PERFORMANCE MONITORING: A SURVEY OF TECHNIQUES UTILIZED BY THE UNITED STATES ARMY

by

THOMAS BRYANT

B. S., Morgan State University, 1963

A MASTER'S REPORT

submitted in partial fulfillment of the

requirements for the degree

MASTER OF SCIENCE

Department of Computer Science

KANSAS STATE UNIVERSITY
Manhattan, Kansas

1976

Approved by:

Major Professor

ACKNOWLEDGMENT

The writer wishes to express his appreciation to Drs. Paul S. Fisher, Myron Calhoun, and Virgil E. Wallentine of the Department of Computer Science, Kansas State University, for the constructive criticism and advice given in the preparation of this report. The writer also wishes to thank his wife, Sandra, for her patience and endurance during the compiling of this report.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

1 Introduction

1.1. Background

The complex applications of the current generation of digital computers have increased the importance of methods used to measure and evaluate a computing system's performance. In the past, computers were designed with a fairly simple objective, which was to be as fast as possible in performing a particular application. However the measurement of speed was often dependent on the area of application.[1]

Early computers were designed for either general purpose or special purpose use. General-purpose computers were usually classified as either scientific or commercial.

Because the principle application of scientific computers was to perform calculations, the speed of scientific computers was based on certain discrete capabilities such as add, multiply, or divide time. Frequently, the computer was used to perform a well defined application which required a repetitive set of calculations. The arithmetic speed of a system would be measured in relation to the calculation process.

The commercial data processing field evaluated computing systems based on input/output rates. Large volumes of data were fed into computing system through use of a card reader or some other type of input device; the data was processed by the system, and such things as payrolls,

---

[1]Mansford E. Drummond, Jr., Evaluation and Measurement Techniques for Digital Computer Systems (New Jersey: Prentice-Hall, Inc., 1973), p. 23.

inventories, accounts, and billings resulted. The computing system's performance was assessed by the rates at which reading and writing occurred. Cards per minute, lines per minute and characters per second were typical performance parameters used in commercial data processing.

As computing systems and their applications changed, the simple but useful evaluation techniques of examining the add time or input/output rates were inadequate when evaluating a computer systems performance. Innovations such as I/O buffering have been devised. This technique allows the slow process of reading or writing of data to be overlapped with some other computing activity. Other innovations developed have been large capacity direct access storage devices and programming packages capable of loading other programs from an input storage medium into main storage. These innovations have increased the complexities of computing systems.[2]

Consequently, as computing systems have become more complex, there has been an increase in the number of elements to be measured in order to evaluate a computing systems performance.

To evaluate the total system's performance, such things as secondary storage, I/O channel and peripheral unit performance must be evaluated. In addition to the hardware performance, the software performance must also be considered. This requires the use of different measurement techniques to properly evaluate a computer system's performance.

Currently there are various measurement tools and techniques available which can be utilized to evaluate both the hardware and software performance of a computer system.

---

[2] Ibid., pp. 32-60.

1.2. Purpose of the Report

The United States Army is a very large user of computer systems and has several Data Processing Installations (DPI) and computing centers located throughout the Continental United States, the Pacific, and Europe, whose function is to process various personnel, finance, and logistic records. The United States Army utilizes various performances monitoring techniques to evaluate the performance of their computer system's hardware and software.

The purpose of the report is to survey the performance monitoring techniques utilized by the United States Army.

1.3. Organization of the Report

Section 2 highlights the current state-of-the-art in performance monitoring. Specifically, this section examines the tools and techniques which are available and used in conducting performance monitoring of computing systems. In this section only the hardware, software, and hybrid monitors are discussed.

Section 3 surveys performance monitoring tools and techniques utilized by the United States Army. Included is a discussion of how these tools and techniques used by the United States Army compare with the tools and techniques available on the commercial market.

Section 4 contains the conclusion and recommendations reached as a result of the survey. This section discusses the "effectiveness" of the United States Army's computer's performance monitoring techniques. Recommendations are made for deficiencies or shortcomings noted.

2 Highlights of Performance Monitoring

Tools and techniques for computer system performance monitoring have been around for several years.[1] The purpose of this section is to highlight the current state-of-the-art in performance monitoring. As such this section will examine the tools and techniques which are available and are used in conducting performance monitoring of computer systems. The ability of these tools in providing assistance in improving the effectiveness or optimization of computing systems will be discussed. General advantages and disadvantages which the users of the performance monitoring tools should be aware of will be listed.

2.1. The Need for Performance Monitoring

The computer equipment presently in use throughout the world represents an investment of billions of dollars.[2] The operating, programming and maintenance expenses of computer hardware are very large. With these tremendous expenses involved, it would seem reasonable to expect that computer systems perform to their maximum or at least optimal capabilities. However, reports have shown that the average large computer

---

[1] Mansford E. Drummond, Jr., Evaluation and Measurement Techniques for Digital Computer Systems (New Jersey: Prentice-Hall, Inc., 1973), p. 265.

[2] James H. Sood, "Shopping the World Computer Market," Info Systems, Vol. 20, No. 3 (Mar. 1973), p. 57.

operates at less than 30 percent efficiency.[3] Furthermore, the reports indicate that many computer systems are poorly coordinated and wasteful; the central processing unit (cpu) is not active much of the time, with very little cpu/peripheral overlap.

These inefficiencies in computer systems have affected the job turn-around time and system throughput rates. This would indicate that such inefficiencies represent a very poor return on the amount of money invested in computer systems and that there exists a need to improve the computer system performance. A possible solution to poor computer system performance could be to purchase a "bigger and better" computing system, however this does not guarantee that the same inefficiencies will not exist in the "bigger and better" system.

A more viable solution would be to correct the inefficiencies that may exist in the present computer systems. The correction will require the optimization or tuning of all available computer systems resources. This would include the hardware which consist of the cpu, memory, tapes, disc, card readers, and other peripherals and the software consisting of the assembler, compiler, operating system, subroutine libraries etc. Optimization of a computer system could mean a change in the number, type, speed, and organization of system components; greater interaction between peripherals; more cpu-I/O overlap, or a redistribution of peak loads and other bottlenecks that degrade a systems performance.

Before any of these changes can be accomplished effectively, a thorough evaluation of the computer system will be required. Questions such as the following will have to be answered. What is the system

---

[3]Dudley Warner, "Monitoring: A Key to Cost Efficiency," Datamation (January, 1971), p. 41.

doing? When is it doing it? Why and what is the priority? To what end is the system utilizing its resources? These questions can be answered through the use of system performance measurement techniques.[4]

2.2. Performance Measurement Defined

What is a performance measurement? A performance measurement can be defined as the process of quantitatively describing the operating characteristics of a digital computer system.[5] It may also be defined as a process which allows the user to obtain quantitative measurements of a computer system's performance while normal processing is underway. These definitions represent a contrast to the intuitive measurement of system performance that might be provided by users of computer systems.

2.3. Tools and Techniques

There are generally two preferred techniques to system performance measurement. Both involve the process of monitoring the computer system to obtain data describing the actual performance of an existing system. The techniques involve the use of hardware and software monitors. These terms refer more to the means used to perform the measurement than to the type of information being collected. However, each monitor has its own areas of operation. These will be explained in the discussion of each monitor.

2.3.1. Hardware Monitor

As previously stated the hardware of a computer system consists of the central processing unit (cpu), memory, tapes, discs, cardreaders, and

---

[4] _Ibid._, p. 40.

[5] Systems Development Corporation, "A Guide to Computer System Measurement," (unpublished, undated), p. 2.

other peripherals. A hardware monitor, sometimes called a "black box," is a device that is connected directly to the computer's circuitry and is used to measure the performance of computer hardware.[6]

Listed below are some of the categories of information that can be measured by a hardware monitor. Many other functions can be monitored and measured at the option of the user.

1. Active time of central processing unit (cpu).
2. Problem vs. supervisor time.
3. Activity by region (e.g. core mapping, data set organization, file structure).
4. Task-switching frequency.
5. Instructions executed.
6. Program and routine timing.
7. I/O organization and selection by
   a. channel.
   b. controller.
   c. device.
   d. component (e.g. disc arm).
8. Overlap of cpu-I/O.
9. Multipath balance (e.g. paths to each device).[7]

2.3.1.1. Classes of Hardware Monitors

Hardware monitors can be divided into two classes of monitors. The summary type has the characteristic that it accumulates the time or occurrence of a particular type event for the total amount of time or

---

[6] Drummond, Evaluation and Measurement Techniques for Digital Computer Systems, p. 207.

[7] Warner, "Monitoring:," p. 42.

total number of occurrences over a specified period. The summary type monitor provides the "how much" or "how often" information as to occurrences of an event. The dynamic monitor is the second type and provides the "when" of event occurrences. The inherent speed and capacity of the device and the speed of the output media are the primary differences between the two classes of hardware monitors.[8]

2.3.1.2. Organization and Function

The hardware monitor can accomplish the measurements previously listed because it is physically connected to the host computer. Figure 2.1 shows a hardware monitor and is an example of the main components of

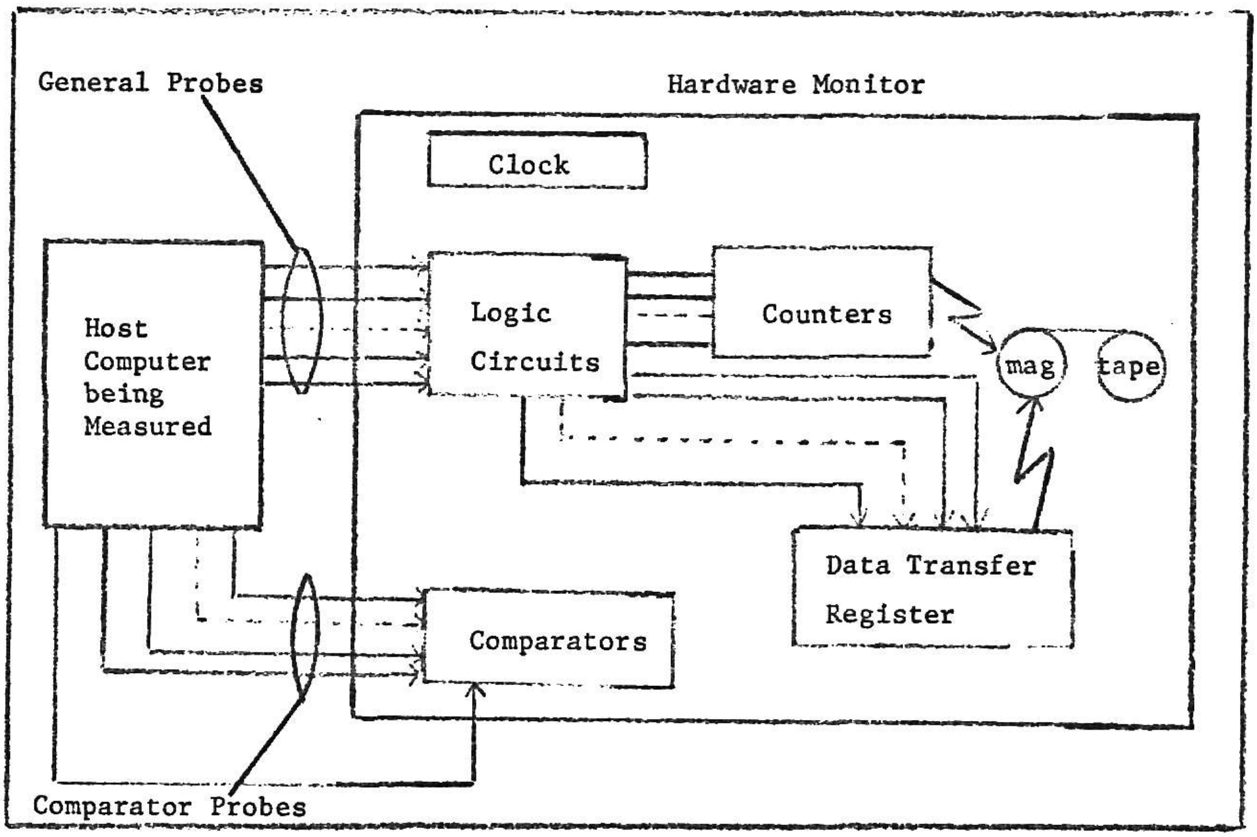

Figure 2.1: Elements of a Hardware Monitor

[8]Drummond, Evaluation and Measurement Techniques of Digital Computer Systems, p. 240.

the hardware monitor.[9] It should be noted that this figure is only one example of the many hardware monitors available on the market today.

The principle elements of a hardware monitor are:

1. General probes are a set of electronic signal signal sensors connected to the host computer system by physically connecting to the wire-wrap pins on the back panels of the computer system. These sensors detect electronic signal fluctuations representing such functions as "cpu wait," "channel busy," supervisor mode, and so forth. The sensor is used to detect the presence or absence of the signal which it is monitoring. This sensing process does not degrade the host system and the sensors are claimed to be "electronically transparent" to the host system.[10]

2. Logic circuits accept signals from the general probes and allow logical combinations of the signals (AND, OR, INVERTS, etc.). Events of interest (e.g. cpu active, channel active etc.) may be defined by combining signals from various pins in the required manner.

3. A group of counters are used to count the occurrences of various events or to measure the time between events by counting the number of intervening clock pulses. An example would be to count the occurrence of disc seeks, and also time the duration of these seeks. Selected information is then placed into buffered storage, or accumulating storage or both for later output. This operation may vary depending on the type of hardware monitor.

4. Comparator probes are similar to the general probes, but are used to

---

[9] Jerre D. Noe, "Acquiring and Using a Hardware Monitor," Datamation (April 1974), p. 89.

[10] Ibid.

sense a number of bits that appear in parallel (e.g. as in an address register).

5. Comparators provide a means for comparing the parallel bits with some value that has been preset by the user and may be used to determine either equality of incoming signals or the relationship of the signals to boundary conditions.

6. The data transfer register provides the means for passing data directly from the host computer to a magnetic tape recorder or to some display depending on the type monitor.

Hardware monitors may vary in function, cost, etc., but certain aspects (e.g. use of sensing probes, etc.) are common to all hardware monitors. Information is displayed and recorded in a number of ways. In an operational environment, some monitors have real time displays available on a small video screen, while for a more permanent record, a magnetic tape recorder is used.[11]

At the high price end of hardware monitors, a dedicated processor such as the Dynaprobe 8000 with a PDP 11-45, represents some of the latest technology. Earlier hardware monitors were specially constructed devices designed to perform the monitoring function. The use of another computer, specifically a minicomputer, represents a response to the rising cost of these specially designed hardware monitors. As the cost of specially designed hardware monitors rises, it approaches that of a minicomputer.[12] The minicomputer can be used to handle the monitoring task as well as a variety of other functions.

---

[11] Richard Slatter, "Sparing the Time to Analyse," Data Processing (July-August, 1974), p. 222.

[12] See Data Pro 70 The EDP Buyers Bible for price comparisons.

2.3.2. Software Monitor

The second technique of monitoring involves the use of a software monitor. The feasibility of a software monitor is derived from the method used by most operating systems to detect conditions that affect resource utilization (e.g. privileged operation mode and the interrupt structure).

A software monitor is a specially written program(s) which when used resides in memory and by an intercept or a sampling concept extracts data on various facts about the computer system environment. By inserting additional code at key points in the system control program or problem programs, intercept points are established which invoke the software monitor routine. Figure 2.2 shows the concept of intercept points to invoke the software monitoring routines. The sampling concept involves the invoking of the software monitor program at a regular interval, or time or when a combination of events occur that meets some criterion.[13]

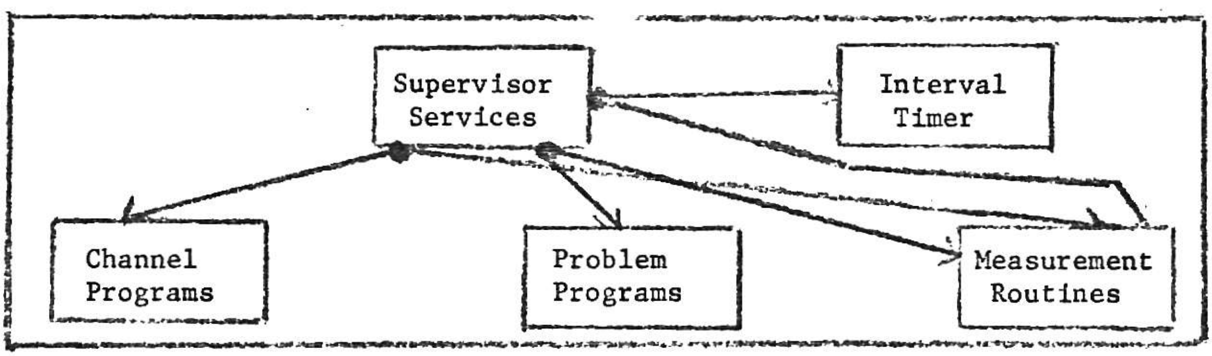

Figure 2.2: Concept of Intercept Points

When the software monitor is executed, it examines and extracts data pertinent to the activity about the system. Periodically this

[13] Drummond, Evaluation and Measurement Techniques of Digital Computer Systems, pp. 219-20.

accumulated information is recorded on suitable media such as tape or disk. A program analyzer, normally supplied with commercially available software monitor, is used for reduction and analysis of data gathered.

The output reports produced by the analyzer display information in the form of lists and charts. Information, identified as system or application related, of the following types is usually included in such reports:

1. Total elapsed system time.
2. CPU active time.
3. CPU wait time.
4. Operating system CPU time.
5. Problem program CPU time.
6. Channel(s) busy time.
7. Channel(s) overlap time.
8. Controller(s) busy time.
9. Device(s) busy time.
10. Total number of interrupts.
11. Total number of seeks on direct access storage devices.
12. Core memory partition usage.
13. Device activity (number of events).

There are two basic types of software monitors available. One type is the system monitor which gets its name because it is a facility used to measure the overall computer system activity and is device resource oriented. A second type is the program monitor which is used to measure program related activity in order to identify inefficiencies or bottlenecks in a particular program. Figure 2.3 illustrates one structure of a software monitor. This structure minimizes core usage by separating the

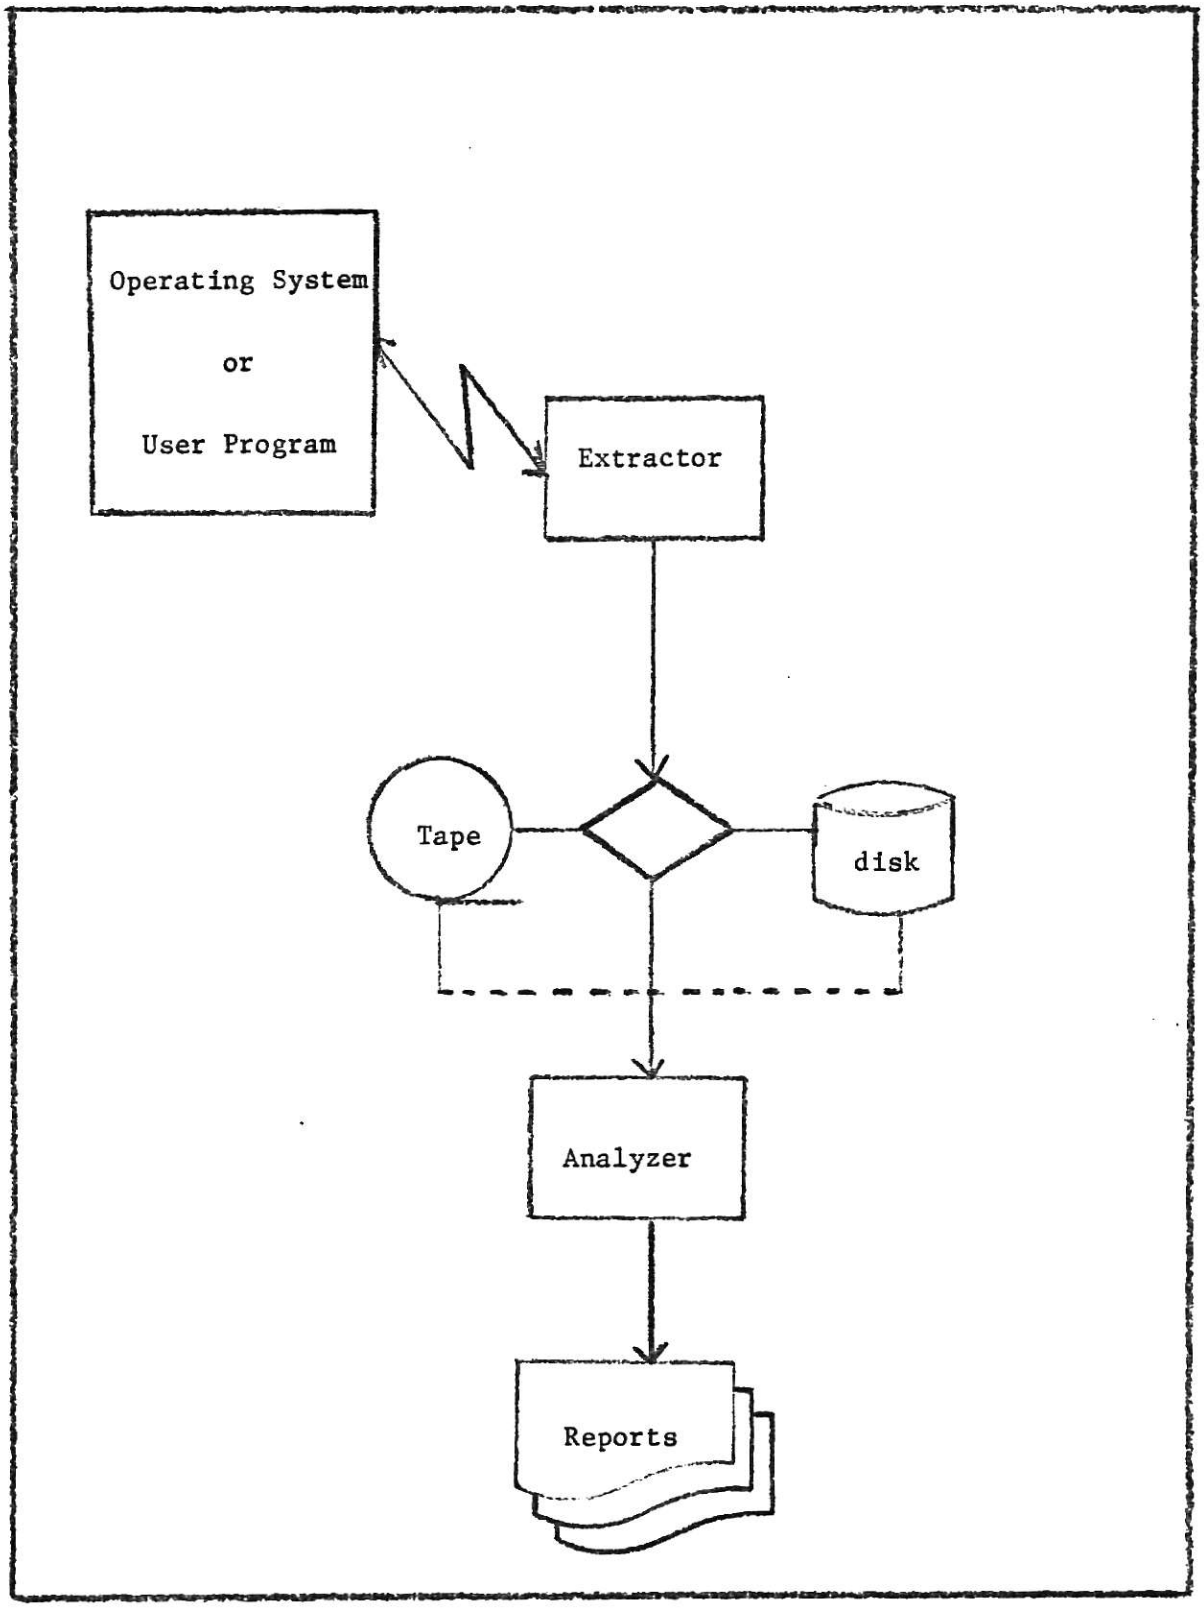

Figure 2.3: Type of Software Monitor

data extractor mechanism from data analysis function. While some software measurement tools combine the extractor and data reduction mechanisms, experience with them tends to strengthen the case for separation. Through various methods, the extractor enters a problem program or the operating system to collect data about the program or system. This data is periodically written on tape or disk. When the analyzer is invoked, the data is processed and various reports are produced.

2.3.3. General Advantages and Disadvantages

Hardware and software monitors each have their relative merits. Each type of monitor has its advantages and disadvantages. Figure 2.4 is a comparison of the hardware monitor versus the software monitor.

	Hardware	Software
1. Cost	High	Low-medium
2. Overhead	None	10-40%
3. Hardware/System dependent	No	Yes
4. Precision	Monitor dependent	Host dependent
5. Qualitative Measurement	Limited	Yes
6. User Training Required	Extensive	Limited
7. Flexibility	Good	Poor
8. Ease of Use	Poor	Good
9. Portability	Fair	Good
10. Device Monitoring	Yes	Maybe
11. Interrupt Recognition	Yes	Yes
12. Lifetime	Good	Fair

Figure 2.4: Comparison of Hardware vs Software Monitors

A software monitor has the following advantages: (1) very easy to use, (2) portable, (3) can monitor entire system and point out bottlenecks, (4) relatively inexpensive ($8-12K), (5) can also monitor individual programs, (6) has a unique advantage in that it can determine queue lengths.

A software monitor has the following disadvantages: (1) distortion of results (Heisenberg uncertainty principle),[14] (2) high overhead while executing, (3) operating system dependent, (4) cannot monitor multiple cpu's, (5) sampling technique prevents gathering all occurrences.

A hardware monitor has the following advantages: (1) no cpu or device overhead on monitored system, (2) no distortion of data, (3) can sample or collect all events, (4) extremely accurate, (5) has special advantages of being able to monitor multiple cpu's and can accurately monitor the activity of the operating system.

The disadvantages of a hardware monitor are: (1) very expensive ($70K), (2) talented user required, (3) high set-up time, (4) can collect data, but provides little intuition.[15]

2.3.4. Hybrid Monitor

The choice between a hardware and software monitor must be weighed in terms of system resource utilization and consequent degradation caused by the type of instrumentation used and the relative degree of difficulty dictated by the technique to be employed. Therefore, the very purpose of

---

[14] While using a software monitor to measure the performance of a computer system, the monitoring may have an effect on the performance of the computer.

[15] The hardware monitor has a major defect, in that it cannot conveniently provide detailed information about where in the system and problem programs and the data sets the various activities are concentrated.

monitoring which was to gain insight into the behavior of the system, to help trace and determine cause/effect relationships suggests a combination of both techniques. Since a computer system should be viewed from an integrated (hardware/software) point of view, a concept was developed using a similar approach in the monitoring field. This approach resulted in the development of the hybrid monitor. The hybrid monitor retains a hardware data acquisition front end while having a software that can interact with the operating system. The HEMI (Hybrid Events Monitoring Instrument) represents an attempt at synthesizing the former hardware and software monitoring techniques. It is an experimental instrumentation system being developed for use with the CYBER 70 and 170 series computers. This effort represents some of the newer technology in the field of monitoring.

2.4. Utility of Performance Monitoring Tools

Although hardware and software monitors have been around for several years and improvements in the technology of the monitors have been accomplished, the widespread use of monitors did not occur until the early 70's.[16] Perhaps the reason for this slowly developing acceptance and use lies in the fact that these monitors are fairly sophisticated scientific tools and in the past there weren't that many skilled people with the knowledge of how these monitors should be used. Perhaps, the use of monitors did not become important until computers and their applications became so complex.

Experience of some users of monitors has shown that the best results have been obtained in those cases where the user had defined a specific

---

[16] Kenneth W. Kolence, "A Software View of Measurement Tools," Datamation (January, 1971), p. 32.

problem area in his shop and used a hardware or software monitor to identify the causes of the problem. The worst results occurred when the user simply monitored his system with no specific goal in mind.

Users of monitors indicate that before monitoring is attempted, one should define a specific objective to be accomplished. This objective could be one of the following: (1) increase throughput (get more work done by the system), (2) reduce turnaround time (get a job, program, or transaction through the system in less time), (3) improve cost performance (get the work done at minimum cost).

With this objective established, the environment in which the monitors are to be used must be identified which involves the determining of the following: (1) representative work cycle (daily, weekly, and monthly), (2) operations scheduling constraints (pre-emptive priorities), (3) equipment schedules (is the equipment configuration going to change in the near future?).

After the objective and the environment have been identified, a general system profile should be obtained by using a hardware or software monitor. The profile should provide a good picture of overall systems performance, from which areas of potential improvement can be identified. If changes in system hardware or software organization are indicated from the profile, these changes should be accomplished and measurements repeated to see if changes had any effect.

An example of a performance measurement that resulted in improved job turnaround, was the measurement conducted at an Allied Chemical Data Center. The center had an IBM 360/50 and a 360/40 with a work load consisting of local batch processing, remote job entry and online systems. Measurements were taking using both hardware and software monitors.

These measurements revealed a cpu utilization on the 360/50 of 37-percent and selector channel utilization low and unbalanced. Consequent system changes and reorganization that were accomplished after the initial monitoring, resulted in a 57-percent cpu utilization and an increase in channel and cpu overlap on the 360/50. All the center's workload was transferred to the 360/50 and 360/40 was released. This report and others have indicated the utility of hardware and software monitors.[17]

---

[17] Philip G. Bookman, Barry A. Brotman, and Kurt L. Schmitt, "Use Measurement Engineering for Better System Performance," Computer Decisions (April, 1972), pp. 27-32.

3 Survey of Performance Monitoring in the United States Army

The United States Army spends more than $378 million annually for computing systems. It manages approximately 950 computers of all types and sizes. As a user of computer systems the United States Army has an interest in insuring that these computer systems perform effectively. One manifestation of this interest is the performance monitoring efforts taken by various Army agencies. This section surveys the performance monitoring tools and techniques utilized by the United States Army. This survey includes a discussion of the need for performance monitoring within the Army and examines the performance monitoring that is conducted by various Army agencies.

3.1. The Need for Performance Monitoring Within the Army

The costly commitment to computing systems by the Army developed during the 1960 decade. An annual growth rate of 15 percent occurred as a result of the "information explosion," and the use of sophisticated weaponry, personnel, and logistics techniques. Although the Army made significant gains in the battle to stem spiraling computing systems costs, the resource commitment remains significant. As a constant goal, the Army has sought to support its various information needs efficiently and at the lowest possible cost.

Various circumstances have retarded progress in attainment of this goal. One of the circumstances has been the fact that over the years, new applications and increasing workloads have imposed greater requirements for sophisticated computing equipment and highly skilled personnel. This resulted in an ever greater system complexity. This complexity has, in turn, challenged management efforts to obtain accurate measurement of hardware performance. Traditional job accounting systems furnished by computer manufacturers for performance measurement have become less and less satisfactory. The simple measurement of job run time and collection of "billing type information" fail to provide information of the kind and in the amount necessary to manage efficiently today's average computer system and workload. A second circumstance related closely to the first is the increasing reliance on sophisticated software and, in particular, the significant rise in software development and maintenance costs in recent years. This attaches increased importance to software costs in relation to other system costs.

Fortunately, various products and techniques to measure hardware performance and to curb and reduce software development costs have appeared over the last few years.

The Army recognized the potential of these new computer performance measurement products and techniques as early as December 1970 when information relative to their use was first disseminated Army-wide. Limited use began in 1972. The acquisition and use of these devices have thus far been accomplished on a decentralized basis. Various Army publications contain verbal policy encouraging individual data processing installations to use the new techniques. None of these publications prescribe a specified performance measurement program. It would seem that a centralized approach for the development and implementation of policy and procedures for performance measurement would more fully tap its available potential and yield a much wider range of benefits.[1]

3.2. Performance Monitoring in the Army

Experience in performance monitor usage has been growing since the first market survey was made by the U.S. Army Computer Systems and Evaluation Agency (CSSEA) in February 1970. In December 1970, the Army published Technical Bulletin, 18-20, Improving Computer System Efficiency. This bulletin contains general information on monitor usage, a list of vendors, their advantages and disadvantages, and represents an initial effort to encourage use.

In May 1972, CSSEA acquired a Dynaprobe 7700 hardware monitor. With the acquisition of this monitor, CSSEA has also had responsibility for maintaining a current and thorough knowledge of performance monitors, providing advice and assistance on monitor usage to all Army activities, and conducting performance measurements with monitors as requested.

In February 1973, a change to Army Regulation 18-1 required Army activities to obtain approval prior to acquisition of monitors. Except for this provision, no other policies, procedures, or standards for monitor usage are mentioned in this or related regulations. The approach to performance measurement largely has been a decentralized program. Noteworthy results with monitors, however, have been achieved at CSSEA, Computer System Command (CSC) and the Automated Logistices Management System Agency (ALMSA).

---

[1]U.S., Department of the Army, Army Audit Agency, "Armywide Audit of Selected Aspects of the Utilization and Management of Automatic Data Processing Equipment" (5 Oct. 1973), p. 116.