Optical Disk Technology

CONTENTS

DEPARTMENTS

FEATURES

COVER:
Optical disk technology is a high-density storage solution that will grow and evolve as more users become aware of its potential. For a general overview of OD technology, turn to Fred Schuff's article on page 35. Cover illustration by David Kramer.

PUBLISHER

Bob Thomas

While reader response to my question, "Mainframe Conference — Is The Time Right?" (posed in the March 1989 issue) was overwhelmingly affirmative, several readers had thoughts on the subject worth pondering. One reader (and a contributing writer in this issue), Jim Johnson, Director of Hallmark Cards' Data Center, sent several observations I would like to share with you.

"From my perspective, a general or combined conference has less chance of succeeding than a focused conference. Each specialty (database, quality, help desk, capacity management, network, etc.) has a loyal following with individuals justifying one or two trips per year based on business interest.

Analyzing the potential audience:

Top MIS personnel prefer direct vendor contact, existing conferences (Nolan Norton, Index, Diebold, Gartner) and relating with peers.

Middle DP management attend GUIDE, SHARE and specialty groups such as Productivity for Systems Development by ACR.
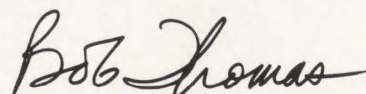
Technical DP personnel focus on SHARE or Computer Measurement Group (CMG) seminars.

Many conferences have organizational support. Examples include Systems Management, CMG, AFCOM, DPMA, etc. The conference proposed has no organizational support."

As evidenced from the strong response we received about the Mainframe Conference, it is clearly apparent that there is a strong need for more information and education among DP professionals. It is further apparent that many DP professionals are not aware of the various user groups, trade associations, organizations, conferences and seminars that already exist. For example, GUIDE International is one of the largest associations for users of IBM mainframes. However, following publication of the article titled, "GUIDE: One User's Perspective," by Pete Clark (MAINFRAME JOURNAL, January 1989), we were surprised by the numerous responses from readers not at all familiar with GUIDE who were seeking more information.

Since the mission of MAINFRAME JOURNAL is to inform and educate DP professionals using IBM mainframe (and compatible) systems, we are going to make a concerted effort to inform you about many of the relevant organizations, trade associations, user groups, conferences and seminars. In fact, we will be attending and promoting several of the conferences starting with National CASEcon (New York, June 20-22) and the inaugural conferences for both the International DB2 Users Group (Chicago, August 6-9) and the Disaster Recovery Symposium (Atlanta, September 11-13).

Is the time right for a Mainframe Conference — maybe not. Before the decision is made to proceed with plans, I think we should first explore the possibility that there already exist ample education and training resources that are unknown to many in our readership. Stay tuned; you will be informed!

Bob Thomas

Dyslexia and DP Professionals

I enjoyed the article on dyslexia ("How Dyslexia May Affect Data Processing Professionals" by Ted C. Keller, March 1989) and would encourage the human side of technology to show forth more often.

W. B. Ratcliff
Gulfstream Aerospace
Savannah, GA

The article on dyslexia was especially enlightening. Data processing is more than hardware or software; it is people.

Robert E. Jones
Southern Pacific
San Francisco, CA

Your article on dyslexia was wonderful. I was nine years old when it was determined that I, too, was dyslexic. In reading your article I found many of the problems Mr. Keller had in myself. While dyslexia is not curable, with work it can be overcome. My congratulations to MAINFRAME JOURNAL for producing a magazine that not only gives good, quality data processing articles, but also is bold enough not to limit itself exclusively to the computer.

Neil Benchell
Computer Associates
McLean, VA

I applaud your article on dyslexia. Working with computers for the last six years has helped my self-image 100 percent. More articles about this subject are necessary to help employers who have dyslexic employees and to give those of us who have dyslexia more confidence in our work environment.

Timothy Hearley
City of Albany
Albany, NY

COBOL Efficiency and Programming Dinosaurs

My husband and I enjoyed the article by Harvey Bookman on COBOL efficiency in the February 1989 issue very much. It was enormously refreshing to find that someone, besides us, believes that efficient code is a desirable objective.

Thanks for the tip on APPLY WRITE-ONLY. I don't believe I have ever seen it used. I knew about the odd number of digits in COMP-3 fields because I have, from time to time, looked at full PMAP to see what silliness COBOL was guilty of. In fact, I reduced run time for one program by 40 percent by changing a whole gang of S9(6) to S9(7) after I saw the arm-waving caused by the attempt to deal with the high-order nibble.

One other great timesaver is as follows: if you are doing random reads in an ISAM file, check to see if you already have the record you want *before* you issue the READ. ISAM is dumb as a post and will meekly trot off and go through his whole bag of tricks when he has already got the record you need. That simple change reduced a four-hour run time to two hours on a file as small as 40,000 records. I don't know if VSAM has the same problem but I will set up a test one of these days to find out.

My husband and I are both programming dinosaurs. He has been at it for 25 years and I for 27. We have both been informed, on multiple occasions, particularly when we were pointing out that some of the *structured* techniques resulted in egregious code, that efficiency is not an issue and that we are hopelessly backward and out-of-date. Doesn't anybody out there realize that the hardware vendors have a genuine vested interest in promoting inefficient software? Why on earth would any of them espouse techniques which will let you run in 65K when *state-of-the-art* ideas will require at least a megabyte?

Margaret C. DeVault
Dallas, TX

Critical Dataset Identifier (CDI)

Having been in data processing for seven years, I find the articles in MAINFRAME JOURNAL interesting and the technical advice useful. In 1988, our corporation conducted three tests at a disaster recovery site. During our planning sessions, we made major changes to the data restoration part of our disaster recovery plan. What follows is a brief summary of what is known to us as the Critical Dataset Identifier (CDI).

We developed and designed the CDI process to address the concept of data restoration at a disaster recovery processing site. Initially, our thoughts were to restore full-volume backups from the weekend followed by daily incremental backups to bring us forward to the point of the recovery. Unfortunately, this approach assumes that a matching configuration (single, dual and tri-density DASD) is installed at the recovery site. This not only proved to be a poor assumption, but also the data restoration process would take many hours to complete. With assistance from an outside consulting firm, we have developed, designed and implemented a process which identifies datasets essential for the recovery of our batch application systems. Using our DASD management system, we execute separate backups of these files on a daily basis. Critical tape datasets are verified against the tape vaulting list to ensure movement to an off-site vault.

In addition to substantially reducing the time required to restore files, this process also satisfies an outstanding audit requirement. This requirement is to ensure that if an application flow was changed, a file was renamed or a new application implemented, the proper recovery file is identified and entered in the appropriate table for movement off-site. In the past, we were at the mercy of an application programmer to supply the correct information to the responsible party and then pray that, once in production, nothing ever changed. Our process, CDI, which executes daily in our environment, generates an exception report identifying critical datasets which are not found in the Critical Master File table.

We have been successful in using this process to restore our critical files at an alternate processing site during our third recovery test of 1988. If you wish, I would like to share with your readers the detailed process and experiences we have encountered.
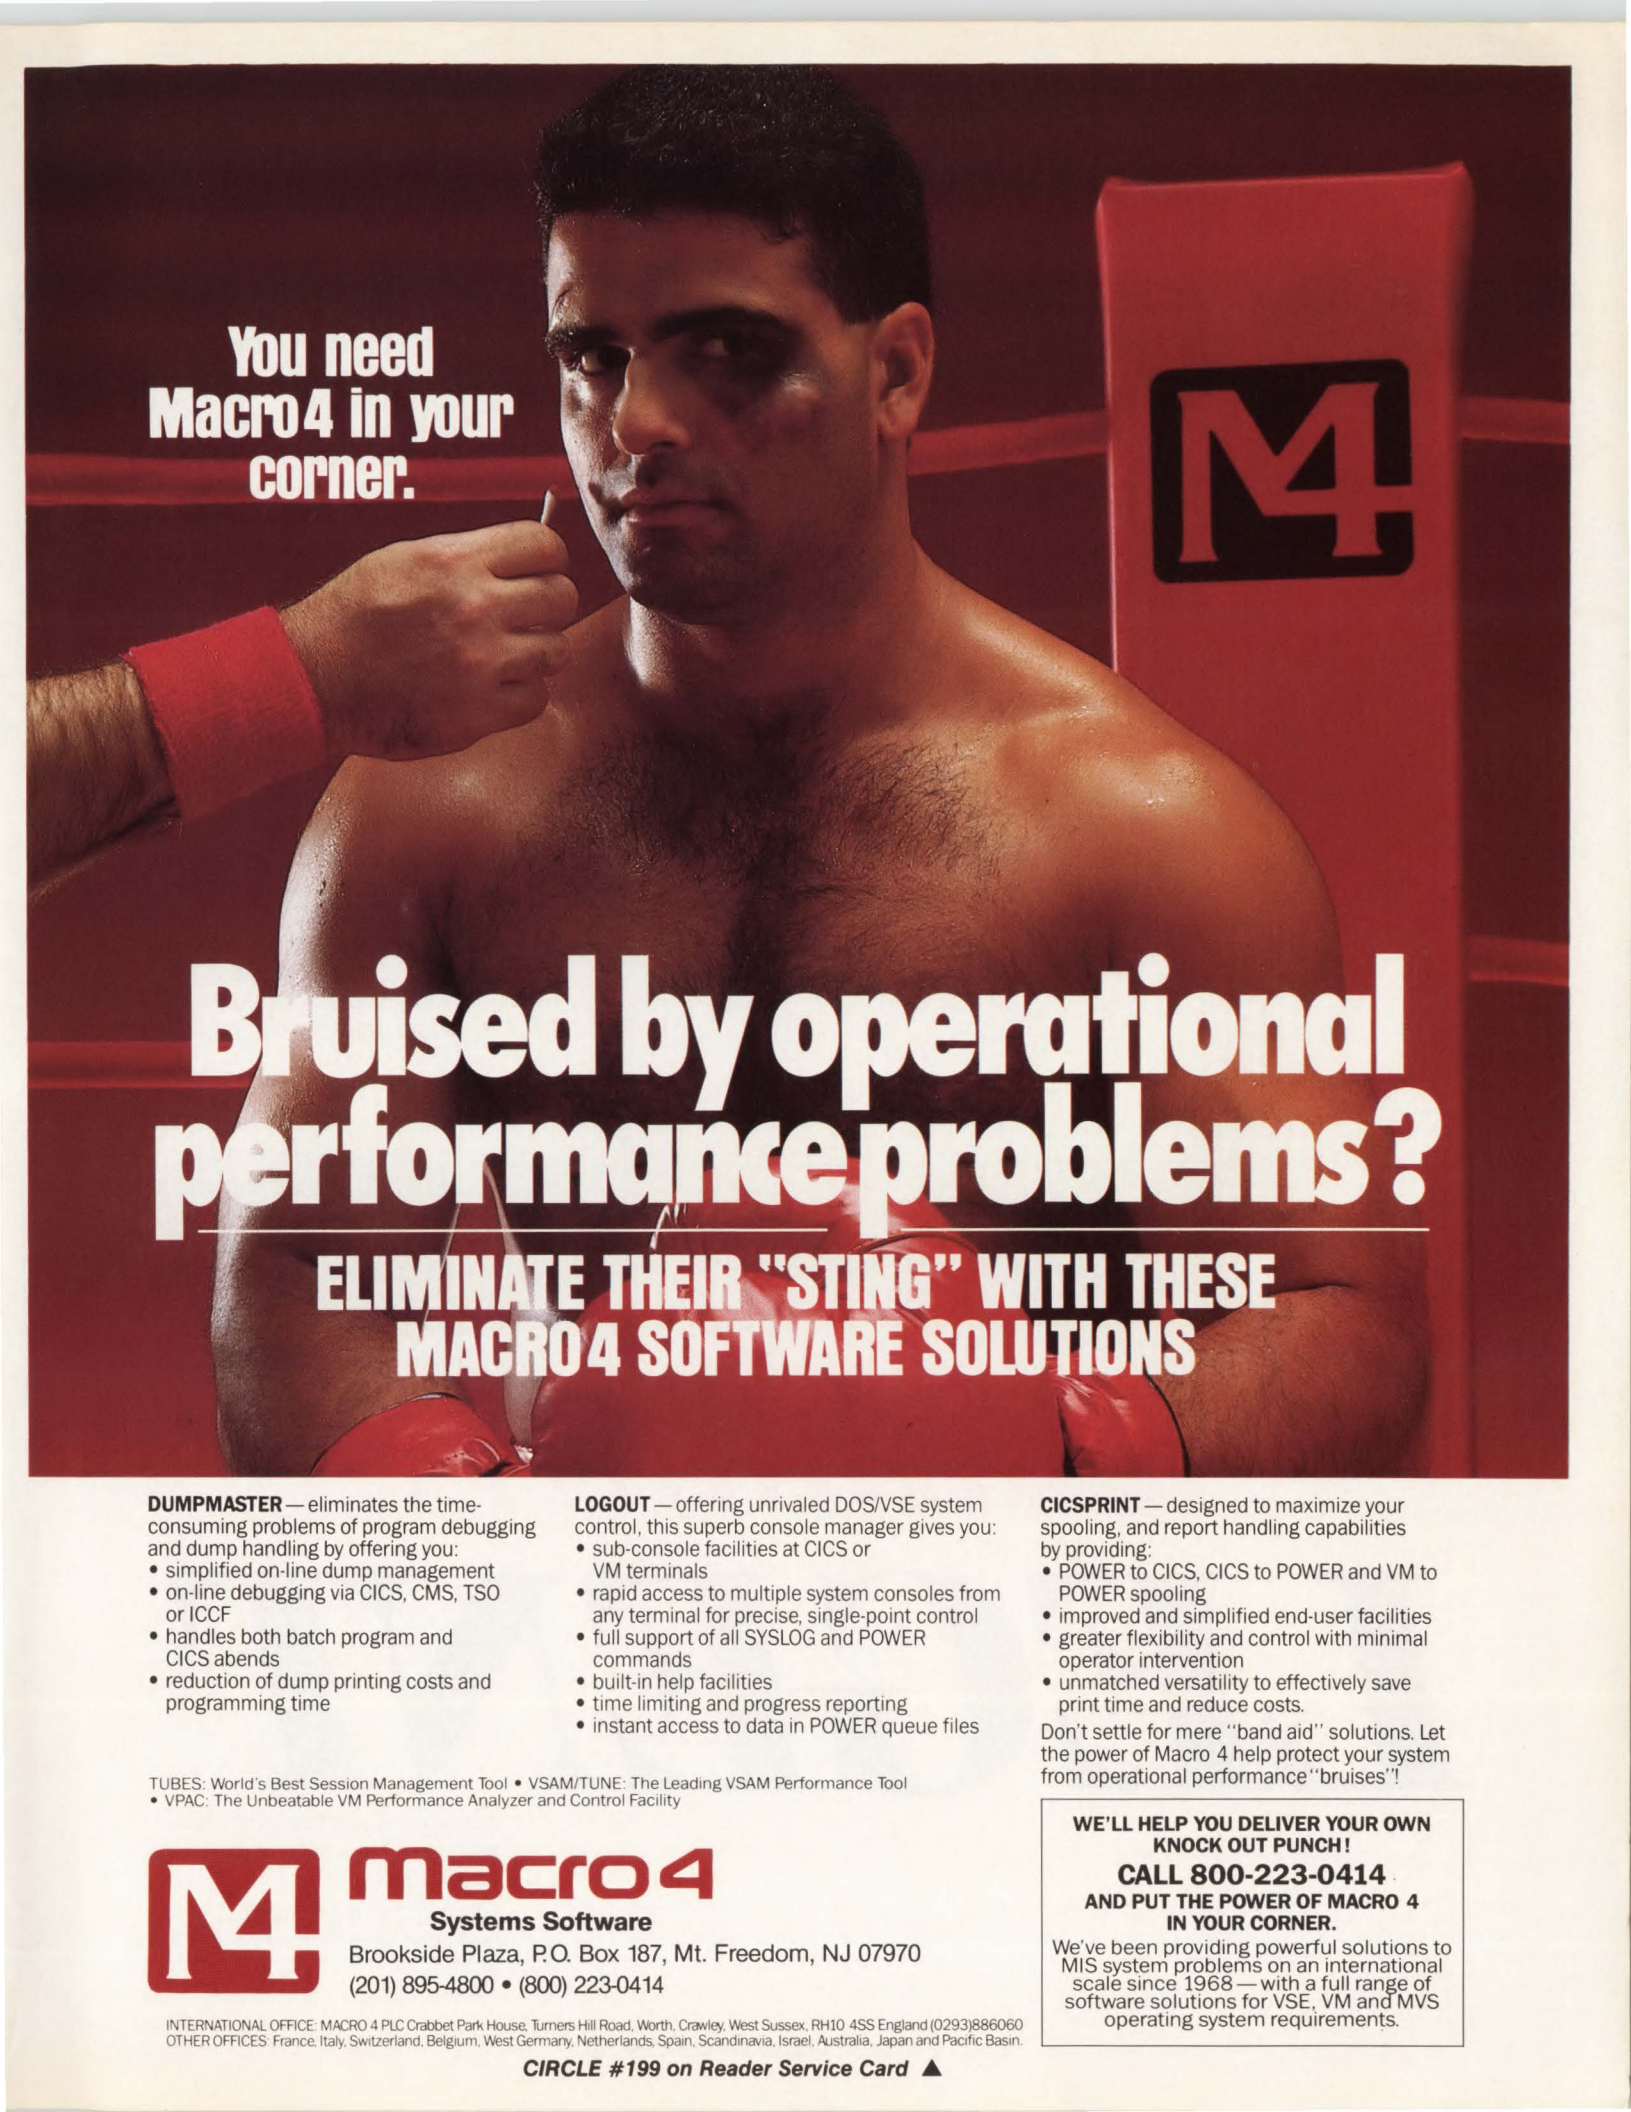
Patrick S. Dolan
National Liberty Corp.
Valley Forge, PA

Editor's Note: We have taken Mr. Dolan up on his offer to provide a more detailed explanation of his experiences in a future article: look for it in our September issue which will focus on Disaster Recovery Planning.

Correction

CA-SPOOLMAN, a spool management facility that operates in the VM environment, is a Computer Associates product. It was incorrectly referred to as being from Compiler Associates on page 90 of the July/August 1988 issue. ◈

You need Macro4 in your corner.

Bruised by operational performance problems?

ELIMINATE THEIR "STING" WITH THESE MACRO4 SOFTWARE SOLUTIONS

DUMPMASTER — eliminates the time-consuming problems of program debugging and dump handling by offering you:
- simplified on-line dump management
- on-line debugging via CICS, CMS, TSO or ICCF
- handles both batch program and CICS abends
- reduction of dump printing costs and programming time

LOGOUT — offering unrivaled DOS/VSE system control, this superb console manager gives you:
- sub-console facilities at CICS or VM terminals
- rapid access to multiple system consoles from any terminal for precise, single-point control
- full support of all SYSLOG and POWER commands
- built-in help facilities
- time limiting and progress reporting
- instant access to data in POWER queue files

CICSPRINT — designed to maximize your spooling, and report handling capabilities by providing:
- POWER to CICS, CICS to POWER and VM to POWER spooling
- improved and simplified end-user facilities
- greater flexibility and control with minimal operator intervention
- unmatched versatility to effectively save print time and reduce costs.

Don't settle for mere "band aid" solutions. Let the power of Macro 4 help protect your system from operational performance "bruises"!

TUBES: World's Best Session Management Tool • VSAM/TUNE: The Leading VSAM Performance Tool
• VPAC: The Unbeatable VM Performance Analyzer and Control Facility

macro4
Systems Software
Brookside Plaza, P.O. Box 187, Mt. Freedom, NJ 07970
(201) 895-4800 • (800) 223-0414

INTERNATIONAL OFFICE: MACRO 4 PLC Crabbet Park House, Turners Hill Road, Worth, Crawley, West Sussex, RH10 4SS England (0293)886060
OTHER OFFICES: France, Italy, Switzerland, Belgium, West Germany, Netherlands, Spain, Scandinavia, Israel, Australia, Japan and Pacific Basin.

WE'LL HELP YOU DELIVER YOUR OWN KNOCK OUT PUNCH!
CALL 800-223-0414
AND PUT THE POWER OF MACRO 4 IN YOUR CORNER.
We've been providing powerful solutions to MIS system problems on an international scale since 1968 — with a full range of software solutions for VSE, VM and MVS operating system requirements.

CIRCLE #199 on Reader Service Card ▲

VM Software, Inc. Expands Focus and Becomes Systems Center, Inc.

Even though VM software products continue to be its roots and expertise, VM Software, Inc. has shed its VM moniker in favor of Systems Center, Inc. By adding such non-VM products as: The Network DataMover products, network administration products and relational database management products, the name change should alleviate confusion. According to Robert E. Cook, Chairman and Chief Executive Officer, "The VM Software name no longer fits the character of our business. Systems Center, Inc. is an excellent descriptive umbrella for our business."

Duquesne + Morino = LEGENT Corporation

The merger of Duquesne Systems, Inc. and Morino Associates, Inc., announced last December, has been completed and the result is LEGENT Corporation (Pittsburgh, PA). With more than $100 million in annual revenue and approximately 600 employees, LEGENT becomes the third largest independent IBM system software vendor. Mario M. Morino is now Chairman of LEGENT and Glen F. Chatfield is Chief Executive Officer. The new company offers a line of approximately 40 products that represent a range of performance, networking, optimization and management support systems. LEGENT will have 12 sales offices in the United States and will market through subsidiaries as well as agents and representatives in two dozen nations. Combining the technological resources, financial strength and quality customer support of two of the industry's fastest growing and most profitable companies will make LEGENT Corporation formidable.

BlueLine Software Continues to Build VM Product Line

BlueLine Software, Inc. (Minneapolis, MN) recently acquired ownership of VLOCK from the product's original developer, Thomas Ericsson. This comes right on the heels of BlueLine's acquisition of the RD/LABS product line (RD/SHARE, RD/CHANGE and RD/COMM) in October 1988. With the addition of VLOCK, BlueLine now provides an extensive collection of VM tools for making both end users and the technical staff more productive. Bill Cecchi, BlueLine's President, comments, "As IBM's Systems Application Architecture (SAA) strategy unfolds for mainframe sites running the VM operating system, BlueLine's goal is to be positioned as a single source for selected tools aimed at applications development, systems management and end-user productivity in a VM environment."

3090 Mainframe Storage Prices Slashed

EMC Corporation (Hopkinton, MA) and Cambex Corporation (Waltham, MA) have announced price reductions for 3090 mainframe storage. EMC Corporation announced a $60,000 reduction on the price of its 64MB expanded storage upgrades for IBM 3090 systems. David Guy, EMC's Vice President of Marketing says, "EMC expanded storage upgrades in 64MB increments were priced at $157,000. Now users can buy the same upgrade for $97,500 and because this pricing is based on each side of the CPU, IBM 3090 users with double-sided machines such as models 280E, 400E, 500E and 600E have a potential savings of $120,000." EMC also provides 3090 central storage at $189,000 per 32MB. Cambex Corporation also announced price reductions for its STOR/9000 central storage and expanded storage memories for IBM 3090 mainframes. Cambex STOR/9000 has been reduced to $190,000 from $210,000 for a 32MB increment and the STOR/9000 expanded storage has been reduced to $145,000 from $175,000 for a 64MB increment.

Annual VM Workshop Scheduled for June 13-16

Kansas State University (Manhattan, KS) will host the annual VM Workshop June 13-16, 1989. Usually held on university campuses, the VM Workshop is an outstanding source of VM information, contacts, tools and camaraderie. The VM Workshop is an all-volunteer, self-supporting organization that provides a forum for individuals and organizations with an interest in VM to exchange information and experiences. The Workshop Tools Tape each attendee receives is almost worth the trip itself! More than 300 VMers attended last year's Workshop. For information and registration information call (913) 532-5575.

Control, Audit and Security of IBM Systems Conference Set

The Ninth Annual Conference on Control, Audit and Security of IBM Systems is scheduled to be held in Chicago, September 18-21, 1989. Keynote speaker Charles B. Wang, Chairman and Chief Executive Officer of Computer Associates International, Inc. (Garden City, NY), will speak on integrated solutions to security in today's complex business applications. The conference is designed to bring EDP auditors, internal auditors, security professionals and data processing professionals up-to-date on control, audit and security of the entire range of IBM mainframes using MVS to IBM PCs. The conference is sponsored by the MIS Training Institute. For information call Patsy Day (508) 879-7999.

Great *Myths* of Capacity Planning

By Jerry L. Rosenberg

Capacity planning, just like any other discipline, has developed a complete inventory of myths during its maturation process. As often happens, some of these myths have replaced intelligent, decision-making approaches and have, in fact, become the capacity planning methodology in some shops.

This article will attempt to expose some of these myths to the light of critical analysis and point out the deficiencies inherent in them. As part of the investigation, this article will also present the available pragmatic alternatives to these mistaken notions and illustrate how the capacity planner can be more successful.

Some of the myths that will be addressed are the following.

- We cannot perform capacity planning for a system or complex that is having performance problems.
- Capacity planning is done once a year.
- Capacity planning is modeling; modeling is capacity planning.

- Capacity planning and performance analysis are totally different.
- Capacity planning can be adequately performed with rules of thumb.
- Capacity planning is a technical exercise; therefore, I only need to consider machine-readable data.
- The technical solution should always be adopted.
- The newest piece of hardware is always the right answer or "the bigger is better syndrome."
- There exists a single measure to which all data can be distilled that answers all capacity questions or the "I only need MIPs complex."
- I can develop one report that captures everything that all levels of management and technical personnel require concerning resource utilization.

Background

All disciplines, on the way to maturity,

develop a mythology or underlying set of premises that are intended as a support for the basic structure of the discipline. This has been historically true for most forms of human activity (from religious to social to economic). The emerging mythology can be as complex as the pantheon of Greek or Roman gods or as simple as the financial axiom — "What's good for General Motors is good for America." In all cases, it would seem that myths are created to explain an otherwise inexplicable set of events.

In the case of the Greeks, they were looking for an explanation for the intricacies of the universe and nature. These were really complex issues requiring a detailed, structured set of rules and beliefs. In the General Motors example, the company was merely looking for a rationale to easily explain and justify some financial controls and political actions that might otherwise have been more difficult to implement.

You.
Westinghouse.
IBM.

In managing your data center, you can't afford to let just anyone come between your users and your IBM mainframe.

At Westinghouse, we help IBM mainframe users make the right connections to productivity. We offer a complete family of productivity enhancement software for MVS, VSE, and VM operating systems. Software that provides all the right features without sacrificing performance...from mainframe to user terminal. More power with less resources and overhead...no matter what the application.

We've been helping corporations manage their data centers for over 18 years with sensible, cost-effective solutions.

So if you're looking for a proven solution to your MIS productivity problems, give us a call. We're a name you can be sure of.

Westinghouse Software Solutions. Because in managing your data center, it's what's between you and your IBM that counts.

Westinghouse
Management Systems Software

P.O. Box 2728
Pittsburgh, PA 15230
(800) 348-3523
(412) 256-2900 in PA

In each of these extreme cases, as well as the whole spectrum of myths, there are several common ground rules.

Each myth generally has conditions associated with it. Unfortunately, over time, the conditions are forgotten and only the "rule" remains. This leads later converts to believe that the myth has universal application under all scenarios.

What makes this even more pernicious is that each myth has within it a grain of truth. If the myth is used under the right conditions, it works superbly. This enhances the belief of the adherents and leads to further universal application.

What should be done, but often is not, is the simple retesting of the myth under changing conditions to verify that it still applies. Once a myth has been accepted, it becomes almost sacrilegious to question it. Data processors, however, should be accustomed to the concept of regression testing. They are always testing old code under new operating systems to ensure compatibility before accepting a new product. So it should be safe to assume that they would not fall prey to myths. Right? Wrong!

Just like the Greeks and the Romans, capacity planners have developed myths to provide quick answers to difficult problems. Some of these myths have not been updated to conform with a changing technology while others have only marginal applicability in the first place. However, like all good myths, they are being used as universal truths. The existence of these myths is, in fact, keeping many well-intentioned capacity planners from developing methodologies that reflect the real requirements of their unique environments.

By investigating some of the sacrosanct beliefs that are still pervasive in the capacity planning community, I hope to point out how they actually inhibit the development of sound planning processes at many data centers. I have found these myths entrenched at major state-of-the-art installations and not just among the aborigines and bushmen of data processing. It is time to exorcise these demons and get on with capacity planning.

Capacity Planning Definition

Before examining the value (or lack thereof) of various myths in the capacity planning process, first agree on a definition of the process and its goals. I like to define capacity planning as the process of ensuring sufficient resources are available to supply satisfactory and consistent serv-ice in a cost-effective manner. This task encompasses a wide range of technical and managerial skills including data gathering and analysis, user liaison and interface functions, workload characterization and forecasting, projection techniques (that is, modeling, simulation) and technical and management reporting.

While all may agree on the noble goal embodied in the above definition, the required action items may be viewed slightly different. Some of these disagreements will stem from differing groups' unique roots. One group of capacity planners have come from the technical systems programming and performance school of thought. Others have been trained as applications developers. A third group has emerged from statistical or mathematical backgrounds. In truth, each of these disciplines will contribute to the successful capacity planning process. It is the manager's task to integrate these diverse skills to yield the optimal product. The effective manager will avoid the easy answers offered by the following myths and get on with supplying the best possible process to satisfy data processing, users and management.

The Myths

I am sure that I have left out some often-used cliches. Please do not be insulted if your personal favorite is not on the list. If you send it to me, I promise to include it in Release 2. (That, by the way, is one of the great promises of data processing, "It will be fixed in the next release.")

To illustrate how dangerous the acceptance of cliches can be, the first myth is a common, non-DP rule that has been used often by computer professionals.

■ *"If it ain't broke, don't fix it."*

This rule definitely has a ring of truth. I can remember using it myself when I encountered my son beginning to disassemble my alarm clock. The cliche actually works in many situations; however, it is not a universal truth.

How comfortable would you be flying with an airline that adhered to this rule as its maintenance philosophy and waited until the engine failed before overhauling it. In fact, this is quite contrary to the data processing approach of preventive maintenance.

The use of this cliche within data processing circles was probably initiated by applications programmers who believe (and rightly so) that any time you make a change to code, you leave yourself open to unanticipated errors. Therefore, let sleeping dogs lie (if the cliche fits, use it). The rule does not have applicability for the airline and it is not a good philosophy for capacity planning.

The basic goal defined above requires anticipating resource needs and acting to avoid any potential problems. Capacity planning necessitates a preventive maintenance mentality. Perhaps a better cliche would be, "Never let it break." Look for more effective ways to anticipate problems, identify solutions and implement those solutions before the failure occurs.

This leads us to a pair of mutually exclusive myths that are both widely held. The reason that both are accepted by such wide audiences is that the truth lies somewhere between the two statements.

■ *Capacity planning and performance analysis are the* "same task."

■ *Capacity planning and performance analysis are* "totally different."

First, I suspect the choice that is adopted in a particular shop is probably political, not technical. It depends on whether or not the same group that does performance analysis does, or wants to do, capacity planning. In examining the tasks involved in each process, it is evident that performance analysis is generally the act of isolating and correcting problems after they occur. Capacity planning should be the process of predicting future problems and avoiding them. Typically, the same set of data may be used in each case (that is, SMF, RMF, CMP, NPM and so on) but the mind set and the approach should, by necessity, be different. In addition, capacity planning must have more rigid user interfaces established and must address a report set that also satisfies management levels. The capability for both capacity planning and performance analysis must exist in your data center (unless you have such total faith in your capacity planning process that you never expect a performance problem).

Another interesting question in this same general area relates to whether or not the same individuals can physically perform both the performance and capacity tasks concurrently. There are many reasons to assume this is true — common data requirements, common data analytic activities, knowledge of the hardware and software environments. However, there is one overriding element that usually spells disaster for this double job specification scenario. The performance analyst is re-

IBM mainframe application software has a new DB2 champion

We're Lawson, the IBM mainframe business application software company with the confidence to challenge the industry's Goliaths. And hit them with new ideas. Like low cost, low maintenance—yet full-featured—business application software.

Lawson: Leader in software technology.

Our new technology is rapidly showing the "big guys" a few things about PINSTRIPE® low cost, low maintenance software. Forget "other" IBM mainframe software. With its high prices, old code, time-consuming installation and expensive upkeep.

New technology made our lower costs possible. It also made us first with application software for the DB2 relational DBMS. Now you get a choice: DB2 or VSAM. With newly developed code that's easy to install. Simple and inexpensive to maintain.

Solid application technology.

Comprehensive, well-thought-out application features reflect the experience of our package designers. Seasoned accountants, CPAs and human resource professionals who know what features and functions your users need.

And features they request. In the last four years, 84 percent of user-group suggested enhancements have been added to Lawson PINSTRIPE® packages.

Service and support — cutting the "big guys" off at the knees.

We're big enough to get the job done—but small enough to be sensitive to our clients' needs. With the best service and support available on the IBM mainframe market today. Including your personal software consultant, an expert who knows your needs *and* your application.

A full range of optional Support-Plus services is available, too. Like implementation consulting, data conversion, technical and application consulting. And a post-sale audit to make sure you're utilizing every feature we've built into PINSTRIPE® packages.
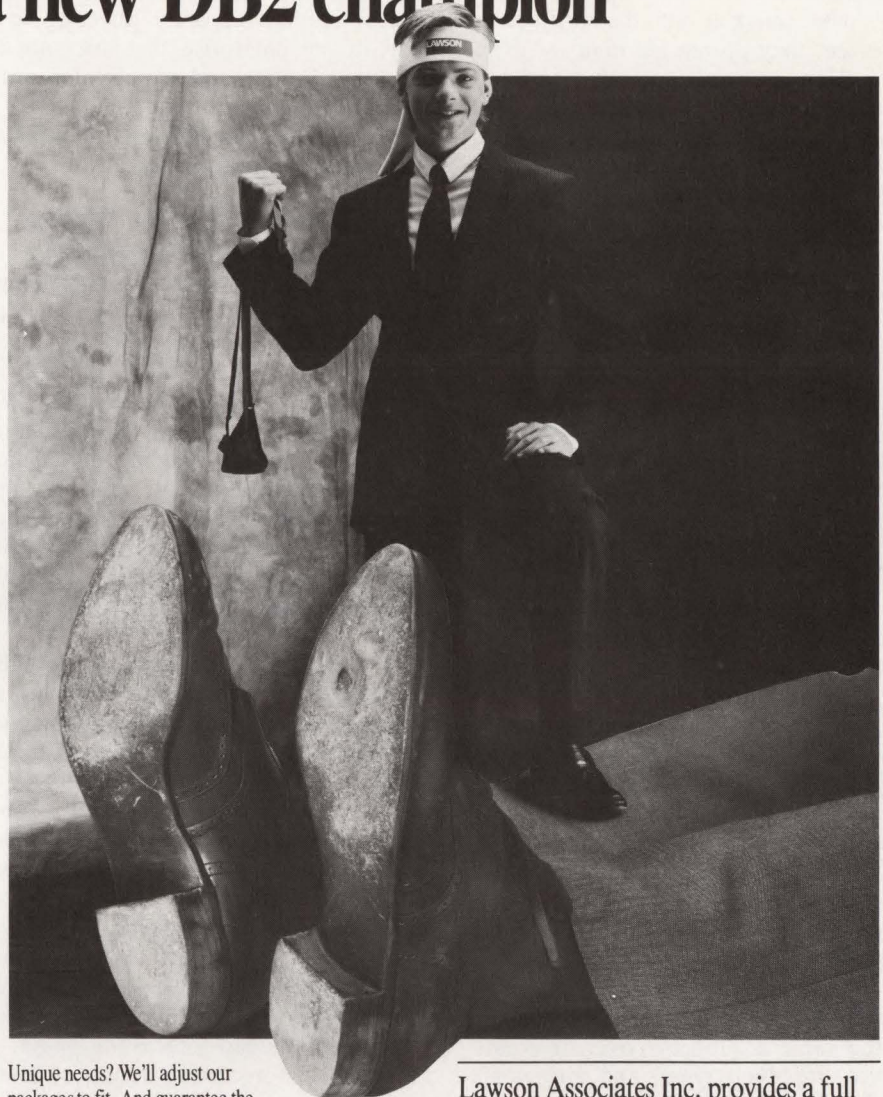
Unique needs? We'll adjust our packages to fit. And guarantee the work. If you'd rather do the job yourself, you'll find Lawson software easy to modify. And we'll still support the rest of the package. Other application software companies won't.

Lawson, IBM mainframe software's new champion.

We've put it all together. New technology with fresh code, including both VSAM and DB2 expertise. Full-featured application technology. Backed up by support that sets a new industry standard.

All in low cost, low maintenance application software from Lawson. No wonder we can take on the "big guys." And win.

Lawson Associates Inc. provides a full range of feature-rich PINSTRIPE® business application software packages for the IBM mainframe hardware environment. Available in the U.S., Canada and overseas.

Accounting
- General Ledger • Accounts Payable
- Accounts Receivable • Fixed Assets
- Project Accounting

Human Resources
- Payroll • Personnel • Benefits

Distribution/Procurement Management
- Order Entry • Inventory Control
- Purchase Order

LAWSON

...WHATEVER IT TAKES!

For more information, call
Sue Weinacht

(612) 379-0258

LAWSON ASSOCIATES INC.
1300 Godward Street, Minneapolis, MN 55413

CIRCLE #32 on Reader Service Card ▲

action driven. If there is a problem, he must respond now. This responsiveness forces other activities onto the back burner and often, they do not get done — ever. It has been my experience that capacity planning becomes a permanent back burner task when assigned to performance personnel and does not get done. This leads to the general feeling that performance is more valuable and that capacity planning does not really accomplish anything. If you actually do want to accomplish something, give capacity planning the support and visibility it deserves. Let it stand alone with sufficient resources.

Since capacity planning and performance analysis are similar but different, you could then incorrectly assume the following myth.

■ Capacity planning is done once a year.

A common belief is that capacity planning is only done to support the annual budget preparation activity. What you really want to do at budget time is develop a budget. As it turns out, this is when many data processing managers wake up to discover that it is not feasible to cost out a future configuration until they actually have some idea of the contents of that configuration. In order to assess the contents, they need to forecast requirements. It is starting to sound a lot like capacity planning.

If you wait until these questions need to be answered before starting the analysis, the game is probably over. It is difficult to effectively amass and analyze all the requisite information in a narrow timeframe. An effective projection of future requirements is based on data analysis and user interfaces that should take place throughout the year. It is built on solid reports that track growth (allowing for predictive trending) and that evaluate the credibility of past volume growth predictions.

Try a different cliche instead, "If you do not understand it now, you will never predict future behavior." The process of adaptation requires that an organism understands its environment and anticipates upcoming changes. It must then make the necessary behavioral changes to conform with the environmental changes. You really cannot wait until the midst of winter to worry about whether there is fuel in the furnace.

The establishment of an ongoing, effective capacity planning process will ensure the validity of the answers needed to support management requirements in a timely fashion. The lack of a continuous process reduces capacity planning to a panic-driven enterprise that does not fit into the definition and actually defies the concept of planning.

While in this general area, look at an offshoot of the above belief.

■ Capacity planning is modeling: modeling is capacity planning.

It is common for otherwise knowledgeable managers to greet me on the first day of a capacity planning consulting study with something like, "Okay, let's roll up our sleeves and build this model." My first task is to gently slow down this runaway locomotive. Modeling, simulation and benchmarking are all valuable tools in the capacity planning arsenal. They are not, in and of themselves, capacity planning. Each one requires significant upfront analysis to ensure a successful effort.

A model is a representation (usually mathematical) of current reality that can be modified parametrically so as to anticipate the performance of a future expected reality. In order to develop this picture of reality, you must first understand it. This requires data collection and analysis, configuration analysis and workload characterization. While there are many tools available to assist in these activities, the upfront analysis will still take up the major portion of the study. A considerable portion of the remaining time will be taken up by the workload forecasting phase.

Workload forecasting requires determining the anticipated changes to each workload for the time span covered by the study. Clearly, an open interface with the user communities is critical. In addition, a method of statistically analyzing historical data is extremely useful.

The tasks of workload characterization and forecasting typically consume 80 to 90 percent of the time spent on a capacity planning effort. So it should be evident that capacity planning is much more than modeling.

It is possible that this myth was created as a result of the fact that many modeling tools and packages do have the added benefit of being usable in the workload characterization process. That is a nice extra but modeling per se does not begin until a significant amount of groundwork has been covered.

Having mentioned the issue of workload forecasting, I will move on to another sacred cow.

■ Capacity planning is a technical exercise. Therefore, I only need to consider machine readable data.

This is the school of thought that is convinced that if you can get one day's SMF and RMF data from a site, you should be able to determine all of its problems and make recommendations for the future. I have to make a confession. After 20 years of performance and capacity work, I am still not that good.

As discussed above, even the technical portion of the capacity planning activity requires some detailed legwork and will depend on several different measurement intervals to ensure the validity of the process. Beyond that, there are major data sources that are not in typical monitor-produced, machine-readable formats. At the very least, projections of anticipated volumes will be required from users, new systems under design and relevant business changes (that is, new lines of business, mergers, expansions and so on). To do a complete job, obtain information on budgetary constraints and concerns, on-order equipment, lease expirations and service levels (explicit or implicit). These are simply not produced and archived by monitors; they require human interfaces.

It is possible (and probably desirable) that this type of data, once obtained, is input into a SAS or LOTUS format. Maybe that is why some people think they are only dealing with machine readable data. I strongly recommend setting up an ongoing database of human-based data that will facilitate analysis and graphing. I have developed several quasi-automated approaches that utilize SAS and/or LOTUS successfully. The existence of this type of facility greatly enhances the capacity planner's ability to understand and predict the changing nature of his environment.

A few paragraphs back, I mentioned service levels. This leads me to consider the next myth that is really an industry standard. Several years ago, IBM attempted to kill this one with little success.

■ There is a single measure to which all data can be distilled that can answer all capacity questions or the "I only need MIPs" complex.

In the mid-70s, I was responsible for computer performance measurement and analysis for a major aircraft manufacturer. Many of the management personnel came from an engineering background and approached data processing problems with a slightly different viewpoint. Most of the
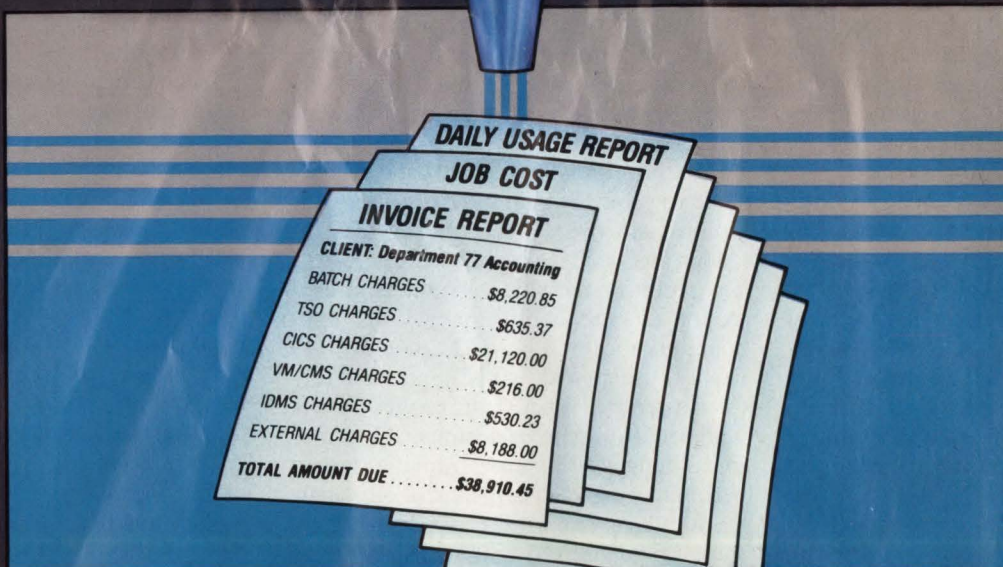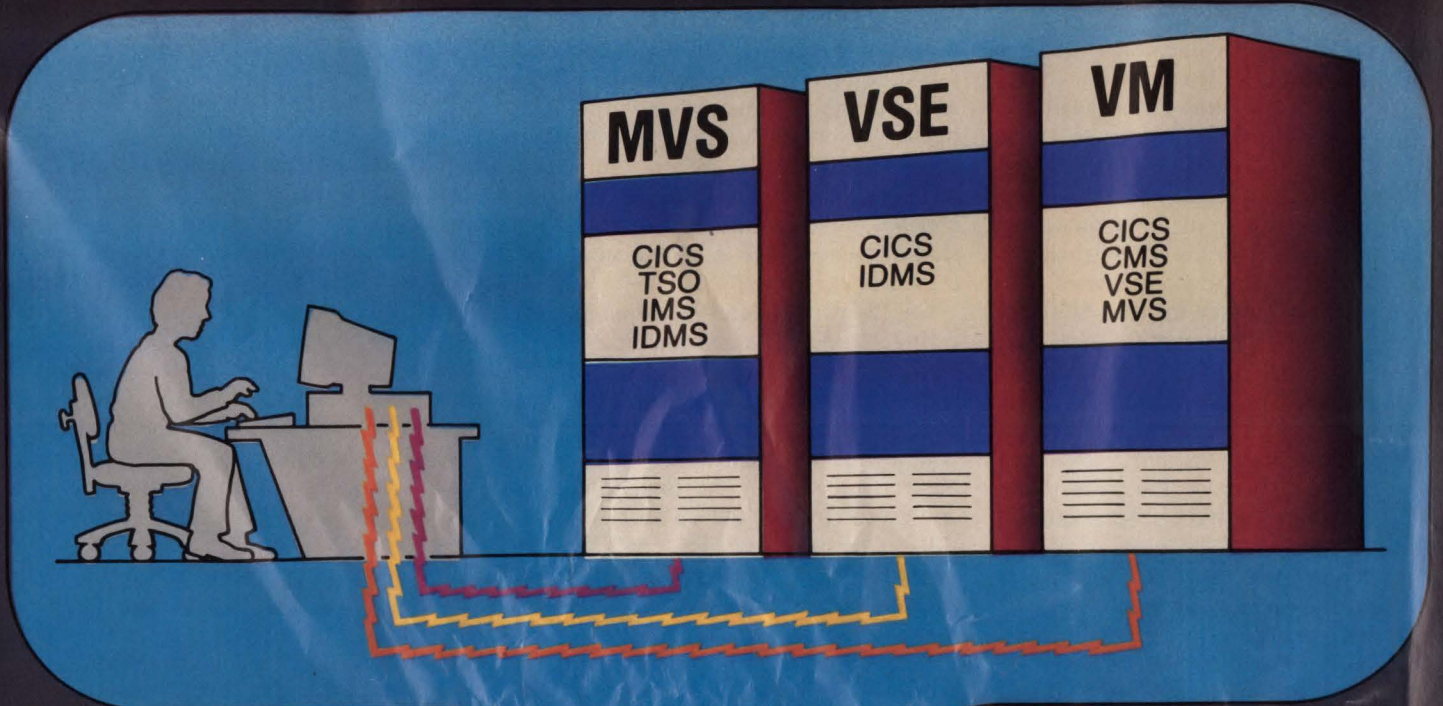
time this resulted in a new process that turned out to be of exceptional use and greatly ahead of its time. In one case, though, I am still waiting for the available instrumentation and techniques to catch up with the concept. It was called the "tailpipe temperature gauge." The name and the idea came from the engineering problem of building a highly efficient engine. Having built the engine, it is possible to develop measurement instruments that trap the exhaust of the engine and analyze it. The resulting information, which is limited to only a handful of values, would yield a wealth of knowledge concerning the strengths and weaknesses of the engine. To prove the point, New York State (and I assume several others) requires a yearly vehicle emission inspection that relies on just such a device. It produces about five numbers that attest to the proper combustion and general overall acceptability of automobile engines.

It seemed logical to attempt to distill the available measurement data and develop a "measure of goodness" for the computer system. Please keep in mind that this was a serious undertaking entered into with great hopes of success. Also, remember that this was just prior to MVS (Who remembers SVS?) and the intricacies were a little less imposing than they are today.

After some detailed analysis (scientific, statistical, intuitive, mystical and so on), a great flaw in the original analogy was discovered. While it was possible to determine from a tailpipe temperature gauge how well the engine performed, it did not supply any insight into such ancillary functions as the electrical system, onboard computers, navigation, passenger comfort, communications and so on. In fact, while getting an overall picture of the engine, a picture of the overall aircraft was not evolving. Total performance was not being measured or predicted. This insight came about at the time of a general awakening within capacity planning and performance analysis circles that they were not really chartered with keeping the CPU (engine) at 100 percent but with supplying some form of acceptable service to a user community (overall performance). As in the aircraft, there were many other subsystems within the computer complex that contributed greatly to the real measure of goodness. I/O rates, page faults, RPS misses, channel utilizations, control unit contention, memory occupancy, path concurrency and many other factors affected total performance. In fact, when they listed

the available data elements that might be needed to capture, archive and analyze, a fair sized pamphlet was published. (If you are a disbeliever, find a shop with a SAS-based performance database and do a PROC CONTENTS. When you see the quantity of values that are trapped, counted, collected and archived, you will start to wonder how any application code ever gets executed.)

So, they considered combining, analyzing, reducing, summarizing and otherwise massaging the "significant" measurement values to obtain a few succinct parameters of systems goodness. The bottom line was that the best measure of the performance of the total system on any given day was the number of logged calls to the help desk. This really was not what they were after but it did substantiate that the prime goal was to satisfy the users.

Since I personally went through this learning experience so long ago, I am often amazed that data processing personnel still cling to the concept of measuring and planning by MIPs alone. Millions of instructions per second is a measure of CPU capability and usage only. It does not relate to the current or future delivery of service. What is missing is the interactions of the operating system and the applications, I/O requirements, memory usage and so on. At best, it is a rough ballpark estimate to limit the magnitude of the required capacity study. The study must include a careful analysis of the current environment (including, but not limited to the CPU), anticipated growth, new systems, new technology, budgetary constraints and overall corporate goals (including service levels and business plans). The result of the study should be a projection of the ability of various total configurations (CPU, I/O, memory) to supply users with consistent, satisfactory response times and service levels while meeting corporate financial guidelines (budget). Modeling, simulation and benchmarking offer a means of anticipating future service level capabilities. They should be employed in the capacity planning process. However, they are not the post-processor, all-purpose monitor that the tailpipe temperature gauge was intended to be.

While I still patiently await the announcement of a true tailpipe temperature gauge, I am certain that MIPs is not the answer. In fact, I have seen many people badly burned by relying on the results of analysis based solely on MIPs. Use MIPs

for sense of direction sizing only. If there is any doubt, do a complete study.

■ *The technical solution should always be adopted.*

This is another common mistake made by evolving capacity planners. I suspect there has been more frustration, employee dissatisfaction and job hopping as a result of this myth than any other single cause. To illustrate how this occurs, consider the capacity planner who has done a textbook study. He (or she) has performed a superb workload characterization, interviewed users, reviewed past data and created reasonable projections of future application and operating system requirements. Using a robust modeling tool, a meaningful assessment of future hardware and software needs have been developed and a concise report to management has been produced. The analyst is rightly proud of a job well done. Management, however, does not choose to implement the suggested optimal configuration. The analyst is crestfallen and believes that: a) management does not understand the capacity planning process; b) his work has been in vain and c) he is really not appreciated. This often leads to a downward spiraling morale problem that culminates in a job change. The corporation has lost a valuable employee and the employee is probably destined to repeat this encounter at his next job.

In fact, management actually liked the report and was able to discern many valuable insights from it. Management's goal (unbeknownst to the capacity planner) was to meet corporate "bottom line" goals while supplying service. It is possible that a different configuration than the one suggested offered a better financial advantage or allowed for business plans that were not known to the analyst. Management may, in fact, have made the right business decision which is not necessarily the optimal technical decision.

Why did this happen? Let me begin by saying that there is fault on both sides. Management should have advised the author of the report that it was a well-thought-out piece of work. If possible, it should have been made known that the final decision was not a result of disfavor with the capacity planning process. Bidirectional communication is a critical component of capacity planning.

As for the analyst, he made the mistake of expecting management to simply act as a rubber stamp for his recommendation. The result of the study is only one input

Managing DASD Performance Can Be a Pleasure Trip... Not a Survival Course

DASDMON will revolutionize the way you manage MVS DASD performance.

Traditional DASD tuning techniques are being outpaced by faster CPUs and increasing user demands for more online data. Although cache controllers, data spaces and hiperspaces will relieve some of the strain, they also create new complexities. Tuning DASD can feel like struggling through a survival course.

DASDMON takes you away from it all

DASDMON offers you system managed performance today. By cutting through the complexities of conventional tuning, managing DASD performance is more automatic and streamlined than ever before. Once you've set response time objectives, just sit back and...

- **DASDMON** will identify problems by continually measuring individual I/O response times and comparing them to your objectives.

- **DASDMON** will determine the cause of problems by gathering and analyzing data. No longer will you spend hours scanning reports.

- **DASDMON** will present solutions to your problems in precise English, and simulate their effects to ensure that a new problem is not introduced.

- **DASDMON** will select the best possible set of solutions so you can implement recommendations with confidence.

And, if you have any questions or concerns along the way, our experienced technical support staff will be glad to guide you—24 hours a day.

To maximize your system's overall performance, use DASDMON with our DASD Performance Optimizers. These products eliminate DASD I/Os, resulting in automatic response time reduction.

For more details, call **(800) 323-2600** or **(412) 323-2600** in Pennsylvania.

Performance and Optimization products from LEGENT. **Enjoy the results.**

LEGENT

The company formed by the merger of Duquesne Systems and Morino.

Two Allegheny Center
Pittsburgh, PA 15212

© 1989 Duquesne Systems Inc.

Only one automated operations product is a perfect fit.

OMEGAMON　　　**AF/OPERATOR**

Service level management began with OMEGAMON®. So we precision-crafted automated operations to fit perfectly with OMEGAMON. Because automation is just another part of performance management.

AF/OPERATOR™ learns about MVS, CICS, IMS, or DB2 problems from OMEGAMON. So as soon as an exception is triggered, AF/OPERATOR can issue commands to pinpoint the problem. Then speed a message directly to your best troubleshooter for fast resolution. It can even take action itself to improve service at machine speed.

Maybe the CICS claims application is way over its service level. AF/OPERATOR can issue OMEGAMON commands to determine which resource is delaying claims. Or if other jobs are causing the delay. It can then automatically track down the user to see if the conflicting job can be cancelled.

Of course, AF/OPERATOR lightens the load of your operations staff with console message management, automatic startup of time-dependent activities, reliable IPLs, and orderly shutdowns. And it simplifies your operations by becoming the single point of control for all your systems — MVS, IMS, and JES to name a few.

For over ten years, data center managers have trusted the performance of their systems to Candle. Now you can automate your data center with the same confidence. Expect AF/OPERATOR to fit your data center perfectly — now and in the future. For more information, call Terry Forbes at **(800) 843-3970**.

!Candle®

CIRCLE #120 on Reader Service Card ▲

into a complex analysis and decision process that ultimately might require the expenditure of millions of dollars. Such an investment must be based on more than just technical criteria. To avoid this situation, the capacity planner must readjust his thinking to accept that the goal of his efforts is to supply management with the information necessary for *it* to make decisions. To enhance the value of his recommendations, he might consider including management-related topics such as finances.

■ *The newest piece of hardware is always the right answer or the "bigger is better" syndrome.*

Many capacity planners and managers also subscribe to this belief. I maintain that this myth has been carefully nurtured over the years by vendor salesmen. A case in point that illustrates the flaw in this particular rule concerns CICS and the recently available IBM CPUs. CICS is a transaction processing program that operates as a single address space under MVS. Therefore, it can only use one engine in a dyadic or dual-dyadic (or six-plex) complex. Think back in recent history to the heydays of the 308X series. (Was it only just a few years ago?) While the 3084-Q was the flagship, top-of-the-line machine, it was actually made up of four engines. Each of these had less speed than the single engine in the 3083-J. It turned out then that, in many cases, the 3083-J was a better choice for CICS performance. Today, there is no benefit for a single region CICS workload on a 3090-600 over a 3900-200. True, some ripples have been imposed by subtasking (for example, VSAM), MRO and ISC implementations and by the fact that many shops are running multiple CICS regions concurrently. But the moral of the tale remains — *before you begin a capacity study, understand the goals.* If you were expected to maximize CICS throughput, bigger was not necessarily better.

The second lesson learned is that in order to perform a complete capacity analysis, the capacity planner must have a solid knowledge of the environment (hardware and software). This knowledge is not acquired through quick, once-a-year methods but is supported by an ongoing process that allows for the analysis of information under various conditions in non-panic mode. This process requires a sound, regular reporting system that provides insights into the system.

■ *I can develop one report that captures everything that all levels of management and technical personnel require concerning resource utilization.*

In developing this reporting system, do not fall prey to this assumption. This is a reinterpretation of the MIPs syndrome discussed previously. This mythical report does not exist. There are many different types of people who will make use of capacity planning reports. They range from technical personnel (such as performance analysts and capacity planners) to users interested in service levels to management personnel who want concise reports on the overall state of the system. You cannot take the same report and simply change the cover letter and ship it to a new audience. The reports must be tailored to the needs of the particular audience that they address.

As you can see, myths can keep you from completing the capacity planning task. The next myth is one of the more dangerous myths that will actually keep you from ever starting the job at all.

■ *We cannot perform capacity planning for a system or complex that is having performance problems.*

I have dubbed this the "tuned system trap." The first reality that confronts this rule is that there are no systems that are ever fully tuned. Does that mean that you never do capacity planning? That is hardly an acceptable answer to either capacity planning professionals or to the corporations that require the answers that they supply.

Fortunately, the answer to this dilemma is that the capacity planning process has within its structure (via modeling or simulation or benchmarking) a means for isolating bottlenecks and test driving the potential solutions to these conditions. This means that you can actually utilize capacity planning to assist in performance analysis and tuning. It also means that the anticipated performance improvements can then be applied to the predictive tool to complete the study. In effect, you have

CIRCLE #149 on Reader Service Card ▲

Carolina Steel Makes The Most of VSE

Like tens of thousands of other IBM shops, Carolina Steel is sticking with VSE. However, they are getting a lot more out of it than most.

By Mary Lou Roberts

For years, VSE shops have been faced with the decision of whether or not to convert to VM or MVS. However, it seems that, for a large percentage of these loyalists, neither positive inducements ("Have I got a deal for you!") or negative fear tactics ("VSE will not be supported much longer") have succeeded. DOS shops — all 30,000 of them — have been sticking to their guns and their operating systems with all of the tenacity of a Garfield clinging to a car window.

Are they wrong, these lovers of batch-oriented processing and limited partitions? The answer is a resounding "NO."

Sticking with VSE is *not* akin to being trapped in the Middle Ages of information processing technology. Or, at least, it does not have to be. When coupled with back-to-basics management and creative approaches to operations, as it is at Carolina Steel, VSE can shine. Some of the frills may be missing; but the work gets done. Users are satisfied and costs are kept to a miminum.

The Carolina Steel Environment

When Charlie Rice joined the staff of Carolina Steel Corporation 26 years ago, there was not much question about which operating system the company should use. Its new IBM 1401, the company's first computer, only allowed for one option. However, life was simpler then.

Today, things are much different. Rice, who is the assistant manager of Information Services, has seen lots of changes take place over the years — both in his company that currently handles more than

From left to right are Don Stonemen, systems programmer, and Charlie Rice, assistant manager of Information Services, at Carolina Steel.

$250 million in sales and in the increasing number of information technologies that are available.

Operating out of its Greensboro, NC corporate headquarters, the Information Services (IS) group of Carolina Steel services six fabricating plants, 13 service and processing centers and two concrete plants across the southeastern part of the country.

Two major types of applications are involved. In supplying steel and concrete

for the construction of buildings and bridges, IS must deal with small volumes of large orders. For this aspect of its business, the data processing applications are the more traditional manufacturing job-shop programs such as Bill of Materials and Costing.

Processing for the other side of the business, the steel service center, has different requirements. As a major distribution system, it must handle large volumes of much smaller orders. Applications

here include Order Entry and Inventory Management.

All of this processing is done on an IBM 4381 that replaced a 4341 about one year ago. The system has an on-line CICS network with 350 terminals that are distributed all over the southeast. The inventory of about 3,000 COBOL and FORTRAN programs is structured in approximately 1,000 job streams.

From 8 a.m. to 6 p.m., the system is primarily dedicated to CICS with no scheduled batch operations. During these hours, Carolina Steel averages 100,000 transactions per day with response times of less than one second except for line delay. All batch is run at night while CICS is still active but handles only a small number of transactions.

The operating system is VSE/SP. Due in large part to Rice's efforts, Carolina Steel has been operating with some version of DOS since the late 1960s — a fact of which the company is proud and which it believes has contributed in large part to the smooth running and cost effectiveness of the corporate information systems operation.

What Is Different At Carolina Steel?

Several people have pointed to Carolina Steel as a model shop for its size: one in which efficiency and user satisfaction are high and costs are low.

What makes the company different? If you ask Rice, he will tell you that Carolina Steel differs from other shops of its size in three major respects. "First, is our commitment to and effective use of VSE. Second, is our completely automated operations and finally, the absence of a database management system."

VSE

Rice notes that rumors about IBM's continued support of VSE (or lack thereof) have been floating around for several years. "But now, the future of VSE has come back to life. IBM simply cannot ignore 30,000 users. IBM has to support them because these users just are not going to convert. Now, IBM has even committed to COBOL II and CICS under VSE," he points out.

Rice, who has been active in IBM's GUIDE users group for many years and is the current Group Manager of GUIDE's VSE Group in the Operating Systems Division, has had the opportunity to make his opinions well known to IBM. He believes that IBM will concentrate its efforts on continuing to make it easier for VSE users to move to MVS. "This means that the longer we wait, the easier it will be. It can only get better for us," he foresees.

Carolina Steel has evaluated the alternatives. At one time the company *did* install VM and used it for six months. "All it did for us was to allow systems programmers the freedom from having to come in on weekends. Otherwise, it was pure overhead. It just wasn't worth it," Rice explains.

Carolina Steel has also looked at what it would take to convert to MVS. "We estimated that it would cost in excess of $1 million in out-of-pocket expenses over one year. In that same year, we would not be able to do any new development. Furthermore, our ongoing operations costs would multiply by a factor of four! You have to gain some real benefits for that. And we just haven't seen those benefits

yet. Conversions are often political. Not many people can make a real business case to make a change like that at such great cost without getting a lot back for your money," Rice states.

Instead, Carolina Steel continues to operate VSE in the most efficient way possible — native mode.

Rice highlights still another way in which Carolina Steel differs from most other VSE shops. "I think that we're one of the few shops in the world that is run-ning VTAM in a private address space in VSE. It allows us to allocate more memory to other address spaces including CICS."

Given the current level of satisfaction and success with VSE, Carolina Steel sees no reason to convert. The company is aware, too, that Software Pursuits (Alameda, CA) sells a VSE replacement product called MVT/VSE. Rice notes, "It has some nice features. They support 16MB of real memory and would give us more partitions. Those things would be nice to have. If we were worried about IBM's continued support, we would certainly take a closer look at it; but, it's pretty clear that we'll stay with IBM as long as they stay with us."

Automated Operations

"One major way in which we differ from other shops is in the degree to which we have automated our computer operations. In fact, I have never seen another shop that is as automated," Rice points out.

In late 1986, Carolina Steel established the objective of achieving "unattended operations." Its goal was for a computer operator's sole duties to be: 1) mount tapes; 2) put paper in the printer; and 3) call people whenever the system told them to. Today, thanks to the implementation of three major programs developed by Don Stoneman, a systems programmer at Carolina Steel, the company is 98 percent of the way toward that goal and is able to run all computer operations with only two people.

The VSE/POWER Scheduling System is a program for controlling the time and order in which jobs are run under VSE/POWER. It provides one of the essential functions of unattended operations.

"When we first took a hard look at our job streams, we found that they had been set up in ways that would make it the easiest for operations to run them — not for greatest system efficiency," Rice says.

However, thanks to the scheduling system, things are done differently today. The users relate POWER jobs together and define a schedule for each of these groups of jobs. A strict syntax is provided by the scheduler to define the order in which the jobs are to be run and the events which must occur, if any, before each job is released.

Once a schedule is defined, the user simply submits the jobs that the schedule controls into the POWER reader queue with a hold disposition. The scheduler will then release or delete these jobs automatically, as required.

The Automated Console Response System, known as RSVP, adds the power of a high-level programming language to control the VSE operating system. Using this system, the user codes an RSVP procedure that is read from SYSIPT by the RSVP program, compiled and then initiated.

Each procedure is made up of one or more tasks. These tasks receive each mes-

Introducing a subsystem that cuts through the brown tape.

IBM mainframe users can now compress 24 reels of data onto one 8mm cassette.

It's called the 6800 Series Cartridge Tape Drive with the F1011 Data Compression Feature. Using IBM's SNA compression algorithm, the F1011 enables you to put a colossal 4.5 gigabytes onto one compact 8mm cartridge. Without operator intervention.

What's more, a fully configured 6860 tape subsystem features up to 7 cassette drives, which essentially replaces 196 reels. Goodbye stuffed storage rooms. Goodbye downloading monotony.

Of course, the IPL 6800 Series F1011 Feature uses no more of your system's resources than a normal tape drive operation. And as with all IPL products, the 6800 connects directly to your IBM system. There's no need for hardware or software alterations.

So if you're looking for a way to cut through the problems associated with reel-to-reel backup and storage, consider the IPL 6800 series.
It's an innovation that turns a reel hassle into a real pleasure.

Call IPL today at 1-800-338-8ipl, in MA (617) 890-6620. Or contact us at 360 Second Avenue, Waltham, MA 02154.
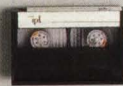
ipl *systems, inc.*

sage as it appears on the system console (SYSLOG) and can then perform any of a number of functions to handle the message. These functions include sending commands to the console or VSE/POWER, scheduling commands by time and executing other RSVP tasks.

Some RSVP procedures are short-lived, running only seconds and then continuing on with the next program in the job stream. Others can be resident for long periods of time, handling various conditions as they arise throughout the day.

A procedure can be run in a partition by itself or it can run concurrently with another program. It can be simple, responding to a single message; or it can be complex, consisting of several tasks and responding to a large number of conditions.

The function of the Report Archival and Distribution System is to handle report distribution and to archive those reports for reprinting or displaying on-line under CICS.

With this system, jobs are extracted selectively from the VSE/POWER List Queue. Then they are compressed and written into an Archive File. Once archived, these jobs can then be printed for normal distribution, reprinted later or displayed on-line. Up to 32 generations of each job can be retained.

"We have really seen the benefits of automated operations. Before we implemented these systems, we were a three-shift operation. Since then, we have added more work, been able to cut back to only two shifts, reduced our staff to only one operator per shift and are usually finished with all of our work by 10 p.m. However, we aren't completely satisfied yet. There's still a little way to go to reach our goal of 100 percent. We're going to keep working on it," Rice notes.

No Database System

"We consider the fact that we have no database management system to be one reason for our cost effectiveness. For much larger shops, it might be necessary. But the only thing it would give us is the advantage of adding fields to records. And we don't need to do that all that often. In fact, of all of the companies I've talked to, not many even try to claim that they have a successful database system," Rice says.

Carolina Steel does have a lot of end-user applications. However, to them that means that the users play a major role in defining the application, designing the system and submitting the jobs once they are in production. The users themselves have no hands-on access to tools like SQL or FOCUS.

"From time-to-time, some of our users have asked for tools that would allow them to have direct access to a database. We've been happy to listen to them and show them what it would entail. However, when we lay out the costs for them, they always want to stick with things the way they are. Data processing, here, is run as a cost center. Everything is charged back. And, so far, our users simply aren't willing to pay the price for databases and end-user tools," Rice explains.

The Computer Operations Expert System

Another project on which Carolina Steel is currently working that will bring even greater efficiency to its operations is the development of a PC-based expert system to assist computer operations.

This is one feeling you'll never have when you buy our software.

Ever happened to you? You buy a piece of software that sounds good, fulfills a real need and generates high hope. Then at the first sign of trouble, you're left standing by the software vendor. Alone. Abandoned. And frustrated.

Well, this simply doesn't happen with DTA Software Products. Because every one of our highly useful products comes with crowd-pleasing documentation. Clear, well organized, readable documentation that is considered to be among the best in the industry.

You also get our exclusive DTA Straight Line. A direct access phone line that puts you one-on-one with DTA's Software Product Developers. So, in the highly unlikely event that you hit a snag, we'll tell you how to overcome it quickly and effectively. No matter what time of day or night it happens.

And all this comes with the assurance you'd expect from software carefully developed in every detail. Then tested under honest, realistic conditions at user sites.

If you are looking for useful software products that provide valuable solutions to complex operation and development issues, join the crowd that has discovered DTA. You'll never be alone again.

DTA®

Software Products Group
550 Waterford Park
505 N. County Road 18
Minneapolis, MN 55441
800-521-6773
612-591-6100

DTA Software...More efficient in every detail.

DejaView.

```
------------- ISPF/PDF PRIMARY OPTION MENU VERSION 2.3 -------------
OPTION ===>
                                              USERID   - XASC21
  0  ISPF PARMS   - Specify terminal and user parameters  PROC     - P75XSP
  1  BROWSE       - Display source data or output listings PF KEYS  - 24
  2  EDIT         - Create or change source data           TERMINAL - 3278
  3  UTILITIES    - Perform utility functions              TIME     - 11:01
  4  FOREGROUND   - Invoke language processors or SCRIPT   JULIAN   - 89.018
  5  BATCH        - Submit job for language processing     DATE     - 89/01/18
  6  COMMAND      - Enter TSO command or clist
  7  DIALOG TEST  - Perform dialog testing
  8  LM UTILITIES- Perform library management utility functions
  C  CHANGES      - Display summary of changes for this release
  T  TUTORIAL     - Display information about ISPF/PDF
  X  EXIT         - Terminate ISPF using log and list defaults

Enter END command to terminate ISPF.
```

```
------------- IMS-XPERT/ISPF PRIMARY OPTION MENU -------------
OPTION --->
                                                   VOLSER   -
  0  ISPF PARMS    - Specify ISPF parameters and IXP PF keys RELEASE  -
  1  BROWSE        - Display data base contents             RELEASE  -
  2  EDIT          - Create or change data base contents    CPU ID   -
  3  UTILITIES     - Perform utility functions              USERID   - XASC21
  4  EXTRACT/LOAD  - Extract or load a data base or subset   TIME     - 10:41
  5  PRINT         - Print an audit trail                   DATE(G)  - 89/01/18
  6  SEL CRITERIA  - Create or change selection criteria    DATE(J)  - 89.018
  7  XREF          - Create or change segment/layout XREF   TERMINAL - 3278
  8  NOT AVAIL     - Reserved for future use                PF KEYS  - 24
  9  APPL REL      - Create/change application relationships LANGUAGE -
  P  IXP PARMS     - Specify IMS-XPERT parameters
  C  CHANGES       - Display summary of IMS-XPERT changes
  T  TUTORIAL      - Display information about IMS-XPERT
  X  EXIT          - Terminate IMS-XPERT and return to ISPF

Enter END command to terminate IMS-XPERT/ISPF

       Copyright (c) 1985, an unpublished work by XA Systems Corporation.
                          All Rights Reserved.
```

Productivity has never been so familiar.

It only seemed natural for us to create our XPERT Series of programmer productivity tools with an ISPF-like design. You'll learn skills faster and produce results immediately—something you can never see enough of.

DATA-XPERT, IMS-XPERT and DB2-XPERT all work the way programmers do. They're compatible with all the commonly used programming languages, file organizations and data base management systems in the IBM MVS environments.

Access to files and data bases like you've never seen.

Without getting tied up by using time-consuming utilities, you can easily access VSAM, ISAM, BDAM, PDS, Sequential files, and IMS and DB2 data bases. Once you're there, you can selectively edit, browse, extract, load, convert, reformat, and print any size file or record to get just the data you need.

See for yourself.

Find out why the XPERT Series is rated #1 over the competition.* Call today for more information or a free 30-day trial of DATA-XPERT, IMS-XPERT or DB2-XPERT.

Call 1-800-344-9223; in Canada, 1-800-344-9224.

The XPERT Series. Things will never seem the same again.
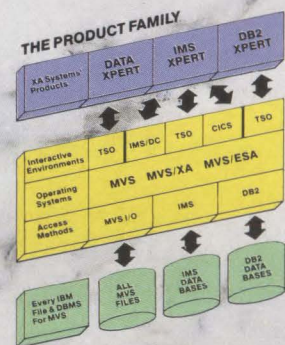
The Essential Software Company

XA SYSTEMS CORPORATION

The Eight Myths

of Measuring Software Development Performance

By James R. Johnson

Productivity equals output divided by resources consumed. Why can't this simplistic equation be used to establish precise productivity figures? If you have studied this area, the obvious problems are known: changing definitions, recording consistent statistics, selecting comparable projects and so on. Our database here at Hallmark Cards was extensive when we started a formal measurement project. Still, the task proved difficult and many myths were discovered as this article documents.

The lessons learned may apply to others, specifically those developing and maintaining software. This is the story of what was learned about productivity measures.

Myths

To start the story, consider eight myths associated with performance measures:

1. Performance measures are intuitive
2. Implementation requires less than three months
3. All activities should be quantified
4. Department measures roll down
5. Improvements automatically follow
6. Zero defects is the objective
7. Competing management methodologies are replaced
8. Corporate benefits are directly linked.

The arguments disputing these myths result from implementing and monitoring nine specific measures in the software factory, an environment where more than 200 professionals develop, install and maintain data processing systems.

Myth 1: Performance Measures Are Intuitive

Definitions of performance measures include four types: productivity, quality,

estimating and other ratios. Until clear definitions were completed (a six-week process), terminology caused considerable confusion. Defined in terms of ratios, the performance measures were as outlined below.

Productivity Measure

An output divided by a resource. The value of this ratio is dependent on how directly the output is related to the resource: the closer the relationship, the better the ratio. The process of producing output is the critical management issue for improving productivity. Examples of outputs for software development: function points, Lines Of Code (LOC), programs. Examples of resources: man-days, dollars, employees. Five ratios were defined in this subset as shown in Figure 1: support productivity, project productivity, function point productivity, LOC productivity, information center productivity.

Quality Measure

Number not meeting a quality standard divided by total number. Examples: defects — total items, program bugs — Total LOC. Two ratios were defined as noted in Figure 1: maintenance UCRs and installation UCRs. The denominator, Unusual Condition Report (UCR), in this context is simply a program bug causing a system termination. In many discussions of productivity, quality is assumed to remain constant. Measures of quality are indirectly related to the productivity measurement itself and must be monitored concurrently.

Estimates

Actual number of resources or outputs meeting estimate (or schedule) divided by the estimate base. Examples: actual hours — budgeted hours, projects on schedule — total projects. Two ratios were defined (see Figure 1): schedule estimate and budget estimate. Estimating is also related to productivity but is based on expectations and not necessarily improving output. It should remain constant or improve as productivity goes up.

Other

Influencing factors related to productivity; actual divided by goal/target. Examples: staff turnover — target, training hours — target. No ratios for this subset were defined.

In the nine performance measures, only tangible data were used. It is, of course, possible to incorporate opinions and other subjective data. Another item avoided was dollarizing the resources. Instead, user man-days or -hours were used. Converting to dollars introduces inflation, an additional complexity. Also, to cover all costs, software and hardware dollars would have to be included. This involves depreciation, leases, rental agreements and

so on and would further complicate the calculation.

All measures were formally recorded yearly; however, UCR counts for maintenance and projects were monitored monthly as they occurred in net amounts versus ratios. In the software factory it is not possible to link all measures to a month-by-month cycle since implementations occur throughout the year.

Myth 2: Implementation Requires Less Than Three Months

After six weeks, definitions existed but it was obvious other steps followed. Over the next four months, specific measures were established, data was analyzed, objectives were set and plans for improving ratios documented. This effort was longer than anticipated; however, implementa-

---

Without consistent output/resource data, an operational system for a major department requires a year or more.

---

tion could be significantly longer if consistent output/resource data is not available. In the software factory, years had been devoted to refining definitions of the two outputs, function points and LOCs.

Historically, monitoring LOC has received considerable criticism; however, if recording rules are consistent and if only one language is used (COBOL in our organization), then LOC is considered a reasonable measure of output. Collection simplicity is one inherent advantage with no separate calculations needed. In fact, the collection of data may be automated. Although the LOC definition was calculated with minor variations for three measures (Numbers A1, 2 and 4), the general intent is as follows: *All source statements (physical LOC) including comments for production programs (excludes conversion and test programs).*

Also, an existing time-reporting system isolated the resource (productivity ratio denominator, expressed in man-days) as project, maintenance or enhancement time. With this data available, defining terms, setting objectives and documenting an improvement plan were possible in

four to five months. Without consistent output/resource data, an operational system for a major department requires a year or more.

A related subject is permanency. Although the management team had confidence in each measure, Performance Measure Number A4 (LOC Productivity) initially did not accomplish its intended purpose of accurately combining project and enhancement LOC outputs. During the first year, deleted and changed LOC were not recorded; thus, the enhancement LOC was understated. After the automatic LOC counting system was enhanced, the measure became meaningful. Performance measures are subject to refinement and change even when organizational functions remain constant.

Myth 3: All Activities Should Be Quantified.

After completing performance ratios, the question was asked, ''What percent of management's responsibilities are measured?'' To answer the question, departmental reponsibilities were grouped into three categories as noted in Figure 2. The first item, Business Systems Planning (BSP) was subjectively valued as one-third of our responsibilities. This involves selecting a portfolio of projects that optimize the corporate contribution. In other words, if a department is not working on the best projects for the company, productivity is a secondary factor.

The next third includes personnel management, introduction of technology and other general management activities. The last third covers implementation activities. The nine performance measures defined are direct measures of this last activity. BSP is not covered and the middle third activities are only indirectly monitored by the performance measures. For example, if there is a strong training program, this should result in more productive project implementations and improved ratios.

Establishing performance measures for the other responsibilities is much more

---

FIGURE 1

Performance Measure Results and Objectives

Performance Measure		General Objective	1988 Specific Objective
A Productivity			
1. Support Productivity =	Total LOC / Equiv Support Staff	Increase	160,000
2. Project Productivity =	Project LOC Implem / Project Man-Years	Maintain or Increase	23,000
3. Function Point Productivity =	Function Points / Project Man-Days	Increase	8
4. LOC Productivity =	LOC Activity / Equiv Enhance/Dev Staff	Increase	27,000
5. Info Center Productivity =	Function Pt / Programming Man-Days	Increase	11.0
B. Quality			
1. Maintenance UCRs =	Total LOC / Tot UCR Cnt — Proj UCRs	Maintain or Increase	13,000
2. Project UCRs =	Project LOC Implem / Project UCRs	Maintain or Increase	6,000
C. Estimating			
1. Schedule Estimate =	Projects On or Ahead Schedule / Total Projects	Increase	85%
2. Budget Estimate =	Projects On or Under Budget / Total Projects	Increase	80%

FIGURE 2

Major Responsibilities

- Business System Planning
  - User Interface
  - Portfolio Selection } 1/3
- Personnel Management
  - Career Planning
  - Recruiting
  - Training
  - Turnover
- Technology Introduction
  - 4GL
  - Expert Systems } 1/3
  - End-User Computing
- Other
  - Planning To-Dos
  - Standards
  - Internal Audits
- Implementations — Direct Measures } 1/3

FIGURE 3

Software Factory Functions (3a)

	MEASUREMENT VALUE	
Process	LOW	HIGH
Structured	Planning meeting, Standards, Audits	Project implementations, Maintenance, Info Center Service
	VERY LOW	LOW
Unstructured	Business Planning, Career Planning, Other Mgt. Activities	System Design process
	FEW	MANY

NUMBER OF OUTPUTS

Corporate Functions (3b)

	MEASUREMENT VALUE	
Process	LOW	HIGH
Structured	Engineering, Finance	MIS, Manufacturing, Distribution
	VERY LOW	LOW
Unstructured	Personnel, Legal	Purchasing, Marketing, Product Management
	FEW	MANY

NUMBER OF OUTPUTS

difficult because data is intangible and not repetitive. Projects, on the other hand, provide a base of comparison, especially when data on 10 or 15 projects is documented. Consider how difficult it would be to establish measures on the BSP function. Who could provide a scale measuring the contribution of this planning activity?

In general, the value of purely quantitative measures depends on two variables: process structure and number of individual outputs. See Figure 3a for the software development activities. In the factory, the implementation process, with high structure and volume, facilitated numeric performance measurement. However, the design process lacks structure, as do other management activities in which the value of quantitative measures is limited.

Activities in different functional areas of the corporation vary considerably in these factors. For example, the manufacturing process is inherently one of high structure and many outputs; a legal department might display low structure and possibly few major projects per year. Figures 3a and 3b link the two variables to functional areas and also predict the value of each combination. Since direct experience only substantiates the software functions, other functional area placement in the matrix might be debated. Also, as with the software factory certain aspects of each functional area could be quantified with performance measurement.

Support versus development is also a distinction that must be made in most operating groups. There is usually some level of activity that goes to maintenance or

support of existing facilities. Output versus resource should be separate for the support activity. Thus, for almost every group, there should be a minimum of two measures. This was not true for the information center, where only one measure was defined: function points divided by programming man-days. A dramatic increase obtained one year (10.2 vs. 6.8 function points per man-day) was due to cloning information requests and diverting resources from consulting/training activities, which were not measured, to the measured activity. Obviously, the second reason was not intended. Although the effort to define more measures can have a diminishing rate of return, in this case, consulting and training activities were monitored.

The value of our measures is in top-down order with the majority of benefits at the director level.

Myth 4: Department Measures Roll Down

All nine performance measures were established to measure department trends, the efforts of over 200 individuals. Assume the organizational hierarchy of programmer, project manager (six to 12 staff), systems development manager (25 to 30 staff), director of systems development (200 + staff). Many people envision using performance measures at the first two levels, programmer and project manager. However, the value of our measures is in top-down order, with the majority of benefits at the director level, some benefits at the systems development manager and project levels and minimal value at the individual level. To substantiate this point, consider the range of values reported by eight similar groups (systems development manager level) in Figure 4.

Remember, these variances were in one organization with the same standards, productivity tools, experience level and working environment. The difference among the groups is project uniqueness and application area; that is, manufacturing, finance, order processing, market-

ing, distribution, purchasing, inventory control and personnel.

The more groups in an average, the more reliable are the statistics. For example, department results were reasonably close to goals because the eight individual areas were averaged and the "law of large numbers" improved the reliability.

However, the lack of correlation to the department average does not necessarily imply lack of correlation from year-to-year of one particular group. With a few exceptions, based on unique projects, the variance for the same group from one year to the next was reasonable. Thus, establishing goals has value with staffs of 25 to 30 completing five to seven projects per year. At lower levels (smaller groups), project uniqueness produces inconsistencies.

Myth 5: Improvements Automatically Follow

Some productivity experts recommend tracking ratios over time to establish trends. It was our philosophy that tracking was only part of the responsibility along with setting objectives and defining how to improve — what tools and/or methodologies will positively impact the ratios. Over the years, an emphasis on productivity resulted in a variety of programming/testing software tools to perform functions, such as CRT screen generation, test data preparation, interactive program debug and librarian functions. The maintenance environment utilizes these tools as well as restructuring code facilities.

Figure 5 lists 10 approaches cross-referenced to the measures impacted. In our

FIGURE 4

Performance Measures

Measure	Department Average	Range for Eight Systems Development Managers
A. Productivity		
1. Support Productivity	170,186	115,700 to 311,600
2. Project Productivity	30,638	17,100 to 111,900
3. Function Point Productivity	94	47 to 2.21
B. Quality		
1. Maintenance UCRs	14,012	6,500 to 33,100
2. Project UCRs	7,984	2,900 to 39,800

FIGURE 5

How to Improve

	Performance Measures								
	Productivity					Qlty		Est	
Approach	1	2	3	4	5	1	2	1	2
1. Monitor for consistent subsecond response time	X	X	X	X	X			X	X
2. Implement the JAD design process		X	X	X					
3. Search for code generator	X	X	X	X	X	X	X	X	X
4. Continue to look for software application packages	X	X	X	X	X			X	X
5. Clone code		X	X	X	X				
6. Install faster terminals	X	X	X	X	X			X	X
7. Review the PC prototyping software		X	X	X					
8. Pursue improved testing tools						X	X		
9. Pursue estimating packages								X	X
10. Increase user system knowledge	X								

situation, the major items to pursue were the first three: monitor for consistent sub-second response time, formally implement the Joint Application Development (JAD) concept and pursue COBOL code generators.

It may be surprising that Fourth-Generation Languages (4GLs) are not on the list for consideration. After years of analysis, 4GLs are recommended exclusively for end-user "programming" of simple nonintegrated systems. Conversely, the systems development staff implements large, tailored, integrated complex corporate systems, not easily programmed or maintained with a 4GL.

Myth 6: Zero Defects Is the Objective

Attempting to maximize a productivity measure is not always logical since a balance of activities may produce a more optimal organizational result. For example, attempting to minimize support on systems could eventually produce an adverse effect, that is, inability to support the system.

Consider the objective for Quality Measure Number 1, maintain or increase the LOC per UCR. For some installations, improving operational quality might be the primary objective of the performance measurement program; however, our system operational quality was considered more than acceptable over the past few years. For example, in 1987, only one UCR occurred per year for every 14,012 LOC installed (not including project installations).

It may sound like poor management to admit an acceptable error level and you might propose an absolute zero defects objective. However, if too much pressure is exerted by a ratio, negative ramifications may result, such as delayed implementation dates because of unnecessary repetitive testing. Obviously, this logic is only appropriate when UCRs (program bugs) are under control. Also, the objectives must be attainable.

Remember, when completing projects on time is the dominant goal, poor corporate decisions may result. For example, delaying known system enhancements may keep the project on schedule but cost significantly more to add after installation.

Myth 7: Competing Management Methodologies Are Replaced

The management process is generally defined as five activities: plan, organize, staff, direct and control. Controlling involves monitoring results, usually through a reporting system. Our nine ratios are obviously directed at the control process since we implemented the concept dealing strictly with numeric information. Other management methodologies, such as Critical Success Factors (CSF) or Management By Objective (MBO), primarily address the control function but are less quantitative than performance measures. They are more adaptable to planning, organizing, staffing and directing activities.

Some would propose performance objectives as the primary tool; however, this is not recommended. Rather, performance measures should be used as an ancillary management tool.

Myth 8: Corporate Benefits Are Directly Linked

One surprising observation: benefits to the corporation were not included in any of the calculations. By improving performance, there is an overall benefit to the company but the value for individual projects is not input to the nine measures. Accumulating benefits for modern on-line systems is dependent on an organization's ability to quantify intangible contribution, a difficult task in a dynamic environment. However, if this is the objective, other approaches relying on management judgment are preferable to performance measurement.

Implementing a performance measurement methodology obviously requires resources and management attention. IBM states that five percent of the project budget should be directed toward measurement activities. In our case, almost all the work was accomplished via line management, those managers responsible for implementing the measures. If an organization is growing rapidly and involved in a number of new projects, it may not be an appropriate time to install a formal measurement program. For us, the basic factors such as LOC, function points and unusual condition reports had been monitored for a number of years. Thus, the advantage was formalizing our approach and communicating it to everyone in the organization. Over time, this formal approach will help document productivity improvements.

Professor Skinner, author of "Productivity Paradox" (*Harvard Business Review,* July-August, 1986) disclosed a potential disadvantage of productivity measures: "Managers under pressure to maximize productivity resist innovation." The objective must be to install the best overall corporate solution (application packages, end-user computing, 4GLs) versus maximizing ratios. The function point measures complement LOC measures, providing a balance and encouraging innovation.

Conclusion

Organizational and environmental factors dictate the value and type of measures. In our situation, a large number of projects provided stability in the ratios when all groups were included. Since most of the data needed were available, the cost of implementing performance measures was minimal; however, years were spent defining function points, collecting unusual condition reports and automating LOC data. The measures contributed to our organizational effectiveness and complement other management activities.

In retrospect, the process of implementing performance measures parallels other management tasks: define terms, set objectives, establish a plan, collect data and monitor results. Line management should drive the process although some staff support may be required. Simplicity in an operational program is important; if data collection is expensive or too many measures are defined, the benefits are diminished.

Managers must decide if their organizations are ready for performance measurement and, if so, how extensive a program is needed. ≜

ABOUT THE AUTHOR

James R. Johnson is Director of Data Center Operations for Hallmark Cards in Kansas City, MO. With more than 20 years data processing experience, Johnson is also a well-known lecturer and frequent contributor to Datamation *and* Computerworld. *This article is excerpted from Johnson's book titled, "The Software Factory: Managing Software Development and Maintenance," published by QED Information Sciences, 170 Linden Street, Wellesley, MA 02181, (800) 343-4848.*

ADABAS 5: more than more performance.

ADABAS 5's leading position as a high-performance DBMS has been verified once again by a series of standard, fully scaled benchmarks. Each was conducted on a National Advanced Systems AS/EX™ 100 (equivalent in power to an IBM 3090 500S). And audited by Coopers & Lybrand.

In the **standardized TP1 debit-credit benchmark,** ADABAS 5 worked with a data base containing **1 million accounts.** The results: a **record-breaking 388 transactions per second** (tps) – 99.3% serviced in under one second!

For the first time, an authentic **ET1 debit-credit benchmark** was conducted with ADABAS 5 managing **10 million accounts.** The results: **167 tps** (from terminal in/to terminal out) – 99.5% serviced in under one second!

These figures represent major savings in time and money. But they're not surprising – at least not to the thousands of organizations which already use ADABAS 5.

They expect more performance. Plus the benefits that come from 18 years' experience in DBMS technology: – relational functionality which allows adaptable data structures and fast information retrieval based on multikey criteria – document management and free text retrieval – 24 hour processing in a fully integrated DB/DC environment – portable applications across various hardware and operating systems – distributed processing – entity relationship data bases with recursive retrieval functions for knowledge-based systems.

Discover how much more profitable your DP organization can be. **Conduct your own ADABAS 5 benchmark,** using your own data and application profile in your own production environment. The facts will speak for themselves.
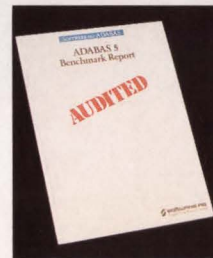
Demand the performance and functionality only ADABAS 5 can offer.

To order your free copy of the ADABAS 5 Benchmark Report, call toll-free: 1-800-843-9534 today.

SOFTWARE AG
Programming Business Success

™ AS/EX is a trademark of National Advanced Systems Corporation.

Optical Disk Technology

OD Technology Comes Of Age

By Fred Schuff

Optical Disks (OD) were introduced about 15 years ago. However, only in the past five years has OD technology matured enough to be applied to the requirements of data processing. *Optical* storage is associated with the concept of large data capacity at costs which are significantly less than standard magnetic media (tape or disk). The OD is a high density storage solution for data that is infrequently used but needs to be kept on-line. The media is reliable (one error in 10 to 12 bits) and cost effective. ODs can be used to off-load large tape libraries and have data in an on-line or near-line (manually retrievable) system that ensures fast, error-free access.

This article is intended to provide a general overview of OD technology with the primary focus on its applicability to mainframe computers rather than micro or mini-computers.
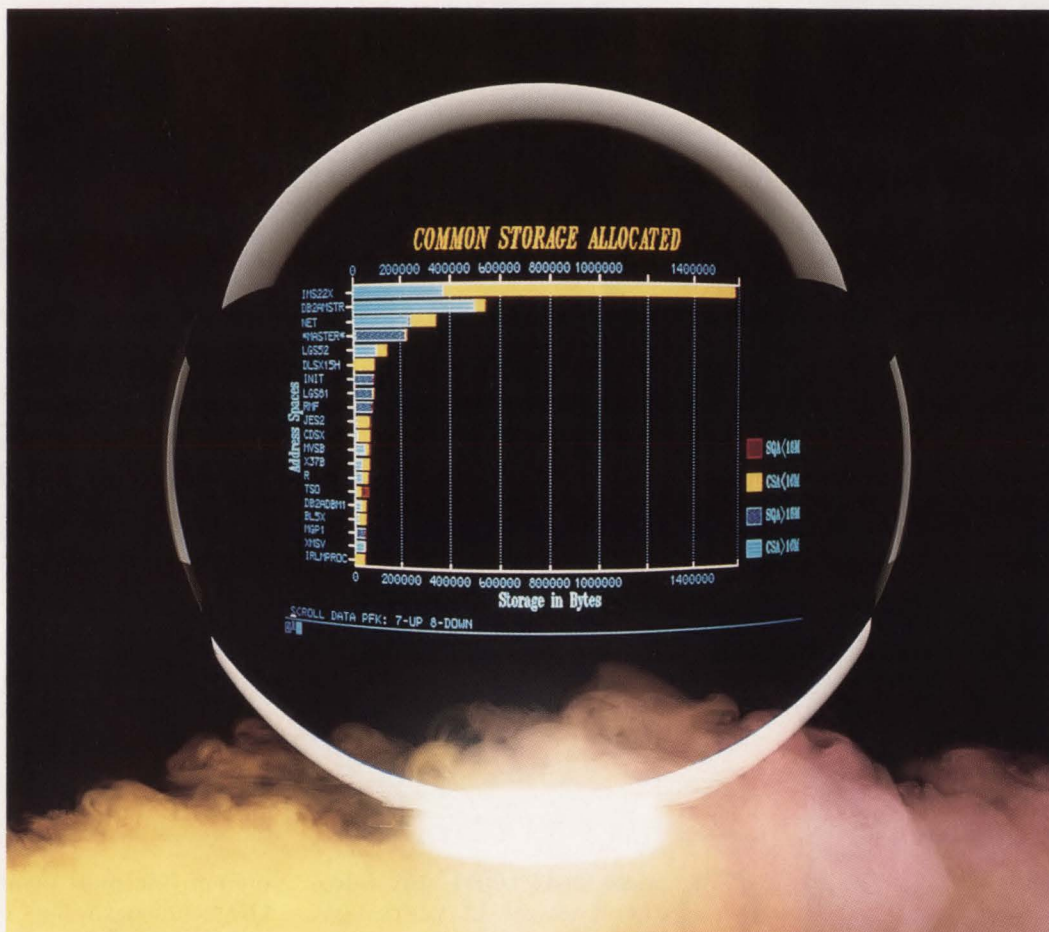
The basis of OD technology refers to using a LASER (a highly concentrated light beam or Light Amplification by Stimulated Emission of Radiation) to read/write data that is recorded on a thin film surface. The recording surface is usually record-shaped and enclosed in a ceramic, plastic or glass casing to form the total recording media platter (see Figure 1). At present, the most common contact with OD technology is the Compact Disk (CD) and Compact Disk — Read Only Memory (CD-ROM). These CDs are used for audio recordings, video disks and static data files (mostly personal computer based).

CD-ROM technology uses a master disk to *press* out data platters much like normal records are mass produced. The laser beam can write data permanently on the OD using what is called WORM (Write Once Read Many Times) methodology. When data is written it is permanently *burned* into the media but it can then be read many times. There are more than 250 companies involved with OD products in either a manufacturing or value-added software capacity. Most OD systems consist of OD components such as OD drives, juke boxes, robotic pickers, OD platters, controllers and so on. They are packaged on a micro- or mini-computer platform and driven by customized software.

New OD recording techniques and enhanced hardware components are being announced constantly. Each year shows a significant change in the state-of-the-art technology. Each month there are new announcements and predictions in the optical storage field. Current improvements include increased capacity of devices and an

erasable media and recording methodology. Even by the time this is read, some of the capacities noted will most probably be out-of-date.

A predominant number of the current OD systems and components use high capacity (2.6GB per 12″ platter or 600-1200 MB per 5¼″ platter) optical storage media (see Table 1).
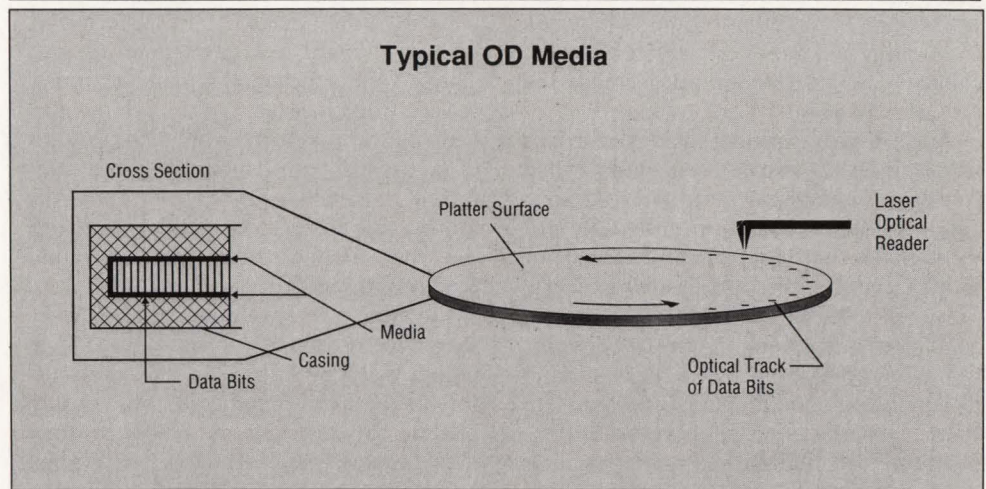
The benefits of *non-erasable* OD technology are:

- The capacity to store large amounts of data on a small surface
- The data is permanent, that is, a permanent track of data written is maintained forever
- Media lifetime is in the 30 to 100 year range rather than the traditional five to seven years for magnetic tape
- Laser technology does not involve the close proximity of read/write heads to storage surface so *head crashes* are eliminated
- The media is removable for greater total capacity and data transportability at a greatly reduced cost.

Recording is commonly done in a helical spiral around the platter surface. The data is usually written at a constant speed referred to as Constant Linear Velocity (CLV) at all locations on the platter surface. The density of the recording mechanism determines the amount of data that can be stored on the surface along with the transfer rate. OD devices are about equal in density of data stored on one track. Much like the magnetic DASD devices, a relatively constant rotational speed is maintained. Many OD devices rotate at 1800 rpm, about one-half the speed of magnetic disks. Tracks are often placed closer together on OD devices by a factor of 10 compared to magnetic disk devices.

Furthermore, OD platters are remova-

FIGURE 1

Typical OD Media

Cross Section · Platter Surface · Laser Optical Reader · Media · Casing · Data Bits · Optical Track of Data Bits

TABLE 3

Approximate Cost/GB Comparison

3380 DASD Double Density (2.4GB)	Direct Attach 4-12″ OD (10.4GB)
$14,500 / GB	$9,600 / GB

Juke Box 40-12″ OD (104GB)
$4,800 / GB

ble from the OD read/write devices. Because there can be many platters per OD reader in an OD system, this allows packaging large storage capacity in a small area. OD readers only require a short time (four to 10 seconds) to spin-up a platter to processing speed. These systems will encourage multi-terabyte (1,000GB or 1,000,000,000,000 bytes) on-line or near-line, library capacities (see Table 2).

Erasable ODs appeal to users who want OD technology for high volume DASD storage devices. They want to use OD devices like traditional magnetic storage devices. The limiting factor in this direct replacement is the slower access speed for

OD devices compared to magnetic media. This gap is expected to remain for the next few years. Additionally, this *erasable* technology eliminates the *permanency* feature from the OD stored data.

The data in Table 2 is being affected by current changes in other technology as well:

- New 3480 type tape cartridges (18 track) are now being announced with recording densities of around 32K BPI which will quadruple the capacity of the Automatic Tape Library Systems (ATLS) and projections are for recording densities around 100K BPI
- New 3380 type DASD is expected with greater storage capacity, faster access times and reduced cost/MB like the IBM 3380J and K will then be followed with the 3390 series that again multiplies capacity and transfer rates
- New OD devices have been announced with 14GB/platter on a 14″ WORM platter (up from a 6GB capacity) and 650MB on a 5½″ erasable device (up from 30MB capacity)
- New ATLS systems offer competitive data access times within enclosed cabinet devices (50-2000 recording devices such as platters or tapes using a similar robotic access mechanism, see Table 3).

With systems using multiple platters per read/write device, several robotic mechanisms exist to reduce the intervention to switch and manage the platters. One of the most common is the juke box for storing a number of platters and has a robotic arm for fetching, loading and unloading the platters into the optical drives (see Figure 2). Most units of this type are designed with 50 to 256 platter capacity. Many units can be expanded by adding more juke box slots or entire juke box

TABLE 1

Equivalent OD Capacity Comparison

OD Device	2400′ ½″ Tape 1600bpi / 6250bpi		3480 Cart. Tape	3380 SD DASD	5½″ Floppy
12″ Platter (2600MB)	68	17	15	4.2	7200

TABLE 2

Terabyte Capacity Comparison

		12″ OD	2400′ ½″ Tape 1600bpi 6250bpi		3480 Cart Tape	3380 SD DASD	5½″ Floppy
1-Terabyte	=	384	26000	6700	5900	1600	2800K

units. Maximum sizes of these systems are currently in the four to six terabyte range. Improvements and innovations in juke box technology will most certainly increase the capacity of these units and decrease the cost per GB of total storage capacity.

Additionally, optical systems/devices usually interface in a different manner than standard mainframe peripheral devices. The physical connection is usually not through the same control unit as DASD or tape devices but through another interface (network, RS232, SCSI or IPI protocol, MVE Bus and so on). A network connection allows one Optical Storage System to be connected to multiple, different hosts to share data. The other connections are implemented to interface to micro or minicomputer systems. More recent products have introduced direct channel-connected systems (using 3480 and 3420 control unit emulation). The newest addition is a DASD control unit emulator. The trend is toward direct mainframe host connections becoming available especially in the 1989 to 1992 time frame.

There are two characteristics that hinder Optical Storage Systems. First, due to the

F I G U R E 2

Sample Robotic Configuration
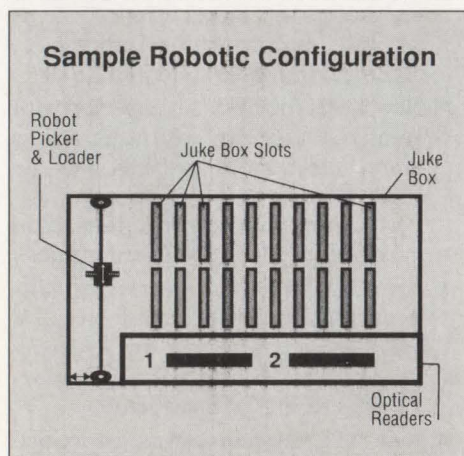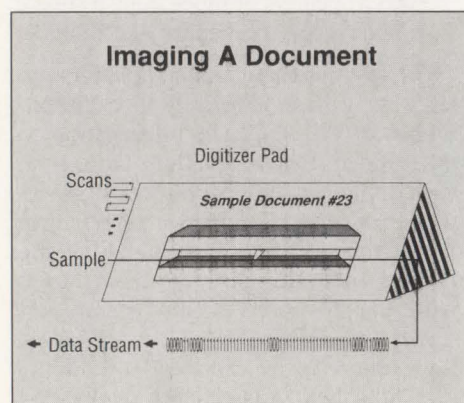
nature of the optical media, there is a 300 to 350KB/second read rate limiting the amount of data that can be transferred to/from the device in a given time interval. The second is the small number of optical readers implemented to support large juke box systems that further limits the volume of data which can concurrently be transferred to/from an OD system. Thus, the more ideal application of the technology must be oriented towards capacity rather than access speed.

Previously, it was noted that OD research has developed optical storage technology that temporarily modifies the structure of the recording media in the optical platter. This provides the ability to erase data because the laser does not permanently *burn* a hole in the surface. Some of this research has reported the ability to multiply recording densities by several orders of magnitude and promises greatly expanded storage capacities for OD devices.

An Optical Disk Application

Document imaging is an ideal application for OD technology. Imaging is storing a *digitized image* of a document for reference or retrieval at some later time. A digitized image is a document that is broken down into a matrix of dots to represent the document. Storage of the matrix requires a large amount of storage space because the density of the matrix required for many documents generates a large volume of data (see Figure 3). The original documents come from many sources such as medical records, insurance records, bank transactions, engineering drawings, X-rays and so on. The increased capacity of the OD technology, along with the long media life, makes the process feasible and economical.

Many imaging systems are developed as stand-alone systems. They are comprised of micro or minicomputers connected to the OD storage system components. Data input and retrieval is isolated to this auxiliary system and exists outside of the mainframe environment. The limitations of speed and connectivity of the OD components have much to do with this segregated environment.

Conclusions

The OD market is in a state of rapid change and has not yet found a firm technology (media, interfaces, format and so on) to implement. Standards have not been set across the various manufacturers. New devices are rapidly coming

onto the market and old ones become obsolete just as fast. The technology for storage devices using magnetic technology is also going through significant changes and improvements.

OD technology has a number of potential enhancements in development:

- Read both sides of the platter simultaneously using a dual-head read/write mechanism to double the transfer rate
- Interface the controller for the optical drive(s) through a standard, mainframe-compatible controller to let the OD system act like a standard DASD device
- Implement four to eight optical drives connected through a control unit with multiple channel paths
- Increase buffered storage area in the optical drive controller and denser recording methods to improve transfer rates.

In my opinion, optical devices and OD systems will become more numerous as the systems evolve and users become aware of this technology. This growth will be spurred by the stabilization of the optical media, standardization of the interfaces and the realization that optical will be around for the long term. Keeping track of the changes and charting the growth path is perhaps the major obstacle.

In its current state of rapid evolution, OD users are both encouraged and confused at the same time. The challenge is to have the insight to get your feet wet and avoid getting soaked. ≜

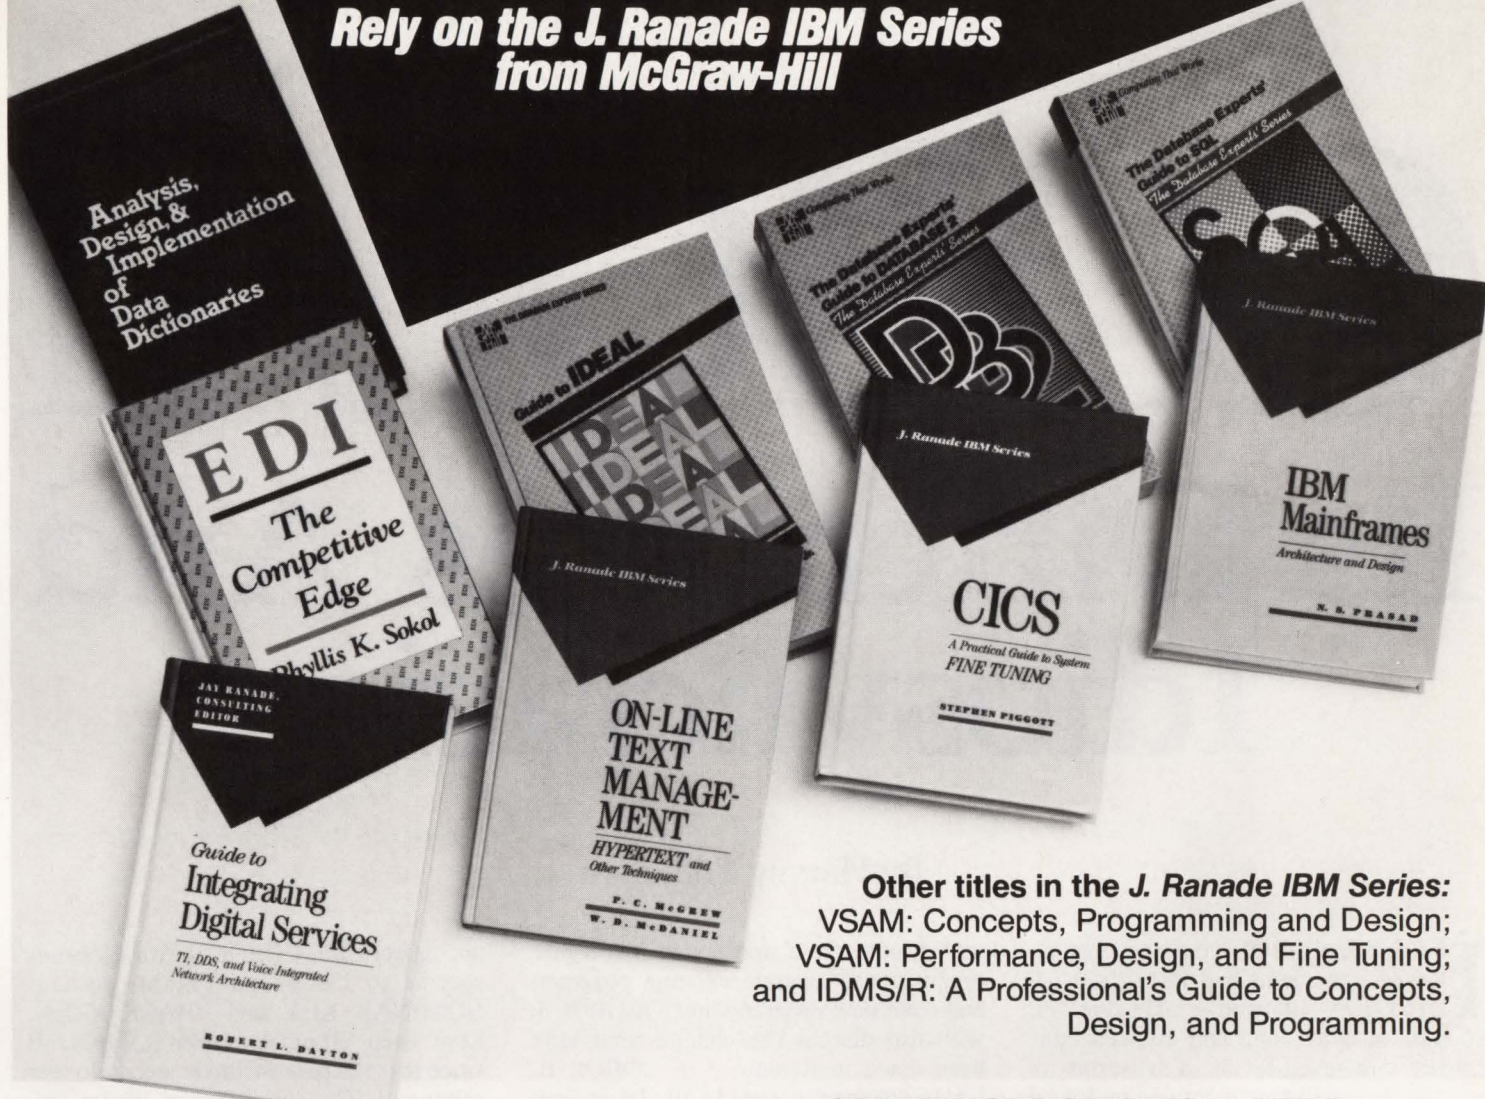*ABOUT THE AUTHOR*

Fred Schuff is a Systems Programming Consultant for Systems/Software Engineering, a systems and management consulting firm in Wayne, PA. He has worked in a number of different industries in systems programming, database support and application design and consulting. S/SE, 940 W. Valley Rd., Suite 1603, Wayne, PA 19087, (215) 341-9017.

F I G U R E 3

Imaging A Document

COBOL II
Its Differences And Idiosyncracies

By Harvey Bookman

Imagine a COBOL compiler that allowed more than one PROCEDURE DIVISION, allowed a table that was 16 million bytes long and allowed table entries with seven levels of subscripts or indexes. It is something that you no longer have to imagine because it actually exists. It is Release 3 of the IBM COBOL II compiler that was made available at the end of 1988.

Future programmers will have plenty of COBOL options to deal with. For present VS COBOL users, until Release 2 of COBOL II, the *COBOL II Application Programming Guide* made certain to explain and to chart the differences between VS COBOL and COBOL II. This was to ease conversions and to assist programmers in quickly finding the new features. Beginning with Release 3, these charts are gone. So many changes were made from Release 2 to Release 3, IBM now needs pages of charts to show the differences between the COBOL II releases themselves.

The COBOL II updates fall into one of three areas: old features that are eliminated, old features that are updated and new features that are added. This article

will discuss these updates so that a programmer can migrate existing programs and code new programs in COBOL II. It will also discuss the changes that have been made in Release 3 of COBOL II. Unless otherwise noted, all discussions will be about Release 2 of COBOL II which is the version most widely used at the current time. Changes that are mainly pertinent to CICS will be discussed in the CICS section rather than in the general enhancement section of this article.

Eliminated Features

After issuing Release 1 of COBOL II, IBM saw that elimination of features is not taken lightly by users. Release 1 eliminated exponentiation, floating point data, complex occurs depending on usage and the APPLY WRITE-ONLY option. A number of features, including the four just mentioned, were reinstated in Release 2 of COBOL II. IBM termed the enhancements as, ''functions identified by early users as key inhibitors to conversion of applications from OS/VS COBOL to VS COBOL II.''

There are, however, a number of features that existed in VS COBOL that are

not supported in COBOL II. Some clauses specific to ISAM and BDAM, such as NOMINAL KEY and TRACK-AREA, have been eliminated from COBOL II since these access methods are no longer supported. The entire Report Writer feature has been removed. While not commonly used, its removal does present a problem for converting those existing programs that do use the feature. Use of the WHEN phrase in the SEARCH statement now has added restrictions. If a table contains a key named KEY-FIELD and you wanted to code a SEARCH statement to check if it was equal to WORKING-STORAGE field MONTH-NAME, VS COBOL allowed either WHEN KEY-FIELD IS EQUAL TO MONTH-NAME or WHEN MONTH-NAME IS EQUAL TO KEY-FIELD. COBOL II only allows the first clause since it restricts the first field named in a WHEN condition to a key field and restricts the field being compared to as a non-key field.

The EXAMINE and TRANSFORM verbs have been removed. The INSPECT statement, which previously existed, now performs the function of the three verbs. To enable the INSPECT statement to han-

STROBE . . .
The Key To
Well-Tuned
Applications

**You know your peformance
is great when:**

- your customers expect and get fast
  online response time
- your management reports regularly
  meet deadlines
- you confidently schedule batch
  processing before online start-up
- you can bring up new applications
  without straining your resources

**Even if you have already tuned your
MVS system, you can assure
greater performance with STROBE,
the premier application
tuning product. STROBE will:**

- furnish detail not available from
  system tuning software
- attribute resource use to specific
  causes within application programs
  and subsystems in any address space
- show CPU use by source program
  statement or system service function,
  and show disk unit access time by
  cylinder within data set or index

**STROBE measures and reports on
programs written in COBOL, PL/1,
Fortran, and Assembler,
operating in all MVS system
environments, and in CICS, IMS,
DB2, IDMS, and other subsystem
environments.**

CIRCLE #33 on Reader Service Card

PROGRAMART
1280 MASSACHUSETTS AVENUE
CAMBRIDGE, MA 02138
617-661-3020 FAX: 617-864-6558

P · R · O · G · R · A · M · A · R · T

dle all of the functions of EXAMINE and TRANSFORM, Release 3 now allows an unlimited number of TALLYING and REPLACING phrases, contains a CONVERTING phrase and allows multiple BEFORE and AFTER phrases.

EXHIBIT has been removed and its function is to be performed by the DISPLAY verb that existed in VS COBOL as well. In Release 3, the DISPLAY statement has been enhanced. It now allows the WITH NO ADVANCING clause that suppresses skipping to the next line when data is written. While the READY TRACE and RESET TRACE statements do not cause compilation errors, they do not serve any function in COBOL II. The MOVE statement options of CURRENT-DATE and TIME-OF-DAY have been removed and equivalent information can be obtained using the DATE and TIME options of the ACCEPT statement. In Release 3, the ACCEPT statement now has a DAY-OF-WEEK option which is self-explanatory. The REMARKS paragraph and NOTE statement have been eliminated; standard comments are to be used instead. The ON statement (for example, ON 1 GO TO FIRST-TIME) is no longer supported.

While not actually an eliminated feature, there are new "reserved words" in COBOL II. This can cause errors in programs that would have compiled correctly. For instance, if a program used a paragraph name of INITIALIZE or a DATA DIVISION field name of EVALUATE, these words will cause compilation errors since they are now COBOL II reserved words.

While the eliminated features do not cause any real restrictions for new programmers, they do cause considerable problems in converting VS COBOL programs to COBOL II. While the changes are not difficult, they are tedious and not preferred work by most programmers. Changes to CICS programs include all standard changes plus those specific to CICS.

There are, however, a number of conversion software products available to perform the conversion function for CICS as well as batch programs. IBM has its own COBOL and CICS/VS Command Level Conversion Aid. Other conversion programs include the CA-Optimizer from Computer Associates (Garden City, NY) and MHtran-2 from Prince Software Products (Mahwah, NJ). Any conversion aid should convert CICS as well as batch programs. It should also handle any software specific statements your company uses such as IDMS code and run with software products you have like VM, Panvalet or Librarian. Additionally, it should clearly flag any code that it is unsure how to convert and should ensure that any program not fully converted will compile with errors. This will prevent inadvertent runs of a partially converted program.

Updated Features

The next area of changes to be discussed will be the COBOL features that have been updated. The FILE STATUS has been expanded with the extended VSAM feedback area. In addition to the two byte file status that was previously used, an additional six byte VSAM status is available. The information consists of three two-byte binary fields and is only filled in when the file status is not zero. The first two bytes contain the VSAM return code (register 15). When the return

code is not zero, the following two fields are filled in. They contain the VSAM function code and the feedback code.

The internal processing of the OCCURS DEPENDING ON clause has been changed which may yield more efficient runs. COBOL II will perform the length computation of a data item defined with the OCCURS DEPENDING ON clause at the time the data item containing the clause is referenced. In VS COBOL, when the value of the object field of the clause was updated, the length of the data item was computed, regardless of whether or not the variable length data item was used. This meant that as you were continuously adding to the number of occurrences in the table, the table length computation would be computed each time. The VS COBOL length computation would be considerably more efficient if a variable length table was created once each time the program ran and was then repeatedly used. The new method is more efficient when the table length is constantly changing. It should be mentioned that while there is no way to increase the efficiency of the COBOL II method, the VS COBOL method could have been programmed efficiently by incrementing a data item that was not the object of the OCCURS DEPENDING ON clause and then moving this item to the object. The length computation would only be done when the move was done, not when the incrementing took place.

Conversions of floating point numbers have been enhanced to be more accurate in COBOL II. Be aware that results may actually differ from VS COBOL when floating point numbers are used. This may make it difficult to use automated procedures to do regression testing to verify that the VS COBOL and COBOL II versions of a program produce the *same* results.

The CALL statement in COBOL II allows the use of the BY CONTENT phrase to allow constants to be passed to another program and the BY CONTENT LENGTH OF phrase to allow the length of a data item to be passed to another program. Release 3 allows the BY CONTENT phrase to be used with data names as well as constants. If any data passed in the BY CONTENT phrase gets modified in the called program, its value is not affected in the calling program.

The PERFORM statement has often caused programmers confusion since conditions in the UNTIL clause were examined before the PERFORM was executed,

FIGURE 1

```
DATA DIVISION.
WORKING-STORAGE SECTION.
01    GENDER                              PIC X(6).
01    AGE                                 PIC S9(3).
PROCEDURE DIVISION.

       . . .
       IF GENDER = 'MALE' AND AGE > 25
           PERFORM RATE-1
       ELSE IF GENDER = 'MALE' AND AGE NOT > 25
                 PERFORM RATE-2
             ELSE IF GENDER NOT = 'MALE' AND AGE > 25
                    PERFORM RATE-3
                 ELSE PERFORM RATE-4.
```

```
DATA DIVISION.
WORKING-STORAGE SECTION.
01    GENDER                              PIC X(6).
01    AGE                                 PIC S9(3).
PROCEDURE DIVISION.

       . . .
       EVALUATE GENDER = 'MALE' ALSO AGE > 25
           WHEN TRUE ALSO TRUE
                PERFORM RATE-1
           WHEN TRUE ALSO FALSE
                PERFORM RATE-2
           WHEN FALSE ALSO TRUE
                PERFORM RATE-3
           WHEN OTHER
                PERFORM RATE-4
       END-EVALUATE.
```

not after it. COBOL II now allows the code to specify either WITH TEST BEFORE (the default if not coded, compatible with VS COBOL) or WITH TEST AFTER (which is easier to follow and will probably be used by the next generation of programmers). The PERFORM statement also allows the performed procedure to be coded in-line; the code to be performed is put directly after the PERFORM statement and the paragraph name to be performed is not coded.

Some error processing done by the compiler has become considerably more user friendly. For example, when you have an erroneous GO TO in a program, you might fall through the code beyond the last line in the program. In VS COBOL, this produced a 519 User ABEND. When the same condition occurs in COBOL II, an error message is produced that informs the programmer that the execution preceded beyond the last line of the program.

Not only have the error messages produced by the compiler changed, but also the letter that indicates the severity has changed as well. The errors used to be, from least to most severe, W (Warning, return code of 4), C (Caution, RC = 8), E (Error, RC = 12) and D (Disaster, RC = 16). The severity levels are now I (Informational, RC = 0), W (Warning, RC = 4), E (Error, RC = 8), S (Severe, RC = 12) and U (Unrecoverable, also called terminating, RC = 16). Notice particularly that the term ''Error'' message has changed in meaning from a severe

error to what used to be called a cautionary error.

A nice advantage in COBOL II is that a number of compiler limits have been expanded. A 77 level or 01 level may now be up to 16MB in length. Table entries may have a maximum length of 8MB. Elementary fields, previously limited to 32K, now also can be up to 16MB. The LINKAGE SECTION which supported up to 255 BLL cells (01 levels) now has no fixed upper limit.

COPY statement processing has been enhanced to allow nested COPYs, that is, a copy member may now contain COPY statements. This is allowed in many other languages and is a welcome addition to COBOL.

New Features

The INITIALIZE statement allows the initialization of an elementary or a group item with a single statement. It does not require that each individual elementary field within a group be listed. Each alphabetic, alphanumeric or alphanumeric-edited field is cleared to spaces while numeric, numeric-edited and floating-point items are cleared to zeros. The statement also allows specific values (if other than spaces and zeros are desired) to be used, as well as allowing only certain types of data (alphanumeric, numeric and so on) to be initialized.

The INITIALIZE statement yields a nice advantage. If a data division copy member or a record description is modi-

fied with new or deleted fields, the INI-TIALIZE statement usually does not have to be updated. When the program is recompiled, the new or deleted field will be recognized by the compiler. The only negative aspect of the statement is that a programmer can code a more efficient initialization routine. However, this would be at the expense of ease of maintenance.

The EVALUATE statement allows enhanced condition testing to simplify decision logic. In many cases it can replace nested IF statements. Figure 1 shows how an IF statement would look when converted into an EVALUATE statement.

The CONTINUE statement is a null statement that does not generate executable code. It can be used wherever the compiler syntax allows an imperative or conditional statement. Examples of its use are in an IF statement instead of a NEXT SENTENCE or in an EVALUATE statement where NEXT SENTENCE is not allowed.

Figure 2 shows the explicit scope terminators which have now been added to the compiler. What they basically do is show the end of a particular statement. While they are often used to improve program structure, they are valuable to turn conditional statements into imperative statements, allowing the statement to be nested. Figure 3 shows how a READ statement with an AT END clause can appear in the middle of nested IF statements — something that was not possible in VS COBOL.

FIGURE 2

Explicit Scope Terminators

```
END-ADD
END-CALL
END-COMPUTE
END-DELETE
END-DIVIDE
END-EVALUATE
END-IF
END-MULTIPLY
END-PERFORM
END-READ
END-RETURN
END-REWRITE
END-SEARCH
END-START
END-STRING
END-SUBTRACT
END-UNSTRING
END-WRITE
```

Compiler Option Changes

The compiler options have changed considerably. The debugging facility has been revamped. This means that the COUNT, STATE, FLOW and SYM-DUMP options no longer exist. The OS-DECK option is not necessary and has been removed. The LISTER feature that included the LSTONLY and LSTCOMP options is no longer supported. While the listing produced by this feature was quite nice, the feature was not widely used. Also, the COBOL II compilation listing is improved over VS COBOL and contains some of the details produced by the LISTER feature.

The compiler listing of the Assembler language expansion of the source code now shows fields by their real names, not names like DNM = 1-245 that VS COBOL used. You, therefore, no longer have to cross reference the field names back to the Data Division Map. The error messages also include actual field names which make it considerably easier to fix compilation errors while working on a terminal and no listing is available. The data name cross reference now shows whether each use of a data field caused a modification of the field or is just a reference of the field. The procedure name cross reference shows whether a reference is due to a PERFORM or a GO TO.

I find it extremely annoying that, for no apparent reason, the names of a number of the COBOL II compiler options have been changed from their VS COBOL names. The SXREF option is now XREF and the VS COBOL XREF option (that produced a non-sorted cross reference) is no longer supported. A table of the main options that have had their names changed can be found in Figure 4.

A number of new compiler options have been added, the most important of which are SSRANGE, RENT, PFDSGN and FASTSRT. The SSRANGE option, which should have been included in COBOL years ago, checks subscripted areas, indexed areas and OCCURS DEPENDING ON areas to ensure that they do not exceed the storage size allotted to them by the compiler. Invalid values can cause storage overlays and/or protection exceptions (0C4s). For instance, if a one-dimensional table entry occurs 10 times and a program attempts to reference the 15th entry, an error message will be displayed. There is also an installation option that will terminate a program that gets this error.

However, the implementation of the SSRANGE option is somewhat bizarre. If it makes no sense to you, you are in excellent company. When the SSRANGE option is used in a compilation, each usage of a subscripted data name generates a call to a COBOL subroutine (an incredibly inefficient method). Furthermore, if a subscript is referenced on multiple statements within a paragraph and is not changed, the COBOL subroutine will be continually called to revalidate the subscript.

If it still makes sense to you, realize that the full extent of the absurdity of the SSRANGE option is yet to be described. When the INITIALIZE statement is used to clear a table, the code generated by the compiler continually calls the subscript range check routine to ensure that it is not

FIGURE 3

```
DATA DIVISION.
WORKING-STORAGE SECTION.
01    WS-RECORDS-READ                        PIC S9(5) COMP-3 VALUE ZERO.
01    WS-GET-REC                             PIC X.
PROCEDURE DIVISION.
      . . .
      IF WS-GET-REC = 'Y'
          READ FILE1 AT END
              IF WS-RECORDS-READ = ZERO
                  PERFORM NO-RECORDS
              ELSE PERFORM SOME-RECORDS
              END-IF
          END-READ
      ELSE PERFORM DONT-READ-FILE
      END-IF.
```

FIGURE 4

VS COBOL Compiler Option	Equivalent COBOL II Compiler Option
SYNTAX	NOCOMPILE
CSYNTAX	NOCOMPILE(S)
FLAGE	FLAG(E)
FLAGW	FLAG(W)
PMAP	LIST
DMAP	MAP
LOAD	OBJECT
CLIST	OFFSET
SXREF	XREF

outside of the table's limits! You must also realize that each subscript in a multi-subscripted table element is not checked for validity, only the resulting address. For instance, if a three dimensional table has a maximum subscript of (12, 10, 3) and you use a subscript that generates a value (1, 1, 6), no error will occur. Although one of the subscripts is outside its proper range, the address computed using all three subscripts will fall well inside the table limits.

The RENT option makes a program reentrant. This allows one copy of a program to be shared by many users, such as is the case with CICS programs. It is implemented by copying the areas of the COBOL program that get modified into a GETMAINed area. This includes the TGT and WORKING-STORAGE section. When the RENT option is used in conjunction with the DATA option, you may specify that program storage be acquired from unrestricted storage, either above or below the 16MB line.

PFDSGN is a new option that can make a program run more efficiently. It instructs the compiler that: your numeric fields have valid signs; and that the signs are F for unsigned fields and C or D for signed fields. Programmers often do not realize that A, C, E and F are positive signs while B and D are negative signs. None of these values will cause data exceptions (0C7s). However, when packed decimal calculations are performed on an IBM mainframe, the preferred signs of C and D are placed into the results. The NOPFDSGN option tells the compiler that numeric fields in a program do not necessarily have the preferred signs of F, C and D.

The PFDSGN option may not seem to be problematic until you realize that neither the PFDSGN nor NOPFDSGN will generate executable code with the same idiosyncracies as the VS COBOL compiler. In Release 3, IBM has removed the PFDSGN and NOPFDSGN options. Their functionally equivalent counterparts in Release 3 are NUMPROC(PFD) and NUMPROC(NOPFD). In addition, as a migration aid, NUMPROC(MIG) has been added to request COBOL II to perform similar sign processing to VS COBOL.

The FASTSRT option allows a program to process a sort faster. This is accomplished by having the sort run external to the program rather than having every record funnel through the COBOL I/O routines. It, therefore, only works when a SORT statement contains either a USING

or a GIVING clause. If the SORT statement contains an input procedure and an output procedure the FASTSRT option does not affect sort processing. In this case, FASTSRT does not cause an error if it is used as a compilation option. However, you do receive an information message that the fast sort was not done.

COBOL II and CICS

COBOL II has lifted a number of restrictions that existed under VS COBOL CICS. The INSPECT and UNSTRING statements may now be used. STOP RUN may be coded and is equivalent to coding an EXEC CICS RETURN END-EXEC. The CALL statement can now be used to issue a static call from a COBOL II program to another COBOL II program or to an Assembler program. Dynamic calls are still not allowed.

The most dramatic change to coding CICS applications has come in BLL cell usage. The BLL cells are no longer coded and initialized and programs converted from VS COBOL to COBOL II must have their BLL structures removed from the LINKAGE SECTION. To enable the handling of BLL cells in a more structured method, enhancements were made to the COBOL II compiler.

Data fields may now be defined as USAGE IS POINTER. If a field is so defined, the field may not have a PICTURE clause. It generates a fullword containing an address and has an implicit PICTURE of S9(9) COMP. Pointer fields may only be used in the SET statement, in relation conditions and as parameters in a CALL statement. They may not have arithmetic performed upon them. If you wish to do arithmetic on the address contained in a pointer field, you can code a COMP field which REDEFINES the pointer and then perform computations on the redefined field.

The ADDRESS OF clause is an addition to the COBOL compiler. It may be used in a SET or CALL statement. The ADDRESS OF clause is used with data names that are 01 or 77 levels in the LINKAGE SECTION. It is basically a fancy way to refer to a BLL cell. As an example, if TABLE1-PTR is defined as USAGE IS POINTER and TABLE1 is an 01 level in the LINKAGE SECTION, SET TABLE1-PTR TO ADDRESS OF TABLE1 will move the value of the BLL cell used for addressability to TABLE1 to field TABLE1-PTR. When used in a CALL statement, the ADDRESS OF clause can be used to pass an address of

data, rather than the data itself, to another program.

The changes mentioned above led to the change in the SET parameter used in CICS commands. The syntax in CICS using VS COBOL was SET(TABLE1-BLL-CELL). The new syntax is SET(ADDRESS OF TABLE1) or SET(TABLE1-PTR). Programmers no longer have to worry about LINKAGE SECTION fields that are greater than 4096 bytes which required more than one BLL cell to be initialized. This is taken care of automatically by the COBOL II compiler. In addition to the updates to the SET options, CICS programs must also eliminate any other references to BLL cells. Statements like ADD 4046 TABLE1-BLL-CELL2, to address an area greater than 4096 bytes, should be deleted.

For efficiency reasons, the COBOL II SET statement should be used instead of the CICS ADDRESS statement whenever possible. For example, SET ADDRESS OF TABLE1 TO TABLE1-PTR should be coded instead of EXEC CICS ADDRESS SET(ADDRESS OF TABLE1) USING (TABLE1-PTR) END-EXEC.

The LENGTH OF clause is a new feature that allows a program to access the length of any data field. Like address fields, lengths have implicit definitions of S9(9) COMP. While the LENGTH OF clause is defined under the CALL statement in the COBOL II manual, it may be used in any statement that allows a numeric data item, except that it cannot be used as a receiving field or as a subscript. When the length of a table entry is requested, a subscript need not be supplied and the length returned will be of a single occurrence.

The LENGTH OF clause has enhanced COBOL/CICS processing. When the FROM and INTO options are used in an EXEC CICS command, the LENGTH, FLENGTH, FROMLENGTH, MAXLENGTH, MAXFLENGTH, DESTIDLENG and VOLUMELENG no longer have to be coded. However, these options can still be coded if you wish to use a length other than the defined length of a data field.

Note that the ADDRESS OF, the USAGE IS POINTER and the LENGTH OF clauses may also be used by batch COBOL II programs.

Other Release 3 Changes

Not only have a number of new features been added to Release 3 of COBOL II, but also there are a number of state-

ments which execute differently between Release 2 and Release 3 of COBOL II. An option of CMPR2 has been added to Release 3 that when used allows a Release 2 program without modifications to execute under Release 3 with the same results. When this option is used, many new features of Release 3 are not available to the program.

I will briefly mention a number of statements that function differently between the two releases. When a field is checked for ALPHABETIC, lower case letters will now make the test true while in the past only upper case letters satisfied the test. Intermediate results on MULTIPLY and DIVIDE statements used to cause an ON SIZE ERROR whereas they do not in Release 3. When a group item has a subordinate data item containing the OCCURS DEPENDING clause and the field defining the number of occurrences is defined within the same group item, the maximum length of the group item may be used for the move. When this is not the case, the compiler will issue a message indicating that the actual rather than the maximum length is being used. In this case, the field defining the number of occurrences must be set prior to the move to the group item.

Another change in Release 3 is in the state of PERFORM ranges in a called program. In Release 3, all PERFORM ranges are re-initialized upon each entry and PERFORMs will execute as if the program was called for the first time. In Release 2, PERFORMs were left in their last-used state when a GOBACK was issued. While this change normally will not cause any changes in execution, poor programming techniques such as exiting PERFORMed paragraphs with a GO TO or GOBACK before they finish execution, could lead to programs that will behave differently.

Among the most interesting features added to Release 3 is the EXTERNAL attribute. If many programs wish to use the same field, such as a count of the number of records read, they may all share the same field without passing it between programs. As an example, if REC-COUNT is to be shared in many programs, all programs would define the field as '01 REC-COUNT PIC S9(3) EXTERNAL. The EXTERNAL attribute gives added flexibility to programs since it may be used for files as well. No longer do you have to write special programs to read and write files shared by many programs. Files may be defined as EXTERNAL and

the READ and WRITE statements in each program will all access the same file.

Among the most enigmatic enhancements to the COBOL II compiler is the concept of nested programs, the ability to code multiple PROCEDURE DIVISIONs within a single program. Most programmers have enough of a problem coding one PROCEDURE DIVISION in a program. The new implementation allows multiple IDENTIFICATION DIVISIONs and END PROGRAM statements indicating where each nested program begins and ends. While a CALL to a nested program will function with the efficiency of a PERFORM rather than a CALL, which is quite good, I see this nested program feature adding to the complexity of COBOL code, not diminishing it. The complexity includes the understanding of COMMON nested programs which may be called by all programs within the nested structure and GLOBAL variables which may be used by all of the nested programs.

It is interesting that the syntax of a compiler statement got updated because so many programmers kept making the same coding mistake. While ADD A B GIVING C was correct syntax, ADD A TO B GIVING C was not. The TO verb may now be used in an ADD statement with the GIVING option. IBM could have left the statement as a syntax error but since it had only caused a warning message in VS COBOL and was causing a severe error in Release 2 of COBOL II, IBM had to make another change.

The symbols $<=$ and $>=$ may now be used to signify LESS THAN OR EQUAL and GREATER THAN OR EQUAL. The colon ':', formally not a valid character in COBOL code, is now allowed as a separator character. It is used in what the COBOL II manual calls *reference modification*. This feature allows part of a data item to be used without defining the breakdown of the data item in the DATA DIVISION. It may only be used with DISPLAY fields and has the format DATA-FIELD (x:y) where x is the beginning character position to be used in DATA-FIELD and y is the number of bytes. For example, MOVE FIELD1 (3:7) to FIELD2 will move the third through ninth bytes of FIELD1 to FIELD2.

As I have previously mentioned, tables may be defined with up to seven levels of subscripting or indexing. Subscripts and indexes may now both be coded in the same statement to access a data item. Also, subscripts may now be coded in the

form of "SUBSCRIPT + x" or "SUBSCRIPT − x", where x is an integer. This form of coding had previously only been allowed for indexes. The VALUE clause may now be specified on a field that contains an OCCURS clause or a field subordinate to a data item containing an OCCURS clause. When the VALUE clause is used on a data item that occurs many times, each occurrence is set to the value coded.

Numeric edited items may now be de-edited, allowing a numeric edited field to be moved to a numeric field. COBOL code may be in both upper and lower case (lower case was only previously allowed in literals) with the compiler making no differentiation between upper and lower case letters (except in literals). Non-numeric literals may now be specified in hexadecimal.

The AWO compiler option is a nice addition to Release 3. It is used for efficient processing of blocked variable length output records and serves the equivalent function of the APPLY WRITE-ONLY clause. A program can run more efficiently by recompiling it with the AWO option rather than having to change the program itself. Restrictions in VS COBOL for files processed with APPLY WRITE-ONLY have been eliminated in COBOL II.

The TRUNC and NOTRUNC options have been removed and are replaced with TRUNC(STD), TRUNC(OPT) and TRUNC(BIN). The TRUNC(STD) option works the same way the NOTRUNC option used to work. The TRUNC(BIN) option is a new feature, allowing a binary field to be truncated based upon the maximum value in the binary field (halfword, fullword, doubleword), not the maximum value based upon the PICTURE clause. The TRUNC(OPT) option was added for performance reasons to generate the most efficient code, either truncating to the length of the PICTURE clause or the binary field. This option should only be used when a program's logic never allows the value of binary fields to exceed the number of digits defined in their PICTURE clauses.

Unexpected And Concealed COBOL II Differences

Most of the COBOL II changes discussed so far are highlighted in charts in the IBM COBOL II manuals. However, the differences between VS COBOL and COBOL II that are not clearly defined will

AT $1,595 THE MAJOR DIFFERENCE BETWEEN OUR 3270 PRINTER AND THEIRS.
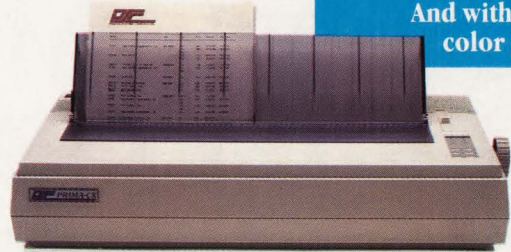
OURS EXISTS.

That's right. Take a minute and flip through this magazine. You won't find another mainframe printer that gives you as many features for anywhere close to this price. Starting at only $1,595, the Prima-CX is designed for industrial-grade applications, not PC-type printing. It attaches to IBM mainframes as well as 4300 and 9370 mid-ranges — via IBM 3174, 3274, and 3276 cluster controllers.

So whether you're producing memos or listings, there's a Prima-CX model (80- or 136-column) that's for you. The IBM-type front panel puts control at your finger-tips. You can feed cut sheets without removing continuous forms. Program tear-off and RPQ options. Cancel print jobs from the printer. Take advantage of the 324 cps draft speed. You'll also enjoy the new optional GDDM and PSS color graphics. All this in the smallest footprint of any coax-compatible printer.

No one else offers a first-year, on-site service contract for only 49 cents.

They'd be afraid to. But, buy a Prima-CX professional printer and we'll give you a comprehensive, *first-year, on-site* service contract for *only* 49 cents!* We can make this incredible offer because our design ensures maximum reliability. Contact us today for complete details.
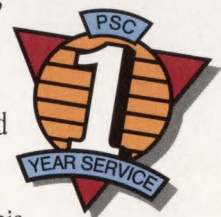
Printer Systems Corporation
9055 Comprint Court
Gaithersburg, Maryland 20877

1-800-638-4041
(1-301-258-5060, in Maryland)

with DB2!

Should DB2 Be Your DBMS Of Choice?

By Joel S. Goldstein

There are presently four times as many DB2 licenses as there were three years ago. Some say it is because DB2 and SQL are major components of the Systems Application Architecture (SAA) and SAA is IBM's strategic direction.

Management has become increasingly aware of the actual value of data and the need to access it without waiting weeks or months for new programs and systems. Relational systems, like DB2, provide the facility for combining and extracting data from many different tables and creating new sets of data to satisfy almost any reporting requirement. SQL provides a high level access to data, and, depending on the requirements and application, may be executed without the need of a surrounding application program. Additionally, one SQL statement can perform the work of hundreds or thousands of database calls for a non-relational DBMS.

The primary benefits of DB2 could be summarized this way:

- Access to data without waiting for application programs to be written
- Significant productivity improvements
- Ease of creating tables and creating new relationships from existing tables
- Ability to modify existing data tables without affecting existing queries
- It is easy for people to visualize a tabular data organization.

Everything has a price and a cost/benefit tradeoff that should not be ignored. In many aspects, this is a bit of deja vu.

Some History

At the beginning of the last decade, IMS was the hot product. Some of IBM's major marketing points for IMS were data independence, increased productivity, elimination of redundant data and the programmer did not have to know the physical structure of the database. All of this was true — if you took the statements with a couple of pounds of salt. The salt was, and still is, the performance implications of logical relationships that eliminated data redundancy; the inefficient processing sequences of programmers that were not properly trained because it was not *necessary;* and the hordes of programmers who immediately became database administrators as soon as they learned how to spell IMS. Similar scenarios now occur daily in the DB2 arena; the basic flaws, such as using the DB2 defaults for Closerule and Validate, occur as often as poor table design and inefficient SQL.

Present and Future

Coding SQL is like writing a sentence using English. Anybody can do it. Right? Well, SQL can be many times more powerful than a DL/1 call to an IMS database. Potentially, one SQL request can retrieve or update *every* row in a table. It can JOIN the data from many tables based on simple or complex criteria. The difference between an incorrectly and a correctly coded SQL statement is the difference between a couple of seconds of CPU usage and many minutes of CPU usage. A simple SQL statement that is properly coded often executes in less than .1 seconds in a transaction processing environment. Obviously, these transactions are not scanning large tables or updating many rows. The elapsed time and resources consumed bears a direct relationship to the amount of work — no matter *which* DBMS we are using.

I do not consider the DB2 Catalog tables to be large when many application data tables contain several million rows. Quite recently I had a call from someone who was executing a query against the DB2 catalog and could not understand why the query needed many minutes to complete. A quick look through an on-line monitor showed that the query had already caused 90,000+ GETPAGES, used almost two minutes of CPU time and was still running. Since this was a fairly busy DB2 system, I cancelled the query. An evaluation of the SQL indicated that it was JOINing five tables and some of them did not have an index. The major concern is not the elapsed processing time for this particular query. The concern is the impact upon the rest of the DB2 user community and the balance of the work on the processor that includes several IMS copies, many TSO users and a wide variety of batch jobs.

COULD YOU USE A LITTLE HELP MANAGING YOUR DL/1 DATABASE?

Is the time consuming management of your databases keeping you from more important tasks, like your conversion to DB2? Are you too busy to monitor those databases as closely as you would like? Then you need Data Base Control (DBC™) from Boole & Babbage.

DBC actually takes over the burdensome, and often neglected, task of ongoing DL/1 database monitoring. It eliminates the need for you to pore through the stacks of statistical reports that often hide critical performance indicators. When there is a problem, DBC pinpoints its location, so you don't waste time looking for problems where none exist. And DBC even saves machine time by reducing SMU runs up to 90%.

DBC provides threshold monitoring, alerting you to performance or space related problems only when they're serious enough to require action... and before performance degrades. DBC even lets you set the warning thresholds, so you can tailor your monitoring to the specific characteristics of your database. Which lets you define what conditions constitute a problem.

In short, DBC is like having full-time assistants who recognize and react to problems as you would. So your database administration is more productive and more efficient. The result is a more stable IMS environment, maximum availability and greater control of your database.

Get the help you need to get control of your database administration. Call to arrange a free demonstration of DBC, today. Call Marty Johnson: in California, 800-624-5566. Outside California, 800-822-6653.

Boole & Babbage, Inc., 510 Oakmead Parkway, Sunnyvale, CA 94086.

Boole & Babbage

DBC is a trademark of Schumann Consulting Group, Inc.

Perception vs. Reality

Similar dilemmas have been faced for almost two decades. A business requirement to retrieve several records from a database of several million — the answer set is small but the processing requirements might be unbearable. The *key to success* for any database application is the *technical knowledge* of how the DBMS works — what is efficient and what is not efficient. Reality is the amount of work (processing cost) to produce the answer and the business requirements must justify the value of the answer based on the processing cost to produce it.

Some points cannot be emphasized strongly enough; so, sometimes it is necessary to be deliberately redundant. The *critical* factors in the success or failure of any system are the application and database design. Proper and efficient design *demands* a high level of technical knowledge. Real knowledge derived by practical experience is the basis for success. Simply reading the manuals and understanding the theory will lead you directly into a black hole.

Everyone has his favorite war stories. Here is one of mine. Over a dozen years ago, I received a contract to investigate performance problems with a large IMS manufacturing application. This was strictly a batch system. How bad can batch performance be? The overnight window was 12 hours. Benchmarks indicated that it would require more than 400 hours to process the nightly batch work. The design problem was classic.

The designers had read the manuals, taken an application programming course and designed the application. They had swallowed one of the selling criteria for IMS. They had *completely eliminated* all data redundancy. All the databases had multiple Logical Relationships. It was an interesting two weeks. This was the most extensive and complete set of application documentation I had ever seen. Every program had a complete call schematic for every function it could perform. Based on the known IMS instruction path lengths for the calls and the path lengths for I/O instructions, a simple pencil and paper calculation could predict the elapsed time for any program within five to nine percent. The solution to the problem? It was shown (via calculations) that creating some data redundancies would allow processing to complete within the overnight window with a couple of hours to spare. Unfortunately, it would have required thousands of man hours to modify the programs and databases. Therefore, the entire project was scrapped.

How Is This Relevant To DB2?

More information was available about IMS internal processing. It was known exactly how data would be accessed and the access path and method could be controlled. There are levels of variables that determine the DB2 path length for a call before data is accessed and significant differences when data is qualified in the RDS instead of the Data Manager. The access path for rows of data can be affected; however, the Optimizer selects the actual access method and criteria. Using the EXPLAIN facility shows how the data will be retrieved. This should be *mandatory* for *all* SQL! Often, restructuring the SQL or using redundant predicates can force the Optimizer to select a better access path. This is especially true when using *static* SQL with *host variables*. When the Optimizer sees host variables, it does not know the actual predicate values and makes default decisions based on catalog statistics and its filter factors regarding the number of pages that may have to be accessed. A query that can be satisfied efficiently through a non-matching index scan might be processed as a tablespace scan at a higher processing cost. A thorough discussion of the Optimizer, how it functions, the criteria it uses and the effects of varying query structures upon its access path decisions could more than fill this entire magazine.

The Cinderella Analogy

Based on the assumption that *everyone* saw this movie in their youth, What was the significance of the glass slipper? Quite simply, if it fit the foot, the scullery maid became the princess. If your next large application uses DB2, will its success turn you into today's equivalent of royalty?

Think back to the movie. Remember the nasty step-sister trying to cram that size-10 foot into the size-five slipper? If it had not popped out, it still would not have been possible for her to walk. This same effect will occur if your application really does not fit a relational implementation under DB2. Some applications fit perfectly, most can work well and occasionally there will be some that do not belong in a relational DBMS.

The Answer?

Should DB2 become your corporate DBMS of choice? It depends. What will this statement mean, what impact will it have and how will it be implemented? The flexibility, ease of access and productivity improvements of DB2 are needed.

There must be hardware cost considerations. The use of DB2 will require *at least* two and one-half times the processor resource as the same application using IMS. It will also require more expanded memory for large buffer pools. The additional hardware costs will, initially, expand at a faster rate than had been expected using older DBMS technology.

Overall, the *critical* factor is how the statement *DBMS of choice* is issued and implemented. It requires two clearly defined understandings for success.

- Every new application should look at DB2 first with the understanding that some applications are not suitable for DB2 implementations.
- People with adequate levels of technical expertise in both DB2 and the alternative DBMS being considered shall have a major voice in the decision. Also, the decision will be based upon *both* technical criteria and business requirements.

So, at last the answer is — *yes*, it should be the *choice*. However, this is dependent on the acceptance and understanding that some new applications really should not use DB2 at this time.

For those who should not use DB2 right now, they can be designed to minimize a future conversion. This can be accomplished using structured programming techniques, not *black box* DBMS interface modules with their additional levels of overhead and processing inefficiencies. ◒

ABOUT THE AUTHOR

Joel S. Goldstein is President of the Responsive Systems Company, an information systems consulting firm that specializes in on-line database systems. He has 26 years of data processing experience. Responsive Systems Company, 281 Highway 79, Morganville, NJ 07751, (201) 591-0911.

Why Isn't A CPU Second Consistent?

By Cheryl Watson

Great quotes heard by the "Why did I get into data processing?" school of thought are, "Your charge-back system stinks — it bills me a different amount every time I run a job," "I want to run all my jobs on the 3090 instead of the 3084 because it costs less," "Your new operating system stinks — now my jobs cost twice as much to run," "You were overcharging me before the new operating system was put in because now my jobs cost half as much," "Can't you get anything right," "It wasn't like this back in the good, old 158 days," "Why can't you just charge in 158 hours?"!!!!

If you have heard comments like these, continue reading! This article discusses the reasons for differing charges based on CPU time used. (The discussion could continue if I also took on I/O charges and storage charges. However, for now I will concentrate on CPU variances.) While I have used MVS for most of my examples, the concepts apply to both VM and VSE.

A fact of life: CPU time varies from one run to the next even when using the same program and the same data. This generally plays havoc with charge-back systems, capacity plans and management reporting systems. Take a look at some CPU variability in two cases: a stand-alone, benchmark situation and a normal multiprogramming environment with other jobs running. In both cases, assume that you want to evaluate a job that uses the same data every run.

Stand-alone Benchmark

Most people do not have the luxury of running a stand-alone benchmark test. However, variabilities still exist due to the hardware microcode and software differences. Hardware variabilities exist due to: instruction speeds, pipelining (based on Translation Lookaside Buffers (TLBs) and CPU cache buffers) and multiprocessing effects. Software differences occur because of the software releases, program location, product changes, software parameters and configuration differences.

Instruction Speeds

Each processor has a given speed for a single *cycle* with the more expensive processors having faster *cycle speeds*. However, this is not a true indication of instruction speed since instructions can take a single cycle or multiple cycles. Additionally, vendors do not use the same microcode for instructions so one vendor may have fast addition and move instructions while another vendor has faster multiplies and branch instructions. Given two programs (like benchmark programs), some will run faster on the first processor while the other will run faster on the sec-

ond processor simply because of the mix of instructions. In fact, some well-known benchmark tests, such as Whetstone and Megaflops, attempt to provide a mix of instructions to test all processors. (Note: beware of benchmark tests provided by vendors — they know which instructions run better on their machines!) Unfortunately, standard benchmarks do not adequately match your environment. As an example, most COBOL applications consist of 50 percent code that moves data to and from memory (rather than in registers). Therefore, COBOL applications do better on processors that have fast *memory move* instructions.

The reason that "MIP" ratings (Millions of Instructions Per Second) are distrusted is that the rate depends on which instructions are being executed. What can you do? Run your own tests based on selected samples of your workload and compare the results. Speaking of MIPs, you have probably noticed that some vendors are refusing to rate their machines in MIPs anymore because of the variability of the instruction speeds and other factors. IBM, however, does publish a factor that is used to determine Service Units (SU). SUs were designed to be machine independent measurements of resource usage: CPU, I/O and memory. The *MVS Initialization and Tuning Guide* from IBM

gives SU rates per CPU second for each processor. For example, the SU rate for a 3084-QX (single-image) is 373.8 SU/second while a 3090-400E (single-image) is 700.0 SU/second (an 87 percent improvement). It is interesting to note that some reports list the 3084-QX as a 29.1 MIP machine and the 3090-400E as a 61.5 MIP machine (a 111 percent improvement).

Pipelining

Hardware vendors have implemented two primary techniques for increasing the effective speed of a given CPU: lookahead processing and buffering. Lookahead processing refers to the simultaneous processing of an instruction and the decoding of the next instruction(s). Some processors are more effective than others with this, partly due to the use of TLBs (buffers of frequently-used or previously-used addresses). Obviously, a larger TLB will allow faster translation of addresses. Some programs take more advantage of this fact than others due to the way they are written. Programs with little branching make effective use of lookahead processing. (''GOTOLESS'' programming, however, still results in many, many branches.)

Cache buffers (also called high speed buffers) are used to store real storage memory during execution. In the older processors, code could be executed from either cache or real memory but many new processors require that all addresses be contained in cache before execution. The difference in speed between an access in cache and real storage can vary from two to 18 times as fast (a recent IBM presentation showed a one-to-16 ratio on a 3090). Cache is an expensive memory storage and varies by the CPU. For example, a 3084-QX has a cache of 64K while a 3090-400 has a cache size of 256K. Obviously if a program has a small ''locality of reference'' (that is, references few 4K blocks), all references could be within cache and take one-sixteenth the time on a 3090. A program with a ''large locality of reference'' might require all references outside of cache and result in much larger CPU time. Therefore, if you run a program on two machines with different size caches, some programs will run better and others will not.

Cache is used by several processes and if cache is needed by another program, the first program's buffers may be overlaid in cache and it will have to go back to referencing real storage. Other programs in the system will cause this, as well as interrupts on the system (such as I/O interrupts and timer interrupts).

Multiprocessing Effects

In a uni-processor, a single CPU executes code from a single operating system. All code is available to the CPU. On a single-image multiprocessor (such as a dyadic or a quad), multiple CPUs use a single copy of the operating system but must single-thread through some parts of the code (you would not want to ''dispatch'' the same job on two processors!). The MVS implementation of single-threading is the use of ''locks,'' essentially a flag that indicates the code is in use. When the lock is held by one CPU, the other CPU typically ''spins'' or loops until the lock is freed. Because of this competition, a portion of the CPU time reported in an MP mode is due to the *spin* code and will result in a higher amount of CPU time recorded. To adjust for this effect, the service rate assigned by IBM is adjusted depending on the number of processors. For example, MVS/XA assumes 100 percent for a single processor, 93 percent for two processors, 88 percent for three, 85 percent for four, 80 percent for five and 76 percent for six. Additionally, machines can be run in *partitioned* mode rather than single-image. In this case, there is often cross system or interprocessor communication that causes interrupts to the CPUs. Because of this effect, IBM has assigned different service unit rates to machines running in single-image versus partitioned (3090-400E single-image is 700.0 SU/second and partitioned is 765.88). Note that these effects are taken into account in the calculation of service units but not in the case of CPU time which will vary, much to your chagrin!

Software Releases

Some changes in CPU measurements occur because of changes in the measurement tools. For example, measurement of overall CPU time in MVS changed from sampling to actual measurements of wait time in one release while other releases captured more of the consumed CPU time in separate fields. In order to understand the CPU measurements, you must first understand the operating system that is in use.

A good example of the result of changes in releases occurs because of the use of Cross-Memory Services (XMS). Prior to MVS/SP 1.3, most communication between address spaces was done via the use of Service Request Blocks (SRBs) and CPU time for this processing is reported as ''SRB time.'' MVS/SP 1.3 introduced XMS which allows a program to access another job without an SRB (so time is now reported under ''TCB'' time). As new releases of the operating system and products running under that operating system such as CICS, IMS and so on convert to using XMS instead of SRBs, the recorded CPU time moves from SRB time to TCB time (it is like magic!).

The changes are not over. For installations moving to MVS/ESA, there is now additional time to account for such as the time spent in addressing hiperspaces (this CPU time is recorded under separate fields in SMF).

Program Location

While the software may not change, the location of the operating system programs and the user's programs may affect CPU time. For example in MVS if a program is located in PLPA, the CPU time is less than if it is located in SYS1.LINKLIB. A new release of a product could move modules from one location to another. Additionally, the size of PLPA or SYS1.LINKLIB and its concatenations could affect CPU time due to the search time of the directories. Use of STEPLIB or JOBLIB can also impact CPU time, even if the program is not found (especially if the program is not found).

Product Changes

New releases of products can impact CPU timing as much as operating system releases as shown before. You should also consider changes within your installation. Have you installed a new release of a product, changed the operational parameters of a product, installed a new product (such as an on-line monitor) or installed installation changes? A rewrite or addition of an SMF exit can impact everyone's CPU time!

Software Parameters

Several parameters specified by the installation can affect the total amount of CPU time used. Without going into detail, I will mention some of the parameters in MVS that can impact the CPU time consumed by jobs or by the entire operating system: RMF cycle time, RMF recording time, SRM invocation time, SRM dispatch algorithms, SMF recording parameters, SRM balancing algorithms, swapping parameters, JES parameters and so on.

Remote printing on a PC without a single compromise.

Until now you've had to rely on a S/36 or AS/400® to deliver remote printing. Now a PC with BARR software and adapter sustains print speeds of 6,000 lines-per-minute and line speeds of 128,000 bits-per-second. Only BARR maintains all this performance with the reliability and ease of use PC users expect.

BARR/SNA RJE or BARR/HASP software drives up to five printers from a single PC. What's more, you can enter data, print, and receive output all simultaneously — without interruption. BARR's advanced multi-tasking software easily manages even the most complicated tasks, including LAN access, tape support, file transfer, and special forms printing. In addition, BARR offers one year of friendly, dedicated customer support with each purchase.

Communications adapters and software are available for the IBM PC, PS/2, and compatibles.

Try BARR for 30 days. We've helped thousands save millions. Call 800-BARR-SYS.

Protocols
SNA RJE with Multiple Sessions
HASP — BSC Multileaving

Host Systems		
JES2	DOS/POWER	CDC NOS
JES3	RSCS	VS1/RES

BARR

BARR SYSTEMS INC.
2830 NW 41 Street, Gainesville, FL 32606
800-BARR-SYS, 904-371-3050, FAX: 904-371-3018

AS/400 is a trademark of IBM Corp.

CIRCLE #91 on Reader Service Card ▲

Configuration

The amount of real storage, expanded storage, solid-state storage, page dataset placement, swap dataset placement, channel balancing and so on can also have a great impact on the amount of CPU time consumed. Although some measurement tools, such as SMF, attempt to separate the *system* CPU time from the *program* CPU time, the tools are not always precise. A change in any of these configuration aspects will cause a change in total CPU time and may impact program time.

Multiprogramming Environment

Assume that you cannot or do not want to run benchmarks. What variation will you see if the same job is run on a machine with multiple jobs? The variations can be due to impact on cache buffers due to other programs or interrupts and queue searching and program differences.

Cache Buffers

If multiple jobs are in the system, while one job is waiting for an I/O or other request, another job is dispatched. While that job is dispatched, its storage references are overlaying cache with its memory. If the first program is then dispatched, few if any references will be from cache and will consume more CPU time to move them from real memory. This means that more programs in the machine will lead to a higher amount of CPU time consumed by a job! Additionally, more jobs produce more I/O and timer interrupts. Each interrupt causes a branch to the interrupt handlers who need to use cache. This means that more programs in the machine will lead to a higher amount of CPU time consumed by a job! Sound familiar?

Queue Searching

Much of the normal processing for a job consists of searching queues of control blocks. Examples of these control blocks include storage control blocks for GETMAINs, enqueue control blocks (although since MVS/SP 1.3, this is now improved by hashing), write-to-operator buffers, VTAM control blocks, JES control blocks, SRM control blocks, device allocation and so on. This means that more programs in the machine will lead to a higher amount of CPU time consumed by a job! Do you start to get the picture? You should realize that many of these queues do not always shrink to their original size. They simply continue to grow. For example, GETMAINs from Common Serv-ice Area (CSA) in MVS require searching some virtual storage control blocks. After a week or two without an IPL, this chain may be large due to fragmentation. In benchmark situations, you should always start with a fresh IPL.

Program Differences

I have mainly discussed CPU variability if the program does not change. But there are many things you can do to a program that will impact CPU time. These include: changes in the compiler version or compiler options; changes in the amount of data processed; changes in the organization or location of files; blocksize of files; number of file buffers; and so on.

Summary

The purpose of this article was to explain why variations occur in CPU timing. It does not answer the question of how to change it. You cannot! The position I have heard from IBM (and I believe it) is that the overhead of adequately measuring CPU time is greater than the benefits of accurate reporting. Their stance has been that charging users on CPU time or any other resource is unfair. You should be charging on the service provided. For example, you know that you pay 10 cents per check to a bank; the user should be able to pay 50 cents per payroll inquiry! Unfortunately, this is easier said than done! ◒

ABOUT THE AUTHOR

During her 23 years in IBM mainframe software, Cheryl Watson has worked for such companies as EDS, Amdahl, Morino and Landmark Systems. She is now president of Watson Resources, Inc. and provides MVS courses including: SRM, SMF, RMF, Capacity Planning, Performance Measurement, IPCS and IBM Products (PR/SM, ES, ESA). Watson Resources, Inc., 4844 Cloister Drive, Rockville, MD 20852, (301) 897-8030.

CIRCLE #190 on Reader Service Card ▲

SMP/E Keeps Third-Party Software Under Control

By David N. Shein

System Modification Program/Extended (SMP/E) is an IBM software product that has the purpose of maintaining the code for an installed operating system and its ancillary features. Although SMP/E was developed to maintain IBM system software, it is readily adapted to other uses. One of SMP/E's strengths is the ability to track the relationships between software components. This is a capability that can be put to use in a variety of ways, one of which is described here.

Recently the Automobile Club of Southern California installed a software product that is closely tied to the internal architecture of a particular operating system component. Thus, the software product has to be updated when changes are made to the system component. What was needed, in effect, was some sort of automated *tickler* to indicate when this situation arose. The technique that was developed (using SMP/E) is the subject of this article.

The remainder of the article consists of four parts. First is a description of the problem that prompted development of the technique described. This is followed by an explanation of the technique itself and the SMP/E features it utilizes. The coding details are then illustrated, after which I conclude with some thoughts on maintenance.

The Problem

Some time ago, VPS from Levi, Ray & Shoup, Inc. (Springfield, IL) was in-stalled. VPS is a software product that provides an interface to remote 327x printers. While installing VPS, Technical Services elected to include several user exit routines, two of which use data generated by JES2, a subsystem of the MVS/XA operating system. These two exits use certain JES2 macros to determine the format of the JES2 data areas they access, so they must be assembled using the JES2 macros. This in turn means that whenever JES2 is updated, the VPS exits may need to be reassembled.

Rather than rely on calendar type ticklers or fallible human memory, my colleagues and I wished to develop a mechanism that would react to JES2 macro changes by generating some sort of non-ignorable reminder to reassemble the affected VPS exits. As it turned out, we were able to go one better: a technique was developed to exploit the features of SMP/E, so that whenever one of the macros used by the VPS exits is modified, the exits are reassembled automatically.

The Solution

In order to make our scheme work, SMP/E had to be made to do two things: store the VPS exits' source code in the SMP/E environment (so SMP/E could assemble it when necessary) and identify the macros used by the exits (so SMP/E could determine when a reassembly was necessary). I will describe the method for accomplishing these two objectives in a moment. First, however, a brief digression is in order to review the key features of SMP/E that I will refer to.

SMP/E maintains a database called the Consolidated Software Inventory (CSI) that is essentially a repository of status information for the software under SMP/E control. The software itself is stored in Partitioned Dataset (PDS) libraries; only the status and control information is kept in the CSI. There are many types of records, called *entries,* in the CSI; three types are of particular concern to us. These are SRC (source) entries, MAC (macro) entries and ASSEM (assembly) entries.

The SRC and MAC entries describe, respectively, source code modules and macros which are stored in libraries controlled by SMP/E. The ASSEM entry will be discussed later.

SMP/E has specific packaging requirements for software that it controls. Software and software modifications *not* produced by IBM are typically introduced into SMP/E in the form of a USERMOD or User Modification. This is the packaging method used for our VPS exits. The actual USERMOD consists of the source code for the exits as well as some JCL to be scanned by SMP/E with appropriate SMP/E instructions embedded at the proper places.

There are three SMP/E commands which are significant for the purpose of this article. They are RECEIVE, APPLY and UCLIN commands.

The RECEIVE command simply tells SMP/E to load the USERMOD into a temporary work area where it is saved. At this time, SMP/E does some rudimen-

tary editing and validation of the USER-MOD to ensure that it is packaged with valid format and syntax. The APPLY command is then used to actually install the USERMOD. The APPLY command is central to the technique because this is what causes SMP/E to create, in the CSI, the control information that makes the technique work.

APPLY processing causes the source code for the exits to be stored in an SMP/E-controlled library. At the same time, SRC entries are built in the CSI specifying where the source code was stored so SMP/E can locate it later.

APPLY also causes SMP/E to look for a ++JCLIN statement in the USER-MOD. Following the ++JCLIN statement is the JCL for a link-edit job step, an IEBCOPY job step and two assembly job steps. This JCL is not actually executed by the system; rather, it is scanned by SMP/E to extract needed information about the structure and contents of the software being installed. The linkage editor JCL tells SMP/E how to link-edit the VPS exits, that is, what to name the load modules and which load library to store them in. SMP/E records this information in the CSI and uses it when performing reassemblies. The IEBCOPY step tells SMP/E which library to store the SRC (source) members in.

The assembly job steps are the real key to the technique because they include a copy of the actual source code for the exits. SMP/E reads the source code and scans it, searching for references to macros. As a result of this processing, SMP/E does two things.

1. It builds an ASSEM entry in the CSI for each exit. The ASSEM entry contains the actual source code for the exit.
2. For each macro detected in the source code, SMP/E locates the appropriate MAC entry in the CSI and adds to it a subentry that identifies the corresponding ASSEM entry. These subentries are extremely important because their presence means that *in the future, whenever a change is made to one of these macros, the code in the ASSEM entry will be automatically reassembled and re-linked.*

Finally, a UCLIN command is used to delete the ASSEM entries from the CSI. At first blush this may sound rather strange — why delete the entries everyone just went to such great pains to create?

The answer lies in the way SMP/E

reassembles source code. When SMP/E decides, due to a macro change, to reassemble one of the exits, it looks for the ASSEM entry which is named by the MAC entry. If the ASSEM entry does not exist, *SMP/E will look for a SRC element by the same name and assemble that instead.*

Remember that the ASSEM entries contain the actual source code to be assembled. This means that if the source code is of any real length (and these exits

are), the entries take up a great deal of space in the CSI. Furthermore, the ASSEM entries are redundant because corresponding SRC modules were installed at the same time. Thus, since the ASSEM entries waste CSI space and are not needed, they are deleted and that completes the installation of the USERMOD.

Coding Details

Figure 1 uses a slightly simplified example (a single exit routine) to illustrate

the USERMOD coding. This is the raw input which SMP/E sees at RECEIVE time. The line numbers do not appear in the actual code; they are included for ease of discussion.

Line one shows the USERMOD statement. The name assigned to the USERMOD is in parentheses and is required. This name is selected according to SMP/E naming conventions and local shop standards. ABC1234 is an arbitrary name used for illustrative purposes only.

The VER statement on line two is required so SMP/E can verify that this USERMOD is being installed in the correct environment. Z038 tells SMP/E that this USERMOD is for MVS or one of its components. The FMID or ''function ID'' (HJE2221 in the example) specifies the particular software component being modified. The example uses the FMID of the current version of JES2.

Line three specifies that a source module is to be added or replaced in its entirety. The module name is EXIT001. The actual source code is inserted immediately following the + +SRC statement. This copy of the source code (two copies are present in the USERMOD) is the one which SMP/E will store in a source library.

The remainder of the USERMOD (after line five) consists of JCL and text to be scanned by SMP/E. The JCL in the example includes three job steps. The first is an assembly step; the source code is inserted (again) at line eight. This copy

			Line
+ +USERMOD (ABC1234).			1
+ +VER (Z038) FMID (HJE2221).			2
+ +SRC (EXIT001).			3
(source code for the exit routine)			4
+ +JCLIN.			5
/ /	EXEC	ASMS,MOD = EXIT001	6
/ / SYSIN	DD *		7
(source code for the exit routine)			8
/ *			9
/ / STEP	EXEC	PGM = HEWL,PARM = 'LIST,LET,XREF,NCAL'	10
/ / SYSLMOD	DD DSN = SYS1.SMPVPS.PROGPROD,DISP = SHR		11
/ / SYSLIN	DD *		12
	INCLUDE AVPSMOD1(EXIT001)		13
	NAME EXITLM1		14
/ *			15
/ / STEP	EXEC	PGM = IEBCOPY	16
/ / SYSIN	DD *		17
COPY	INDD = AVPSSRC,OUTDD = VPSSRC	TYPE = SRC	18
S	MEMBER = EXIT001		19
/ *			20

FIGURE 1

of the source code is the one which SMP/E will scan to determine macro dependencies and which will be stored in an ASSEM entry in the CSI.

The link-edit step beginning at line 10 tells SMP/E where to store the assembled and linked load module and what to name it. SMP/E uses the IEBCOPY step starting at line 16 to determine which library to store the SRC module in.

Three SMP/E commands are needed to install the USERMOD. First is the RECEIVE command, which looks like this:

RECEIVE SOURCEID(USR0001).

There are a number of operands which can be coded on the RECEIVE command. All of them are optional, however, as long as the USERMOD is the only input. The

SOURCEID operand serves to *stamp* the USERMOD with the label USR0001 for control and audit purposes.

Once the RECEIVE is run successfully, the APPLY command is issued:

APPLY S(ABC1234).

This command tells SMP/E to install USERMOD ABC1234. This triggers the processing described earlier (scanning the JCLIN and storing the SRC members in the appropriate library).

The final step is the UCLIN command which consists of three lines:

UCLIN.
DEL ASSEM (EXIT001).
ENDUCL.

SMP/E now knows which macros are used by the exit. Whenever one of the macros is updated, SMP/E will assemble and link the affected exit automatically.

Ongoing Maintenance

As is always the case, installation is only part of the job. Provision must be made for ongoing maintenance. While the technique described here definitely helps, it solves only one-half of the maintenance problem. Changes to the macros will automatically update the exits, true, but SMP/E has no way of automatically knowing when the exits themselves change! Thus, it is still necessary to remember to reinstall the USERMOD when new versions of the exits are implemented.

The reinstallation process is fairly straightforward. Use the RESTORE command of SMP/E (not illustrated here) to remove or *deinstall* the USERMOD. Then, repackage the USERMOD with fresh copies of the exits' source code and

repeat the installation process described above.

Concluding Thoughts

The SMP/E commands used in this article offer numerous features and options which were deliberately omitted. The purpose of this article was to describe a concept, not to explore SMP/E's many intricate coding possibilities. Thus, the coding details have been kept simple for illustrative purposes. When constructing input for SMP/E, the IBM manuals should always be used as the "Bible." (The basic SMP/E reference manual is IBM order number SC28-1107, *System Modification Program Extended Reference*.)

Similarly, the technique illustrated in this article is only one of many ways in which SMP/E can be put to use managing non-IBM software. In fact, this technique is not the only way to achieve the limited objective stated in this article. Another method would be to load the entire software product (not just a few exit routines) into the SMP/E environment. Doing this, as one might expect, yields significantly greater control and increased flexibility in managing installed software at the price of requiring a bit more SMP/E expertise.

SMP/E is a complex tool with a rich repertoire of features and capabilities. Properly exploited, it can give an installation a degree of control over critical software that is rarely achieved otherwise. ≜

EDITORIAL EVALUATION

Please circle the appropriate numbers on the Reader Service Card.
1.) This article was:
   275 (Interesting/Helpful), 276 (Too Technical), 277 (Too Basic)
2.) Would you like more articles on the same subject?
   278 (Yes), 279 (No)

ABOUT THE AUTHOR

David N. Shein is a systems programmer with the Automobile Club of Southern California. The technical staff to which he belongs provides system software support and maintenance, capacity planning, performance management and DASD management. The environment they support includes multiple IBM mainframes, MVS/XA, CICS and a 3000-terminal network. 3333 Fairview Rd., Costa Mesa, CA 92626, (714) 850-5374.

Workload Trend Analysis

———•———

A capacity planning technique for projecting resource requirements in a CICS environment

By Ted C. Keller

It is almost impossible to perform any kind of capacity planning task without using some form of trend analysis. In recent years much has been written about business element forecasting, natural business units and performance modeling and planning. These are all essential components of capacity planning and without them it is likely that any capacity or performance plan will fail. However, these tools will be inadequate unless current workloads and growth patterns are understood. There is almost always more growth in resource utilization than that which can be specifically identified, even with the best of planning. Workload trend analysis can help provide an additional dimension to the complex task of capacity planning by identifying some of the hidden sources of growth.

When most people think of trend analysis they tend to think of system-wide statistics summarizing recent history (perhaps the past two or three years) with projections based on historical trends. Graphs are typically used with solid lines showing actual history and dotted lines showing projections or ranges of projections. Sometimes graphs are produced comparing projections to what actually occurred. This type of planning can be useful; however, effective analysis of workloads requires more depth. Only when the true nature of the workloads involved and how they are growing are understood, can trending techniques be used intelligently with any reliability.

In this article, I will discuss some techniques for analyzing workloads in a CICS environment. The data used in the anal-

yses discussed in this article has been extracted from CICS transaction history provided by the CICS CMF facility. Other sources could have been used for much of this information. All of the analysis has been done using simple SAS (SAS Institute, Cary, NC) procedures. While this study is drawn from a CICS environment and references data meaningful for CICS systems, many of the ideas presented can also be applied to systems with other dominant workloads.

What Is Workload Trend Analysis?

Workload Trend Analysis (WTA) is the study of growth patterns of important workloads and the various components of their growth. WTA is *not* the simple plotting of statistics like "total CPU used"

For CICS/VSAM Data Recovery, Integrity Solutions is the Global Choice.

These leading IBM mainframe companies trust the integrity of their corporate data to Integrity Solutions. They selected Integrity Solutions' DATA RECOVERY SYSTEM (DRS) because it provides the most comprehensive forward and backward recovery for files accessed by CICS/VS and VSAM batch processing.

Integrity Solutions introduced DRS more than 8 years ago, and with over 1000 product installations worldwide, Integrity Solutions will continue setting the data integrity standard. By providing leading-edge products and the highest level of support.... to serve your data recovery needs now and in the future.

Guided by what users say they need next, Integrity Solutions is setting the data integrity standard for the emerging electronic vaulting and mirror processing marketplace as well.

The DATA RECOVERY SYSTEM from Integrity Solutions, because your data is worth it.

Integrity Solutions, Inc.
1-800-289-9900
1-303-794-5505

© 1988

Integrity Solutions... Setting the Data Integrity Standard.

DASH QUICKSILVER™

Dash Quicksilver Well-known, handsome software hero from another universe; raised from infancy by the Quicksilver family. Dash's wrist band gives him unique and powerful telecommunications abilities.

Penny D. Quicksilver Dash's step-sister; ambitious and smart. As manager of the data communications department, she has always been able to help the users and stay within the budget. She has stubbornly refused to call Dash for help.

Aaron S. Scowl Manager of the data entry division. Powerful, sour and resentful towards others, he is particularly out to get Penny.

Data Entry Personnel Once highly productive and motivated employees, they are daily dropping off as response times grow slower and slower.

Penny is faced with a critical situation. Productivity is at an all-time low. Scowl is seeing his opportunity to make her look bad. What will she do?
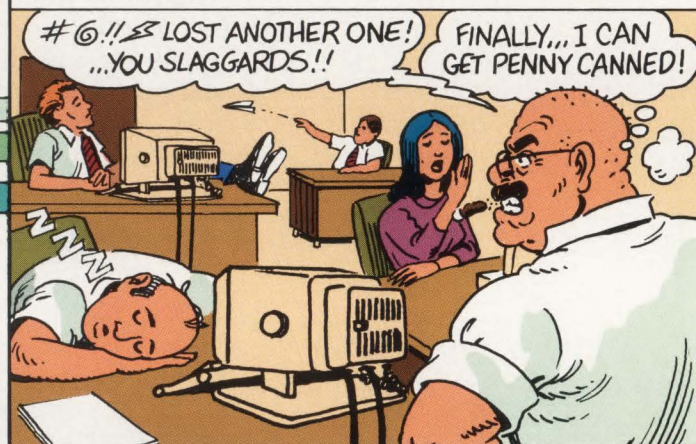
Dramatic Response Time Improvement!

Responses returned faster than quicksilver—that's what you get with our 3270 SUPEROPTIMIZER™* products, which utilize the patented technologies of IMAGING™* and Input Suppression*. You are assured the maximum optimization of outbound and inbound data streams on both local and remote networks. And you can postpone indefinitely or totally eliminate the costs of hardware or communications facilities upgrades.

BMC's years of experience in working with 3270 data streams and the systems that use them, has produced an extensive family of optimization products, each customized to its environment.

The 3270 SUPEROPTIMIZER products offer:

- Improved terminal response times and end user productivity.
- Reduced line utilization and telecommunications costs.
- Eliminated or delayed hardware and telecommunications upgrades.
- Statistics for network planning and reports for data stream diagnosis.
- Optimization of printer data streams.
- Erase Input Key Allowed feature.

or "peak CPU usage" on a month-by-month basis. WTA goes beyond this. It involves coming to an understanding of which workloads are significant, how much they are growing and, most important, what is causing them to change. WTA is not a static statement of what once was, but it is the continuous process of studying the nature of systems and their patterns of change.

Identifying Workloads

Perhaps the first step in WTA is to identify periods of interest and the significant workloads within those periods. This is a basic function that is probably already being done as part of any organized CP process. Any good performance analyst or capacity planner should be able to tell you quickly that a certain hour on a given day of the week or at a certain time of month will typically represent some measure of peak utilization. Because data processing systems are commonly involved with business functions, there is usually a well-known correlation between most data processing peaks and some aspect of business cycles. Since this information is so basic to the capacity planning process, it is assumed that it is a common practice and that it is not necessary to spend any more time discussing this.

Once the periods of interest have been established, the next basic step is to identify the significant workloads in each period. This will normally be done by analyzing RMF data to determine which performance groups are consuming which resources. Once a profile of activity is developed for each CICS region, the next step is to analyze transactions or systems (logical collections of transactions) within each region.

The analysis of workloads within CICS can be a somewhat cumbersome process. Since a separate record is written for each CICS transaction, the amount of data to be analyzed can be considerable. It is usually beneficial to have some kind of summarization routine (as provided by MICS from LEGENT Corp.) pre-process the data before any detailed analysis is done using a statistically-based language. In fact, some of the analyses discussed later would be almost impossible unless detailed data had been summarized and preserved over time in a summarized format. Additional complications lie in the fact that most systems consist of more than one transaction code. The analyst must determine these relationships and adjust his or her studies to reflect this. This problem is often sim-

FIGURE 1

Components Of CPU Growth
For Selected Applications
SYSTEM=BILLING PERIOD=EVENING
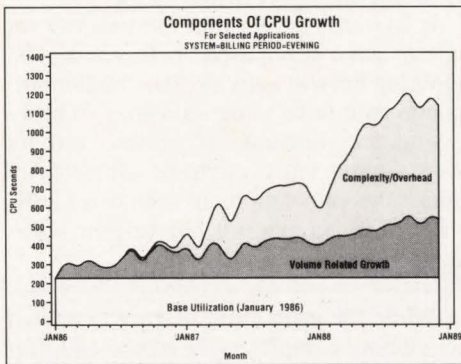
plified if the first two or three digits of the transaction code uniquely identify individual systems.

A second caveat when analyzing CICS workloads is that peak processing loads may occur at different times in different regions. In fact, peak levels of activity in some regions may occur at times totally unrelated to overall system activity. It is worthwhile to know what the peak periods of activity are for each CICS region, especially if any region will use more than 50 percent of a processor. Since most of the work in each CICS region is scheduled under a single MVS TCB, measurable degradation will begin to occur whenever a region attempts to use more than about 60 percent of a single processor. Special attention should be given to changes in workloads which could drive any CICS region beyond this level of activity, regardless of whether those changes will occur during overall peak periods.

Workload Trend Analysis

After determining the periods of interest and significant workloads, the next step is to start gathering data on the historical growth of each workload. This is a fairly common practice and it is not unusual to see graphs tracking the peak utilization or peak numbers of transactions within each CICS region for each month or quarter over a two- to four-year period. After establishing this base, the real work of WTA begins. At this point you should begin to see trends developing and start to ask questions about the nature and causes of these trends.

Borrowing some concepts from cost accounting, you can start to break growth into its composite parts. Figure 1 illustrates how this might be accomplished. The vertical axis represents the average CPU seconds per hour utilized by a single system during its typically busiest periods. For this system, the hours selected

were from early evening in the last business week of each month. Adjustments have been made for processor upgrades, restating previous utilization levels in terms of current processor power.

Three components of growth are shown. The first is base utilization. In this example, January 1986 is the starting point of the analysis. January is convenient not only because it is the start of our business year, but also because there is usually less activity for this system in January than in most other months. This allows me to plot seasonal increases and other growth without having to deal with negative quantities. The base level of CPU utilization for this application was about 230 CPU seconds per hour — about 6.4 percent utilization of a single 3090 E-class processor.

Three components of growth are base utilization, volume and changes related to complexity or overhead.

The second component of growth shown is for volume. This component represents the growth in CPU utilization related to increases in transaction volume. It is calculated by first multiplying the average CPU used per transaction during the base period (January 1986) by the total number of transactions actually executed in the current period. The difference between this number and the CPU utilized in the base period is the volume-related component.

The volume-related component reflects changes in utilization of the system by the user. It may, in many cases, be correlated to changes in business volume. In the case of our sample shown in Figure 1, the total number of transactions is generally related to changes in levels of business activity.

The third component of growth shows

FIGURE 2

Average CPU Per Transaction
For Selected Applications
SYSTEM=BILLING PERIOD=EVENING
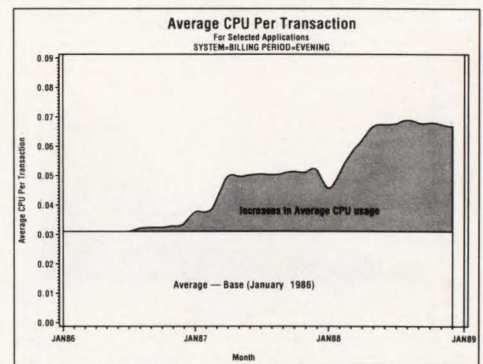
changes related to complexity or overhead. It is the difference between total CPU utilized and CPU usage associated with total volume. It represents the total amount of CPU consumed beyond that used in the base period and adjusted for volume increases. As we can see in Figure 1, most of the CPU currently being utilized by this system is attributable to either complexity or overhead not present during the base period. In essence, it reflects the impact of increases in the amount of CPU used per transaction over time.

There is no way to determine from the graph whether the increases were caused by additional functionality or added overhead. This determination cannot be made automatically. It requires that complexity figures be reviewed periodically and that significant deviations be researched on a timely basis. In theory, a separate database could be developed summarizing the nature of changes in complexity or overhead and this could be factored into the graph. However, this has not been done here.

Figure 2 shows a history of average CPU per transaction. This is a partial indicator of complexity/overhead. It is only meaningful if the system is dominated by a single transaction (as this system happens to be) or the mix of transactions is stable from month-to-month. If properly researched, this data can be extremely valuable to the capacity analyst. In an environment in which new functionality is constantly being added to existing systems but few truly new systems are being developed, the information derived from this analysis can be used to help justify hardware upgrades. When all that management sees is increased costs of new equipment without the addition of major new systems, it can be handy to have a list of major functions that have been added to existing systems and their costs in terms of processor utilization.

FIGURE 3

Components Of CPU Growth
All CICS Systems
PERIOD=MORNING
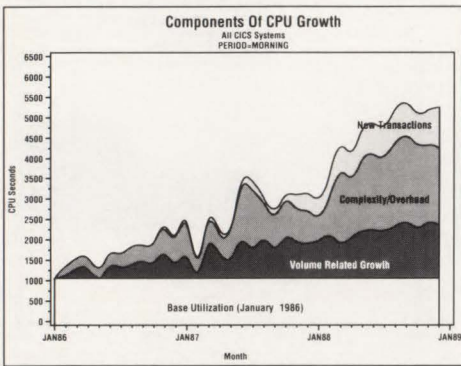
In order for this type of data to be useful and effective, it needs to be researched quickly. I do a detailed study of various types of growth each month and attempt to research which changes may have contributed to growth. I have found that even with a good change control system, the nature of many applications changes are difficult to track as little as a month later. The sooner I can determine that change has occurred, the easier it is to study it.

Unfortunately, not all changes in complexity/overhead are related to increased functionality. At times, relatively minor program modifications with little change in functionality can have a major impact on resource utilization. In theory it would be nice to know about such things before they showed up in production. However, it is almost impossible to conduct performance studies on the numerous minor changes occurring almost constantly in most large CICS environments. When faced with a disproportionately expensive change (in terms of resource utilization), the capacity planner could be faced with the politically sensitive task of confronting the responsible programming group or user or figuring out how to explain such changes without embarrassing anyone. Of course, if you were not doing this kind of analysis, you would never need to face this type of challenge.

System-Wide Workload Analysis

Figures 3 and 4 present the same basic kind of workload analysis discussed above but on a system-wide basis. The base level CPU utilization for each transaction is calculated, as are the volume and complexity components for each. The graphs show the sum of the separately calculated components of each transaction. A new component is added for transactions or systems which did not exist during the base period. The results are shown for two periods of interest — in this case,

early morning and early evening.

As can be seen from the graphs in Figures 3 and 4, increases in complexity or overhead represent the largest portion of current CICS CPU utilization during both periods of interest. It appears that little processing during the evening hours is due to new systems but that 20 percent of all activity in the morning is associated with systems added within the past 3 years. From these graphs, I can easily see that most growth in processor utilization comes from changes to existing workloads. If growth patterns continue as they have, this will probably be the source of most additional processor utilization over the next year or two. Effective analysis of these workloads will be critical to any successful capacity planning process.

Figure 5 shows another interesting system-wide trend — *logical I/O density.* Logical I/O density represents the number of *logical* random I/Os requested in CICS per CPU second directly consumed by transaction processing. If logical I/O requests are indicators of real work being done (and I believe a case can be built for this), logical I/O density can provide a measure of the effectiveness of the work being accomplished. Each *logical* random I/O request represents the retrieval of data or the accomplishment of some function. The larger the number of logical I/O requests achieved per application CPU second, the less processing required for each unit of work accomplished. A lower level of logical I/O density is an indicator of either increased processing complexity or expanded overhead. Lower I/O density is more commonly a function of overhead than increased complexity since increased complexity is usually accompanied by additional logical I/O events.

The reason I have selected logical I/Os requested instead of physical I/Os performed is because expanded use of LSR and other internal caching techniques over the past several years has eliminated many physical I/O requests. If RMF or SMF statistics were to be used (instead of CICS transaction statistics), the results could be misleading. In fact, from an operating system perspective, the more work that can be accomplished without resorting to the services of the I/O subsystem, the more efficient the process. From an operating systems perspective, you should hope to see an ever decreasing level of *physical* I/O density. However, from an internal perspective the reverse is true. The decrease in logical I/O density shown in Figure 5 probably indicates some increase

in application complexity along with considerable increases in overhead.

It is worth noting that random I/O requests have been used in this analysis, ignoring browse activity since random requests tend to be more indicative of processing requirements. If browse activity were significant, it might be necessary to adjust the calculations to include an estimate of the number of blocks browsed.

Summary And Conclusions

While the proportion of growth related to volume, complexity/overhead and new transactions will be different in each environment, each of these components will be present to some extent. The examples shown in this article have demonstrated how rates of growth can be examined. Of course, many graphs similar to Figure 1 might be used to track the history of significant systems. What is not shown in these graphs is the research done each month to trace causes in growth. In the installation from which this data was derived, there is and has been a long backlog of enhancements scheduled for most existing systems. Many of these systems are old and have not been restructured for more than ten years. With this information, I can project that much of the growth associated with increased functionality for

FIGURE 4

Components Of CPU Growth
All CICS Systems
PERIOD=EVENING

FIGURE 5

Logical I/O Density
All CICS Systems
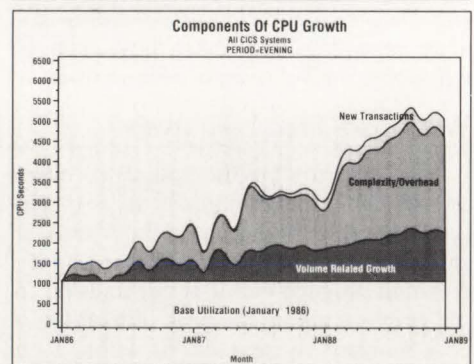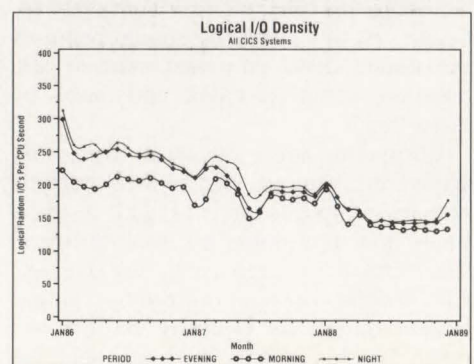
Making Automated Operations a
W-Win
Proposition

DP shops are using A/O as a means of providing better and more productive operations. It is not a diabolical plot to trim operator headcount.

By Bill Carico

The attention being given recently to the topic of automated or unattended operations is well deserved and long overdue. Without a doubt there are tremendous benefits awaiting the DP shop that will begin to automate operational functions and procedures. Those individuals who are strong proponents for data center automation exude tremendous confidence in their approach which is *not* just another case of unfounded optimism. Their certainty is based on the proven results of the numerous shops that have already embarked down the Automated Operations (A/O) path. Figure 1 summarizes some of the benefits associated with A/O.

For many, the most difficult area to deal with will be people's resistance to A/O, especially from within the operations department. Initially the majority of people within the operations group will very likely view A/O as a win-lose proposition in which the data center "wins" (becomes more efficient) and the operators "lose" (become expendable). The eventual acceptance and ongoing success of an A/O project will depend on your ability to overcome objections and get everyone pulling in the same direction.

The fact of the matter is that an A/O project should be viewed as a win-win situation for *both* the data center and the operators. The win-lose mentality has to be overcome but it should be dealt with using extreme care. One manager who recently attended a seminar asked, "Don't you really mean that the operations management wins and the systems management wins but the operators really are going to lose?" I was happy to field the

question and have another opportunity to clarify my position: A/O can definitely benefit the operator by making the job less monotonous and more rewarding.

I spoke with another data center manager where the A/O project has been implemented totally from the operations side with minimal assistance from systems personnel. His operators are extremely pleased because not only do they have more challenging jobs now, but also their positions (salary brackets) have been upgraded due to their increased level of responsibility. This is a prime example of a win-win proposition.

Definition Of Terms

In order to make sure that I mean what I say and say what I mean, it is always a

Summary Of Benefits Associated With A/O

- Improve system availability
- Better utilization of computer resources
- Opportunity to absorb growth without adding resources
- Improves quality of life for employees:
  - No need for off-shift, shift rotation, or holidays spent at work
  - Monotony decreased
- Improving the productivity of the data center:
  - Error-free operations around the clock
  - Critical procedures performed accurately and on time
  - No actions forgotten or taken too late
  - Actions occur at machine speed
  - Console operations never impacted by operator absence/turnover
  - Eliminates premium pay for overtime, holidays and so on.

Conclusion: A/O smooths out the data center.

good idea to agree on terminology beforehand.

Refer to Figure 2 for a summary of the terms that pertain to various automation efforts. Basically, A/O includes any efforts directed at accomplishing human operational functions by using machines instead. It is any attempt that moves the data center closer to being able to run on auto-pilot.

Running The Obstacle Course

As is true with so many new endeavors, there will be numerous obstacles to overcome in implementing an A/O project in your data center. Some you will see ahead of time, some you may not. Following some common sense guidelines can make a significant difference in how much opposition you encounter. Here are several ideas that you should consider when putting together your A/O project plan and for implementing its various phases.

Getting Started

"People don't plan to fail, they fail to plan" is a widely quoted one-liner that carries an important message. There will be no substitute for having a well thought out and well documented plan that explains what it is you are going to accomplish with A/O. This means making an accurate assessment of such areas as current needs, long term objectives and project resource requirements. While it is imperative that you plan carefully, it is also important to make sure you do not drag your feet! Therefore, a "cautiously aggressive" approach is recommended. It has been demonstrated time and again that automating minor functions can save an organization major bucks! Make no mis-

take about it, getting started with A/O can provide immediate gains.

To drive home the point, consider the story of one large data center with a sizable community of CICS users. CICS would infrequently issue a message that always required operator action but on occasion the message would roll off the console screen before an operator would see it. When the message was overlooked, the result was an eventual CICS crash. Following the crash it took about five minutes to bring CICS back up again. Using an automation tool, the data center implemented a trap that would intercept the message and generate an automatic reply, thus avoiding CICS crashes. By automating this one message reply, the organization has demonstrated productivity savings of $1.5 million annually.

Diplomacy Has Its Virtues

Be sure to clearly state your objectives and avoid using terminology that threatens operators. You must communicate your commitment to people (assuming, of course, that you are committed to your people) and the wording of your stated goals and objectives should reflect this commitment, not obscure it. If you are not already a politician, then consider the benefits of using a little diplomacy. There is minimal benefit to emphasizing terms and phrases such as "unattended" and "reduction in head-count" when you are actually trying to upgrade jobs within operations, increase productivity and provide opportunities for lateral transfers into other areas.

DP shops are using A/O as a means of providing better and more productive operations. It is not a diabolical plot to trim operator head-count. In many instances, A/O has meant that shops have been able to absorb and manage data center growth without increasing existing head-count.

Clarify Your Goals

Automation goals and objectives may differ greatly from one installation to the next. Perhaps you already are running a tight ship and have implemented effective operational standards. In this instance an A/O project may not be viewed as a pressing issue. On the other hand, maybe your data center is out of control!

Resource constraints keep you in fire fighting mode! If you are in this boat, then let me suggest that you can get ahead in the long run if you will quit fighting fires long enough to go find the arsonist. In either case, regardless of the urgency

involved, a successful A/O project will save the organization many times more dollars than it would cost to implement it.

Many shops have been working for years on automating different aspects of their operations. These shops already have achieved unattended operations in the sense that they run without operators at night, on weekends and during holidays. Other shops have established aggressive plans for becoming an unattended data center in the not-too-distant future.

On the other end of the spectrum, there are still numerous shops that are just beginning to grasp the potential of A/O. Due to a lack of understanding about A/O, it is not uncommon to find management confused about what A/O means. For example, some individuals may suggest that message suppression is a major component of A/O. If your shop has been writing and testing user exits to run under the MVS Message Processing Facility (MPF), then you are right to consider the chore somewhat laborious.

However, message suppression can be viewed as a somewhat trivial and simple aspect of A/O activity if you are using an automation management software package. Shops that have invested in this type of software report that the console message rate drops by 50 percent or greater

Definition Of Terms

Automated Operations — *using software to perform operational tasks previously done by operators.*

Automated Console Operations (ACO) (IBM's definition) — *the use of techniques and facilities to perform a large subset of those tasks that are traditionally performed by console operators in each of the functional areas. ACO helps you centralize operations and establish remote operations.*

Unattended Operations — *using hardware and software to manage a data center with minimal intervention or participation by personnel. In most companies today, this term is used to describe an environment where the system can run unattended during non-prime shifts, weekends and holidays.*

Principally Unattended — *where only one or two operators are needed to monitor the system.*

Lights-Out Operations — *The physical separation of data center personnel from the machine room and its hardware. There is no need for personnel at any of the sites in the installation.*

Dim-Bulb Operations — *basically the same concept as Principally Unattended.*

Centralized Operations — *All application, system or network operations are channeled to one system that serves as a focal point for operations.*

Remote Operations — *No operations personnel are on the remote site. Console operations are eliminated at the remote site and a system, systems or operators at another site control the systems at the remote site.*

Continuous Operations — *The system is available 24-hours-a-day and seven days-a-week without any disruptive changes or disruptive maintenance.*

Unintentionally Unattended — *People are around but...*

just by installing the A/O product. With a little bit of effort and by using the simple rules provided by these A/O products, it is not uncommon to reduce a message rate of eight to 10 messages per second down to one message every few minutes.

Today's automation technology is just beginning to sprout its wings. The role of automation in the data center reaches far beyond message suppression. In addition to console operation, areas already affected by A/O include job scheduling, job restart, system monitors, network management, DASD management, tape handling, print handling, report balancing, report distribution, help desk, environmental monitors and security controls.

Broaden Your Vision Of Operations

For most DP shops, the best results for automating the data center will be achieved if you let the operations department participate as much as possible. Actually, the most productive A/O projects that I have encountered are those in which operations personnel are running with the ball. Naturally, systems personnel should contribute as necessary to install products and perhaps handle complex automation procedures such as IPL.

The dilemma facing A/O planners is trying to put together the right team by recruiting from rival leagues. The long-standing feud between operations and systems personnel has been waged for years. The problem centers around the fact that while operators may respect members of the system's staff, this does not mean that they like them. On the other side of the fence, systems personnel often view operators as some primitive life form and interact with them in a condescending manner.

So where does this leave you? The fact of the matter is that when moving forward with an A/O plan, the operators fully understand the problems but do not fully understand the solutions. Systems staff understands the solutions but does not fully understand the problems. Getting cooperation from both groups is essential.

The A/O game is best played as a team sport and your job is to pick the right team members. The best candidate from operations may well be the hot shot who is always getting into trouble. The best candidate from systems is your basic human being: someone who can speak English to "non-techies" and who can get along well with others.

Increase your end users' productivity without tying you up.

If your biggest hang up is time lost while solving users' problems by phone, you need to consider our MULTITERM family of session managers for MVS, VSE, and VM systems.

Now problems that users experience with systems or applications can be brought straight to your own terminal! That's because our terminal session managers can display a user's screen on your own terminal screen using a comprehensive set of Help Desk facilities.

But our session managers don't stop there! Automatic log-on, "hot key" session switching, menu driven user profile definition, GIVE and TAKE session facilities, and

message broadcasting are but a few of the many features specifically designed to help you and the end user increase productivity. A virtual channel-to-channel link is also included to allow VTAM users access to CMS or for VM users to access VTAM applications.

All these features, combined with superior worldwide support make our MULTITERM family the ideal session managers for your MVS, VSE, and VM systems.

Call us today at **800-826-0313** and see for yourself how our family of terminal session managers can help you keep your end users productive without impacting your own productivitiy.

The MULTITERM Product Family

BlueLine™
SOFTWARE INC.

New Advances in Performance, Productivity and Planning.

1500 South Lilac Drive, Suite 340, Minneapolis, Minnesota 55416
800-826-0313 • 612-542-1072 in Minnesota

CIRCLE #48 on Reader Service Card ▲

ICP MILLION DOLLAR AWARDS WINNER

Tomorrow's Data Center

The operator's function is already extremely critical; however, it clearly has periods where it can get repetitious and monotonous. A/O removes the repetitious and monotonous and allows operations personnel to design, test and implement A/O solutions, making their job more challenging. Some data centers have created a new position within operations just for this purpose.

How broad is your vision of what operations is now versus what it can become tomorrow? The potential of the A/O opportunity needs to be grasped and you must begin to share your vision with others. A/O should improve the role of tomorrow's operator, making the function of operator into a more challenging and satisfying position. Operators will be challenged to develop their technical, analytical and interactive skills in order to contribute in tomorrow's data center.

Note that the individuals involved with A/O should have a "programming mentality" in order to develop A/O solutions using CLISTs, REXX and/or rules languages. Technical and analytical skills will be required in order to research problem areas and design solutions to the problems. Interactive skills are necessary in order to participate effectively as the A/O project team charts its course, sets priorities, conducts design reviews and keeps others informed. The use of CLISTs or REXX will be required, so make sure that whoever is participating in the A/O project receives proper training.

Do Not Try To Convert The World

As you begin to move forward with A/O, do not make the mistake of trying to get everyone to agree ahead of time that A/O is a good idea. It is okay to have people disagree with the stated objectives as long as they are willing to support the effort. Stated another way, strive for total participation rather than total agreement. Let the people come around in their own time and let the A/O results stand on their own merits. Your main goal is not to get everyone to share your point of view, it is to get the job done.

Communicate Well And Score Early

It is also important to note that A/O projects that run smoothly are also those committed to keeping interested parties well informed. This includes regular status meetings to review plans and frequent memos to keep people abreast of activities.

In setting priorities for what you want to tackle first, try to pick an area that is easy to automate, is highly visible and will help prove the merits of A/O. This may mean automating a reply to a repetitious WTOR message and/or saving CICS crashes as mentioned previously. Quantify the savings and communicate the results.

Get The Right Tools

In spite of the fact that most shops already have NetView, I strongly urge that careful consideration be given to adding an automation management software package to your arsenal. There are obviously many factors to consider. However, do not overlook how critical it is to have the proper tools. The right product will help you centralize A/O controls, reduce the time frames necessary for the development and implementation of A/O controls and make testing and ongoing maintenance easier. As a word of caution, before you go looking for a new tool make sure you are using what you already have. Make sure that justifications of new products are realistic and do not overlap extensively with capabilities that you already have bought and paid for.

When your needs do exceed the capabilities of your current tools, then an A/O management package may be just the ticket. Being able to develop both simple and complex automation solutions quicker also means that you will establish momentum quicker. It also means that you will be consuming fewer development resources along the way and allowing the benefits meter to start running up savings earlier.

Although there are several A/O management products on the market, at the time of this writing there are three products leading the pack: AutoMate/MVS from Duquesne Systems (now called LEGENT Corp., Pittsburgh, PA), AF/OPERATOR from Candle Corp. (Los Angeles, CA), and OPS/MVS from MVS Software, Inc. (Houston, TX). These three products have two major points in common. They each have comprehensive functions and features and an established and happy customer base. All three of these products offer basic functions for managing message and command traffic and time dependent activities via rule-based controls. Each of these vendors also offer Personal Computer (PC) support, either integrally or as an add-on product, for handling remote console operations including IML and IPL and for contacting support personnel via a pager/beeper device.

Automate/MVS is System Application Architecture (SAA) compliant, supporting ISPF and the REXX language. It also offers some noteworthy features such as the Open Interface for VTAM which is available with Release 2 of the product. This interface allows complex rules to be written to monitor full-screen VTAM sessions. CLISTs may also be written to issue commands to VTAM sessions and receive responses back. The strategic importance of this feature is that it dramatically increases the sources of automation data, opening up a whole new realm of possibilities as to new areas that can be targeted for automation.

The latest release of AF/OPERATOR, Version 200, offers a new interface to OMEGAMON that will capture exception conditions, issue commands directly to OMEGAMON and interrogate the responses. AF/OPERATOR is SAA compatible in some aspects and Candle has a goal for 1989 to make the product fully SAA compliant.

MVS Software is the new vendor on the block and is busy carving out its own customer base. Their OPS/MVS appears to be solid technically and also has several unique features of its own. For example, consider the fact that today's automation tools are vulnerable to changes that may occur in the format of messages with a new release of a system. OPS/MVS offers a Variable Data Extractor (VDE) component that allows automation rules to remain one step removed from the actual message text. The VDE keeps message formats in a pattern database and allows the rules to reference message contents using specially-named variables. This increases the likelihood that automation rules will still trap necessary conditions in the event that a particular message's format changes.

What About NetView?

In the A/O seminars that I have presented to date, I have been surprised at the number of shops that have NetView but are still not sure just what to do with it.

This reminds me of the story about the two hunters. One hunter agreed to go out and bring back a grizzly bear if the other hunter would skin it. The hunter was sharpening his knife as his friend came

running through the door being chased by a grizzly and said, "You skin this un and I'll go get you another un!"

I think many shops may look at NetView in this light. While there are many redundant functions in NetView and A/O management software, shops by the hundreds are acquiring A/O products from Independent Software Vendors (ISVs).

NetView is first and foremost a network management tool. It has greatly simplified the network management function, allowing centralization of controls and allowing automated responses to certain network problems and activities. After much interest/pressure from SHARE and GUIDE members, IBM needed to respond to the requirements of A/O. NetView Release 2 was a natural delivery vehicle for providing A/O functionality and included several features targeted at A/O.

NetView works in conjunction with a standard facility in MVS called the Message Processing Facility (MPF). NetView Release 2 added a message automation table but still relies on MPF to handle MVS console traffic. NetView Release 3 contains still more A/O bells and whistles. Make no mistake about it; it has been clearly demonstrated that NetView and MPF can indeed be used to accomplish much A/O activity. The question you must answer is whether or not NetView is the best alternative based on the cost of both human and machine resources that you are willing to invest in your A/O project.

Skeptics are not convinced that NetView is really suited for the long term needs of A/O. This type of speculation is further fueled by rumors that IBM's own advanced work on unattended operations does not include the use of NetView. Furthermore, IBM has suggested that a lot of their future contributions to the A/O effort will be hardware oriented. Putting speculation aside, the biggest indictment against using NetView as your primary A/O tool is that most of the companies I know that have NetView have already purchased or are considering purchasing another A/O management package.

NetView is also getting some competitive heat from another network management package. A recent Datapro survey compared NetView with Cincom's (Cincinnati, OH) NET/MASTER. Based on criteria such as ease of use, ease of installation, power and enhanced functionality, the survey concluded that NET/MASTER out performs NetView.

I have only highlighted a few of the more prominent A/O management packages. There are numerous other A/O related products being offered from a variety of ISVs. Since things can change rapidly in this industry, prudent software evaluators should make sure they are working with the most up-to-date information available when comparing products. However, since it may become too time consuming to watch competitive products play the age old game I call "features leap frog," my recommendation is that you carefully follow the proven approach of analyzing needs versus benefits.

1) Select candidate products being offered by vendors that you can trust and make a thorough product evaluation.

2) Compare your shop's requirements and priorities against the benefits offered by the product.

While a long list of features may look impressive, the real value of a product is only determined by how well it measures up to your specific needs.

Low Or High Profile

At any point of the project, before a new area is addressed it is important to choose between two basic approaches: the high-profile approach versus the low-profile approach. For each phase in a project, your chosen approach is an important tactical decision that may mean the difference between smooth sailing or encountering extreme turbulence. When upper management is already heavily involved with the A/O project, then high-profile activity is your only safe alternative because you are expected to keep everyone well informed of progress.

On the other hand, when you are dealing with individuals who are either openly opposed to automation or who may be fence sitting about a particular issue, then a low-profile position may be warranted. For example, the systems staff in a shop that had been involved in automation for some time saw the need for acquiring A/O management software but management did not. The systems staff evaluated automation tools on the market and selected one that met their requirements. After some specific rules were set up to automate some areas, a demonstration of the capabilities and benefits of the package won management's approval to move forward. The need had been there and was real but it was not easily communicated until the solution to the problems were fully demonstrated.

In another case, systems personnel wanted to implement an electronic report distribution system but the end users would have nothing to do with the idea. Systems people moved ahead and brought in an electronic report distribution software package and gently coaxed an influential end user to help evaluate the package for two months. At the end of the period, the end user that did the evaluation was so excited with the capabilities of the software that he quickly rallied the others to give total support to the idea. The moral of the story here is that people are more likely to accept change when it is their own idea as opposed to having it forced on them.

Summary

While general guidelines will always leave some readers sucking a lot of wind, the suggestions offered here should apply to more than 80 percent of the DP shops considering A/O. It is important to know ahead of time that the biggest obstacles you will encounter may be political rather than technical. Implementing a well-thought-out A/O plan requires that careful consideration be given to the concerns and needs of those most impacted by change. A/O is a new opportunity for the data center and the role of its operators. The vision of what tomorrow's data center should look like needs to be expanded. If properly handled, A/O will be a win-win proposition for all parties involved.

EDITORIAL EVALUATION

Please circle the appropriate numbers on the Reader Service Card.
1.) This article was:
   285 (Interesting/Helpful), 286 (Too Technical). 287 (Too Basic)
2.) Would you like more articles on the same subject?
   288 (Yes), 289 (No)

ABOUT THE AUTHOR
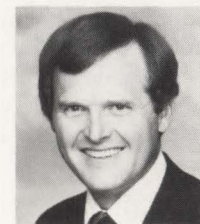
Bill Carico is President of ACTS Corporation that offers a one-day seminar for managers titled, "Automated Operations: Issues and Answers." Excerpts from this seminar served as the basis for this article. For further information contact ACTS Corporation at 1-800-950-ACTS.

Integrated Operations Architecture®
The Only Sensible Solution

Job Scheduling

Job Restart

Report Distribution Archiving and Viewing
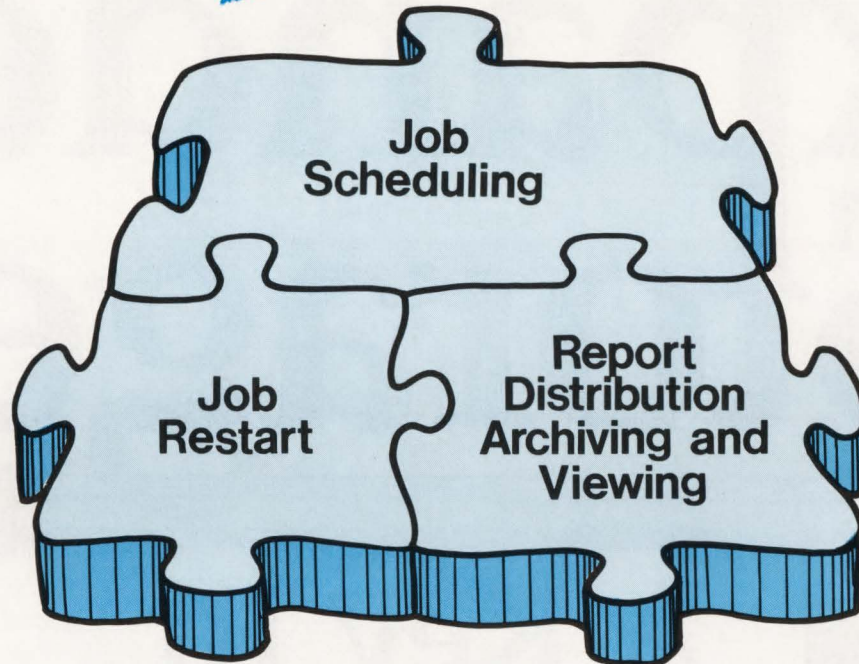
To The Automated Operations Puzzle!

Meet The CONTROL Team

CONTROL-M

- State-of-the-Art Job Scheduler
- NO system/JES hooks, SMF exits or SVCs
- 2 hour Installation
- ISPF, ROSCOE or CICS On-Line Facility
- Forecasting/Simulation
- Automatic Date/Control Card changes
- Fastest schedule implementation
- Automatically Open/Close CICS files
- Can be integrated with CONTROL-D and CONTROL-R

CONTROL-R

- Automated JOB Restart System
- Eliminates manual intervention
- Automatic catalog/GDG adjustment
- Modifies JCL as required
- Eliminates lost time and the errors associated with reruns
- NO system/JES hooks, SMF exits or SVCs
- Integrated with CONTROL-M

CONTROL-D

- Automated Report Distribution, Viewing, and Archival System
- NO system/JES hooks, SMF exits or SVCs
- Easy On-Line report definition and viewing using ISPF, ROSCOE or CICS
- NO permanent database required
- True laser printer technology
- Printer workload balancing
- Can run independently or integrated with CONTROL-M

Products Designed To Work Together

TONE Software Corp.®

Expanded Storage

A Key To MVS/ESA Performance?

By J. William Mullen

In a previous article, I discussed the implementation of expanded storage with the 3090 processors and the significant enhancement it provided to MVS performance as a high speed paging and swapping mechanism. The MVS/ESA operating system will further exploit the expanded storage technology with the dataspace, hiperspace and data windowing capabilities. Each of these new capabilities utilizes expanded storage technology to enhance the MVS operating system facilities and directly, or indirectly, to enhance the performance of the MVS/ESA operating system. While excitement exists over the potential of the new technology and its use of expanded storage, installations should be aware of the potential for performance degradation from over commitment of memory resources. This could be especially significant for

expanded storage use when the new technology must share use of the resource with existing system functions such as paging, swapping and Virtual I/O (VIO). I will review the use of expanded storage by the new technology and the necessity for activation of the mechanism for controlling use of the facilities to prevent system performance degradation.

Dataspaces

Dataspaces in the MVS/ESA operating environment are created through the DSPSERV Assembler macro facility with the allocation of memory as virtual addressable space. The default limit for virtual addressability is 239 4K blocks or approximately 956K of virtual storage. The initial impact of dataspace creation and use is on the availability of central storage frames.

Dataspace virtual pages are not backed by central storage frames until the virtual page is actually used. An unauthorized user may have up to 256 dataspaces of which dataspaces 0, 1 and 2 belong to MVS/ESA and 253 can be created by the user. This allocation would allow a user to address approximately 253MB of virtual storage which, if referenced, would require central storage frame backing. If the creating user is swappable, the frames backing the dataspace become part of the user's working set. If the user is logically swapped, these frames backing the dataspace would be retained in central storage. If the user subsequently transitions to expanded storage, the central storage frames would be transferred to expanded storage.

For most installations, the frame backing requirement should not be a problem

Announcing 3090 Central and Expanded Storage from EMC Corporation.

As a 3090 user you now have the choice to make a smart business decision. Our 3090 Central and Expanded Storage upgrades offer you the reliability, performance, local and remote support that you demand at a price 30% lower than that of IBM.

Actually, this good news is old news. For 10 years we have been offering users cost-effective, state-of-the-art storage solutions. Our strong background in mainframe storage solutions has been recognized by Fortune 500 companies worldwide. These companies have realized that we provide a better business decision for their mainframe storage needs. Now, we have built this mainframe experience into our 3090 product offerings.

For additional good news on how EMC Corporation's 3090 storage upgrades can help you reach your business objectives call or write: EMC Corporation, Hopkinton, MA 01748-9103 **1-800-222-EMC2.** (In MA, 508-435-1000.)

CIRCLE #95 on Reader Service Card ▲

EMC² The System Enhancement Company.

IBM is a registered trademark of International Business Machines Corporation.

CIRCLE #95 on Reader Service Card ▲

since most users that utilize the dataspace facility will not create more than one or two dataspaces and then not use all of allocated virtual storage. The expanded addressability provided by even this limited use of the dataspace capability could result in significantly larger swap working sets as well as total page frame movement for swapping operations. Absorption of the increased swap set load by expanded storage will allow for the increase in swap working set sizes and potential increase in page transfer operations without impact to the performance of the user or the system as in total.

The primary exposure in dataspace creation is the increase in swap set sizes and the potential for exhausting the available expanded storage frames which would result in expanded storage page migration to auxiliary storage. Performance enhancements gained through expanded storage usage will be lost through the additional overhead of transferring the migrated pages to auxiliary storage and subsequently transferring the pages back into central storage when the user is ready to execute.

An additional exposure is the depletion of available auxiliary storage frame slots. As expanded storage migration rates increase and the available auxiliary storage slots are depleted below the MVS operating system's low threshold, creation of new address spaces in the MVS system will cease until the necessary auxiliary slots required for address space backing are available. The auxiliary slot backing requirement for MVS address spaces was reduced with MVS/XA Release 2.2 but was not eliminated.

Thus, performance improvements available through high speed retrieval of data stored in dataspaces whose virtual storage pages are backed by central or expanded storage can be negated with over commitment of the resources. Without the availability of expanded storage, the expected page transfers from use of the dataspace facility would have to be absorbed by either a significant increase in the amount of available central storage or by auxiliary storage devices. Both of these alternatives would be costly and not very attractive for most installations.

Disabled Reference (DREF) Dataspaces

Some MVS operating system routines operate in the disabled state preventing them from taking certain types of inter-

FIGURE 1
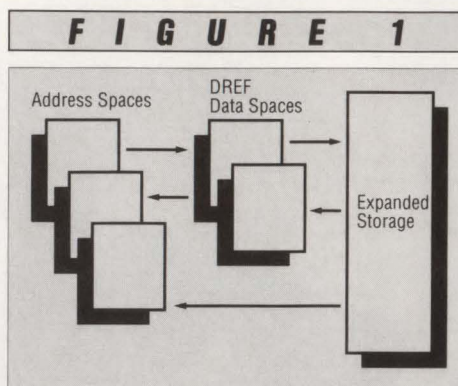
rupts such as page faults which would possibly need to be resolved by bringing a page frame into central storage from auxiliary storage. To prevent this possibility, the routines use page frame fixing to assure that, if needed, the page frames would be resident in central storage. Many of the frames that require fixing reside below the 16MB line. Address spaces also have requirements for page frame residency below the 16MB line.

As the number of concurrently active address spaces in MVS systems increases with the availability of larger processors and larger memories, page frames below the 16MB line have become a critical resource. In many cases, task swap-in failures or delays resulted from the shortage of frames below the 16MB line. Addition of central storage to the processor will not alleviate this shortage. Regardless of the number of megabytes of central storage that are added to the processor, there are only 4,000 frames available below the 16MB line. This number of available frames is further reduced since some portion of the frames are taken by the MVS system.

MVS/ESA addresses this problem through the use of Disabled Reference (DREF) dataspaces. DREF dataspaces can only be created by authorized users. The primary reason for this restriction is the handling of page frames backing DREF dataspaces. Initial frame backing for DREF dataspaces is the same as for dataspaces created by non-authorized users. The virtual storage frames are backed by central storage frames. If the central storage frames backing the DREF dataspace are needed, they are transferred to expanded storage before returning to the central storage available frame queue. Here the similarity ends. Expanded storage page frames that are backing DREF dataspaces *are not eligible for migration to auxiliary storage* (see Figure 1).

The DREF dataspace thus eliminates

much of the page fixing requirements of MVS operating system routines since these pages are assured of residency in central and expanded storage but never on auxiliary storage. The DREF facility also allows the MVS/ESA operating system to reduce frame residency requirements for frames below the 16MB line.

The dark side of the cloud for DREF dataspaces is the permanent residence status of their expanded storage pages. Given creation of many DREF dataspaces, a serious depletion of expanded storage frames could occur. While gaining performance benefits from storage constraint below the 16MB line, this gain could be negated through higher expanded storage migration rates. IBM has recommended that DREF-type dataspaces be used only where necessary and that dataspace pagable storage is a better alternative.

Additional relief for frame availability below the 16MB line is provided by expanded storage. Logically swapped users holding frames below the 16MB line will have these frames stolen and moved to expanded storage in MVS/ESA. If the stolen frames were sent to auxiliary storage, performance benefits gained from the use of the logical swap function would be lost in subsequent retrieval of these pages. The speed of transfer of pages to and from expanded storage allows movement of these stolen pages to expanded storage without negating performance achieved from the logical swap of MVS tasks.

Hiperspaces

The hiperspace facility in MVS/ESA can be used as both a high speed DASD record caching and a data storage and retrieval mechanism. The creation of hiperspaces is available through use of the DSPSERV Assembler macro and the TYPE = HIPERSPACE option or through the CSRIDAC and CSRVIEW data windowing subroutines. The latter requires that the Data Facility Product (DFP) Release 3.1 be installed on the MVS/ESA system.

Access to expanded storage through DASD record caching or data storage and retrieval is not byte addressable. Pages are moved in and out of expanded storage hiperspaces in 4K increments. Access to hiperspaces does not require the use of the ALESERV Assembler macro or Access Registers (ARs). At the Assembler language level, the HSPSERV macro is used to transfer pages to and from hiperspaces in expanded storage. For non-

VPS

Tames The Jungle of MVS Network Printer Support

VPS has been used to replace other products, such as: IBM's 328x/ADMPRINT/DSPrint, CICS supported printing, SASWTR®, RJE and many others, with a single task to drive all your 3270 family printers directly from the JES spool (including cross domain VTAM printers).

1700 Sites use VPS as their shop standard 3270 family printer driver.

✔ Printers supported include the full array of 3174/3274 attached printers, including IPDS support for 42x4's, HP and Xerox lasers, plotters and PC printers.

✔ VPS runs as a VTAM application. NO system modifications. NO JES maintenance.

✔ Automatic forms control, full FCB support, dial-up PC printers, printer pooling and "hot" printers are all supported.

✔ Full screen "ISPF-like" command interfaces for CICS, TSO and ROSCOE® permit end-user control of printers with totally menu driven command entry.

Call or write for more information, or to arrange for a FREE TRIAL
ATTENTION: Marketing Department

LRS *Levi, Ray & Shoup, Inc.*
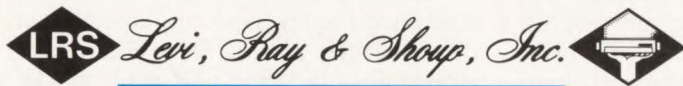
Specializing in Computer Systems Software
2387 West Monroe
Springfield, Illinois 62704
(217) 793-3800 ◆ Fax (217) 787-3286

® SASWTR is a registered trademark of SAS Institute, Inc., Cary, N.C.
® ROSCOE is a registered trademark of Applied Data Research, Inc.
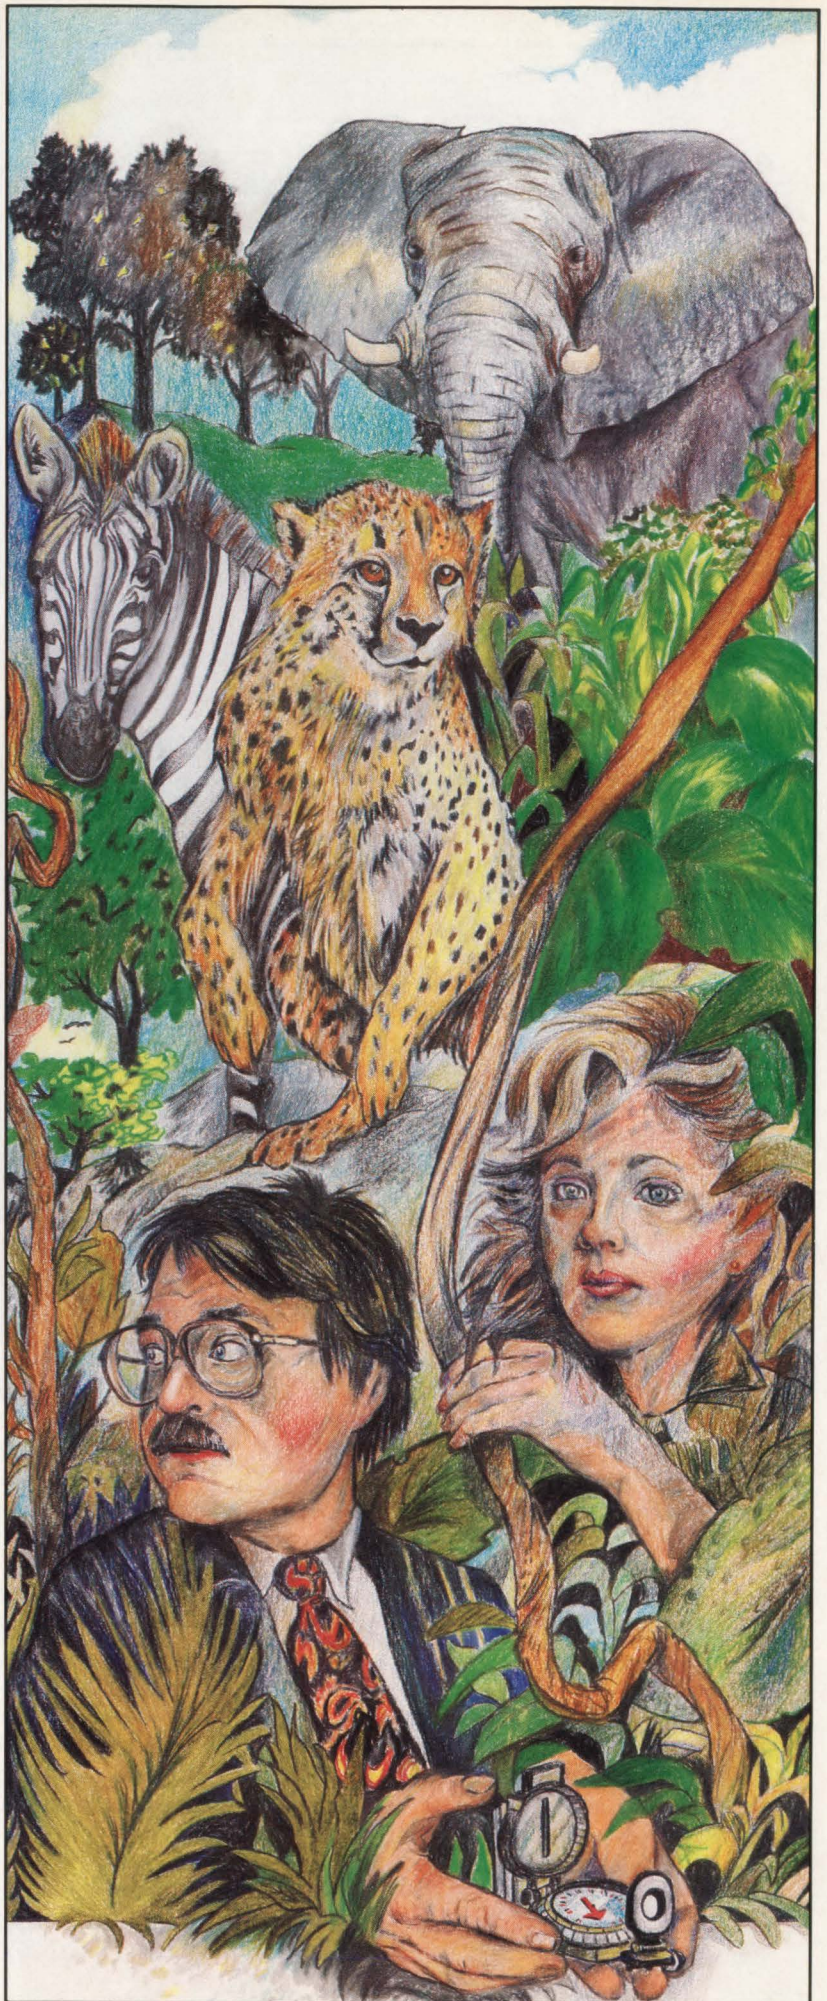
authorized users, the SREAD or SWRITE option on the HSPSERVE macro determines the direction of page movement (see Figure 2).

With data windowing callable subroutines, the CSRIDAC subroutine may be used to create a temporary, scrollable data object (hiperspace) in expanded storage. This use of hiperspaces provides the user with the ability to read and write directly to expanded storage. Frequently referenced data requiring high speed access, such as data to satisfy on-line systems requests, can be stored in a temporary, scrollable hiperspace for the life of the MVS task.

Use of expanded storage provides the MVS task with the ability to store and retrieve data in 4K pages at an average rate of 75 microseconds per page versus the milliseconds required to retrieve the same data from auxiliary storage. A permanent DASD home for temporary, scrollable hiperspace data is not required though one may exist. Access to scrollable hiperspaces is available in the same manner as if the space were a mapped Data-In-Virtual (DIV) object. The option exists for retrieval of data from a VSAM linear dataset for storing in the hiperspace or storing of task *created* data in the hiperspace.

The user may also place data from sequential access method files in the hiperspace. This option is not directly supported. Data from sequential access method files, such as QSAM and BSAM, must first be read into the user's address space and then moved into the hiperspace. It is the user's responsibility to manage this data movement.

The previously discussed limit on the number of dataspaces a user could have currently active includes hiperspaces. A non-authorized user can thus create up to 253 dataspaces or hiperspaces in any combination. The default limit on the number of pages that can be allocated for a hiperspace is 239 or approximately 956K of addressable storage. Expanded storage frames are not actually allocated to the hiperspace until the addressable storage is referenced.

Temporary or scrollable hiperspaces can and will have their pages migrated to auxiliary storage if expanded storage frame shortages occur. The difference with dataspaces and central storage frame backing is use of expanded storage pages might be delayed reducing the possibility of increased expanded storage migration rates. With hiperspaces and direct

FIGURE 2

Hiperspaces

- Authorized Address Spaces 'Cache' Type Hiperspace
- Scroll Type Hiperspaces

FIGURE 3

Data In Virtual (DIV)

reading and writing of expanded storage frames, migration requirements due to expanded storage frame shortages can be immediate.

I/O Elimination

Dataspaces and hiperspaces can provide signifiant reduction in DASD I/O operations with retention of data in storage. The I/O elimination is based upon use of DIV datasets commonly referred to as VSAM linear datasets.

With dataspaces, the user supplies the token for the dataspace when the linear dataset is mapped (opened). Multiple linear files may be mapped into a single dataspace. In this case, the user must supply the IOON parameter with the DSPSERV CREATE Assembler macro request along with the dataspace token, the starting 4K page for the data and the size of the I/O area in 4K blocks.

As discussed earlier, a hiperspace created as a temporary space may have data from a linear space mapped into the space. A temporary hiperspace would be created

with the DSPSERV CREATE,TYPE = HIPERSPACE,HSTYPE = SCROLL .... option of the macro. The token for the hiperspace is supplied in the DIV IDENTIFY for the linear dataset. The size of the *window* and the hiperspace for accessing the data would be defined in the DIV MAP request. Data accessed from the linear dataset will be automatically transferred into the hiperspace and made available to the user through the window (see Figure 3).
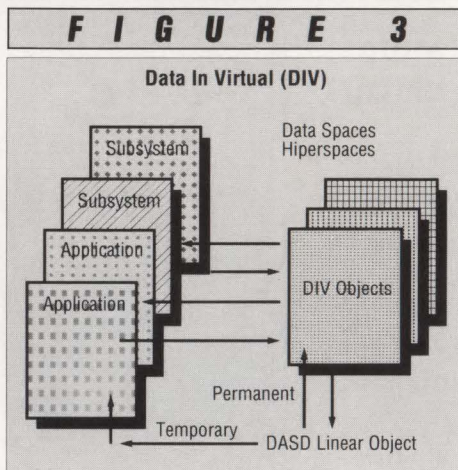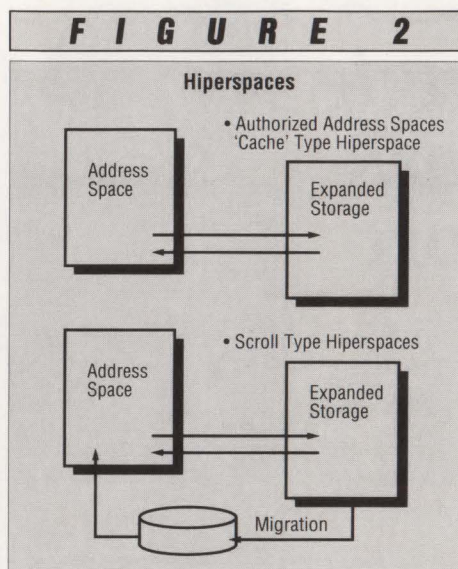
Cache type hiperspaces, referred to as Expanded Storage Only (ESO) hiperspaces, are reserved for authorized users only. Except in cases of critical expanded storage frame shortages, ESO hiperspaces will not have their pages migrated to auxiliary storage. The ESO hiperspace is created much the same as the scrollable or temporary hiperspace through the DSPSERV CREATE,TYPE = HIPERSPACE,HSTYPE = CACHE.... parameters of the Assembler macro. The authorized user then goes through the same identify and map process for the linear dataset as the non-authorized user creating a scrollable hiperspace.

Storing of large amounts of data in hiperspaces will provide significant performance enhancements, particularly for CICS/VS and IMS/VS on-line systems with transactions that have large I/O access requests requirements. Use of expanded storage can reduce or eliminate much of the DASD I/O requirements without the increase in central storage paging rates often seen with on-line systems when trying to buffer large amounts of data in storage.

The second option available for I/O reduction is through Large Virtual Buffering and the creation of Local Shared Resource (LSR) buffer pools for VSAM datasets in dataspaces or hiperspaces. With MVS/XA, VSAM LSR buffer pools are created with the BLDVRP Assembler macro and are constructed in virtual storage in the user's address space above the 16MB line. With MVS/ESA, the buffer pools may be above the 16MB line in the user's address space or directed to a dataspace or hiperspace as an option of the BLDVRP macro. Data accessed is written directly to the LSR pools in 64K increments.

Building the buffer pools above the 16MB line in the user's address space or in a dataspace constitutes an increase in the user's page working set size. If the user is swappable, these pages will subsequently be transferred to expanded stor-

age, if available, and possibly migrated to auxiliary storage if expanded storage frame shortages occur. By directing the buffer pool construction to a hiperspace in expanded storage, a reduction in the number of pages that must be handled by MVS when swapping the user will be gained. Also, only the expanded storage pages of data referenced by the user will be brought into the user's address space. Use of expanded storage for LSR buffer pools would potentially reduce system demand paging by reducing the number of frames in the address space working set while allowing the user to retain large amounts of data in storage (see Figure 4).
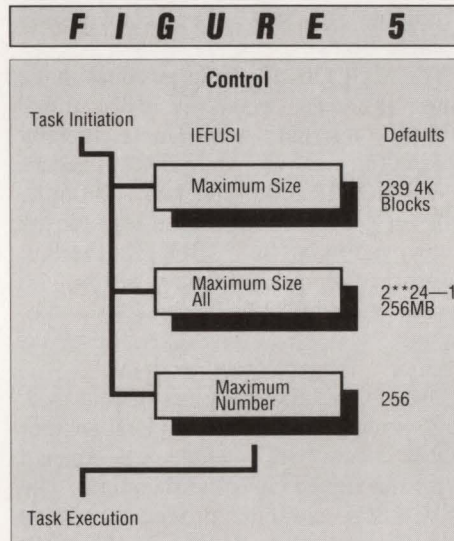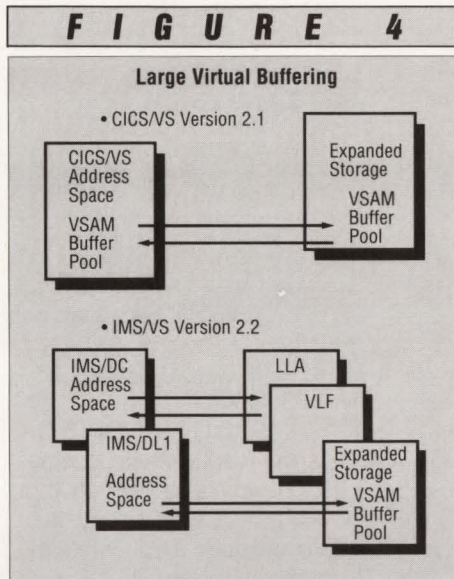
The only restriction is on the size of LSR pools. As of this writing, the author assumes that the limit on the number of buffers in an LSR pool, 32,767, and on the size of the pool, 16MB-1, is still in effect when directing the buffer pool construction to dataspaces and hiperspaces. Of course, with subsystems such as CICS/VS, this restriction could be circumvented through creation of multiple LSR pools.

Expanded storage pages in the LSR pools are eligible for migration to auxiliary storage. Creation of excessive pools could result in extensive expanded storage page movement and migration activity. Again, this would result in data being brought in and out of storage through paging activity which could negate any performance benefits gained through the use of the LSR facilities.

Control

There have been several articles written on MVS/ESA dataspaces and hiperspaces and the expectations for performance improvement through use of the new facilities. An important subject that also needs to be addressed is how an installation protects its system environment from over commitment of memory resources from use of the new facilities.

Control over of the number of concurrent dataspaces and hiperspaces that can be created by a non-authorized user is available through the standard System Management Facility (SMF) step initiation (IEFUSI) exit. The maximum size for addressable storage in a dataspace or hiperspace is two gigabytes. The default limit for a non-authorized user in MVS/ESA is 956K. The default for the number of dataspaces and hiperspaces concurrently active for a user is 256. Each of these default values may be changed by an installation-written IEFUSI exit. When

FIGURE 4

Large Virtual Buffering

- CICS/VS Version 2.1

CICS/VS Address Space

VSAM Buffer Pool

Expanded Storage

VSAM Buffer Pool

- IMS/VS Version 2.2

IMS/DC Address Space

IMS/DL1 Address Space

LLA

VLF

Expanded Storage

VSAM Buffer Pool

FIGURE 5

Control

Task Initiation

IEFUSI

Defaults

Maximum Size — 239 4K Blocks

Maximum Size All — 2**24—1 256MB

Maximum Number — 256

Task Execution

the exit is executed at a task step invocation (batch job step, TSO LOGON and so on), the exit will be passed a three-word parameter list. The installation may alter the values in this parameter list to limit both the number and size of dataspaces and hiperspaces (see Figure 5).

Though most installations do not currently use the step initiation exit, they should consider generating an IEFUSI exit and limiting at least the number of dataspaces and hiperspaces a non-authorized task could create. Though neither are backed by real storage frames until the addressable area is used, a non-authorized task has the ability to create and reference up to 253MB of addressable storage in dataspaces or hiperspaces. It is obvious that a non-authorized user that created 100 scrollable hiperspaces and subsequently referenced the addressable area would have a negative effect on system performance. It should be noted that the limit of

256 dataspaces or hiperspaces is not applied to authorized tasks.

Summary

Expanded storage in 3090 systems has provided the technology which makes the dataspace and hiperspace facilities in MVS/ESA feasible. Without the high speed page transfer rates available from expanded storage, implementation of the dataspace capability would probably have been the only feasible addition to MVS capabilities and then the requirements for additional central storage would have prevented many installations from taking advantage of the new technology.

As subsystems and application systems are upgraded, use of dataspaces and hiperspaces will become a standard mechanism in many installations as a method of satisfying needs for high speed access to data. Use of dataspaces and hiperspaces by system services, such as the Virtual Lookaside Facility (VLF) for library data object and catalog object in-storage management, will alleviate many of the performance bottlenecks present in the current MVS architecture.

The two primary areas addressed by most installations will be making effective use of the new technology while being cognizant of both the performance improvements as well as the performance degradation that can occur with the new technology.

EDITORIAL EVALUATION

Please circle the appropriate numbers on the Reader Service Card.
1.) This article was:
   290 (Interesting/Helpful), 291 (Too Technical), 292 (Too Basic)
2.) Would you like more articles on the same subject?
   293 (Yes), 294 (No)

ABOUT THE AUTHOR
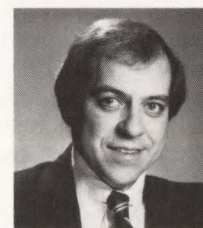
J. William Mullen is a senior consultant for BGS Systems, Inc. He has more than 20 years of data processing experience and has concentrated on the performance measurement and capacity planning area for the last 15 years. Mullen is a past officer and director of SHARE, Inc. and CMG, serving for the last eight years as editor of the CMG Transactions.

cause the greatest problems. Some differences are obvious when seen in a compilation. VS COBOL was more lenient in a number of areas in which COBOL II has decided to strictly enforce the rules. For instance if the file name in an FD statement began in column 11 rather than column 12, it was accepted and no warning was given. The same statement will cause an 'E' level error (condition code 8) in Release 2 of COBOL II. Another example of enforcing restrictions is in values given to an 88 level for an alphanumeric field. When numeric values were given to an alphanumeric field in VS COBOL they did not have to be enclosed in quotes. Now the COBOL II compiler enforces the rule that the values in an alphanumeric field must be defined within quotes or apostrophes. The code that previously compiled with no errors now gets an 'S' level (condition code 12) error.

While the differences just discussed are annoying, even more problematic are those differences that may not be found until an appropriate condition occurs which may not appear in regression testing. For example, if a field defined as PIC 99 was moved to a field defined as PIC S99, VS COBOL moved the field and then inserted a positive sign in the receiving field. This may have prevented an 0C7 program check that would have occurred later in the program. Depending on options used in a compilation, COBOL II may just do a straight move character which is faster but may generate different results. Until Release 3, many did not realize that none of the compiler options instruct COBOL II to create code that exactly matches numeric processing in VS COBOL.

As far as efficiency is concerned, there have been changes that affect it in both directions. Loading of BL cells has been improved and many places in the code that previously had to load a cell no longer do so. The IBM optimizer creates some interesting efficiencies. For instance, if you PERFORM a paragraph in a program, the optimizer will copy the code to make it an in-line PERFORM. If the paragraph is small and is PERFORMed a number of times, the optimizer will change a few of the PERFORMs to in-line. After a limit is reached (the program size cannot increase by more than 50 percent), the PERFORMs are no longer changed to inline. Executable code for the PERFORMed paragraph is always produced in the paragraph's original position.

On the other hand, some code pro-

F I G U R E 5

```
        DATA DIVISION.
        WORKING-STORAGE SECTION.
        01    WS-SIGNED-COMP-DBLWORD          PIC S9(16) COMP VALUE +100.
        PROCEDURE DIVISION.
10            MULTIPLY +1 BY WS-SIGNED-COMP-DBLWORD.
20            MULTIPLY WS-SIGNED-COMP-DBLWORD BY WS-SIGNED-COMP-DBLWORD.
```

	VS COBOL	COBOL II
Line 10	2	0
Line 20	4	1

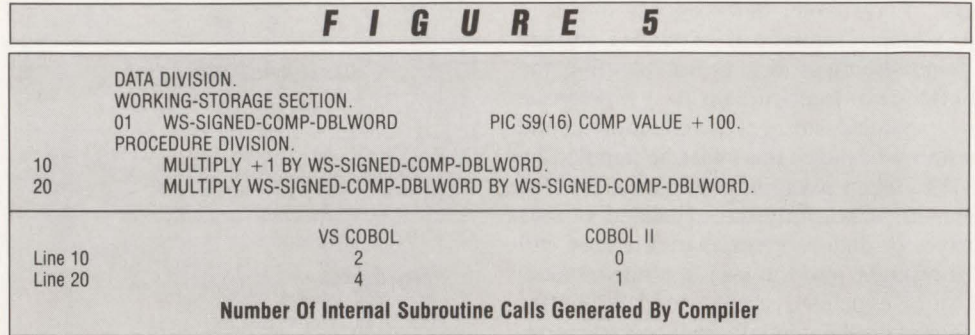**Number Of Internal Subroutine Calls Generated By Compiler**

---

duced in COBOL II is considerably less efficient than VS COBOL. In the statement IF FIELD1 (SUB1) = 1 OR 2 OR 3 OR 4 OR 5, the displacement computation for FIELD1 was done only once in VS COBOL while in COBOL II it may be recomputed five times. Even when the COBOL II optimizer is used, it still may recompute the same displacement five separate times!

The COBOL II compiler does have some instances when it calls fewer COBOL internal subroutines, thereby making the code more efficient. In Figure 5, VS COBOL executed two subroutine calls for the first multiply and four for the second multiply. In COBOL II, the first multiply is done completely in-line and the second with only one call to a subroutine. The efficiency does, however, produce a longer object program.

One change that may seem quite baffling occurs when a constant or another COMP-3 data field is added or subtracted to/from a signed COMP-3 data field. The VS COBOL compiler produced only an Add Packed (AP) or Subtract Packed (SP) instruction. COBOL II still produces the same instruction but then issues a Zero and Add Packed (ZAP) of the resulting field into itself. The ZAP is executed for a subtle reason. When two negative numbers are added together and an overflow occurs, the overflow flag in the condition code is turned on while the resulting sign is set as if overflow did not occur. This means that if two numbers each defined as PIC S9(3) COMP-3 were added together and they had values of -1 and -999 before the addition, the resulting field would be zero with its sign negative. The ZAP will preserve the sign of all other computations but will change a zero with a negative sign to a zero with a positive sign.

Ensuring a positive sign in a zero result was not necessary in the VS COBOL compiler. This is because the VS COBOL compiler used a Compare Packed (CP) instruction to compare numeric fields. It is

interesting that COBOL II now usually produces a Compare Logical Character (CLC) instruction to compare numeric fields. While this is a faster instruction than CP, it requires the same sign value for an equal condition to occur. For example, a CP instruction will find fields with hexadecimal values of '0F' and '0C' equal while the CLC will not.

The ZAP added to arithmetic instructions causes an add to an unsigned field to often be faster than an add to a signed field (the unsigned add will generate an OI instruction to strip the sign, an instruction faster than the ZAP). One idiosyncracy of the compiler is that when the ON SIZE ERROR clause is used in the ADD or SUBTRACT statement, the code produced does not contain the ZAP to correct the sign.

The code produced to clear short fields to spaces has changed. The compiler used to first move a space to the first position of the field with a Move Immediate (MVI), then clear the field by using an overlapping move. This did not require any spaces to be stored in the literal pool of the program. The COBOL II compiler stores spaces and clears short fields to spaces with a single move. The code produced for subscripting and indexing has improved and is now shorter.

Conclusion

Where do all of the COBOL II differences leave us? I personally find many of the new changes fascinating and complex. Had I designed specifications for COBOL II, I certainly would have included more useful options than de-editing and allowing code to be written in lower case characters. Regardless, I enjoy being a technician and will have fun with the new features and their idiosyncracies. However, from the standpoint of the hype that claims that COBOL II will "improve programmer productivity, improve efficiency and ease maintenance," I am not impressed. Complexity and additional features leave more for a programmer to

---

understand, making it harder for a programmer to maintain code. There is now a greater possibility that a programmer will not be familiar with the subset of COBOL features that another programmer has decided to use in code. As always, judicious choices of the available coding options will produce more efficient, easily maintained code while, as always, injudicious choices will send the program in the opposite direction.

The compilation time of a program is certainly slower. The difference in speed is considerable enough that you can literally feel the difference in time when you do an on-line compile. What about the speed of program execution? IBM claims that programs will now run considerably faster. From the people I know that have done minimal studies in execution time comparisons, one company claims that the COBOL II programs run between five percent slower and five percent faster when compared with VS COBOL while another company says that they run from slightly slower to about the same speed.

Regardless of whether or not COBOL II is easier to use and maintain and more efficient, it is here to stay. IBM has created a product that we must all convert to, whether we like it or not. What I cannot understand is why it takes a company like IBM three releases of a compiler before realizing that something as fundamental as compatible numeric processing in a compiler is needed so that programs can be easily migrated? ≜

EDITORIAL EVALUATION

Please circle the appropriate numbers on the Reader Service Card.
1.) This article was:
   260 (Interesting/Helpful), 261 (Too Technical), 262 (Too Basic)
2.) Would you like more articles on the same subject?
   263 (Yes), 264 (No)
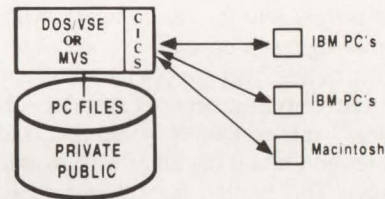

ABOUT THE AUTHOR

Harvey Bookman is President of Bookman Consulting, Inc., a software development company specializing in programmer proficiency testing. He is responsible for the development of the expert system testing service, TECKCHEK™. Bookman Consulting, Inc., 67 Wall Street, Suite 2411, New York, NY 10005, (212) 819-1955.

CIRCLE #167 on Reader Service Card ▲

tuned the system before doing the study (actually, as part of the study). The tuned system problem is solved.

That should help you get started with the task. But what about those people who do not believe that a process is necessary at all? They adhere to the next myth.

■ Capacity planning can be adequately performed with rules of thumb.

This is actually far more dangerous than the MIPs concept. At least with MIPs bigots there are measurements that relate to a specific value that does supply some limited insight into the environment. Rules of thumb suggest a dependency on pseudo-scientific beliefs that are not based on data from the environment. Consider the "medical" rule of thumb — starve a cold, feed a fever. Or is it the other way around; no matter. This implies that all fevers have a common cause which can be corrected by a single cure. Clearly, you need more data specific to the individual who is feverish.

The same is true for data processing rules of thumb. For example, the rule is that an on-line system should not run on a CPU with more than 80 percent utilization. Is this daily average, hourly peak, five-minute interval? The rule of thumb leaves room for error. In addition, I have seen many systems with adequate service that run consistently in the 90 percent range. It depends on the application, the system, the ancillary subsystems, the desired service levels and so on. The same holds true for the most closely cherished rule of thumb of all — DASD utilization should not exceed 35 percent.

Rules of thumb are okay only until you have examined your own system and determined what is right for you. Then you will have specific rules of thumb that work for your environment. You will still need to recheck these values periodically to ensure their continued validity. Sounds to me like what we need is an ongoing process.

Before you go off thinking that there are no rules that you can base your capacity planning process on, let me replace the ones above with a new set. These will, I believe, serve you well for a long time.

New Set Of Rules

■ Know Your Customer

In a broader sense, know what your company does for a living. Unless you have this understanding, it will be diffi-cult to relate to the needs of each segment of the company. This may require that the analyst occasionally review trade publications and establish relationships with peers in non-DP functions within the organization.

Optimally, this will be a two-way process in which you will learn about the "business" and will help to educate others of the value of DP. It is also important to understand the organizational structure and the corporate decision-making process. This will assist the capacity planner in determining the best point in the company to present information for optimal results.

■ Satisfy Your Customer

This is an age-old axiom in business that still holds true. In the process of knowing your customer, it is key that you gain insight into his needs and requirements. Success is tied to satisfying those needs. As a part of satisfying the customer, concerns such as timeliness, accuracy, reliability and cost must be addressed.

■ Gain the Customer's Confidence Through Positive Achievements

Nothing succeeds like success. Take a positive attitude and approach when performing your tasks. Exert some effort in trying to help the organization derive the maximum utilization from available resources. Develop an effective reporting system. Suggest alternatives that eliminate bottlenecks. Prove that the capacity management function is worth the investment. Do it in a positive cooperative manner.

■ Present Multiple Solutions

Whenever possible present more than one alternative for consideration. This is especially important when dealing with higher levels of management. Each option should be clearly delineated with benefits and impacts, costs and timeframes. This will help to get others personally involved and committed to the resulting decision. (Please read on.)

■ Take a Chance (Within Reason)

While you are advised to present multiple options, do not be afraid to support the one that you think is best. Also, if you feel strongly, present original ideas and solutions. Gambles that are based on sound foundations come to be known as investments.

■ Establish Defined Plans

Never appear haphazard and disorgan-ized. Keep your tasks in perspective and document each milestone. The matrix approach presented in this article has this concept as one of its cornerstones.

■ Make Decisions

Management is not impressed by indecisiveness. If you hope to advance, you must exhibit leadership qualities.

■ Constantly Review

Always remember that any process left to its own will eventually atrophy. Set up a viable feedback system that points out the strengths and weaknesses of the process. Correct the weaknesses and publish the strengths. Continually improve the process.

To sum it all up, the basic rule should be — take time to understand what is right for your environment. That means that you will need to analyze your data processing configuration (hardware, software and workloads), your users, your management and your corporation. Since capacity planning, by definition, is tailored to service all of these areas, you should be able to utilize the valuable inputs that each has to offer. Do not neglect any of them and do not look for the mythical easy way out. ◈

ABOUT THE AUTHOR

Jerry L. Rosenberg, co-founder of SRM Associates, has 20 years of experience in the areas of capacity management, performance analysis and systems engineering and support. His background includes hardware evaluations, physical site planning and systems management. As an independent consultant, Rosenberg has conducted numerous consulting studies involving performance and planning decisions for hardware, software and application subsystems. He is also a frequent speaker on the subject of capacity management addressing the concerns of both technical and managerial personnel. SRM Associates, 12 S. Oaks Blvd., Plainview, NY 11803, (212) 484-2601.
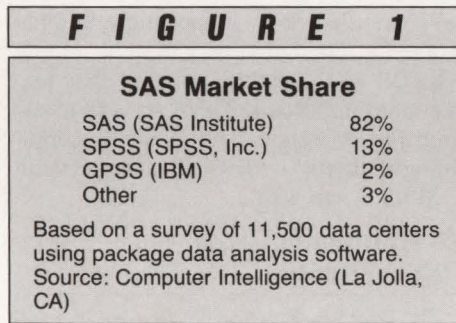---

SAS System
A Program for All Reasons

By John Kador

In the course of preparing these product reviews for MAINFRAME JOURNAL, I have talked to literally hundreds of data centers. One question I always ask is, "What other software packages do you use in your data center?" In the area of data management and analysis, the product that is most often mentioned is the subject of this month's product review — the SAS System from SAS Institute Inc., Cary, NC.

The SAS System is, indeed, one of the most widely-installed data management packages in the computer industry today (see Figure 1). It started out as a statistical analysis support package, grew into a graphics presentation system and evolved in a modular fashion into the big system it is today. Today, the SAS System includes most of the desired facilities normally associated with database management, query and report writing, decision support, graphics and statistics. It is the flagship product of SAS Institute Inc., the nation's largest privately-held computer software company.

My experience is consistent with its market share. SAS is installed in more than 10,000 computer centers around the world, including more than 90 percent of the Fortune 1000 and 99 of the nation's largest data processors. With the possible exception of SyncSort (SyncSort, Inc., Woodcliff Lake, NJ), there is no one system from an independent software company that comes as close as SAS to saturating its market. For that reason SAS Institute, which originally offered SAS for IBM mainframes only, has now migrated it to the Digital Equipment line of VAX minicomputers and workstations, minicomputers made by Data General and Prime IBM PCs and UNIX-based work-

FIGURE 1

SAS Market Share

SAS (SAS Institute)	82%
SPSS (SPSS, Inc.)	13%
GPSS (IBM)	2%
Other	3%

Based on a survey of 11,500 data centers using package data analysis software. Source: Computer Intelligence (La Jolla, CA)

stations. The Institute is also putting a lot of effort into marketing SAS as a capacity planning tool.

SAS is a prime example of a product tailored under Multivendor Architecture (MVA). First conceived in 1984, MVA provides a software architecture that maximizes the SAS System's ease of migration from one operating environment to another. With MVA, SAS is vendor-independent, running identically across all environment's SAS supports. The Institute's decision to rewrite the entire SAS System in the C programming language helped initiate its move to MVA.

Competition

There is no one product that competes with SAS on a function-by-function basis. Rather, many products compete with SAS in specific operational areas. In the area of reporting, statistics and business analysis, SAS's closest competition is SPSS (from SPSS Chicago, IL). The SPSS system provides many of the same capabilities as SAS and serves a broader range of machines but is not as portable. Each SPSS system operates in a unique environment.

In the graphics area, such products as DISPLAA and Tel-A-Graf (both from

Computer Associates International, Garden City, NY) provide IBM mainframe users with graphics capabilities that exceed those offered with SAS/GRAPH.

In the area of data management and report writing, SAS competes against such entries as EasyTrieve Plus (Pansophic Systems, Oak Brook, IL). Both products offer powerful data management facilities. SAS, however, tends to be more portable for use in third-party software environments. The Pansophic product is primarily designed to support the report writing capabilities of other Pansophic products.

For some applications, SAS also competes with Information Center DBMSs such as FOCUS (Information Builders, New York City), RAMIS II (On-Line Software International, Ft. Lee, NJ) and NOMAD (Must Software, Norwalk, CT). Over the years, these DBMSs have expanded to incorporate sophisticated data manipulation, statistical analysis, reporting and graphic presentation functions. These tools can offer the user similar functionality, but typically at a higher cost in terms of licenses and learning curves. The internal data manipulation techniques employed in these DBMSs are more sophisticated than those in the SAS System.

Design

In design, the SAS System is modular. It consists of base software and add-on software products that extend and enhance basic system functions. The software in the base version of SAS revolves around two fundamental operations: a DATA step that enables users to create the dataset to be worked on and the PROC step that defines the application that will operate on the dataset. Some of the pro-

cedures available in base SAS software are a sort routine, a routine to compute frequency and a routine to develop means and correlation data. It also offers a full complement of tools for statistical analysis including procedures for regression analysis, analysis of variance, categorical data analysis, multivariate and discriminate analysis, clustering and scoring.

The optional components of SAS, available under a complicated licensing schedule, provide users with high-quality graphics capabilities, economic analysis features, interactive query and editing functions, letter writing capabilities and spreadsheet functions. Other optional add-ons include an operations research module, a component to help in developing statistical quality control applications and a function that allows users to share data through multiple access and concurrent update capabilities.

The system also provides interfaces to most popular DBMS' including IBM's IMS/VS, DB2 and SQL/DS, Oracle Corporation's Oracle, Cullinet Software's IDMS/R, Computer Associates' DATACOM/DB and a direct link to SAS Institute's own SYSTEM 2000 DBMS. One of the chief attributes of SAS is that it is easy to use. It is widely praised by non-technical end users for its user friendliness; yet, it is accepted by professional programmers, as well. The SAS language is a powerful tool to build applications that translate data into meaningful information.

University Presence

SAS is well represented among the colleges and universities of the country. It is one of the most popular software packages at the Homewood Computing Facility that supports the faculty and staff of Johns Hopkins University in Baltimore, MD. In addition to SAS, the Homewood facility, which serves as an information center for hundreds of researchers and administrators, supports SPSS-X, the IMSL Library of FORTRAN routines, and dozens of other mainframe and microcomputer products. The center runs three mainframes: an IBM 3081 under VM/CMS, a VAX computer under VMS and an AT&T 3B4000 under UNIX.

According to Susan J. Slaughter, Sciences Computing Coordinator, SAS is used primarily as a tool for scientific research, statistical analyses, data management and graphics. In addition to the base system, Johns Hopkins licenses SAS/ETS, a library of 15 procedures designed for use by statisticians involved in business planning and modeling. It also licenses SAS/FSP (Full Screen Product), that provides procedures for entering, editing and browsing for data. It also has SAS/GRAPH, an option that provides color graphics on a variety of terminals, plotters and printers. The University has plans to implement SAS/AF, an application development facility for building front-end menus for SAS applications.

Since the Homewood Computing Facility supports both SAS and SPSS-X, Slaughter is often asked which is better. Her reply is that SPSS-X is a good statistical package for people who already know SPSS. "For people starting from scratch, however, I suggest starting with SAS. It is bigger and has features that are potentially useful such as supporting relational databases. The University may be adding SQL/DS in the future, making that feature more important. There are a few features that are easy in SAS — such as combining datasets — that are more difficult in SPSS," she notes.

No Automatic Cross Field Validation

The Full Screen Product feature is good but Slaughter wishes SAS would provide for a way to do automatic cross field validation. "Although it's easy enough to write a SAS program that checks cross field validation after the data entry is completed, I would like to have it flag invalid entries at the time the data is entered," she says. According to a spokesperson at SAS Institute, this capability is scheduled to be included in the next major release for mainframes and is already implemented in the PC and UNIX versions of the SAS System (Version 6).

SAS plays a prominent role at the California Department of Labor and other state agencies. Raymond G. Fredericks, a technical consultant with the Employment Development Department in Sacramento, reports that SAS is used by more than 300 people in a variety of applications, often using SAS/AF, the interactive application building feature of the product. Several hundred users at the California Bureau of Labor Statistics use SAS on an Amdahl 5090 under MVS and an IBM 3090 under VM/CMS.

Statisticians use SAS to extract data from a range of other systems, run statistical analyses, produce tables and, with SAS/GRAPH, generate charts. SAS/GRAPH includes map datasets which the department uses to generate maps of state job service offices. The department also takes advantage of the SAS/SHARE option that allows dozens of users concurrent access to centralized information.

Many companies use the statistics and graphics capabilities of SAS in support of their capacity planning efforts. Most of the leading capacity planning and modeling programs have hooks to SAS. With the release of the SAS/CPE Starter Set, SAS Institute offers a more direct way to use SAS for capacity planning and evaluation. SAS/CPE, available for OS/MVS, is a menu-driven system for evaluating SMF and RMF data. It combines many of SAS's data management, analysis and reporting tools.

The SAS System and its options are available on an annual license basis only. After the first year, renewal licenses are discounted 25 to 50 percent. Deep discounts for educational organizations are offered. Licenses are graduated according to machine classes for mainframes and minicomputers. Pricing is complicated. There are six categories of IBM mainframes — from the 9370-20 to the 3090-600E — and similar classes for minicomputers, as well. A license for a base system for an IBM 3090 600E is $15,000. Each option is then priced similarly. Contact the vendor for specific prices for specific configurations because not all options are available for all configurations. Versions of SAS are available for IBM and compatible mainframes (running under DOS/VSE, VM/CMS, MVS/SP and MVS/XA), IBM microcomputers and compatibles (DOS and PS/2), and minicomputers from DEC VAX (under VMS), Prime (under PRIMOS), Data General (under AOS/VS), Sun Microsystems (under Sun OS) and Hewlett-Packard (under HP-UX).

For more information, contact SAS Institute Inc., SAS Circle, Box 8000, Cary, NC 27512-8000, (919) 467-8000. ⊜

EDITORIAL EVALUATION

Please circle the appropriate numbers on the Reader Service Card.
1.) This article was:
  295 (Interesting/Helpful), 296 (Too Technical), 297 (Too Basic)
2.) Would you like more articles on the same subject?
  298 (Yes), 299 (No)

ABOUT THE AUTHOR

John Kador is a free-lance writer and a frequent contributor to MAINFRAME JOURNAL.

most existing systems will continue. In some cases, the rate of change will probably be accelerated since additional staffing has been added to work on the backlog of changes. The overhead factor will probably also continue to increase as already complex programs continue to be modified.

While none of these factors is absolute or totally convincing, each provides a small indicator of what to expect in the future. While I still need to plan for major new systems that are being developed, I can see that in this particular environment, enhancements and expansion of existing systems constitute the largest share of my growth.

Capacity planning involves more than just planning new hardware requirements. The effective capacity planner will understand the nature of the work in his or her system, how it is growing and why. Capacity plans need to evolve from a background founded not only on major new workloads but on growth patterns built around the nature of the organization, its structure and plans. The need to work with business forecasts and estimates for newly designed systems will continue. When this is done, remember that you are always adding these elements of growth to established systems. Unless you understand the underlying nature of these systems and the historical causes of their growth, you are likely to make assumptions about the projected environment that will not be realistic. ◓

ABOUT THE AUTHOR
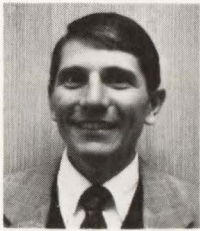
Ted C. Keller is the manager of a group responsible for CICS systems support, performance management and capacity planning at Yellow Freight System, Inc., 10990 Roe Ave., Overland Park, KS 66211, (913) 345-3274. He has worked in various data processing jobs for more than 21 years.

"This system will capture all of the knowledge about our operations and decrease the number of questions that we have to answer on a ongoing basis. By giving our staff easy access to that knowledge, we hope to increase our efficiency and effectiveness even more," Rice says.

This project is underway now, but is not expected to be completed for at least another year.

Sticktuitivity Pays Off

As a group, DOS users represent IBM's largest single user base. There are few, if any, indications that their number is decreasing — some even say it is growing.

Due to their collective rational stubborness, it now appears that IBM has little choice but to continue to offer these users a fairly high level of support.

Will that mean lots of new and exciting enhancements for VSE? Probably not; but, who really cares? VSE offers a stable, tested environment and there is a lot to be said for that.

With the type of solid, innovative, common-sense management that exists at Carolina Steel, VSE users should be able to continue to offer their organizations efficient and cost-effective information systems processing for years to come. ◓

ABOUT THE AUTHOR

Mary Lou Roberts is an independent consultant and free-lance writer. With more than 19 years of experience in the data processing industry, she has been involved extensively in both the technical and marketing aspects of the business.

Q *Please advise on dataset naming conventions for an MVS/XA environment with no tape or DASD management system in place. The datasets (VSAM, BDAM, Partition and Sequential) will be from several vendors. I would like to keep the number of ALIAS and ICF user catalogs down to a minimum.*

A I assume you are chiefly concerned with high-order qualifiers. If you are new to MVS, forget everything you ever learned about arbitrary 44-character dataset names. All datasets, other than temporaries with system created names, must be cataloged and must have systematic names: a one-to-eight character, high-level qualifier ("prefix") chosen from a restricted set of names maintained in the Master Catalog, followed by at least one additional one-to-eight character strings, separated by periods. The first character of each chunk must be alphabetic; the others may be a mixture of alphabetics, numerics and "national characters." The U.S. national characters are "$", "#" and "@". The Master Catalog entry for the high-level qualifier is an ALIAS pointer to the user catalog in which the actual dataset pointers are maintained. Datasets with the SYS1 prefix, catalogs and a few other system-defined datasets are the only ones cataloged directly in the Master Catalog.

If your installation has existing "natural" codes denoting organizational units (such as AR for accounts receivable, FAC for facilities and so on) or three-or-four character department numbers that begin with scattered letters, these might form the nucleus of your naming convention. A system I have seen that worked well combined the department code (one letter, two digits), the person's initials (three letters) and one character for multiple LOGONs or for uniqueness, giving a seven-character user name for TSO, the JCL USER parameter and as a high-level qualifier. Thus, all USER datasets had numerics in positions two and three. Production datasets were distinguished by having alphabetics in at least one of these positions. For example, J. P. Wilson in accounts payable (department F23) might have a USER-ID of F23JPW1. A production dataset in billing (department F47) using the BILLPAK application package might be called QBF47.PROD.BILLPAK.NAMELIST (the QB is arbitrary, assigned to billing's applications to ensure a uniform distribution throughout the alphabet of high-level qualifiers; the PROD qualifier obviously should *not* come first).

These considerations come to mind.

• Tape dataset names should be restricted to 17 characters and should not contain national characters if they are to be ANSI-interchangable.

• Be positioned for system-managed storage by having key attributes associated with high-level qualifiers. Read the SMS overview manuals to understand the class structure and follow it as much as makes sense.

• The person responsible for catalog maintenance should be the only person allowed update access to the Master Catalog (of course, that person should have at least two backups from different departments). That person should have a good understanding of the distribution of high-level qualifiers to user catalogs and be the one to assign the arbitrary qualifier names for new ALIASes.

• If you are using RACF or another access control system based on dataset entries, make doubly sure to avoid clustered names if generic profiles are not used. There is a natural tendency to load up on the letter "S" because of the SYS1 prefix. You should thus avoid naming your SMF archive datasets something like SMFXXXYY.abcd. Concentrated name-space in a busy catalog might lead to excessive CI and CA splits as well.

• Keep your Master Catalog as compact as possible with only essential system datasets and ALIAS pointers in it. The more catalogs, the smaller risk of major problems if any one breaks; and the less need for ugly protracted splits under-the-gun when everyone agrees that a catalog is too big. At any rate, follow pooling principles to keep each catalog with its datasets. Watch for I/O hot spots on catalog volumes and be ready to split whenever a catalog appears to be contributing to key applications' delay.

Qualifiers below the high level for non-production datasets should follow default TSO naming conventions in the absence of a good reason to the contrary. See the JCL Reference under the DD statement for your level of MVS for further information on dataset names.

(Answer provided by Steve Samson, Candle Corporation, Los Angeles, CA)

Q *My shop runs VM/HPO 5 with three VSE/SP guests, all 40MB VAE with CICS1 on VSE1 and CICS2 on VSE2. CICS1 and CICS2 have sub-second response time with VSE1 having a higher priority over VSE2. We attempted to remove VSE2 and ran CICS1 and CICS2 on VSE1 but had to put CICS2 back to VSE2 because CICS2 response time was two minutes on average. I thought the fewer guest machines I had on VM the better performance I'd get. CICS2 has SQL on it. Is it better to have two guests or one? The CPU is a 4381-14.*

A Lots of possibilities. The 4381-14 is a dual processor but you will have both CICS' on the same processor when using one VSE but you could be using different processors with separate VSE guests. CICS2 could have been inadvertently placed below other partitions (besides CICS1) in the single VSE1 case. If memory, as viewed by the VSE page manager, as opposed to the real memory known to VM, is over committed, the VSE guest could be paging even to the extreme of deactivating the CICS2. It is better to have fewer guests than more, especially when the guests are the same operating system (as in this case) because you will have fewer copies of the VSE Supervisor, VTAM, POWER and so on occupying memory. In other words, the real storage working set will be less.

(Answer provided by Ben Moyle, BI Moyle Associates, Inc., Minneapolis, MN and Steve Huggins, BI Moyle Associates, Inc., Washington, D.C.)

Q *What is a good method of cancelling or aborting batch programs in VSE?*

A A job can always be cancelled by an operator with a "CANCEL xx" or "F xx" (POWER spooling flush command) where 'xx' is the partition ID. If you want an application program to cancel itself, these Assembler language macros are available to do this and could be easily packaged into callable subroutines for other languages:

CANCEL — Cancel DOS/VSE job without a memory dump
DUMP — Cancel the job step only with a dump
JDUMP — Cancel the DOS/VSE job step with a dump and bypass additional steps in the job.

Note that the CANCEL command and macro and the JDUMP macro do not affect other DOS/VSE jobs which might be packaged in the same POWER spooling $JOB. This may or may not be suitable in each situation. There is no macro equivalent of the "F" operator command.

(Answer provided by Ben Moyle, BI Moyle Associates, Inc., Minneapolis, MN) ▤

If...

you spend a lot of time tuning VSAM

Then...

you will find this is the best five hours you've ever invested.

Announcing Goal Systems' Spring VSAM Class Series

This class concentrates on VSAM (including DL/I and IMS VSAM) files, dataspace allocation and monitoring techniques which will assist in optimizing overall system performance. Techniques covered allow accurate tracking and effective tuning of VSAM files. This is a "how to tune" class designed for the practitioner. A brief Goal Systems' VSAM product line overview will follow the instruction.

San Diego, CA • Tuesday, April 4 • 8-12:30	**Detroit, MI** • Tuesday, May 16 • 12:30-5:30
St. Louis, MO • Wednesday, April 5 • 12:30-5:30	**Rochester, NY** • Wednesday, May 17 • 12:30-5:30
Ft. Lauderdale, FL • Thursday, April 6 • 12:30-5:30	**Pittsburgh, PA** • Thursday, May 18 • 12:30-5:30
Dallas, TX • Tuesday, April 18 • 8-12:30	**Houston, TX** • Tuesday, May 23 • 8-12:30
Cincinnati, OH • Wednesday, April 19 • 12:30-5:30	**Hartford, CT** • Wednesday, May 24 • 12:30-5:30
Charlotte, NC • Thursday, April 20 • 12:30-5:30	**Cleveland, OH** • Thursday, May 25 • 12:30-5:30
Phoenix, AZ • Tuesday, May 2 • 8-12:30	**Anaheim, CA** • Tuesday, June 6 • 8-12:30
Baltimore, MD • Wednesday, May 3 • 12:30-5:30	**Denver, CO** • Wednesday, June 7 • 12:30-5:30
Philadelphia, PA • Thursday, May 4 • 12:30-5:30	**Milwaukee, WI** • Thursday, June 8 • 12:30-5:30

(800) 537-4891 East/West • (800) 848-4640 Central

Classes are provided at no charge

Goal Systems

Goal Systems International • 7965 N. High Street • Columbus, Ohio 43235 • 800-848-4640
Goal Systems International S.A.R.L. • 88 avenue de Wagram • 75017 Paris, France • Phone: (1)42 67 55 55 • Telex: 641.094

BLOCKADE Provides New Layer Of Network Access Security

Xenos Computer Systems, Inc. (Markham, Ontario, Canada, 800-387-9781) has released BLOCKADE, a software solution to the security problems inherent with dial-up lines. BLOCKADE is a MVS/VTAM Network access security program that provides extended user verification for non-secured terminals (dial-up or dedicated). It also provides security administrators with complete access control facilities directly from their TSO sessions as well as an exhaustive audit trail through the use of SMF facilities.

For more information
CIRCLE #200 on the Reader Service Card

OPMAN Now Supports Unattended Operations For MVS/ESA

Compucept (Morgan Hill, CA, 408-778-2011) recently announed MVS/ESA support in Release 2.5 of its operations management product OPMAN. OPMAN statistics on MVS operator activity allow managers to set and track performance standards, justify tape-to-cartridge conversion, remove operational bottlenecks, increase throughput and more. In addition to ESA support, the new release introduces an ISPF interface for ease of use, new management reports and a link to automation software products. OPMAN's self-automating features enable it to comply with unattended project requirements.

For more information
CIRCLE #201 on the Reader Service Card

DB/AUDITOR Enables DBAs to Easily Produce Audit Reports

Systems Center, Inc., (formerly VM Software, Inc., Reston, VA, 703-264-8000) recently announced DB/AUDITOR. DB/AUDITOR is said to be a complete auditing tool for DB2 that enables DBAs, security analysts and auditors to quickly and easily produce comprehensive audit reports. It simplifies the process of extracting audit data from Systems Management Facilities (SMF) files to produce a series of standard audit reports that consist of: unauthorized access attempts, attempted read/writes, authorization change activity, GRANT/REVOKE activity, audit summary, plan management and utility access.

For more information
CIRCLE #202 on the Reader Service Card

CICS DUMP BUSTER Introduced

International Business Information Systems (New Orleans, LA, 504-897-3367) has just introduced CICS DUMP BUSTER, a product to manage and analyze CICS transaction and storage violation dumps. It provides three basic functions: the reduction of printed output, the reduction or complete elimination of debugging time and the archival and retrieval of dumps. CICS DUMP BUSTER is currently available for MVS.

For more information
CIRCLE #203 on the Reader Service Card

BARR/SNA RJE (TRN) Is A First

Barr Systems, Inc. (Gainesville, FL, 800-227-7797) unveiled what is said to be the only product for PC-to-mainframe communications using SNA multiple session RJE on a direct Token Ring connection. The product, BARR/SNA RJE (TRN), uses the Token Ring connection for fast throughput at 4MB or 16MB per second and high-speed printing up to 6,000 lines per minute. The PC becomes a convenient departmental JES workstation for local printing, remote job entry and batch submission of files from a LAN to the mainframe.

For more information
CIRCLE #204 on the Reader Service Card

Impact Analysis Pinpoints Workload Contention

Candle Corp. (Los Angeles, CA, 213-207-1400) announced the availability of Version 695 of OMEGAMON for MVS and DEXAN for MVS. This version of the performance monitor introduces Impact Analysis for MVS which determines the specific contentions between workloads. It determines to what degree various workloads on a system are interfering with each other and it shows end users specifically which workloads are using the computer resources the impacted workload needs. With a few simple keystrokes, users can quickly see who is specifically causing the impact and why.

For more information
CIRCLE #205 on the Reader Service Card

AutoMate/MVS & AutoMate/XC Updated

LEGENT Corp. (formerly Duquesne Systems, Inc. and Morino Associates, Inc., Pittsburgh, PA, 412-323-2600) announced the availability of AutoMate/MVS and AutoMate/XC Release 2.0. AutoMate/MVS host-based automated operations software and AutoMate/XC, its PC component, are rule-based products which simplify console operations allowing large MVS data processing centers to automate operating procedures, enhance the operator/system interface, monitor and control VTAM applications and manage remote systems. Release 2.0 of AutoMate/MVS includes the Open Interface for VTAM which provides a bi-directional interface for monitoring, managing and reacting to any information supplied in any full-screen VTAM-based application regardless of vendor.

For more information
CIRCLE #206 on the Reader Service Card

MHtran-2 Translator Supports COBOL '85

Prince Software Products' (Mahwah, NJ, 201-934-0022) latest release of MHtran-2 is designed to support ANSI COBOL '85 as implemented in Release 3 of IBM's VS COBOL II compiler. MHtran-2 is a systems software product that automatically converts OS/VS COBOL ('68 & '74) to VS COBOL II. It is designed to operate in all MVS environments and supports Panvalet, Librarian, IDMS, IMS and more. Ease-of-use features like on-line access and menu-driven operation is said to ensure maximum programmer productivity.

For more information
CIRCLE #207 on the Reader Service Card

DCSS/Manager Simplifies VM DCSS Management

The Adesse Corp. (Danbury, CT, 203-790-9473) has announced the availability of Release 1.6 of its DCSS/Manager. DCSS/Manager is said to dramatically simplify the management of discontiguous saved segments by addressing the major weaknesses of their current implementation thereby eliminating SYSGENs and SHUTDOWNs, enhancing control and management functions as well as human factors and user friendliness. DCSS/Manager now supports AUTOFIT and is operational under VM/SP Release 6.

For more information
CIRCLE #208 on the Reader Service Card

Sterling Standards Monitor Analyzes DASD Management

Sterling Software's (Rancho Cordova, CA, 916-635-5535) Sterling Standards Monitor aids in reviewing and analyzing DASD storage environments. It is a menu-driven, interactive system that provides the storage administrator with a complete DASD environmental analysis including recommendations for dataset naming conventions, DASD pool structure and catalog organization.

For more information
CIRCLE #209 on the Reader Service Card

VSAM-LITE Supports Managed-SAM Files

Universal Software, Inc. (Brookfield, CT, 203-792-5100) has announced Release 1.7 of VSAM-LITE, its automatic VSAM file compression package for the VSE/SP operating environment. Release 1.7 of VSAM-LITE provides the capability to compress Managed-SAM files in addition to ESDS and KSDS datasets. It also includes an improved method for managing VSAM-LITE files in multiple CPU and VM/VSE environments.

For more information
CIRCLE #210 on the Reader Service Card

SAS Institute Inc.

SAS Institute Inc., the nation's largest privately owned software company, is looking for creative people who want to contribute their ideas and talents to our development team. We have the following positions open in our Cary headquarters:

VM Systems Support Manager

You will lead and contribute expert systems programming skills to the VM systems support group. Duties include establishing group objectives, priorities, and training; planning, scheduling, and coordinating the implementation of VM hardware and software changes; evaluating, installing, maintaining, enhancing, modifying, monitoring and tuning the performance of IBM® and OEM vendor software; and assisting users with technical problems and maintaining working relationships with the user community. Applicants must have a bachelor's degree in computer science or a related field, or equivalent experience; five to seven years' VM systems programming experience, including system problem diagnosis, system internals, system generation, program product installation, and system monitoring and tuning; experience supervising and developing technical professionals; and the ability to plan, schedule and coordinate software and hardware changes for a complex VM environment. Experience with VM/XA, computer networking, other operating systems, and the SAS® System desirable. (#998A)

Senior Systems Developer

You will migrate micro-to-host communication services from the DOS to the OS/2 environment and research, design, and implement LAN services to support distributed applications. Applicants must have a bachelor's degree in a related field or equivalent experience; five years' experience with communications software support and development; experience in the following areas: IBM 3270 data stream, 3270 emulation interfaces, asynchronous communication protocols, and protocol converter knowledge; and experience or interest in the following: OS/2, NETBIOS, LAN operating systems, and IEEE 802.2/802.3/802.5 protocols. (#1036)

If you are looking for a challenge, send your resume, including the position number to

SAS Institute Inc.
Department MFR050189
Box 8000
Cary, NC 27512

We also have other development positions available. For details, write to the above address.

SAS is a registered trademark of SAS Institute Inc., Cary, NC 27512.
IBM is a registered trademark of International Business Machines.

SAS Institute is an Equal Opportunity/Affirmative Action Employer   EOE M/F/H/V

CICS PROFESSIONALS · MAJOR DEVELOPMENT

O'Connor O'Connor Lordi, Ltd., an Executive Search Firm, has been retained by a major Pittsburgh-based Corporation to carry out an extensive talent search for top calibre CICS professionals. Our Client's Executive Committee has appropriated a multi-million dollar capital improvement budget structured to assure the successful development of this centerpiece organization.

Due to their rapidly expanding business base and the scope of this major development, our Client is seeking leading professionals who will provide technical expertise for the design and development required to implement the architecture of a futuristic Trust Accounting System. These individuals must desire a challenge and the opportunity to work in a creative team environment.

Qualified candidates must possess a minimum of 2 years experience with CICS Command Level (COBOL) programming and 2 years analysis/design on a large scale CICS application development in an online environment. Your candidacy will be greatly enhanced with proven skill in some of the following: VSAM; Basic Mapping Support; TSO/ISPF; and OS JCL/Utilities.

Knowledge and/or experience in the following is a plus: MVS; DB2; BAL; Financial Systems Development; Librarian; DOS JCL; IBM Operating Systems; and FOCUS.

Our Client offers competitive compensation and benefits package, relocation assistance, and most importantly, the opportunity to create a piece of tomorrow. Positions are located in Pittsburgh, PA.

For confidential consideration, please send resume or directly contact:

Kathy Englert, Program Manager
O'CONNOR O'CONNOR LORDI, LTD.
600 N. Bell Avenue, Suite 216
Carnegie, PA 15106
(412) 276-5070
Fax #: (412) 276-8538

For your convenience, phone calls will be accepted:
Monday-Friday 9:00 a.m.-5:00 p.m.
and Monday evening until 6:00 p.m.

All fees are assumed by Client Company.

Compuware Corporation

Compuware Corporation is an international leader in the data processing industry with professional services and software products in use at more than 4,500 data centers worldwide. One of the largest privately owned data processing companies, Compuware ranks among the world's largest systems software providers.

Compuware Corporation was founded in 1973 when its three founders set out to provide professional services to help people do productive things with their computers. Peter Karmanos, Jr., Tom Thewes and Al Cutting began by marketing computer professionals from an office in Southfield, MI. Their growing technical expertise soon provided a staffing resource skilled in most programming languages and capable of helping clients utilize hardware and software systems. Gross revenues for that first fiscal year were $300,000.

Today, Compuware's Professional Services Division offers a range of services including total project management, analysis and design, systems software support, systems and applications programming, database design and program conversion implementation.

As its productivity expertise continued to grow, Compuware developed a family of software products. The first software product, Abend-AID, introduced in 1977, was a breakthrough in program abend analysis, telling the programmer what happened, where in the program it happened and providing essential supporting information in a concise, accessible format. In the past ten years, Abend-AID has become one of the most widely-used programmer productivity tools in the IBM mainframe environment.

File-AID is a comprehensive, menu-driven product that allows quick, secure access to data without programming, regardless of the access method or file organization. CICS Abend-AID frees programmers from the tedious, repetitive task of dump analysis. It is designed to intercept, analyze and provide expert system diagnostics for all CICS transaction abends. CICS Abend-AID analyzes the

From left to right are: Thomas Thewes, Vice Chairman, Executive Vice President-Finance; Peter Karmanos, Jr., Chairman of the Board, Chief Executive Officer; and Allen Cutting, Executive Vice President-Marketing.

problem, determines the cause and recommends solutions to the problem. All diagnostics are immediately available, online, in a concise and easy-to-understand format.

CICS dBUG-AID is an interactive testing and debugging system that combines full screen source-editor characteristics with the comprehensive testing facilities required by the unique CICS environment. It can also improve programmer productivity, reduce testing time, improve production maintenance and increase overall CICS availability and integrity.

CICS Playback eliminates the problems associated with putting CICS systems into production. CICS Playback is a CICS transaction driver that allows repeatable, objective testing (at both low and high volumes) of any form of CICS application — without the use of a terminal network. It provides facilities to recover production transactions, estimate the impact of changes in transaction volumes or terminal network size and enhance Help Desk facilities.

In 1978, Compuware established an Educational Resources Division to provide a resource for high quality technical education. The company offers instruction for data processing professionals at all levels of technical expertise and now

trains more than 4,500 professionals each year. Compuware employs full-time instructors to train at the company's Training Center, on the customer's premises or by using an extensive series of video modules.

In 1985, Compuware introduced two vertical market products: Printer's Plus and Compuware's Law Office Information System. Printer's Plus was developed as a productivity tool for the quick printing industry and handles routine tasks such as sales reports, accounts receivable, job estimating, customer identification and miscellaneous bookkeeping.

Compuware's Law Office Information System was designed to support the legal profession and is made up of five software modules that may be purchased separately or as a group. The modules currently available are Timing and Billing, Accounts Receivable, Accounts Payable, General Ledger and Docket Control.

In Compuware's 15 years of existence, the company has become an organization of more than 850 people. Compuware now has its headquarters in Farmington Hills, MI, seven branch offices in the United States, operations in Europe and sales agents throughout the world.

The company's gross revenues for fiscal year 1988 were $86 million with projections of $115 million for fiscal year 1989. The steady growth of Compuware has been the result of realistic planning and a careful balancing of resources to assure the highest levels of customer service.

Karmanos believes that Compuware's growth and profit are the measure of how well the company serves its clients.

"We are in the business to serve clients. We are competitive but we are not in business to put other companies out of business. We produce satisfied customers by providing expertise and quality services and products. That is why Compuware is successful," Karmanos concludes.●

Vendor Profile is a regular forum whereby a vendor is given the opportunity to introduce the company and its products to MAINFRAME JOURNAL readers.

Space, the Final Frontier

If any of your users need to crunch big spreadsheets, they will soon find themselves exhausting the normal complement of PC memory and coveting one or more expansion memory boards. The recent run-up in RAM chip prices, unfortunately, makes scratching this particular itch a good deal more expensive than it once was. Each such expansion memory board, if it is fully populated with RAM chips, can easily cost more than the basic machine it fits in.

Picking Up the Check

When the totals are all run, just the hardware and software cost of supporting personal productivity applications on mainframe-linked micros turns out to be a step function with an increment of roughly $4,000 per user. As I have pointed out, it is always possible to spend more.

Headcount

However, hardware and software aren't the end of the story. PCs are not mainframes and the PC hardware and software environments are not especially similar to those on the big iron. Anyone bringing in lots of PCs, therefore, also takes on the necessary burden of either training some existing staff to be PC people or hiring some existing PC people to provide the support the user community will expect and deserve.

I should point out here that these needed new troops are not to be confused with trainers. All users, on whatever size machine, have to be trained to use personal productivity tools when first introduced to them. PCs and mainframes come out even in this respect.

New Faces, New Jobs

What I refer to, instead, are genuine PC technical support staffers. These people have many duties in PC-intensive organizations. A partial list includes: assembling new PCs, installing add-in hardware in new and existing PCs, formatting and configuring hard disks, installing PC software packages and subsequent updates, establishing and monitoring standards for file subdirectory organization, establishing and monitoring backup and data security procedures, diagnosing hardware failures, assisting users who suffer hardware (especially hard disk) failures, recovering inadvertently deleted data files and improving user interfaces by writing PC-DOS batch command files or comparable instructions for a commercially available piece of "shell" software.

This personnel component of corporate PC use is actually the submerged part of the iceberg in terms of total cost. The PC hardware costs I explored earlier at least have the saving grace that they do not need to be paid out again each year. People, though, have ongoing salary requirements and the clock keeps running.

Quantity Discount

What is the mainframe-based alternative? If your users already have mainframe workstations, you can provide them with equal or better personal productivity functions for just the price of the appropriate mainframe packages. This sum will be equivalent to just a handful of the micro-based setups previously described. Your mainframe, of course, can handle far more than a handful of users.

Increments

True, the hardware bill cannot be said to be *precisely* zero. Users will consume additional mainframe resources (CPU time, I/O operations, main memory space) when using personal productivity tools and these things do not come free. As this usage adds itself to the background represented by your normal processing load, it will make its own contribution to your occasional need for system enhancements and upgrades.

But even after factoring in a pro rata share of such costs, mainframe personal productivity tools are still a bargain. From a performance standpoint, the use pattern for such tools is that they are invoked sporadically, at need. You may provide hundreds of users with access to these tools without needing to worry that everyone will be using them intensively every hour of your business day.

Staff

Support will be similarly "time-shared" as an incremental duty for existing staff. End users won't have to be acclimated to unfamiliar keyboards and operating system interfaces. Product upgrades can be installed once, centrally, by staffers who already know how to do mainframe software installs and upgrades. There will be no need to chase around updating software levels, one at a time, on dozens or hundreds of PCs.

Conclusion

For the small user, even many medium users, mainframe computing power and cost may not be justifiable for *any* application. In these cases, microcomputers may be the only platforms that can support needed personal productivity tools.

Where a mainframe already exists, however, the advantages of employing it as the platform of choice for personal productivity applications are overwhelming by all economic measures. ⊜

CIRCLE #99 on Reader Service Card ▲

Paying The PC Piper

The advantages of employing a mainframe as the platform of choice for personal productivity applications are overwhelming economically.

By F. Thomas Cox
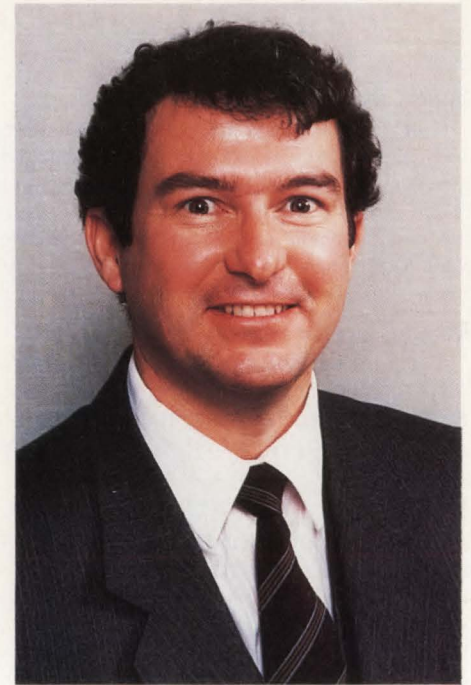
*Y*ou never get a second chance to make a first impression. This old chestnut assumes that good first impressions are an unmixed blessing. Actors know better than most that even a good first impression, if it resists change *too* much, can eventually be a curse — it's called typecasting.

Second Impressions

However, nothing, not even the mind of a Hollywood producer, is entirely immovable. Ann-Margret, Tuesday Weld and Farrah Fawcett each, in her turn, graduated from the ranks of sweet young things to the much more exclusive status of serious actress. In like fashion, Japanese automakers segued, over time, from an image of ultra-low price to one of ultra-high quality — which commands a *premium* price.

Coming now to the point of all this impressionism, it is fair to state that personal computers are in a sort of early Japanese car period — they embody the popular notion of low-cost computing. There is such a strong tendency to believe in the cheapness of PC-based computing solutions, in fact, that they do not always receive proper economic scrutiny. Mainframe users, in particular, should always run the numbers to be sure certain computing applications aren't being "typecast" in costly ways.

The Players

No one can realistically claim that one type of computer is *always* better than another for some purpose. But PCs, nonetheless, are strongly identified with certain personal productivity applications. Word processing, spreadsheet computation, database management, asynchronous telecommunication and pop-up personal organizers are the five that seem to make up the vast majority of PC work in corporate settings. So strongly are all but the database management application identified with PCs, that many user organizations fail to appreciate that there even *are* other implementation options. This applies, in particular, to the use of *mainframe* software for personal productivity applications.

If introducing the notion of mainframes into this discussion sounds peculiar, it shouldn't. I can show that mainframes offer many advantages, especially economic advantages, in the personal productivity sphere. It doesn't damage my case to acknowledge that no one is likely to purchase a mainframe *exclusively* for the purpose of supporting such applications. The crux of my argument is that it applies to the established mainframe user organization rather than the mom-and-pop store.

Marginal Notes

In general, the economic case for mainframe personal productivity tools is one of demonstrating a basis of marginal cost for the acquisition of more than marginal utility. As a mainframe user, you already have all the hardware needed to support these additional functions without new outlays for micros, interface cards and other paraphernalia. You can add personal productivity tools as "icing" to a mainframe "cake" that you originally "baked" for other reasons.

Down the Slippery Slope

What happens if impressions overrule economics and your mainframe shop elects to use microcomputers for personal productivity applications? First, you have to buy micros for every user who is to benefit. Even bare-bones Taiwanese PC-XT clones will set you back $1,000 apiece.

F. Thomas Cox, Vice President of Marketing, Trax Softworks, Inc.

Corners Not To Cut

It *is* possible to buy cheaper clones but only by leaving out the hard disk drives. No one with any regard for a continued career should consider making users run five different sizable applications on floppy-only PCs. Constantly swapping floppy disks in and out gets old very fast.

However, the use of a hard disk also demands that each user be equipped with a good hard disk backup program. This requirement adds another $100 to the running total.

Yesterday, Today and Tomorrow

Of course, what I have said so far assumes that yesterday's PC-DOS technology will be adequate for your needs over the next several years. The capability of graduating to the upcoming OS/2 applications, when they appear, may also be important to you. Most large organizations seem to think so. In this case, the basic tariff becomes the $1,500 or more per unit needed to acquire PC-AT clones. These are the least expensive units having the more advanced 80286 microprocessor needed to run the new generation OS/2 software.

Soft Ware, Hard Price

Then there is the cost of micro-based personal productivity software. Even at discount prices, it can easily cost $1,000 per user to outfit each of them with copies of the leading spreadsheet, database and word processing programs. If asynchronous telecommunications and personal organizers are also on your to-do list, the tab goes up another $100 to $200 per head.
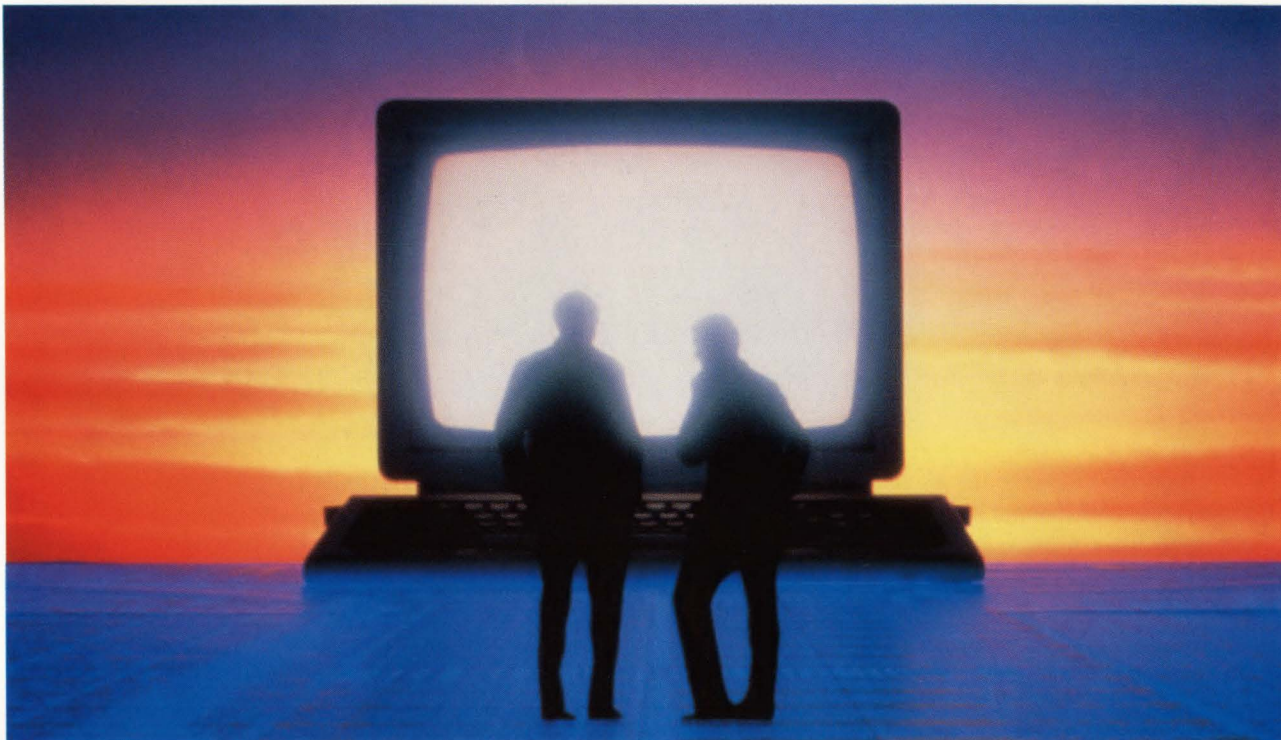
High Price for Cheap Seats

For an investment at or above $2,700 per supported user, then, you can offer basic functionality in all five major personal productivity application areas. Even so, many of your users will find that their PCs come in well short of any reasonable mark.

In particular, your users will be doing their PC computing entirely separate from their mainframe computing. This means you can't even sell off your existing mainframe workstations to defray part of the PC expense. It also means that your existing mainframe users will now have *both* mainframe workstations and PCs gobbling up the surfaces of their desks.