Z personal computer. This has to do with the copy protection scheme used by REALLINK. Also, IBM has installed a different disk controller on this PC model. Our Engineering staff has assigned top priority to creating a patch to resolve this problem. You will be notified when this patch is available.

• Bryan Glassick

### Application Overlays

The following matrix provides you with the release level of Application Overlays required by each supported Series and OS. It is important that you know which Overlay you should obtain prior to a planned upgrade. For example, if you are upgrading a Series 9000 from 5.1 to 5.3 which uses REALCALC, then you will need to obtain the corresponding Overlay release (REALCALC 2.1C) before upgrading.

APPLICATION OVERLAY	SERIES 4700	SERIES 6000	SERIES 9000	SERIES 18
A*L*L 1.1	4.2,4.3	1.1	1.3	N/A
A*L*L 1.2	N/A	2.3	5.3	6.0
PCmicroREALITY 2.0B	4.2	1.1	N/A	N/A
PCmicroREALITY 2.1	4.3	2.3	1.3,5.3	6.0
REALCALC 2.1	4.2	N/A	N/A	N/A
REALCALC 2.1C	4.3	1.1,2.3	1.3,5.3	6.0
REALGRAPH 1.0	4.2	N/A	N/A	N/A
REALGRAPH 1.0C	4.3	1.1,2.3	1.3,5.3	6.0
REALISM	N/A	2.3	5.3	6.0
REALLINK 1.0B	4.2	N/A	N/A	N/A
REALLINK 2.0	4.3	1.1,2.3	1.3,5.3	6.0
REALITY Integrated Office 2.3	4.3	1.1,2.3	1.3,5.3	6.0
WORDLINK 1.4	N/A	1.1	N/A	N/A
WORDLINK 1.4C	N/A	2.3	1.3,5.3	6.0
WORDMATE 1.4D	4.2	N/A	N/A	N/A
WORDMATE 2.1C	4.3	1.1,2.3	1.3,5.3	6.0
TRANSACTION LOGGING 1.2	N/A	2.3	5.3	6.0



• Janet Altman

# PROC DOC



## Keeping ENGLISH in the primary output buffer

Have you ever written a PROC that 'ought to' work, but had it terminate with the error...

### [2] UNEVEN NUMBER OF SINGLE OR DOUBLE QUOTE-SIGNS (' "').

Dollars to donuts you did a SELECT, SSELECT, SEARCH, ESEARCH, GET-LIST or FORM-LIST in the primary output buffer and then loaded an ENGLISH sentence in the secondary output buffer. For example...

```
001 PQN
002 HGET-LIST INVOICES
003 STON
004 HSORT filename WITH criteria BY sort.parameters
005 H HEADING "lots of stuff... . . ."
006 P
```

The problem you encountered is that the secondary output buffer (referenced by turning the 'stack on' with STON) is limited to 140 characters. If you put a large enough ENGLISH sentence there, it will be truncated at 140 characters. \*(See note below.) If this happens in the middle of the HEADING, you'll be missing the closing quotes and will get ERRMSG 2.

There is a very simple alternative. The size of the primary output buffer is limited only by the remaining disc you have! Inserting just two statements will allow you to keep your ENGLISH sentence in the primary output buffer. Observe...

```
001 PQN
002 HGET-LIST INVOICES
003 STON
004 H<
005 P
006 HSORT filename WITH criteria BY sort.parameters
007 H HEADING "lots of stuff... . . ."
008 P
```

The effect of inserting lines 004 and 005 is to execute a carriage return after the GET-LIST and before the SORT. The less-than sign (<) on line 004 is the convention for getting a <RETURN> into the output buffer. This 'slight of hand' merely takes advantage of the fact that a <RETURN> executed at the select list prompt (>) has no effect. It does, however, allow you to finagle this finesse!

\* The PROC Programming Manual recommends using the stack continuation characters (<<) if your command is more than 140 characters. It will terminate execution if no items are found in the primary output buffer.

## SERIES 7000

The following suggestions are offered for your consideration when designing and writing BASIC and KeyBASIC programs:

- Ensure that programs are easily maintainable, clear and concise.
- Use meaningful variable-names. Write comment statements at the beginning of the program to describe each variable.
- Write complex expressions which will reduce the number of statements required.

- Write only one statement per line. You gain nothing by combining statements and it makes the program harder to read.
- Keep all subroutines at the top of your program; this prevents directing execution to subsequent lines of code when using the GOSUB statement.
- Avoid the use of GOTO references, especially forward references.
- When using a conditional statement ensure that the 'normal' exit is not the branch. Always make the exception condition the branch.
- Don't use the ALIGN statement indiscriminately. One ALIGN statement before the main loop of the program and one before the common subroutines should be more than enough.

- If a piece of source code containing over 80 bytes of object code is to be executed more than once during a program, it is best to place the code in a subroutine and call it using GOSUB.
- Avoid using the CAT operator when building long strings. If long strings must be built, set up a variable at the beginning of the program long enough to accommodate the entire string. Use substrings to build up the string. The reason for this is when a variable is set up it is allocated enough memory to accommodate the original size. If you make the variable longer, using CAT, memory must be reorganized to accommodate the new variable size.

You will find additional coding recommendations in the next issue of *ON-LINE*. Your comments and suggestions are welcome.

• Helen James

## COMMS

### Multiplexers (Part 1)

Over the past few years, the rapid developments in the multiplexer market has made it almost impossible to provide a simple definition of a **multiplexer** (also called a "MUX"). All multiplexers perform a primary function: they combine the datastreams from many individual low-speed channels and transmit a combined stream of data over one high-speed communications link. Multiplexers maximize the efficient use of communications links in a network because users can lease one high-speed line for much less than it would cost to lease many low-speed lines. Thus, through multiplexing, users realize significant cost savings.

With the advent of intelligent terminals and distributed processing in the 1970s came the demand for more cost-effective, high-speed multiplexers, and this demand led to the development of **time division multiplexers** (TDMs), and later, **statistical time division multiplexers** (STDMs).

When conventional time division multiplexing is used, each inputting terminal takes its turn at transmitting and receiving data; the order in which the multiplexer serves the terminals is fixed. The device accepts only one bit, byte, or packet of data from each input line and puts it into a specifically allocated time slot on the high-speed terminal in the sequence.

Statistical time division multiplexers (STDMs) take this operating concept one step further: the STDM senses which terminals are currently active, and serves only those units. The STDM provides vari-

able-length, rather than fixed-length, time slots so that more than one byte of data can be accepted from a given terminal before the STDM moves on to the next active terminal. The basic difference, then, between a TDM and an STDM is that while the TDM solicits data from each input terminal in a fixed succession whether the terminal is active or idle, the STDM ignores idle terminals and accepts data only from active devices.

When the multiplexer accepts mixed speeds, the higher speed input lines are transmitted more frequently than the lower speed lines.

For example, if one line is three times as fast as another line, the multiplexer will accept three bytes in any sequence for each one taken from the lower of the two lines. A statistical multiplexer provides more flexibility in line configuration; the aggregate of the inputs can exceed the output as long as not all input channels act simultaneously.

If transmission on the high-speed side is over a conventional analog facility, the TDM must direct its output through modems on the line. If AT&T's Dataphone Digital Service (DDS) is the means of transmission, a data service unit, rather than a modem, is required to provide the bipolar conversion necessary to operate on the digital line.

While conventional TDMs cannot handle an aggregate input speed that exceeds the composite link speed, the STDM has a buffer which makes this possible. In fact, intelligent flow control is one of the most important features of STDMs, which have buffer space for both transmitted and received data. The buffer holds data that exceeds the aggregate input or output rate.

When an STDM is operating at peak activity, i.e., all of the inputting devices are sending simultaneous data at the maximum rate, the buffer will begin to fill up. To prevent data loss, the STDM can signal the transmitting computer or terminal to suspend transmission.

When storage is again available, the multiplexer signals the transmitting device to resume transmission. Flow control on STDMs is generally achieved through one of two methods: X-ON/X-OFF or by dropping **Clear To Send** (CTS). In the X-ON/X-OFF procedure, an X-ON code initiates data flow; X-OFF stops data flow. In the Clear To Send method, the signal is turned on to start data transmission and off to terminate it.

Dynamic buffer allocation, another STDM feature, permits buffer pool sharing among many channels. This feature improves buffer efficiency and allows the multiplexer to operate with an apparent efficiency that is several times greater than the peak channel data rate.

Statistical multiplexers provide a full range of diagnostics, including local and remote, digital and analog loopback tests, self-tests on power-up, port buffer over-flow statistics, and end-to-end diagnostic messages.

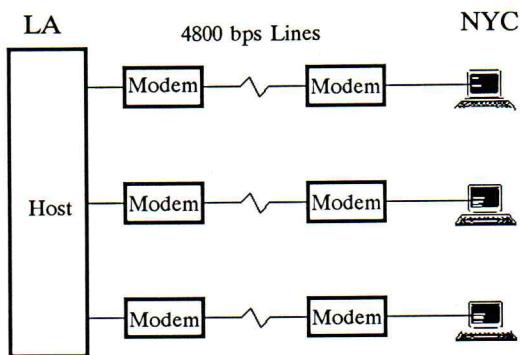
Local loopback tests are internal, testing the transmitting unit only. When this function is performed, the multiplexer is isolated from the data set and composite line, and its internal circuitry is checked for proper operation.

System loopback tests are external; they test the transmitting unit and all other parts of the network through which data passes. This test, performed after the local loopback test has been completed, loops composite link data back to the original multiplexer from a remote point. Channel tests check individual input/output channels, either end-to-end or remotely.

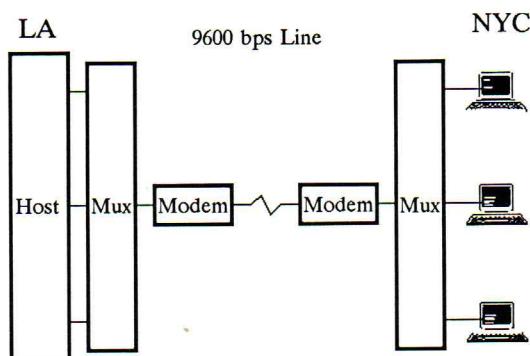
We hope this gives you the basic understanding of what a statistical multiplexer does and some of its features. In the following issues we will cover the specific models of MDCSC multiplexer products and some guiding principles of how to select the most appropriate parameters to optimise your network operations. An incorrect configuration can not only cause poor performance, but can also cause a perfectly good set of hardware and software to fail.

• Niki Jhaveri

#### ***Before statistical multiplexing***



#### ***After statistical multiplexing***



## **STAR POWER**

Consistent with the intent of this column to provide you with information which will assist you immediately in performing your job responsibilities more effectively, this issue's topic concerns insuring that AF-RESTOREs are successful.

During the past quarter, we have handled numerous calls on problems dealing with AF-RESTOREs. Many of these problems resulted from using defective tapes or not verifying the tapes before performing the AF-RESTORE. The usual recovery from an aborted AF-RESTORE requires restoring from a previous file-save. Then you must perform the tedious task of either re-entering all the data since that save, or selectively restoring specific files and data.

All of this unnecessary "hassle" can be greatly reduced if you check the file-save tape(s) you intend to use for the AF-RESTORE **before you restore**. There are two ways to do this, depending on your system and its OS level.

On the Series 6000 2.3, 9000 5.3, and Series 18 6.0 systems, you can use the **VERIFY-SAVE\*** utility which is fully documented in your Programmer's Reference Manual.

On other systems, you can use a "dummy" SEL-RESTORE of a nonexistent Account and File from your file-save. Although you are not restoring any data, you are still reading data from the entire file-save while searching for the "dummy" file. This process will stop with PARITY ERROR or other error messages if either the tape or its data is corrupted. Encountering this situation prior to doing an AF-RESTORE will allow you to produce a replacement file-save for your AF-RESTORE, avoiding recovery from an aborted AF-RESTORE.

\* If you use VERIFY-SAVE to check daily file-saves, be sure to do so during non-peak operational hours, since this process is system-intensive.

• L.W. Abel

#### **FOR THE RECORD**

\* PROC DOC article in Volume 1, Issue 2, page 6:

The following misprints appeared:

**MCL** should be **MCL** (no space)

**MCX** should be **MCXD**

\* STAR POWER article Volume 1, Issue 2, page 9:  
The following misprint appeared:

**.R/RECEIV/PAY 004 READ REC FROM PAY**  
**ABLE, ID THEN** should have read:  
**.R/RECEIV/PAY**  
**004 READ REC FROM PAYABLE, ID THEN**

## FEEDBACK



### McDonnell Douglas Computer Systems Co.:

I would like to thank the following people for their help and assistance in the resolution of the many problems the Hartford Housing Authority (HHA) has experienced over the past weeks. **Jim Smith, Teri Garrett and Mike Bingman** of McDonnell Douglas were extremely helpful and understanding while helping me work on the system. They took the time to explain how things worked. They all were polite, courteous and showed a very professional attitude.

Steve Lefcheck  
Manager Information Systems  
◦ HARTFORD HOUSING AUTHORITY

• The following comments are excerpts from the Software Service Survey conducted by Christine Hardig:

"I have always considered McDonnell Douglas Software Support to be excellent."

Kelly J. Kramer  
MIS Manager  
◦ BOBBI GEE

"In the past, software support was somewhat of a chore when a critical problem arose. Now, with Central Dispatch and your staff of very knowledgeable personnel, I feel very confident whenever a problem occurs. Also, whenever I am not in the office, I know someone in my department will be able to receive support through the "800" number.

Les Gainous  
Data Processing Manager  
◦ AREO HARDWARE & SUPPLY

"We are very pleased with McDonnell Douglas services -- Engineers and Software Support."

Elaine Lillard  
Assistant Office Manager  
◦ NORTHERN KENTUCKY A.D.D.

"The attitude of the Analysts is very good. I feel like I'm part of the team now. Centralized Software Support is great! Keep up the good work."

Jack W. Armstrong  
Manager of Data Processing  
◦ McCREEARY TIRE & RUBBER CO.



(L-R) Central Support Analysts Mike Bingman, Jim Smith, Teri Garrett

## Goofies



"HIT ANY KEY TO CONTINUE"

## IN QUEUE

In our next issue, Niki Jhaveri will continue his discussion about Multiplexers, emphasizing their impact on your daily operations. We will also feature topics on the effective use of KeyBASIC, improving system performance, and discovering several helpful system utilities. The PROC DOC will review using Indirect References. And, as always, we hope to receive your comments, questions, and suggestions for future articles.

## MCDONNELL DOUGLAS CUSTOMER EDUCATION SCHEDULE

COURSES OFFERED	NOVEMBER				DECEMBER			
	7	14	21	28	5	12	19	26
INTRODUCTION TO REALITY OPERATING SYSTEM \$675/Person				NJ IR	DC			
DATA/BASIC PROGRAMMING \$800/Person	DC NJ				IR			
ADVANCED REALITY OPERATING SYSTEM \$850/Person	IR							
SYSTEM THEORY and TROUBLESHOOTING \$900/Person					DC			
ADVANCED DATA/BASIC PROGRAMMING \$850/Person		IR DC						

### LOCATION KEY

IR - Irvine, CA  
DC - Washington, D.C.  
NJ - Union, N.J.

For registration or additional information please contact:

BARBARA STEURER  
MANAGER OF  
CORPORATE EDUCATION  
4000 MacArthur Blvd.  
Newport Beach, CA 92660  
(714) 250-1000 Ext. 2099

McDonnell Douglass Field Service Co. publishes this document without warranty of any kind, including, but not limited to the implied warranties of merchantability and fitness for a particular purpose. McDonnell Douglas Field Service Co. shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this document. This disclaimer may not apply in certain states, provinces or countries.

REALITY®, A\*L\*L®, ENGLISH®, PCmicroREALITY®, REALCALC®, REALGRAPH®, REALLINK®, WORDMATE® are registered trademarks of McDonnell Douglas Computer Systems Company. DATA/BASIC™ and WORDLINK™ are trademarks of McDonnell Douglas Computer Systems Company. IBM® is a registered trademark of International Business Machines. MS-DOS® is a registered trademark of Microsoft Corporation.

### Central Support

### ON-LINE

Published for System Software Users

Published quarterly by McDonnell Douglas Field Service Company Central Support Department for users of McDonnell Douglas computer systems.

L.W. Abel, Editor  
Will Edwards, Technical Advisor  
Pat Dwight, Production Consultant  
Cheryl L. Bouwens, Production

Please write to the Editor at: Central Support,  
MDFSCO, 2351 McGaw Avenue, Irvine, CA. 92713.