111121

COMPUTER PERFORMANCE EVALUATION [CPE]

111121

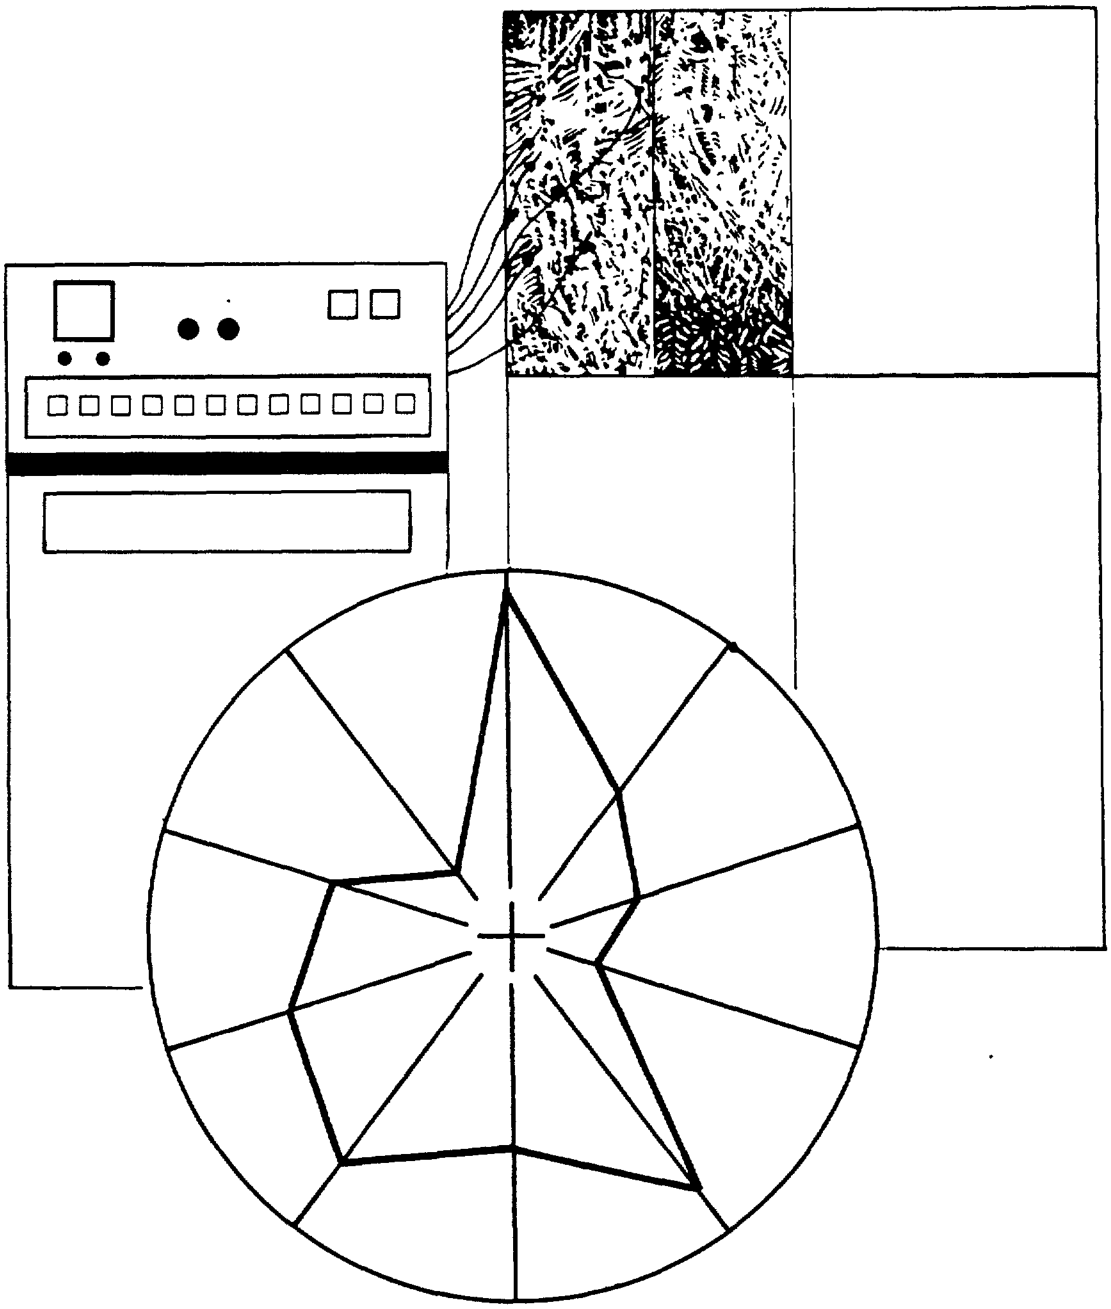

AN AUDITOR'S INTRODUCTION

008103

U.S. GENERAL ACCOUNTING OFFICE • NOVEMBER 1979

111121

OTHER

FOREWORD

This booklet is an auditor's introduction to computer performance evaluation, or CPE. CPE is a specialty of the computer profession that concerns itself with the efficient use of computer resources. It can be valuable to the auditor because it provides a means for controlling data processing costs through measurement and evaluation of computer resource usage. Tools or techniques used to conduct these measurements include accounting data reduction programs, software monitors, program analyzers/optimizers, hardware monitors, benchmarks, and simulation.

The primary purpose of this booklet is to acquaint the auditor with these tools and techniques and with methods of presenting technical CPE data to management so that the data clearly and convincingly supports recommended changes for improving computer performance. For the auditor or team leader whose job includes computer performance evaluation, more detailed descriptions of individual CPE tools are planned as supplements to this booklet.

Some of the material in this booklet is taken, with permission, from the manuscripts of Michael F. Morris' soon-to-be-published book on computer performance evaluation. Mr. Morris served as a consultant to GAO for 2 years.

Because this material is written for the auditor who has acquired the specialized skills needed to work effectively in the computer environment, technical terminology is included without detailed explanation. Comments are welcome and should be addressed to the Director, Financial and General Management Studies Division, U.S. General Accounting Office, Washington, D.C., 20548.

Elmer B. Staats
Comptroller General
of the United States

FOREWORD

This booklet is an auditor's introduction to computer performance evaluation, or CPE. CPE is a specialty of the computer profession that concerns itself with the efficient use of computer resources. It can be valuable to the auditor because it provides a means for controlling data processing costs through measurement and evaluation of computer resource usage. Tools or techniques used to conduct these measurements include accounting data reduction programs, software monitors, program analyzers/optimizers, hardware monitors, benchmarks, and simulation.

The primary purpose of this booklet is to acquaint the auditor with these tools and techniques and with methods of presenting technical CPE data to management so that the data clearly and convincingly supports recommended changes for improving computer performance. For the auditor or team leader whose job includes computer performance evaluation, more detailed descriptions of individual CPE tools are planned as supplements to this booklet.

Some of the material in this booklet is taken, with permission, from the manuscripts of Michael F. Morris' soon-to-be-published book on computer performance evaluation. Mr. Morris served as a consultant to GAO for 2 years.

Because this material is written for the auditor who has acquired the specialized skills needed to work effectively in the computer environment, technical terminology is included without detailed explanation. Comments are welcome and should be addressed to the Director, Financial and General Management Studies Division, U.S. General Accounting Office, Washington, D.C., 20548.

Elmer B. Staats
Comptroller General
of the United States

Contents

ABBREVIATIONS

CPE	computer performance evaluation
CPU	central processing unit

INTRODUCTION

The complexity of a modern Government computer installation makes it difficult for the data processing manager to consistently meet user requirements with an optimum amount of computer resource capacity. Consequently, as an easy solution to the dilemma, the computer installation manager frequently obtains additional computer resources.

To help the installation manager better manage his resources, computer performance evaluation (CPE) includes a group of tools and techniques that can be used to evaluate the use of computer resources. The auditor can use these same CPE tools to identify computer installations that either use resources inefficiently or have more computer resources than necessary. For example, the auditor may measure current utilization of computer resources, and then compare it with potential utilization to get an indication of whether resources are used efficiently.

A permanent CPE effort may not be needed in all computer installations; the cost of CPE should be carefully considered in relation to the benefits derived. Also, the auditor should remember that the installation manager's primary mission is to support those who use the computer. To accomplish this mission, the installation manager may not find it possible or practical to obtain maximum efficiency from the installation computer system.

As a general rule, an auditor is not expected to be able to use all the tools discussed in this booklet. Even so, the auditor should be aware of these tools and their potential for helping to manage the computer reviews referred to in General Accounting Office audit policy. The GAO Comprehensive Audit Manual (part I, ch. 11) states

> "Whenever ADP resources are significant in terms of size, cost, or dependence placed on output products, and whenever there are indications that the resource investment is excessive or results are unsatisfactory, we should consider scheduling the installation for separate review. * * *
>
> "This type of work requires the ability to deal with highly complex ADP technical questions, especially when evaluating current utilization and practical capacity of the computer center configuration. * * * "

SECTION I

DESCRIPTION OF COMPUTER PERFORMANCE EVALUATION

Computer performance evaluation, as we know it today, probably began between 1961 and 1962 when IBM developed the first channel analyzer to measure the IBM 7080 computer performance. The real push toward performance monitoring began between 1967 and 1969.

Earlier generations of computers worked on one job at a time, which made it relatively easy to tell how efficiently each component was being used. However, later generation computers used multi-programming, and efficiency was not apparent because the computer could be working on several jobs at the same time, and each job could be competing for the same resources. This method of operation could result in an idle central processing unit (CPU) or in serious imbalances in loads imposed on peripherals. As a result, these later generation computers actually fostered a need for computer performance evaluation to help balance workloads and improve operating efficiency.

Computer performance evaluation is a specialty of the computer profession and concerns itself with the efficient use of computer resources (e.g., the central processing unit, tape drives, disk drives, and memory). Using selected tools to take measurements over time, the CPE specialist is able to determine computer resource usage. This measurement data provides the CPE specialist with information relative to component capacity, system hardware, software, operating procedure inefficiencies, and workload. Thus, CPE provides the information to answer such questions as:

--Can we decrease the workload or make the application programs more efficient?

--Do we need an additional or a more powerful computer?

--Do we need additional components (memory, tape, disk, etc.)?

--Can we eliminate certain components or replace them at a lower cost?

--Can we improve certain aspects of computer service (response time, turnaround time)?

SECTION II

COMPUTER PERFORMANCE EVALUATION TOOLS

Computer performance evaluation specialists have several tools at their disposal. These tools fall into two broad categories: measurement and predictive. Measurement tools measure or report on measurements of computer resource usage. These tools include accounting data reduction programs, software monitors, program analyzers/optimizers, and hardware monitors. Predictive tools are models which represent the impact of a particular workload (jobs, runs, programs) on a set of computing resources, (cpu, tapes, memory, disk) to produce a planned or required level of output or service (response time, throughput). Simulation is an example of a predictive tool.

Benchmarking is a CPE tool that falls in both of the above categories. A benchmark is a standard or point of reference in measuring or judging quality, value, etc. A computer benchmark is, therefore, a set of computer programs that represents a workload. When the programs are run on an existing computer system to establish a "benchmark" time, and then run on other computer systems to determine how well these other systems perform in comparison with the existing system, the benchmark programs are regarded as measurement tools. On the other hand, when benchmark programs are models of a proposed application, and they are run on a computer system to predict timings or to estimate the impact on current workloads, they are regarded as predictive tools.

These computer performance evaluation tools can be useful in conducting evaluations throughout the life cycle of a computer system. CPE tools that are suitable for use in different phases of a computer's life cycle are shown in the following chart.

ACCOUNTING DATA REDUCTION PROGRAMS

Throughout this booklet, the term "accounting data" means data that describes the amount of computing resources consumed by or in support of each application program that is run on a computer system. Examples include the time a central processing unit is busy, the number of tape and disk drives used, or the amount of memory used. Some facility for generating or collecting this kind of information about the use of various computer resources is now included in nearly every computer above the mini-computer level. Typically, this information is available for each user program that generates computer system activity. The main reason for providing this resource accounting data has been for purposes of billing computer users on a program-by-program basis.

Most general purpose computers above the mini-computer level now collect accounting data automatically. This concept is illustrated below.

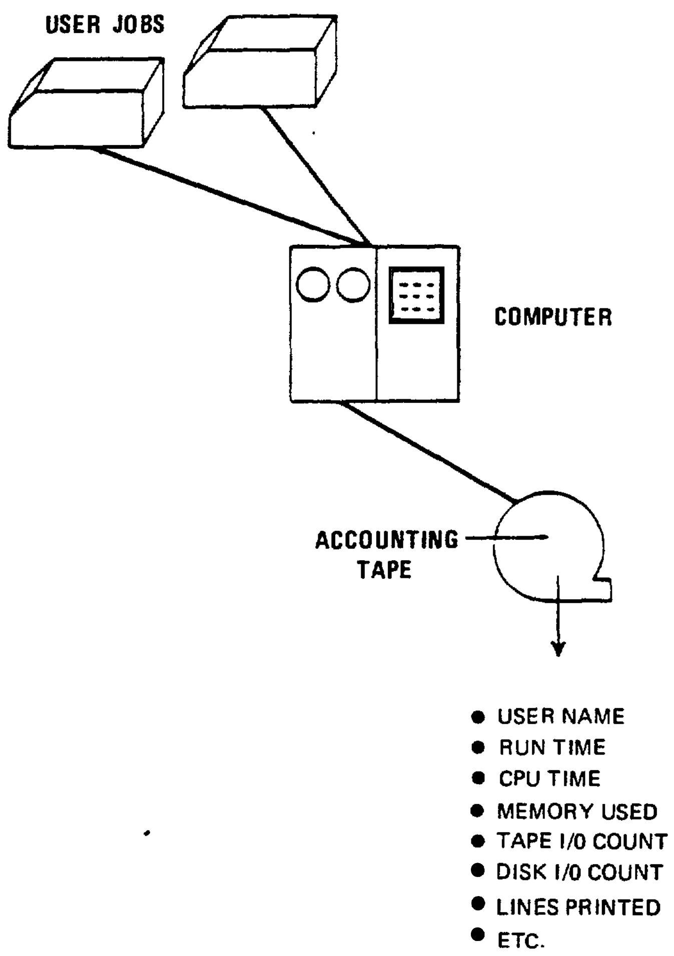

Because accounting data is often produced automatically, it is readily available for any user. As a result, numerous special purpose data reduction programs have been developed to process and analyze this data. These data reduction programs are routinely supplied by most computer manufacturers as a part of the computer control programs. They are generally regarded as "free" resources at the installation because no extra charge is involved.

The easy availability of accounting data coupled with limitations of manufacturer-supplied data reduction programs has fostered the development of special programs that extract and analyze accounting information of great interest to many installations. This commercial development spawned a wide variety of products that, along with manufacturer-supplied programs, is collectively referred to here as accounting data reduction packages. These commercial packages serve such diverse needs as configuration capacity management, job scheduling, library control, standards enforcement, job billing, and numerous other management-oriented functions. Accounting data reduction packages are the primary CPE tool at many installations although they are widely available only for IBM systems.

The use of a data reduction package to extract and analyze accounting data is illustrated in the diagram below.

ACCOUNTING DATA REDUCTION

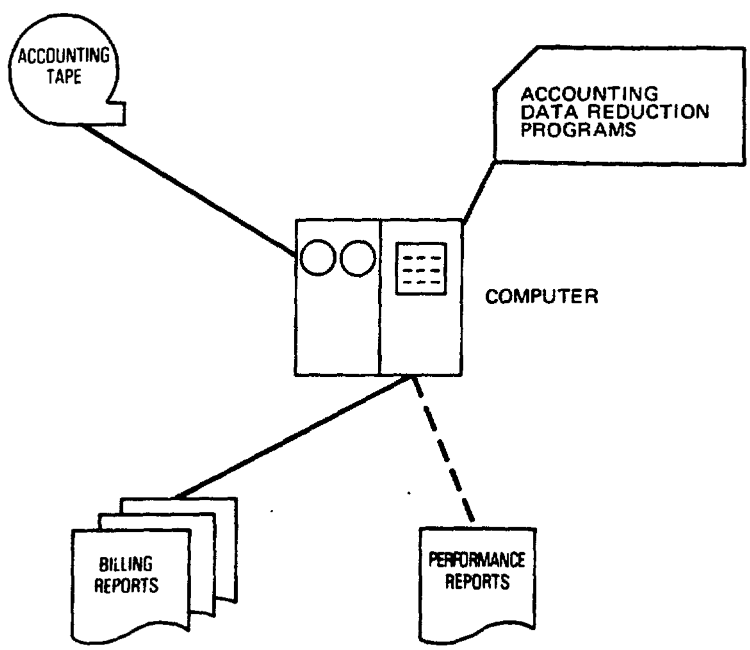

Accounting data reduction packages evolved from the "check flags" inserted in programs by users of the earliest computers. The evolution continued through manual logging and billing, to automated "trace routines," to the current comprehensive accounting data collection programs. Each step in the evolution of accounting data reduction packages was initiated by computer users. In general, the manufacturers simply adopted and refined the more popular developments. Each step in the evolution was taken to solve or understand some real problem in the operation or management of a computer installation.

Because accounting data evolved to solve real management problems, it is probably the richest, single source of information providing high pay-off for most computer performance evaluation or improvement projects. The use of accounting data reduction packages is integral to all continuing CPE efforts and should be the starting point for auditors interested in understanding the detailed characteristics of the workload at a computer installation.

SOFTWARE MONITORS

Another descendant of a "check flag" is the software monitor. Software monitors are specialized sets of computer programs that are usually (but not always) made a part of the computer's operating system. That monitor collects statistical information about the distribution of activity caused by particular programs or sets of programs. This concept is illustrated below.

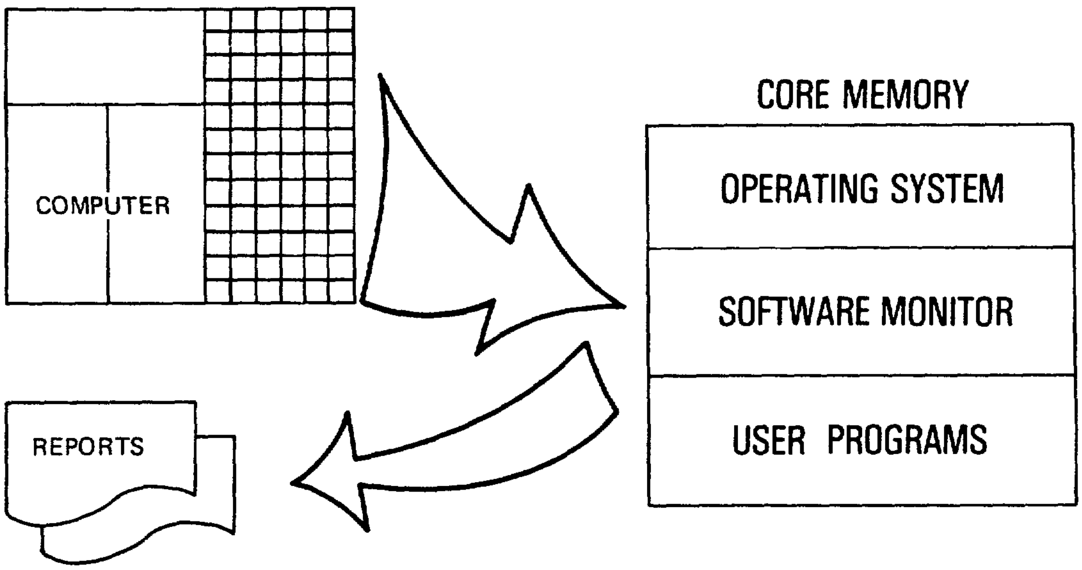

The major difference between software monitors and accounting packages is the level of detail that each is capable of examining. Software monitors can examine the step-by-step execution of instructions within computer programs much more closely than accounting packages can.

Software monitors, like accounting packages, are sold commercially but are available primarily for larger IBM systems. Also like accounting packages, software monitors have been developed by individual computer installations for nearly every brand of computer. Information on the availability of both commercial and user-developed monitors is generally obtainable through "users groups" for the computer brand of interest.

PROGRAM ANALYZERS/OPTIMIZERS

CPE tools that could be considered subsets of either accounting packages or software monitors are program analyzers/optimizers. Program analyzers are specialized computer programs that are usually written in the language of the program that is to be analyzed. They are run along with the application program of interest to collect information on the execution characteristics of that program when it is run with real or test data. Program analyzers are commercially available for major high-level languages like FORTRAN and COBOL.

Although they are tools used in computer program optimization efforts, program analyzers do not, themselves, optimize programs. They produce reports that indicate the areas of a program where a programmer might decrease running time or computer resource usage by employing alternative techniques in the program. A major difference between these analyzers and most software monitors or accounting packages is the type of information collected. For example, analyzers can collect information to determine which parts of a computer program are not used in processing data. In a test environment, knowing that some portion of a new program has not been executed may be more important than knowing that the rest of the code did execute.

In general, optimizers are computer programs which examine a program's object or source code and, without manual intervention, make changes directly to the code. One such optimizer works by eliminating redundant machine instructions from the object code while leaving actual processing logic unchanged.

In a typical installation, most computer system resources are consumed by relatively few of the computer application programs. For this reason, computer performance improvement efforts using program analyzers and optimizers should follow the "80/20 rule." This rule of thumb states that 80 percent of the resources consumed by production programs in a computer installation will normally be accounted for by 20 percent of the computer programs. Therefore, the performance improvement effort should identify programs which are big resource users and concentrate analysis on these programs.

HARDWARE MONITORS

A very different kind of CPE tool that is a descendant of such electronic devices as oscilloscopes is the hardware monitor. Hardware monitors are electronic instruments that may be attached to the internal circuitry of a computer system to count electronic pulses or signals at various connection points. The monitors then record or display information on the number and duration of signals that occur at each connection point. This information may either be displayed and examined immediately, or may be saved for later analysis using a special computer program. An overview of this process is included below.

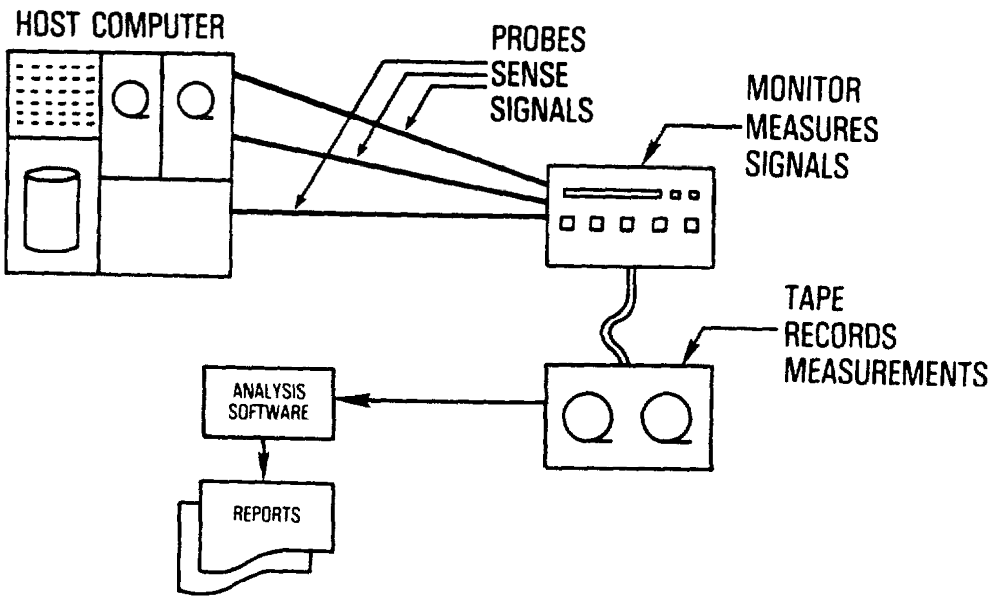

There are three distinct types of hardware monitors--basic, mapping, and intelligent.

Basic monitors can count and time signals or electrical pulses. They usually have some provision for recording that information on magnetic tape. The following diagram shows how a basic monitor counts a signal whenever the signal strength reaches a certain threshold.

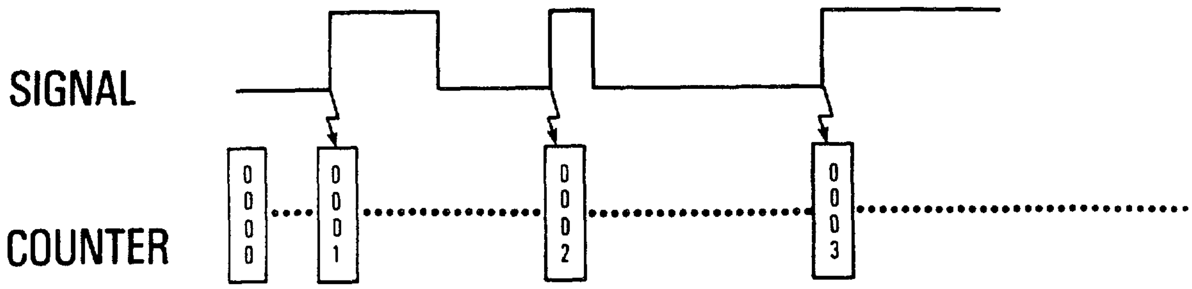

Mapping monitors incorporated memories and facilities to connect large numbers of probes. (A probe is a small clamp-like device which is attached to some part of a computer and transmits signals from that computer part to the hardware monitor.) It became feasible not only to count and time signals but also to determine the values or contents of certain locations in the computer's memory or registers. Mapping monitors made it possible to create maps or distributions of activity as they took place in the computer system being monitored. An example involves collecting data on how often each instruction is used in the computer processor. Collecting the values or contents in a computer's register is illustrated below.

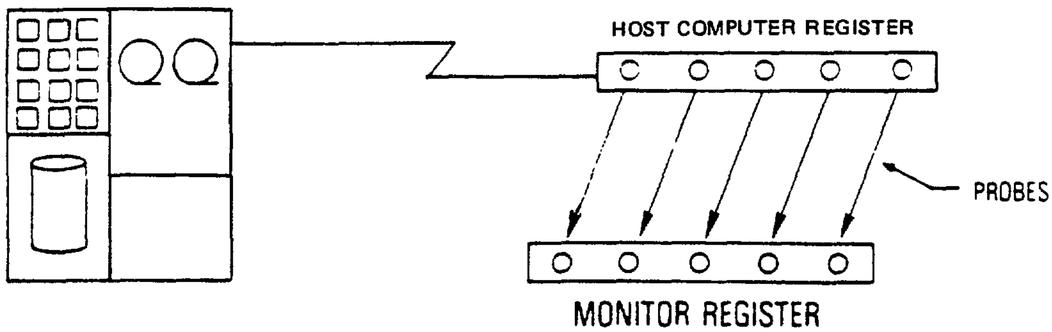

The intelligent monitor contains its own processor which permits the monitor to be programmed for any number of different experiments. In fact, intelligent monitors can be used as small stand-alone computer systems. The intelligent monitor concept is shown below.

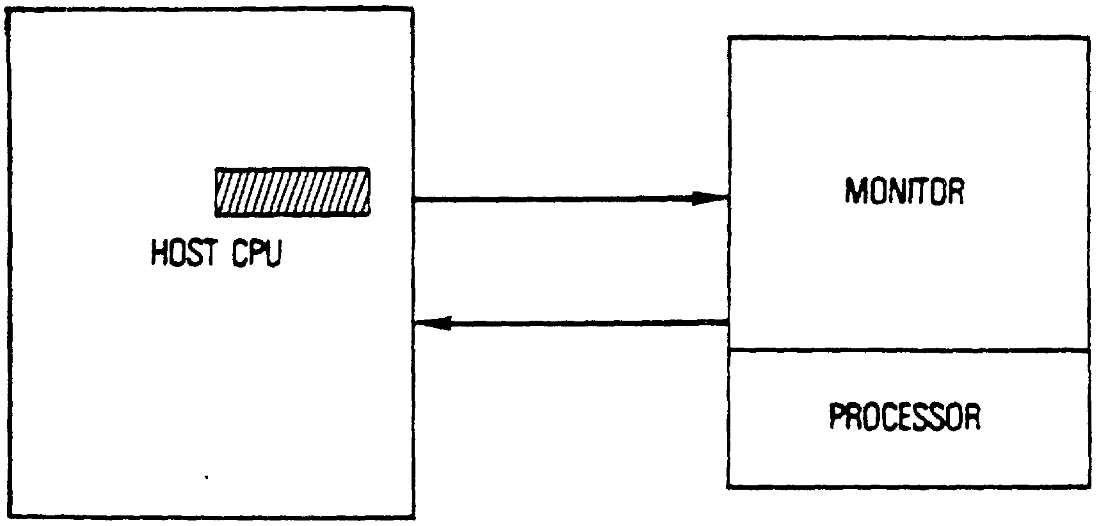

In contrast to accounting data reduction packages, software monitors, and program optimizers, the hardware monitor is not a good tool for the beginner or the inexperienced computer performance evaluator. Recent developments have greatly simplified the use of hardware monitors. However, useful information is not normally collected unless the monitor user (1) thoroughly understands the computer's workload, (2) is formally trained and experienced in using the monitor, and (3) understands the architecture of the computer system to be monitored. Further, hardware monitors can generate data in such large amounts that inexperienced users often find themselves inundated with data shortly after the monitor is installed. However, in the hands of experienced technicians, hardware monitors can be extremely valuable when some piece of otherwise unobtainable information must be measured. Without experienced people, hardware monitors should generally be regarded as the measurement tool of last resort. This does not mean that hardware monitors will not be rewarding. It means simply that hardware monitors are much more difficult to use than most other CPE tools.

BENCHMARKS

Benchmarks are programs or sets of executable instructions that are used to represent a real computer workload on an existing computer, or a workload that is planned for an existing or proposed computer. Benchmarks are typically used to establish the relative capabilities of different computer systems or the alternative configurations that will process a certain workload represented by the benchmarks. Benchmarks are also useful to validate or verify the results produced by other CPE tools.

A small set of benchmark programs that accurately represents the total workload of an installation is always valuable because periodic runs of benchmarks permit the CPE group and the installation managers to assess changes made to the configuration or operating system. Such benchmarks also permit the CPE group to determine when their own actions have degraded system performance.

SIMULATION

Often, the performance of a computer system or set of programs should be examined in detail before the computer system is installed or before programs are written. Such examinations are the most important role of simulation. Simulation involves creating and exercising mathematical descriptions or models of the system's parts to determine or predict the characteristics of the system as it should operate over time. Simulation may also be used to study existing systems when performing experiments directly on the real system is too expensive or time-consuming, or when it might adversely affect the existing system.

Simulation is a discipline in its own right. It is widely used in most scientific and engineering fields with each field developing its own specific simulation tools. In the CPE field, these simulation tools are computer programs or sets of programs that are used as packages that can simulate computer systems. In addition, special simulation languages that are tailored for creating detailed models of computer systems are also widely used in CPE projects.

The differences between computer simulation languages and computer simulation packages are similar to the differences between software monitors and accounting packages. That is, simulation languages (like software monitors) are generally used when more detailed problems are to be studied, while simulation packages (like accounting packages) are most useful when overall systems are under examination.

The major advantage of including simulation as a CPE tool is that it provides an indepth understanding of the entire installation which is invaluable to management in making sound decisions. Simulation permits management to examine a system's total performance or any part of its performance before actually acquiring or constructing a computer system.

COMMERCIAL AVAILABILITY, COST, AND OVERHEAD

CPE tools and their commercial availability, cost, and overhead are listed on the following page. The starting point for computer performance evaluation for nearly every installation is the accounting data reduction package. The second step is typically to acquire a software monitor. Many CPE groups never go beyond these two tools but still produce timely and cost-effective projects on a continuing basis. Next, an installation generally adds a program analyzer or a hardware monitor. Simulation capability is usually the last CPE tool obtained. Benchmarks tend to evolve when they are needed to validate some particular point or when a major system acquisition is planned. Many CPE groups never have to use either simulation or benchmarks.

Computer Performance Evaluation Tools

	Tool	Commercial availability	Cost range (note a)	Computer system overhead
1.	Accounting data reduction package	Yes (note b)	Low to medium	Small
2.	Software monitor	Yes (note b)	Low to medium	Small to moderate
3.	Program analyzers/ optimizers	Yes	Low	Moderate (program)
4.	Hardware monitors:			
	Basic	Yes	Low to medium	None (note c)
	Mapping	Yes	Medium to high	do.
	Intelligent	Yes	High to very high	Slight (note c)
5.	Benchmarks	No	Usually very high	Usually 100%
6.	Simulation:			
	Languages	Yes	Medium to high	Small
	Packages	Yes	High to very high	Small

a/Costs often change dramatically on these types of products. They are shown here in relative terms for most products in each category solely for comparative purposes. According to a December 1978 computer report: low = $5,000; medium = $25,000; high = $100,000.

b/These products are widely available for larger IBM systems but on a very limited basis for other systems.

c/Hardware monitors which produce data that must be reduced after the fact on the subject computer create some overhead at that time. Only intelligent monitors which communicate with the subject computer cause overhead during the monitoring session.

SECTION III

A FLOW DIAGRAM OF COMPUTER PERFORMANCE EVALUATION APPLICATIONS

For the reader interested both in an overview of computer performance evaluation applications for existing and proposed computer systems, and in CPE tools that are best suited to particular tasks, a flow diagram is included below.

A FLOW DIAGRAM OF COMPUTER PERFORMANCE EVALUATION APPLICATIONS

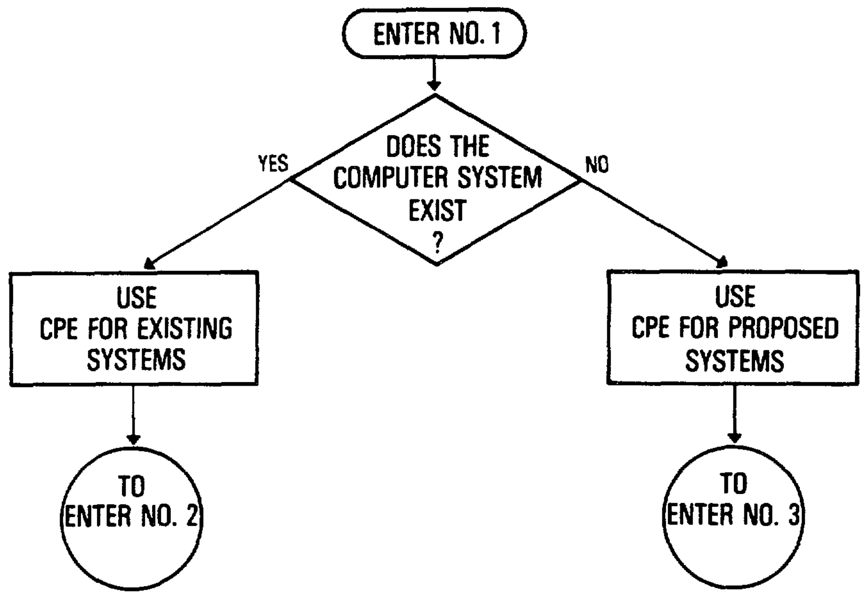

COMPUTER PERFORMANCE EVALUATION FOR EXISTING SYSTEMS

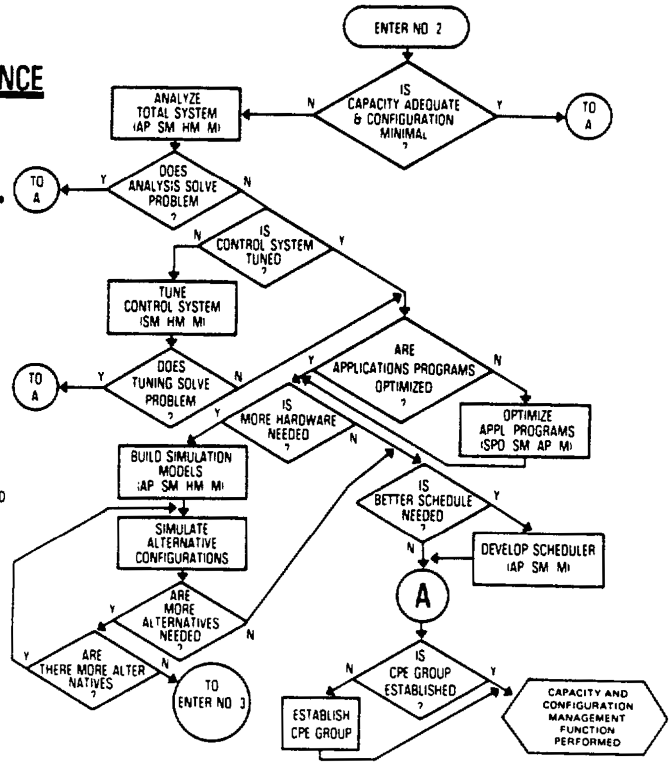

NOTE PARENTHETICAL REFERENCES TO CPE TOOLS INDICATE THE TOOLS BEST SUITED TO THE PARTICULAR TASK AND THE DESCENDING ORDER OF USEFULNESS OF EACH TOOL NAMED

ABBREVIATIONS	
AP	ACCOUNTING PACKAGE
HM	HARDWARE MONITOR
M	MANUAL
SM	SOFTWARE MONITOR
SPO	SOURCE PROGRAM OPTIMIZER

NOTE SOME FORM OF CAPACITY PLANNING IS ALWAYS NEEDED A PERMANENT CPE GROUP SHOULD BE ESTABLISHED WHEN IDENTIFIABLE SAVINGS EXCEED THE COST OF THE CPE EFFORT

COMPUTER PERFORMANCE EVALUATION FOR PROPOSED SYSTEMS

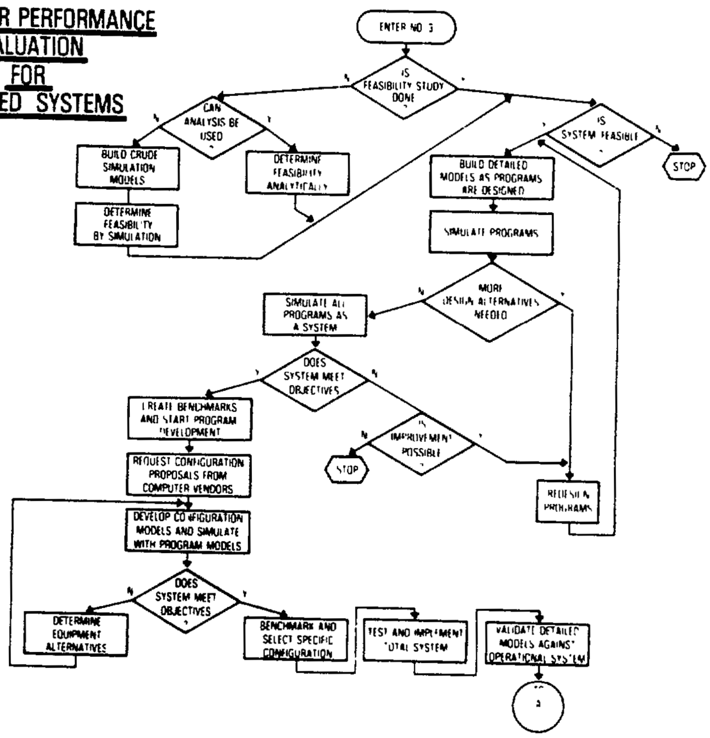

SECTION IV

PRESENTING COMPUTER PERFORMANCE EVALUATION DATA TO MANAGEMENT

An audit finding that cannot be understood by management is of little value. Understanding the workings of computer systems usually decreases dramatically at each level of management above the computer center. Because upper level managers often do not understand the details of a computer's operation, the auditor must put the information into terms that can have meaning for all report recipients.

TERMINOLOGY

On occasion, high-level managers specify the data they wish to have reported to them. In such cases the words used in the reports are familiar management terms like productivity, availability, capacity, service satisfaction, control, waste, and timeliness. Managers always want to be aware of trends. That is, is their agency doing better, worse, or the same as it has in the past?

Generally, first-level managers and technicians tend to track computer performance in terms like CPU busy time, megabytes of memory available, percent busy for channels and peripheral devices, throughput, EXCPs per CPU second, ABEND rates, and so forth. Operational-level people know what these terms mean, but few higher level managers have the time or interest to understand such terms.

CONSISTENCY

It should be recognized that no matter how high computer costs may seem, they are a small part of the total expenditures of most agencies. As a result, higher level management is generally uninterested in computer performance unless poor management of the computer installation is strongly suspected. One thing that causes managers to suspect that all is not well in the computer room is a performance report in computer jargon that cannot be quickly understood. The high level of concern in the use of computers in many industries may be traced to reports that are not understandable. This is especially true when the formats and terms of the reports are changed to suit variations in reporting periods. Reports should be consistent. The auditor is responsible for making certain that audit reports are consistent and can be understood by management.

PERFORMANCE INDICATORS

The best reports are graphical. A "picture" that conveys today's performance and relates it to past performance highs, lows, and averages is easily and quickly grasped. If such a picture can be given with just a few numbers (as in a stock market report), then numbers can be a graphical report. Unfortunately, reducing the performance of a computer installation to a few numbers is, except in very special cases, very difficult.

In industries with essentially one product, a good way to report computer performance is in computer support cost-per-product. In an insurance company, the reporting unit could be computer cost in cents per policy in force or per policy serviced. An aircraft company might report in dollars of computer cost per airplane produced. Along with these present indicators, the reports should state the indicator's previous high, low, and average. When the current number is substantially different from average (whether higher or lower), a short explanation should accompany the report.

Many major computer users cannot create such reports either because their products are so many and so diverse or because they have no identifiable product (as in a government agency). In these cases, more imaginative reports are needed.

KIVIAT GRAPHS

One type of report that has drawn considerable attention in the CPE community is the "Kiviat Graph." 1/ Currently, two forms of the Kiviat Graph are in use. The first and most widely used graph requires the creator to select an even number of performance indicators that are important at the installation where the graphical report is used. Half of these indicators are taken to be "good" when they increase in numerical value. These indicators are numbered alternatively starting with a "good" indicator, next a "bad" indicator, and so forth. The indicators are plotted on the radii of a circular diagram, and each radius, which represents a range from 0 to 100 percent, is evenly spaced around the circle. The radius in the topmost vertical position is numbered "1." The goal is for a Kiviat Graph to approach a star-shaped figure--no matter what parameters are plotted. Following these simple conventions, an eight-axis Kiviat Graph is

---

1/Morris, M.F., "Kiviat Graphs and Single-Figure Measures Evolving," COMPUTERWORLD, Newton, MA 02160; Part 1, Feb. 9, 1976; Part 2, Feb. 16, 1976.

depicted below for a particular IBM 360 installation. The graph shows four performance indicators that were considered good and four that were considered bad.

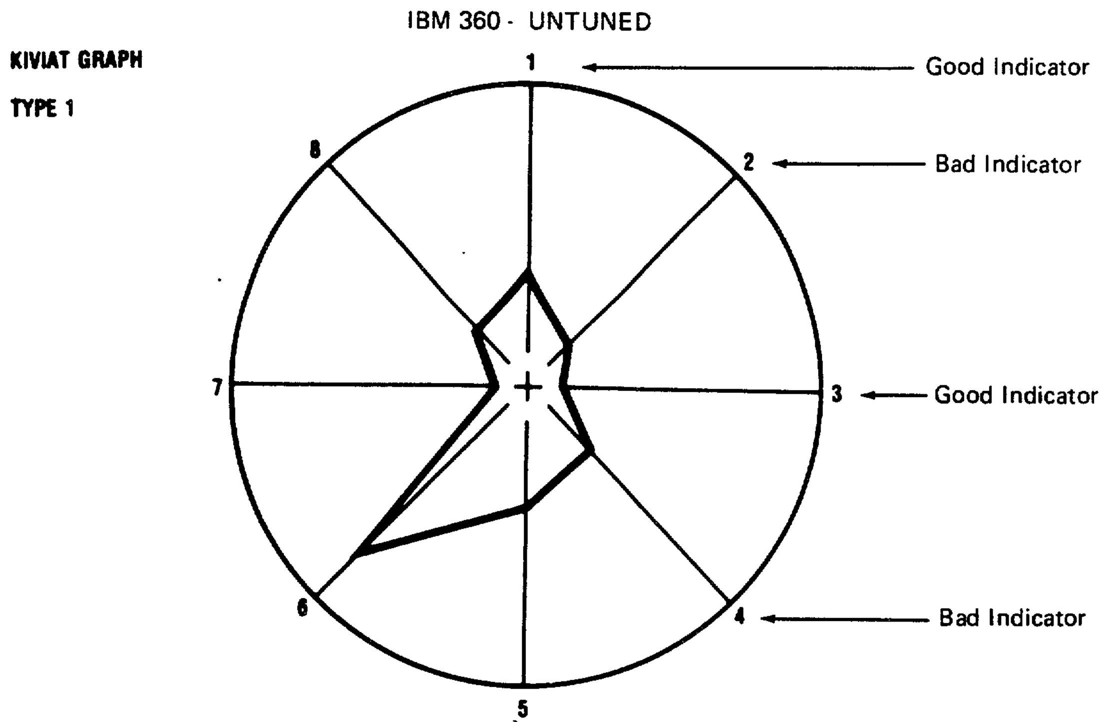

	ACTIVITY	%		ACTIVITY	%
1	CPU Active	28	2	CPU Only	17
3	CPU/Channel Overlap	11	4	Channel Only	21
5	Any Channel Busy	32	6	CPU Wait	71
7.	Problem State	8	8	Supervisor State	20

The indicators shown in this figure are resource usage indicators that might be of interest in an installation concerned with matching the total equipment configuration to the demands placed on it by the total workload. In this example, the Kiviat Graph does not approach a star-shaped figure, therefore the workload does not match the equipment configuration very closely.

The next figure shows the same eight indicators on a system where the workload does match the equipment very closely. As expected, the Kiviat Graph is star-shaped.

IBM 360 — TUNED

KIVIAT GRAPH

TYPE 1

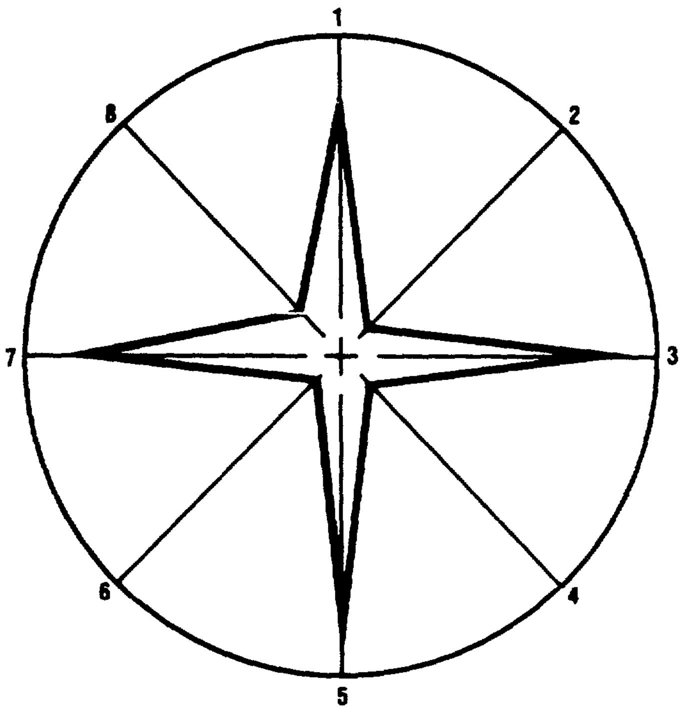

	ACTIVITY	%		ACTIVITY	%
1.	CPU Active	91	2	CPU Only	6
3.	CPU/Channel Overlap	85	4	Channel Only	7
5.	Any Channel Busy	92	6.	CPU Wait	9
7.	Problem State	78	8	Supervisor State	13

The next figure shows ten indicators for a CDC 6400 computer system that are very different than those in the previous figures. It shows the Kiviat Graph of this installation before adding "extended core storage" (ECS) to the system. (ECS provides additional main memory capacity.)

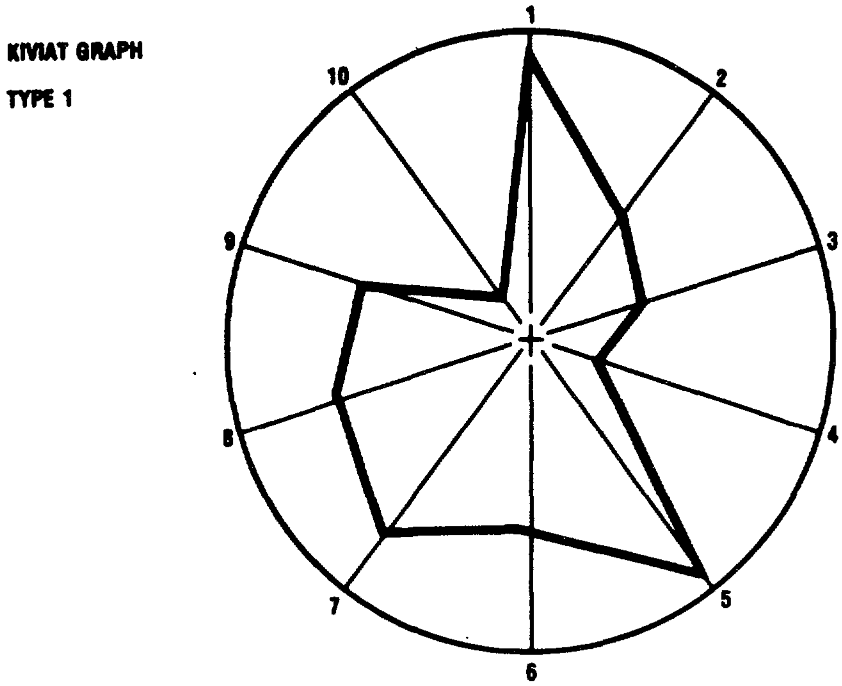

	ACTIVITY	%		ACTIVITY	%
1.	ACTIVE PP'S	93	2	INPUT QUEUE WAIT	51
3.	CP USAGE	39	4.	CONTROL POINT DWELL	23
5	CM USAGE	93	6	AVERAGE TURNAROUND	59
7	ACTIVE CONTROL POINTS	80	8	INPUT QUEUE LENGTH	67
9.	JOBS COMPLETED VS STANDARD	58	10	TIME PP'S ENQUEUE	17

CDC 6400 — BEFORE ECS ADDITION

The following figure shows the Kiviat Graph after extended core storage is added to the above system. The tendency toward a star-shaped graph should make it clear to any level of management that the addition of ECS permitted a better workload and equipment match.

KIVIAT GRAPH

TYPE 1

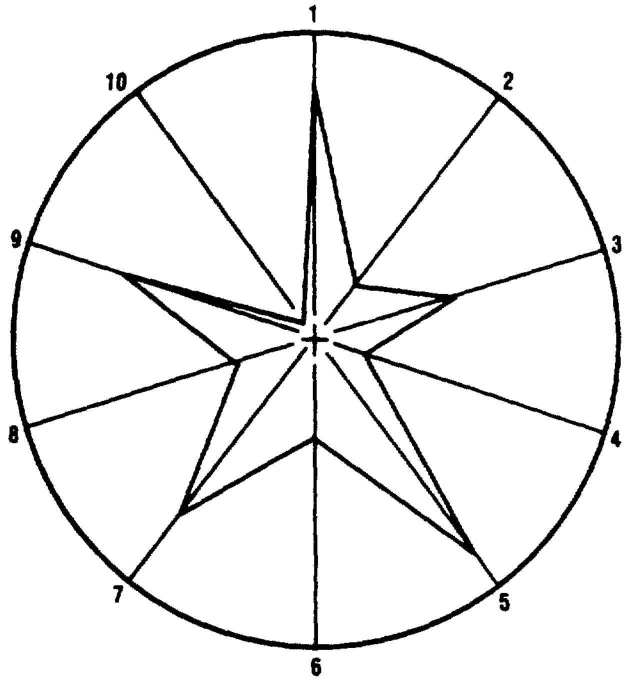

	ACTIVITY	%		ACTIVITY	%
1.	ACTIVE PP'S	87	2	INPUT QUEUE WAIT	22
3.	CP USAGE	47	4	CONTROL POINT DWELL	18
5	CM USAGE	84	6	AVERAGE TURNAROUND	31
7	ACTIVE CONTROL POINTS	76	8.	INPUT QUEUE LENGTH	28
9.	JOB COMPLETED VS. STANDARD	68	10.	TIME PP'S ENQUEUE	2

CDC 6400 — AFTER ECS ADDITION

Trends have been shown in this type of Kiviat Graph by plotting previous highs and lows for each indicator on the axis against a current reporting period's indicator value. A better approach is one which indicates the volatility of each plotted parameter. In that approach, the statistical standard deviations above and below the mean are shown for the observed indicator for the report period. Unfortunately, either of these approaches tends to make the Kiviat Graph visually "busy" and harder to understand at a single glance.

A second form of the Kiviat Graph (illustrated below) incorporates the statistical history of the indicator into the graph itself.

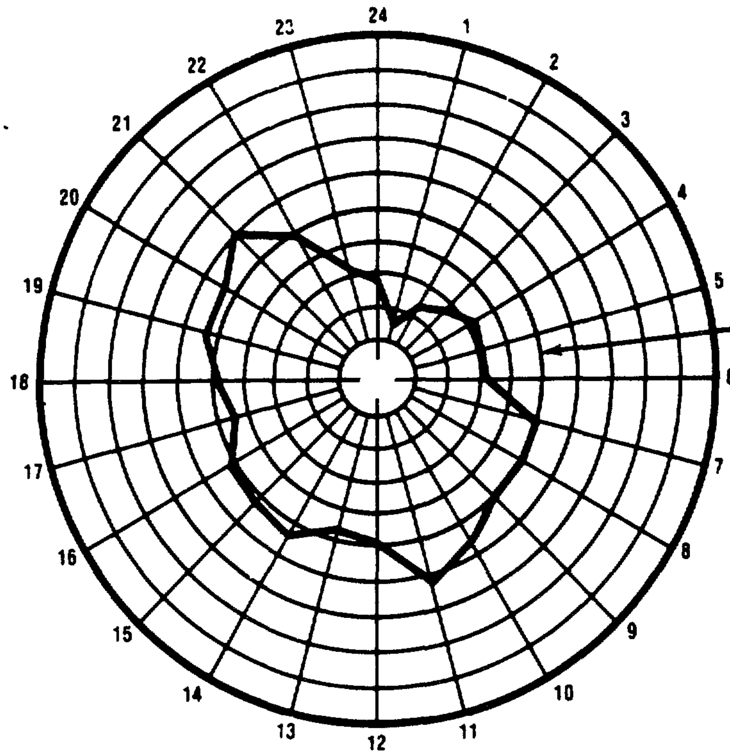

The figure above shows 24 axes whose plots are based on calculations of all previous observations contained in the graph. The axis show the latest period as deviations from the format established by past performance. The center point of each radius in this figure is the statistical mean of all past observations. Outward from the circle, which connects all of the means, are circles which are one-half standard deviation above the mean. Inward from the mean, each ring represents a half of a standard deviation below the mean. That is, beginning from the circle's center, which represents all observations at or below 2-1/2 standard deviations below the mean, the smallest ring is -2 standard deviations, next is -1 1/2, -1, -1/2, 0 (the mean), +1/2, +1, +1 1/2, +2. The circumference of the circle, or outer ring, represents all observations at, or above, 2-1/2 standard deviations above the mean. The plots on this figure show the most recent observations as they relate to all past observations included in the calculation of the variable plots.

The entire area of the above figure between the -1 and +1 standard deviation rings may be regarded as "normal" performance for the indicators plotted. Statistically, this would be nearly 70 percent of all previous observations. If this entire area can be regarded as "normal," and if management is only interested in unusual performance variations, then the figure may be redrawn with a circle at the midpoint of each radius that represents all plots from -1 through +1 standard deviations from the mean.

The following Kiviat Graph includes the above data which is replotted so that any observations between -1 and +1 standard deviations appear on the central, "normal" ring. Those few observations that are not "normal," (see axes 1, 2, and 3 of the figure) would be the only ones that required a narrative explanation.

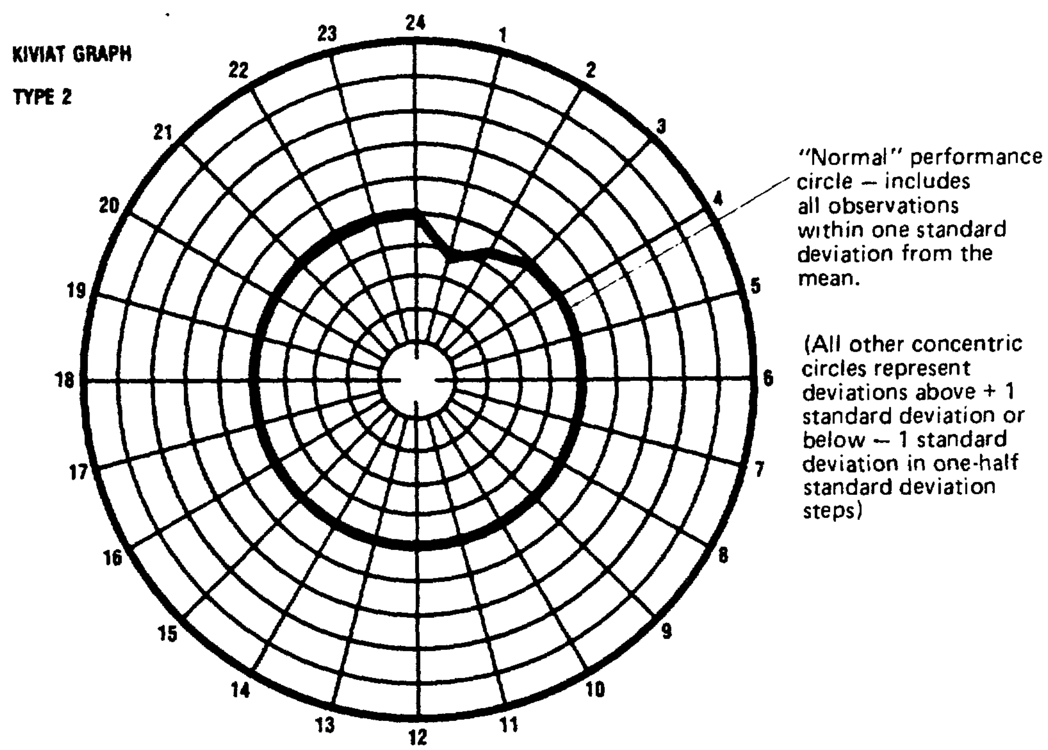

When this second form of Kiviat Graph is used, the conventions regarding even numbers of indicators (half good and half bad) do not apply. Any number of indicators that are of interest at a particular installation may be plotted on the graph. The important point when using this type of Kiviat Graph is that long-term trend information must be maintained for each variable.

SCHUMACHER CHARTS

Another graphical management report that can show essentially the same information as the second type of Kiviat

Graph is called the Schumacher Chart. 1/ The following chart is a Schumacher Chart with a few performance indicators that can be reported weekly.

The structure of the Schumacher Chart incorporates management terminology (effectiveness, availability) as major headings along with more technical computer terms as subheadings. Generally, the most useful method is for each installation to select the important management-level topics and then develop the specific technical indicators which will support those topics.

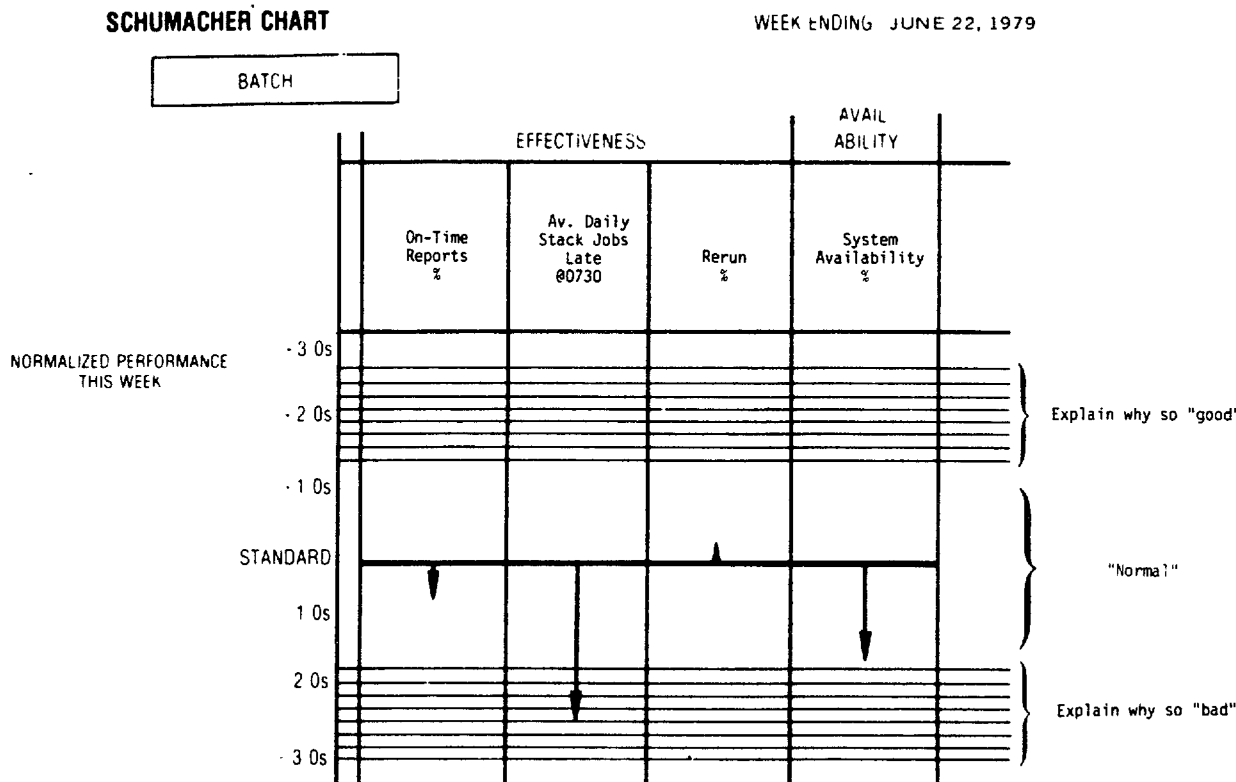

	On-Time Reports %	Av. Daily Stack Jobs Late @0730	Rerun %	System Availability %	
STANDARD	YTD AV	YTD AV	YTD AV	YTD AV	
GOAL					
YTD AV	95 9	6 3	6 7	91 7	
1979 HIGH	98 6	18 8	11 9	96 2	
THIS WEEK	94 6	18 8	6 6	88 1	
1979 LOW	90 4	0 0	4 2	86 8	
BEST VALUE OR DIRECTION	100	0	0	100	

1/Schumacher, David, "A Graphical Computer Performance Report for Management," Proceedings of the BBUG-V Meeting, Oct. 2, 1974.

LINE AND BAR GRAPHIC

Many other graphical techniques are used to report performance to management. Line graphs and bar graphs are examples. Although indicators tend to include technical terminology, the continuity portrayed in these graphs permits overall trends to be seen. Such trends are generally meaningful to management whether the specific indicators are well understood or not. Examples of line graphs and bar graphs follow.

EQUIPMENT USAGE INDICATORS

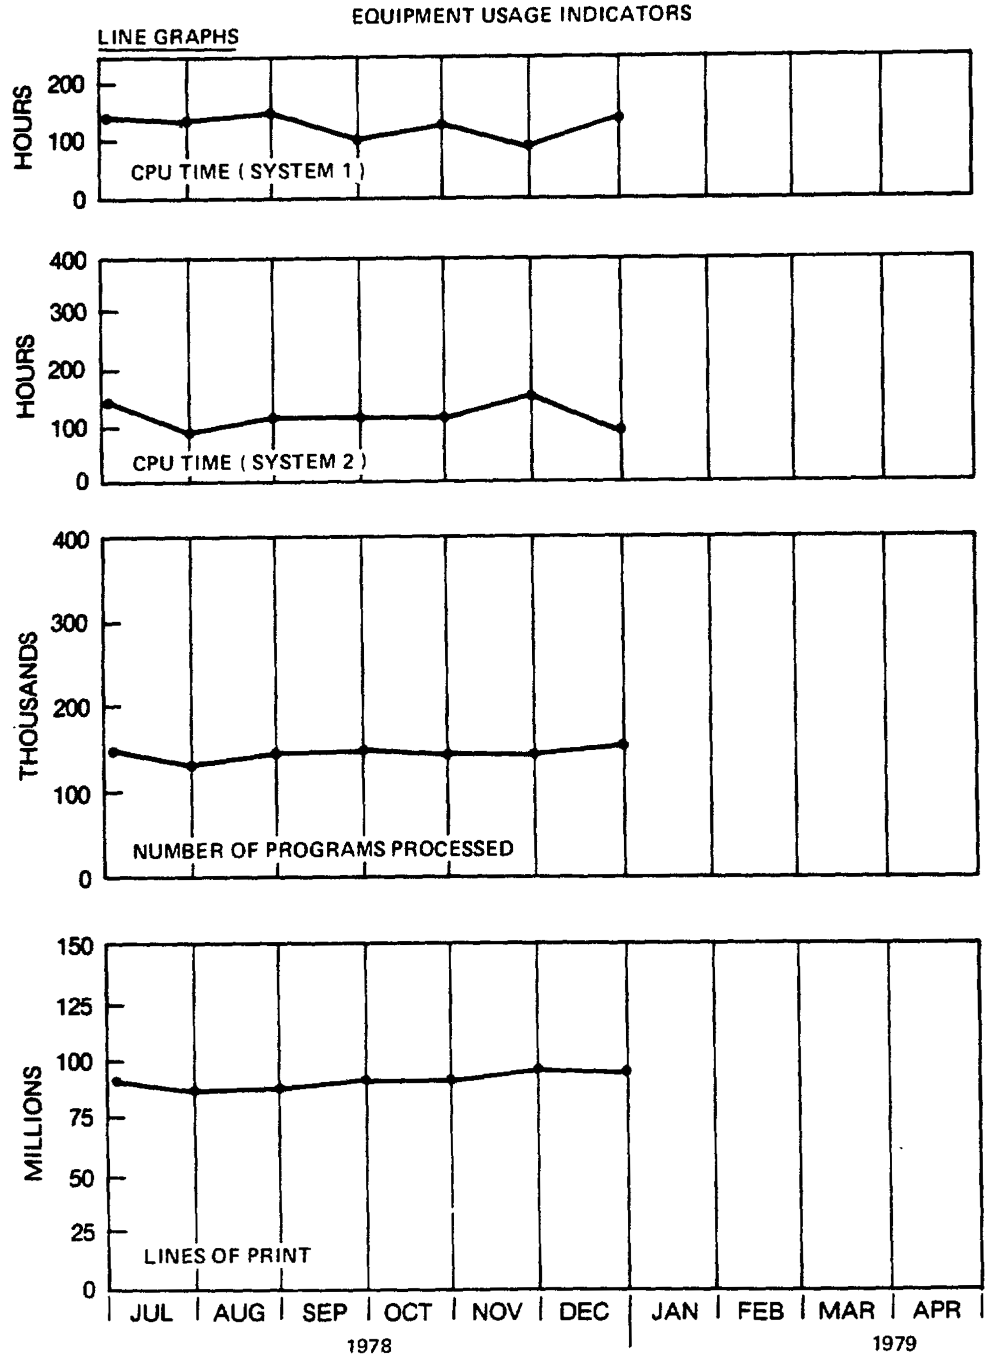

COMPUTER TIME USAGE BY CATEGORY

BAR GRAPH

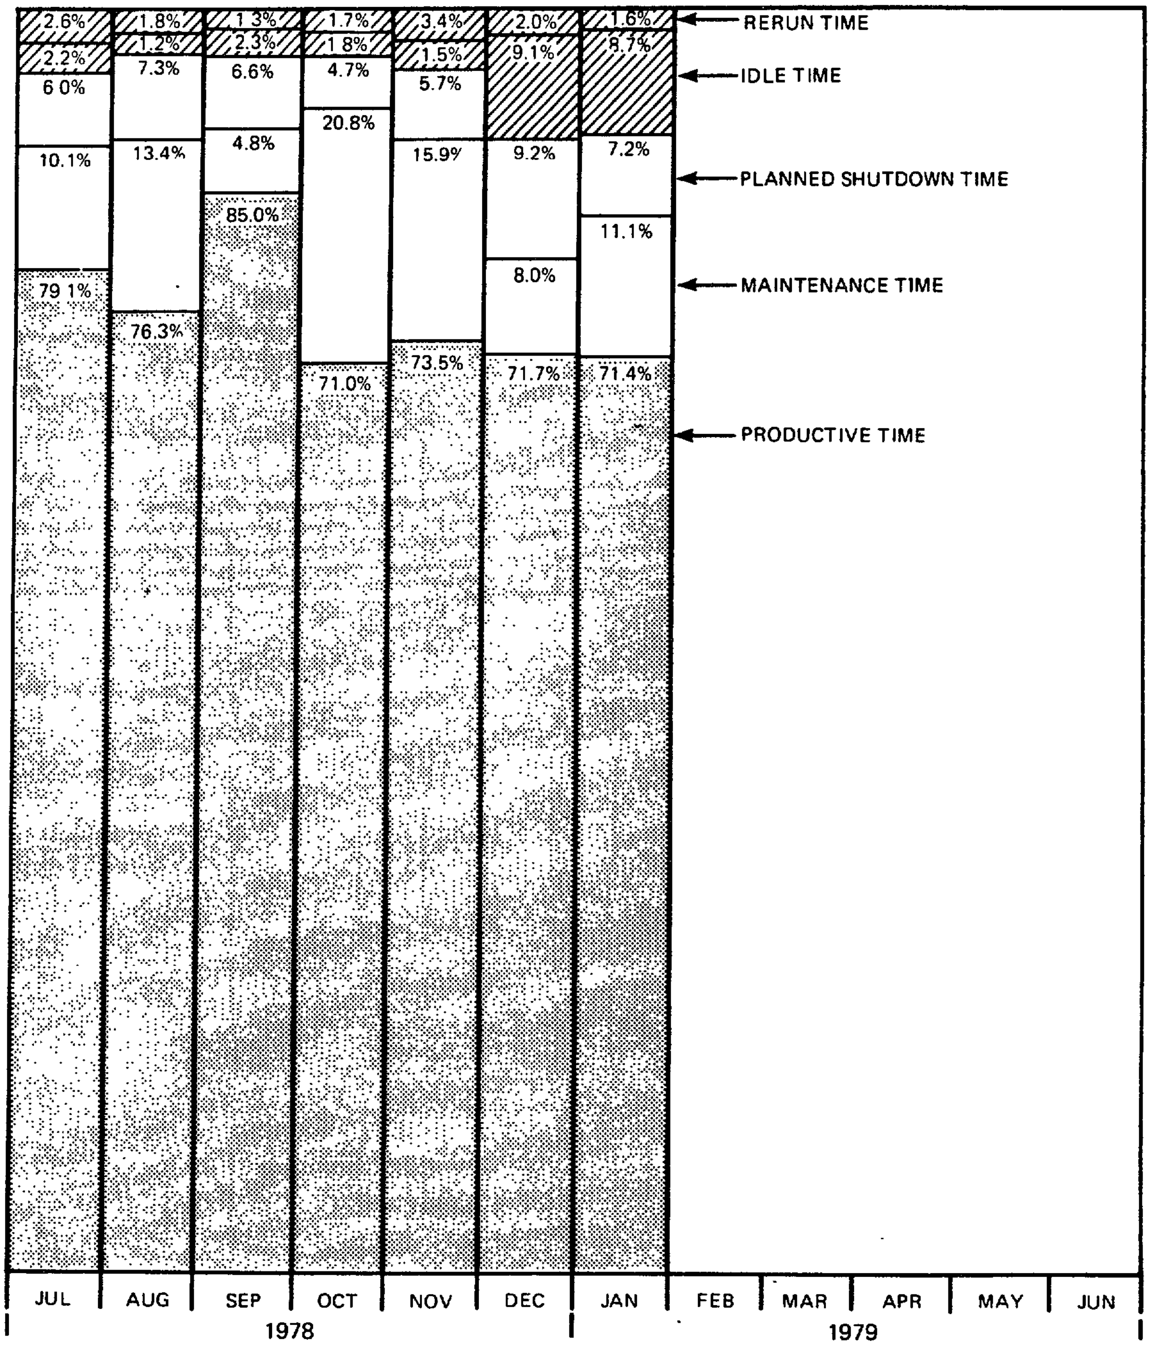

PRODUCTIVE TIME BY MAJOR SYSTEMS

BAR GRAPH

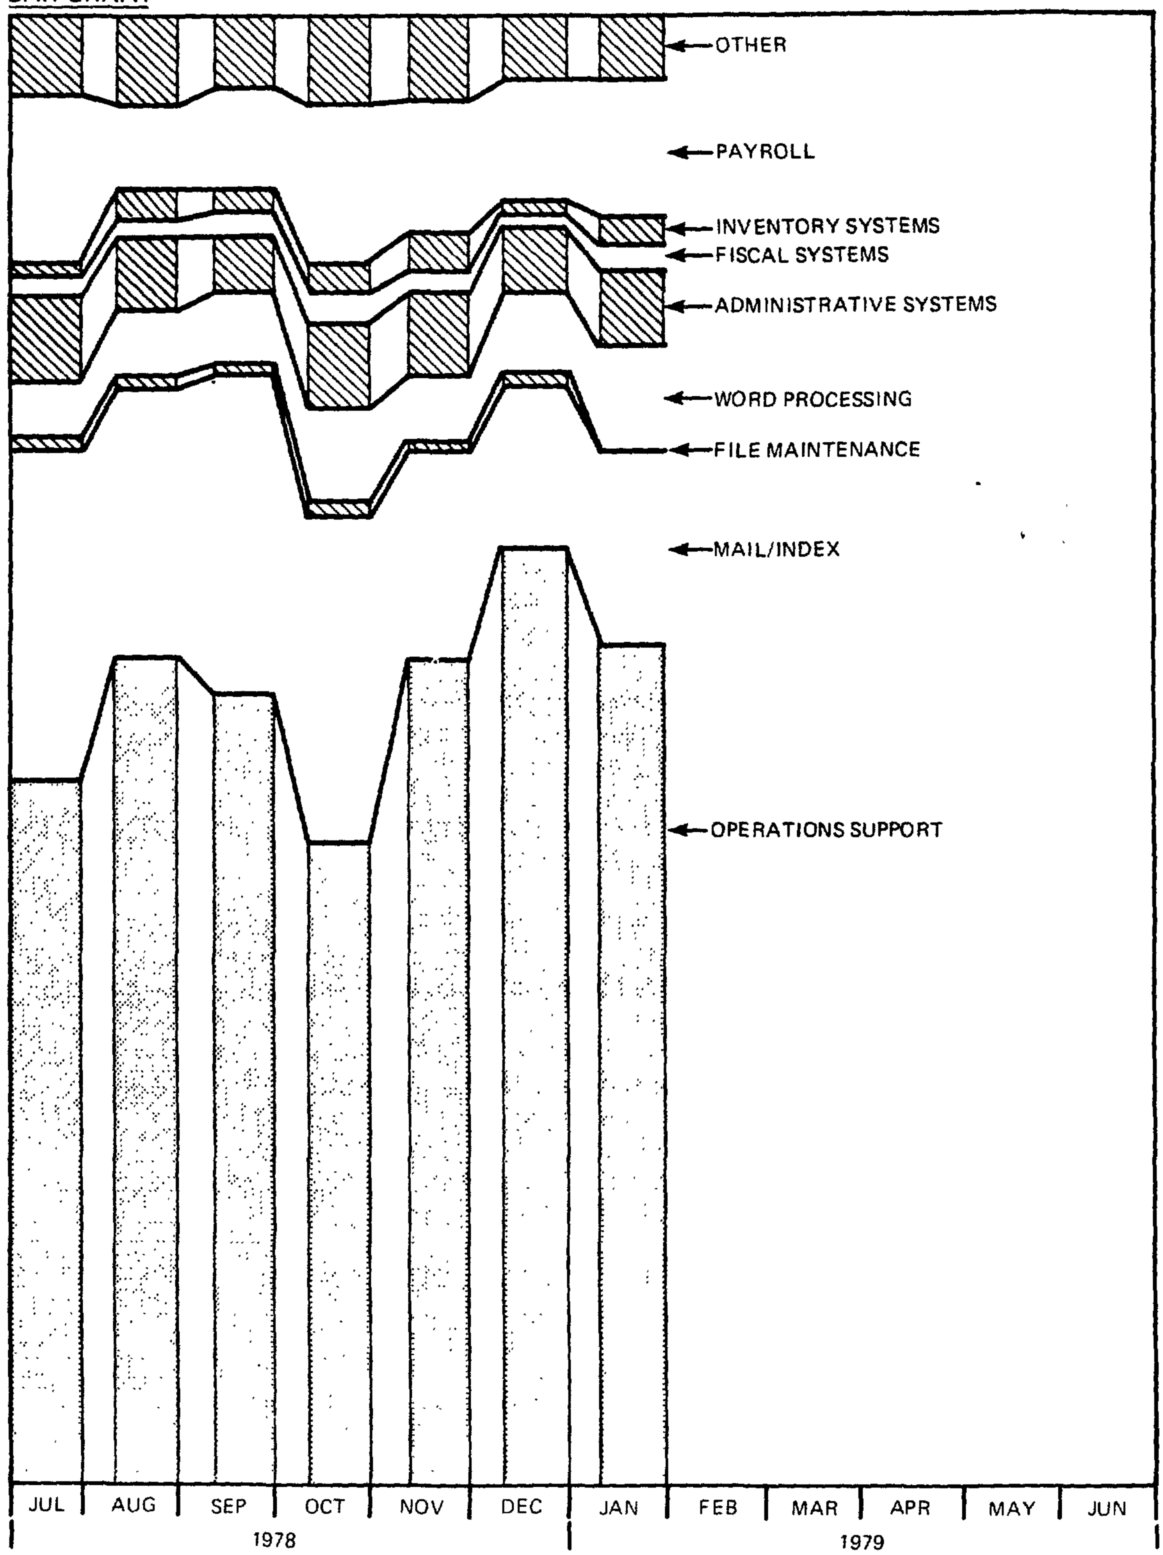

PERFORMANCE GOALS

The need for some goal or standard of system performance is implicit in the above discussion. Many installations have no performance goals. By setting practical goals and comparing them with actual performance, installation managers can evaluate how well their installation is performing. For example, goals can be set for the workload to be accomplished by the computer system (jobs completed, hours of operation, multiprogramming level), the efficiency of the system (CPU usage, memory usage, I/O activity), or the reliability and availability of the system. In addition, installation managers can establish service-level objectives (response time, turnaround time, system availability) with users based on the installation's computer system performance goals.

To establish practical goals, an installation should develop a history of utilization and performance data over 4 to 6 months. Next, the installation should analyze the data to determine correlations between equipment-usage statistics and service-level statistics, e.g., CPU usage vs. response time. Based on this analysis, practical performance goals can be established.

The installation manager can use the performance goals in managing computer performance to help provide answers to such questions as:

--Can we decrease the workload or make the application programs more efficient?

--Do we need an additional or a more powerful computer?

--Do we need additional components (memory, tape, disk, etc.)?

--Can we eliminate certain components or replace them at a lower cost?

--Can we improve certain aspects of computer service (response time, turnaround time)?

Performance goals are important to the auditor performing an efficiency and effectiveness audit of a computer installation because they provide the criteria for measuring actual performance.

SELECTED BIBLIOGRAPHY

ACCOUNTING DATA

A Guide to Major Job Accounting Systems: The Logger System of the UNIVAC 1100 Series Operating System; National Bureau of Standards Special Publication 500-43; Dec. 1978.

"Charging for DP Services: Two Methods Analyzed," Dan Bernard; Government Data Systems; pp. 23-24; Jan./Feb. 1978.

"Cost of Using Excess Capacity," EDWIN L. HEARD; Journal of Systems Management; pp. 24-27; Sept. 1976.

Guideline on Major Job Accounting Systems: The System Management Facilities (SMF) for IBM Systems Under OS/MVT; National Bureau of Standards Special Publication 500-40; Oct. 1978.

"Management Use of Job Accounting Data for Performance Evaluation," V. BANNA; Computer Measurement Group Transactions; Mar. 1977.

"Managing Computer Performance With Control Limits," THOMAS E. BELL; Performance Evaluation Review; pp. 21-28; Jan. 1975.

"The Computer Pricing Process," C.W. Axelrod; Computer Measurement Group Transactions; Dec. 1977.

HARDWARE/SOFTWARE MONITORS

"A Guide to the Use of Hardware Monitors," GARY CARLSON; EDP Performance Review; pp. 1-8, Sept. 1976; pp. 1-7, Oct. 1976.

"Effectively Using Your Hardware Monitor," ROBERT BISHOP; Computer Measurement Group Transactions; Sept. 1976.

"Evaluation and Comparison of Software Monitors," EDP Performance Review; pp. 1-9; Feb. 1976.

"Fallacies of Intuition in Performance Tuning," GARY CARLSON; EDP Performance Review; pp. 1-6; Apr. 1976.

"Making the Most of Performance Monitors," WILLIAM B. ENGLE; Computer Decision; pp. 50; Nov. 1978.

"Software Monitoring: Principles of Design and Use," DUANE RAY BALL; Proceedings of the Computer Performance Evaluation Users Group 1976; pp. 215-219; 1976.

"Techniques for Software Monitoring", T.B. PINKERTON; System Tuning; pp. 347-358; 1977.

"The Design of Software Monitors," I.L. RAMAGE; System Tuning; pp. 359-372; 1977.

"Tutorial: Computer System Monitors," GARY J. NUTT; Performance Evaluation Review; pp. 41-51; Jan. 1976.

PROGRAM ANALYSIS/OPTIMIZATION

"COBOL Optimization and Flowcharting," M.M. HEADRICK; National Technical Information Service; Y/CSD/INF-77/1; 1977.

"Gaining An Awareness of the Performance of COBOL Programs," PAUL J. JALICS; Computer Measurement Group Proceedings; pp. 61-65; 1978.

"Improving Performance the Easy Way," PAUL J. JALICS; Datamation; pp. 135-137; Apr. 1977.

"Measuring Programming Quality and Productivity," T.C. JONES; IBM Systems Journal; Vol. 17, No. 1, pp. 39-63; 1978.

"Performance Evaluation of Programs in a VS Environment," JAMES R. WALKER and THOMAS E. MCKITTRICK; Computer Measurement Group Proceedings; pp. 62-78; 1976.

BENCHMARKS

"Benchmarking Distributed Systems: Objectives and Techniques," THOMAS F. WYRICK; Performance of Computer Installations; pp. 83-101; (D. Ferrari, Ed.) North-Holland; 1978.

"Benchmarking in Selection of Timesharing Systems," D.J.M. DAVIES; Computer Performance Evaluation User Group Proceedings; pp. 27-36; 1978.

Guidelines for Benchmarking ADP Systems in the Competitive Procurement Environment; National Technical Information Service; FIPS-PUB 42-1; 1977.

"Remote Terminal Emulation in the Procurement of Teleprocessing Systems," SHIRLEY WARD WATKINS and MARSHALL D. ABRAMS; AFIPS Conference Proceedings - National Computer Conference; pp. 723-727; 1977.

"Validation--All Important in Benchmarking," ARNOLD L. JOHNSON; Computer Performance Evaluation Users Group; pp. 75-83; 1977.

SIMULATION

"A Case Study of Simulation as a Computer System Design Tool," GARY J. NUTT; Computer; pp. 31-36; Oct. 1978.

A General Computer Network Simulation Model, HOYT M. WARREN; National Technical Information Services; Mar. 1977.

"Computer Simulation: A Tutorial," GERALD ADKINS and UDO W. POOCH; Computer (Magazine of the IEEE Computer Society); pp. 12-17; Apr. 1977.

Simulation with GPSS and GPSS V, FRANK H. SPECKHART and WALTER L. GREEN; Prentice-Hall; 1976.

"The Use of Simulation in the Evaluation of Software," N.F. SCHNEIDEWIND; Computer (Magazine of the IEEE Society); pp. 47-53; Apr. 1977.

GENERAL

"A Survey of EDP Performance Measures," JOSEPH R. MATTHEWS; Government Data Systems; pp. 29-32; Jul./Aug. 1978.

"Computer Performance Analysis: Framework and Initial Phases for a Performance Improvement Effort," T.E. BELL, B.W. BOEHM, and R.A. WATSON; The Rand Corp.; R-549-1-PR; Nov. 1972.

"EDP Effectiveness Evaluation," CORYDON D. HURTADO; Journal of Systems Management; pp. 18-21; Jan. 1978.

EDP Performance Management Handbook, Applied Computer Research, 1978.

"Getting Started in Computer Performance Evaluation," PHILIP J. KIVIAT and MICHAEL F. MORRIS; CMG Transactions; pp. 3-2, 3-19; Dec. 1975.

"Guidance for Sizing ADP Systems," DENNIS M. GILBERT, JAMES O. MULFORD and MITCHELL G. SPIEGEL; Computer Performance Evaluation Users Group; pp. 305-30; 1978.

Guideline on Computer Performance Management: An Introduction, National Technical Information Service; FIPS-PUB-49; 1977.

"Kiviat Graphs and Single-Figure Measures Evolving," MICHAEL F. MORRIS; Computerworld; pp. 17-18, Feb. 9, 1976; pp. 19-20, Feb. 16, 1976.

"The State of the Art In Analytical Modeling," RICHARD F. DUNLAVEY; EDP Performance Review; pp. 1-7; Nov. 1978.

"The Value of Performance Measurement," IVAN BERTI; Computer Management; pp. 57-58; Sept. 1978.

"What are We Measuring?" ROBERT B. FOREST; Infosystems; pp. 92-94; May 1978.

Page

☆U.S. GOVERNMENT PRINTING OFFICE: 1979 O—627-885/1934 REGION 3-1