A11103 089809
NAT'L INST OF STANDARDS & TECH R.I.C.
A11103089809
Nutt, Gary J/A survey of remote monitori
QC100 .U57 NO.500-, 42, 1979 C.1 NBS-PUB

C CE & TECHNOLOGY:

A SURVEY OF REMOTE MONITORING

500-42
9

NBS Special Publication 500-42

U.S. DEPARTMENT OF COMMERCE
National Bureau of Standards

NATIONAL BUREAU OF STANDARDS

The National Bureau of Standards[1] was established by an act of Congress March 3, 1901. The Bureau's overall goal is to strengthen and advance the Nation's science and technology and facilitate their effective application for public benefit. To this end, the Bureau conducts research and provides: (1) a basis for the Nation's physical measurement system, (2) scientific and technological services for industry and government, (3) a technical basis for equity in trade, and (4) technical services to promote public safety. The Bureau's technical work is performed by the National Measurement Laboratory, the National Engineering Laboratory, and the Institute for Computer Sciences and Technology.

THE NATIONAL MEASUREMENT LABORATORY provides the national system of physical and chemical and materials measurement; coordinates the system with measurement systems of other nations and furnishes essential services leading to accurate and uniform physical and chemical measurement throughout the Nation's scientific community, industry, and commerce; conducts materials research leading to improved methods of measurement, standards, and data on the properties of materials needed by industry, commerce, educational institutions, and Government; provides advisory and research services to other Government Agencies; develops, produces, and distributes Standard Reference Materials; and provides calibration services. The Laboratory consists of the following centers:

Absolute Physical Quantities[2] — Radiation Research — Thermodynamics and Molecular Science — Analytical Chemistry — Materials Science.

THE NATIONAL ENGINEERING LABORATORY provides technology and technical services to users in the public and private sectors to address national needs and to solve national problems in the public interest; conducts research in engineering and applied science in support of objectives in these efforts; builds and maintains competence in the necessary disciplines required to carry out this research and technical service; develops engineering data and measurement capabilities; provides engineering measurement traceability services; develops test methods and proposes engineering standards and code changes; develops and proposes new engineering practices; and develops and improves mechanisms to transfer results of its research to the utlimate user. The Laboratory consists of the following centers:

Applied Mathematics — Electronics and Electrical Engineering[2] — Mechanical Engineering and Process Technology[2] — Building Technology — Fire Research — Consumer Product Technology — Field Methods.

THE INSTITUTE FOR COMPUTER SCIENCES AND TECHNOLOGY conducts research and provides scientific and technical services to aid Federal Agencies in the selection, acquisition, application, and use of computer technology to improve effectiveness and economy in Government operations in accordance with Public Law 89-306 (40 U.S.C. 759), relevant Executive Orders, and other directives; carries out this mission by managing the Federal Information Processing Standards Program, developing Federal ADP standards guidelines, and managing Federal participation in ADP voluntary standardization activities; provides scientific and technological advisory services and assistance to Federal Agencies; and provides the technical foundation for computer-related policies of the Federal Government. The Institute consists of the following divisions:

Systems and Software — Computer Systems Engineering — Information Technology.

[1]Headquarters and Laboratories at Gaithersburg, Maryland, unless otherwise noted; mailing address Washington,D.C. 20234.
[2]Some divisions within the center are located at Boulder, Colorado, 80303.

The National Bureau of Standards was reorganized, effective April 9, 1978.

JAN 25 1979
not in circ
QC100
.U57
no. 500-42
1979
c. 2

COMPUTER SCIENCE & TECHNOLOGY:

A Survey of Remote Monitoring

Gary J. Nutt, Ph.D.
Consultant

Institute for Computer Sciences and Technology
National Bureau of Standards
Washington, D.C. 20234

U.S. DEPARTMENT OF COMMERCE, Juanita M. Kreps, Secretary

Jordan J. Baruch, Assistant Secretary for Science and Technology

NATIONAL BUREAU OF STANDARDS, Ernest Ambler, Director

Issued January 1979

Reports on Computer Science and Technology

The National Bureau of Standards has a special responsibility within the Federal Government for computer science and technology activities. The programs of the NBS Institute for Computer Sciences and Technology are designed to provide ADP standards, guidelines, and technical advisory services to improve the effectiveness of computer utilization in the Federal sector, and to perform appropriate research and development efforts as foundation for such activities and programs. This publication series will report these NBS efforts to the Federal computer community as well as to interested specialists in the academic and private sectors. Those wishing to receive notices of publications in this series should complete and return the form at the end of this publication.

National Bureau of Standards Special Publication 500-42
Nat. Bur. Stand. (U.S.), Spec. Publ. 500-42, 34 pages (Jan. 1979)
CODEN: XNBSAV

Library of Congress Cataloging in Publication Data

Nutt, Gary J
A survey of remote monitoring.
(Computer science & technology) (National Bureau of Standards special publication ; 500-42)
Supt. of Docs. no.: C13.10:500-42
1. Electronic digital computers--Evaluation. I. Title. II. Series. III. Series: United States. National Bureau of Standards. Special publication ; 500-42.
QC100.U57 no. 500-42 [QA76.9.E94] 602'.1s [621.38] 78-26313

U.S. GOVERNMENT PRINTING OFFICE
WASHINGTON: 1978

For sale by the Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402
Stock No. 003-003-02013-3 Price $1.50
(Add 25 percent additional for other than U.S. mailing).

TABLE OF CONTENTS

	Page
Abstract	1
Introduction	2
Background	4
Remote Monitoring Techniques	7
Hybrid Monitors	14
Network Monitors	16
Fault Diagnosis Monitors	20
Intelligent and Extended Consoles	21
Summary and Conclusions	23
Acknowledgement	26
References	27

A SURVEY OF REMOTE MONITORING

Abstract

This report describes remote monitoring in the application areas of performance evaluation, diagnostic testing, performance assurance and system security testing. The evolution of remote monitoring is briefly reviewed and, then, remote monitors are categorized into seven classes. Several example systems are discussed for each classification, along with their capabilities in each application area. The views presented in this report represent only those of the author, an independent consultant, and should not be construed as a policy statement of NBS or any other organization.

Key Words:

Diagnostic testing; performance assurance; performance evaluation; remote monitoring; system security testing

NOTE:

Reports on commercial developments included in the publication were obtained from the open literature, supplemented by documentation and correspondence, referenced by permission. Inclusion or exclusion of specific companies, developments, or products cannot be construed as recommendation or endorsement by the National Bureau of Standards.

Paul F. Roth
Systems and Software Division
Institute for Computer
Sciences and Technology

INTRODUCTION

The general area of remote monitoring of computer systems encompasses a broad spectrum of mechanisms for a wide variety of purposes. In this report, the discussion is restricted to monitoring systems or studies where a mechanism is used to measure or observe the performance of a computer system, and that mechanism can be controlled by another device or a human from some geographically distinct location. In most cases, it is expected that the monitoring device itself is designed to collect data about the host system, perform at least preliminary filtering of the raw data, and then either store the filtered data for retrieval by the central controller or immediately transmit the filtered data to the central monitor controller. The nomenclature used for the various constituents, then, is as follows: The host system is the installation being monitored; a monitor that is local to the host is referred to as a remote monitor. The remote monitor is ultimately controlled from a central site by the central monitor controller. The host computer is considered to be the remote facility, while the measurement control and analysis take place at the central site.

This classification of remote monitors admits such approaches as: those implemented purely in software which can be interrogated from an external terminal, programmable hardware monitors, hardware monitors distributed over different portions of the host machine, hybrid monitors, monitors used in distributed computer networks, fault diagnosis monitors, and extended consoles for a computer system. Each of these categories will be discussed in detail in a later section of this report. The classification excludes classic hardware monitors that require plugboard alterations to change the logical combination of probe signals. It also excludes pure software-implemented monitors which use the normal operating system facilities for "triggering," reporting and recording.

Remote monitors are being used in a number of ways that earlier, locally controlled monitors were not used. The most obvious use of the monitors is for gathering performance data. Remote performance monitor facilities are frequently divided into a number of remote data gathering mechanisms plus a single, shared facility to analyze data and prepare reports; the Tesdata facilities are examples of this type.[60] Distributed monitors, perhaps best exemplified by the PARTNER package for Control Data 6000 series machines,[65] are also frequently used for performance measurements; the idea here is to dedicate certain hardware facilities of the host system to the measurement function. Programmable hardware monitors are merely a refinement of earlier hardware monitors, and also are primarily used for performance measurement.

A newer application of remote monitors is for computer system diagnosis and remote exercising of a computer system. A number of computer manufacturers have included this capability in their current product line.[5,32,47,58,59] The basic idea is to replace the conventional operator's console with an intelligent device such as a minicomputer. The intelligent console can be used to inspect any of a number of conditions that

Superscript numbers indicate literature references at end of the report.

exist in the host machine, allowing the observed condition to be analyzed, recorded or transmitted on a telecommunications link to a remote controller. The remote controller may be a human operator or another computer system. Although this is apparently a new concept to most machine manufacturers, it should be noted that Control Data 6000 series computer systems have used this approach to implement their consoles for a number of years.[57]

A new area for which remotely controlled monitors might be employed is that of performance assurance and safeguard studies. The goal is to monitor the workload of a computer system in order to either assure a given level of performance, or to assure that a computer system is not being used for tasks that were not intended to be executed on that system.

Although it would be satisfying to be able to monitor a processor's program counter to determine what program the processor is executing, this is obviously impossible in the general case. It is easy to construct an example that shows that if one could write an algorithm that inspects the program counter locus and identifies a corresponding algorithm, then one ought to be able to write a similar algorithm that inspects the program counter locus and indicates whether the corresponding algorithm will ever terminate or not. The latter algorithm has been proven to be impossible to construct.[19] Nevertheless, there are other activities in the computer system that can be observed with a monitor, e.g. resouce utilization. One can easily compute the ratio of input/output time to central processor time for a given job. This will allow one to partition heavy computer jobs from input/output bound jobs.

Although it is impossible to identify arbitrary programs in execution, it may be possible to recognize a small set of programs when they are executing on the host system. For example, suppose that an installation is intended to only execute programs P_1, P_2... P_n (on arbitrary data). It may be possible to employ heuristic techniques to recognize exactly when one program from that set executes, while any unrecognizable program is declared to be illegal. In this case, remote monitoring techniques can be used to recognize the "signature" of each of the n acceptable programs.

Finally, remote monitors may be used to enhance system security or to provide a mechanism for checking the security of a system. It is clear that the presence of monitors of any type are a threat to the overall security of a computer system, e.g. see references 6 and 13. Whenever a mechanism (i.e. a monitor) is provided the capability of observing critical portions of the operating system, then that same device can be maliciously employed to penetrate the conventional security mechanisms of that system. By partitioning a monitor into a local internal component and a remote external component, system security has a much better chance of being effective. The internal monitor can be written as an internal portion of the operating system itself, subject to the same design constraints (such as proof of correctness, restricted entry points, authorized access, etc.) as other modules. The attendant software is essentially data-gathering code, which is simpler and easier to make secure than a full software monitor. The external portion of the monitor is allowed to access the internal portion through normal, secure paths, thus

allowing authorization checks and entries into predefined procedures of the operating system.[40] Although this approach is not totally secure, it offers a much more effective security policy than undisciplined monitoring of the host system.

Another variant of remote monitors can be used to audit a computer system's security state. The basic idea is to distribute a monitor of internal and external components, as above. The external portion is used interactively by a human that is responsible for system security to audit various portions of the machine with the aid of the internal portion of the monitor. This approach is used in the WWMCCS computer systems,[29] and will be discussed at length in the body of this report. The Rand Corporation has also investigated the use of monitors to detect data bank intrusions and to delay the intruder until other protective action can be taken.[53]

In the remainder of this report the background of remote monitoring will first be examined. The evolution of present-day monitoring systems will be traced from early performance monitoring work. The main body of this report is the next section; seven categories of remote monitors are defined, and a number of examples of each category are discussed. The final section draws some conclusions about capabilities and limitations for various application areas and looks briefly at future trends in remote monitoring.

BACKGROUND

In this section of the report, the evolution of remote monitors is discussed beginning with hardware and software monitors of the 1960-1970 era. In the early 1970's monitoring techniques and tools became substantially more sophisticated, leading to the development of mechanisms that could be construed as remote monitors. This section will briefly describe this evolution into current remote monitoring technology.

Computer system monitoring has become a primary component of system design, manufacture, and maintenance because of its application to performance evaluation. Although testing instruments (e.g. oscilloscopes) were frequently used to monitor the hardware at a very low level, system monitoring did not really begin to be needed until the mid 1960's. In the early part of that decade, computer systems began to reach a level of sophistication where resources were shared among a set of users. Once resource sharing was introduced, then resource utilization became an important metric for that system. If utilization was too high, then the resource represented a bottleneck to system progress; if utilization was too low, then the resource was either over-configured or, perhaps, was being prevented from being used effectively by bottlenecks elsewhere in the system. The result was frantic activity in the areas of hardware and software monitor development.

Software Monitors

Software monitors have not changed significantly in the last ten to fifteen years, although the designs have become more intricate. The predominante idea behind a software monitor is that the operating system of the host machine is modified so that it will collect and save measurement data about the performance of the machine. There are three major techniques for implementing a software monitor (discussed at length elsewhere).[11,36] The first technique is to use the system log as a storage medium for recording the occurrence of events. The system log may subsequently be analyzed to determine the activity of the machine at a level dictated by the post processor and the raw data recorded on the system log. This approach has proven to be a cost-effective mechanism for performance monitoring at a gross level.[35,46,55,64]

For more detailed studies, a more sophisticated data-gathering tool must be employed. The system log cannot normally be used to trace the program counter locus. In these cases, more extreme modifications must be made to the host operating system in order to invoke specially written monitoring and recording code. One technique to do this is the interrupt-intercept method. This technique assumes that the only times at which measurements should be taken are when the system changes states. In an interrupt-driven operating system, (e.g. IBM's OS) state changes are initiated by an interrupt (or a trap). At each such occurrence, the CPU is restarted on a handler, usually as a function of the type of interrupt. The interrupt-intercept monitor modifies this interrupt address so that pertinent interrupts are directed to monitoring code rather than to the handler. The monitoring code records the event and then branches to the handler. The raw data is subsequently recorded on a mass storage device and analyzed offline. (IBM has successfully used the technique in OS measurements,[22] as did GE in the GCOS operating system,[8,9] as well as others.)

While interrupt-intercept monitors can be used successfully to monitor events correlated with an interrupt or a trap, they cannot be successfully used in a machine which does not incorporate an interrupt-driven operating system (e.g. Control Data Cyber machines). An alternative is to use random time intervals for invoking the measurement routines; this technique has been used in many different cases, e.g. see references 24 and 54.

Software monitors have several assets, including their ability to be implemented without modifying existing hardware. Unfortunately, they do tend to add time and space artifact to the host system, producing a halo effect to the measurement experiment. The resolution of a software monitor is always limited by the instruction repertoire and cycle time of the host machine. A software monitor may also be corrupted or bypassed by processes that do not wish to be monitored.

Hardware Monitors

The development of hardware monitors lagged that of software monitors by

a few years, at least in the public domain. Drummond discusses several hardware monitors that were used internally at IBM on the 7094 systems,[11] while customers were still making crude attempts at monitoring in the late 1960's.[48] Independent vendors began constructing special-purpose hardware monitors in 1967-1968, propagating a large number of such firms. Currently, only two such firms still exist (i.e. have any significant portion of the market)--Comten and Tesdata.

Early hardware monitors consisted of three major components: a signal filter and combination unit, a time and count unit, and a data recording unit. The signal filter and combination unit was usually made up of a set of high impedance probes which could be attached to the logic boards of the host machine without disturbing the existing electronic signals on those logic boards. The signals sensed by the probes were passed to a logical combination unit, which could be used to detect the simultaneous occurrence of two or more probe signals, the exclusive occurrence of one signal from among a set of signals, etc. The particular logical combinations chosen were "programmed into" the hardware monitor by conventional plugboard logic. Once an event, or a logical combination of events, had been detected, the occurrence or duration of that event was temporarily recorded in the time and count unit. This unit was ordinarily implemented as a set of addressable registers, although more sophisticated monitors began to incorporate content addressable memories.[3,14,21] Inasmuch as these registers were accumulating data, they were subject to overflow; therefore, their contents were frequently stored on the data recording unit (usually a tape drive).

Hardware monitors of the generation described above had the ability to take measurements of a computer system at a much finer resolution than any software monitor; furthermore, they created no time and space artifact. The monitors were also not reachable from any software in the host machine; hence, they were protected from corruption and/or bypassing. However, it was difficult to draw a correspondence between observed monitor data and a particular task or procedure executing on the host system, i.e. causal relationships were lost. The amount of raw data collected also tended to be too large for reasonable offline analysis (and online storage). Thus, the strengths of pure hardware monitors were the weaknesses of software monitors, and vice versa.

Intelligent Monitors

The first significant step in the direction of improving the chasm between software and hardware monitoring techniques turns in the general direction of remote monitoring (as discussed in this report). The monitor might be distributed across special-purpose measurement hardware and software internal to the host operating system. In 1967, Estrin et al. published a paper describing the SNUPER COMPUTER monitoring system proposed for facilities at UCLA.[12] The SNUPER COMPUTER system was a complete system based around a SDS (XDS) Sigma 7 processor with 16K of 32-bit memory. The input to the processor was from a "sensory system" via normal data channels and high-speed I/O paths. The sensory system design would ultimately include a filtering processor to analyze raw monitor

data, generate event counts and measure the duration of events. Details of the design for SNUPER COMPUTER components were never published; and, apparently, the entire instrumentation system was never operational. Nevertheless, this paper was the first published report of a sophisticated monitoring system in which the instrumentation facility was centered around a programmable processing unit, allowing the hardware monitor to dynamically control a measurement session. Drummond indicates that IBM was using one computer to measure another in the early 1960's with the Direct Couple (7040-7090) system.[11 (p.277)]

The general state-of-the-art in the late 1960's began to see intelligent hardware monitors which employed a simplified version of the SNUPER COMPUTER approach. These monitors were still pure hardware monitors in which the static measurement experiment was determined a priori. However, a processor (usually a minicomputer) was employed to logically combine event signals, to manage event counters and timers in software, to provide a programmable filter for raw performance data, and to control the data recording function. Additionally, such hardware monitors typically provided online displays to describe the state of the host machine in real time. (Notice that this point of evolution essentially corresponds to the use of a remote monitor as a system console as mentioned in the Introduction.) During this period of development, most of the hardware components had been incorporated into monitors so that they could interact with host system software; but, for a few years, the monitors were not used in that manner. Once a programmable monitor is used to query the status of the host in order to dynamically control the monitoring function, that monitoring system has, in effect, become a remote monitoring system, The control element corresponds to the minicomputer-based monitor, and the remote portion of the monitor is that part of the host system which passes information to the external monitor. These applications of minicomputer-based hardware monitors are discussed in detail in the next section of this report.

REMOTE MONITORING TECHNIQUES

In the late 1960's and early 1970's, techniques for remote monitoring of computer systems were beginning to be explored, many of them based on existing hardware and software technology. The state-of-the-art at that time was discussed in the previous section of this report. In the current state of development, there are approximately seven classifications of remote monitors:

- Remotely controlled software monitors
- Internally distributed monitors
- Programmable monitors
- Hybrid monitors
- Computer network monitors
- Fault diagnosis monitors
- Intelligent and extended consoles

Remotely controlled software monitors are pure software monitors in terms of their techniques for obtaining measurement data; they differ from conventional software monitors in their ability to be controlled from an external source (e.g. an interactive terminal). Internally distributed monitors use a combination of software and hardware resources to take measurements, where all hardware resources are parts of a conventional (uninstrumented) computer system. That is, certain built-in hardware facilities are used for monitoring. Programmable hardware monitors are extensions of those discussed in the background section, and hybrid monitors are the logical conclusion of the programmable hardware monitors. The monitor components are distributed across remote internal (software) and remote external (programmable hardware) portions under the control of a local facility. Computer network monitors are incorporated into a conventional network (such as the ARPANET) in order to monitor communication and node performance. Fault diagnosis monitors are characterized as any devices used to check hardware machine state for circuit consistency, etc.; a degenerate example of a fault diagnosis monitor is a logic analyzer or oscilloscope. Intelligent and extended consoles are constructed from programmable devices which serve as a conventional operator's console under "normal" operation and as a pseudo hybrid monitor during monitoring sessions.

This characterization of remote monitors admits to imprecision in the sense that many remote monitors could be classified into two or more of the named divisions; and, in fact, several particular studies will be mentioned under more than one category. The remainder of this section is made up of more detailed discussions of the above-mentioned categories.

Remotely Controlled Software Monitors

A remotely controlled software monitor can be implemented as simply as a system log monitor. The purpose of the monitor is limited only by the ingenuity of the implementers, e.g. it may monitor the CPU utilization, inspect resource utilization, etc. In a pure software monitor, the monitor is instigated, for a particular run, by an operator, a system clock, or it is set for cyclic invocation at system initiation time. A remotely controlled software monitor is triggered by a central monitor controller. There are two prominent examples of actual implementations which employ this technique: the first example is the WWMCCS ADP System Security Officer (WASSO) station,[29,42] and the second example is a technique used on Control Data computers.[15,25]

The WASSO station potentially could be implemented as a true remotely controlled software monitor, or as a monitor distributed between a software tool and an intelligent terminal; the former case is discussed in this section.

Most sites in the WWMCCS are centered around a Honeywell 6000 series computer under the GCOS III operating system. Each site is potentially processing sensitive data in a multiprogramming environment; thus, there is a need for a secure operating policy in the system. Each such site

has at least one ADP System Security Officer whose mission is to monitor the physical protocol of the site, as well as the internal operation of the system. The WASSO station is under development to aid the officer in the latter part of his mission by providing a facility to implement collection, reduction, and analysis of information necessary to assess the site's security posture.

The current status of the WASSO study is that a prototype system is being built, and this prototype will partially rely on the standard software tools that exist in the operating system. The GCOS III operating system has been declared to be inadequate to support the desired security policies necessary for WWMCCS; thus, a new operating system will ultimately be designed and, presumably, that operating system will incorporate new software monitoring tools. The current software tools include facilities to designate one time sharing terminal as a "master terminal" to be used by the WASSO. From this master terminal, one can monitor system status information, line and terminal control tables, as well as inspect all interactions between the computer and an interactive user. The WASSO can also inspect memory utilization, mass storage utilization and file organization via the batch stream (i.e. offline analysis similar to system log analysis). A number of other standard measurement tools are available to the WASSO, all of which are implemented as conventional software monitors. The proposed WASSO station distributes some of the monitoring function onto a tailor-made intelligent terminal. This extension to the facilities of the WASSO clearly puts the mechanism into the class of hybrid monitors; however, it will be discussed in this section for the sake of continuity (and then cross-referenced in later sections).

The WASSO station incorporates a set of monitor data I/O buffers, a processor to analyze data and a full set of peripherals used to report and record monitor data. The station is expected to be implemented on a Honeywell Level 6 minicomputer system, including 48K words of memory, diskettes, magnetic tape, card reader, line printer, a cartridge disk unit, and the console. Functionally, the station will maintain system access controls, alter job priorities, explicitly interconnect the Honeywell 6000 node into the WWMCCS network and handle detected security breaches. The critical observation here is that, again, host software is used to actually <u>collect</u> the data used for determining the activity of the host computer system.

Control Data employs pure software techniques in a manner that is unique due to the architecture of their Cyber series systems. A detailed description of the technique perhaps best belongs in the section on internally distributed monitors, since it employs a peripheral processor to monitor the state of the remainder of the machine. The aspect of the work that makes it appropriate for this section is that software monitoring of a host system can be initiated and controlled from a central site, while the data collection and analysis are performed at the remote site by (a portion of) the host system itself. See the next section for the technical description.

Another example of remotely controlled software monitoring was implemented at MIT in the late 1960's.[17] The monitor probes, and parts of the filtering and analysis, were implemented in the host system software, while online displays, further analysis, and data recording were implemented on a central site minicomputer. A discussion of this study introduces the subsection on hybrid monitors.

Remotely controlled software monitoring can be a cost-effective method for implementing a remote monitoring facility in certain situations. If the environment of the host system is "friendly," then the likelihood of a successful implementation for performance evaluation or diagnostic testing is good. However, if the goal of monitoring is performance assurance or system security, pure software techniques may be unsuccessful due to the well-known complexity of software in the present day and age. Even in an apparantly friendly environment (i.e. there is no need for performance assurance nor security enforcement), it is almost certain that the software monitor can be bypassed and/or violated. (Although there is no proof that the above statement is true, this author knows of no completely secure software system.) In a potentially unfriendly environment, the probability of successfully implementing a completely safe remotely controlled software monitor is vanishing, no matter what the reason for monitoring.

Internally Distributed Monitors

An internally distributed monitor employs portions of the host's hardware in conjunction with (possibly) special-purpose software to monitor the operation of the host itself. If means are provided by which a user can remotely stimulate the internally distributed monitor, then it can be characterized as a remote monitor. (The distinction between remotely controlled software monitors and internally distributed monitors is not well-defined; similarly, hybrid monitors and some internally distributed monitors bear strong similarities.) Ordinarily, only machines with multiple processors can provide a hardware environment for implementing internally distributed monitors; however, note that the multiple processors need not be homogeneous. A number of different domestic companies incorporate multiple processors into their product line (e.g. Digital Equipment PDP 10 series computers, Control Data Cyber series computers, Univac 1100 series computers.) There have also been some machines modified by individual research organizations so that they can support internal distributed monitoring, e.g. see references 20 and 34.

The approach used for internally distributing a monitor is a simple one: One processor of the host system is temporarily reserved for use as a hardware monitor processor, while the remaining processors are employed by the host in a conventional manner. The detached processor simulates the action of the signal filter and combination unit, the time and count unit, and the analysis and recording unit of a conventional hardware monitor. The primary failing of the simulation is that hardware probes are ordinarily not used to gather data for the simulated hardware monitor. Instead, software executes in other parts of the host, as well as in the

detached processor, to perform the data gathering function. The approach can be understood in its entirety by considering an example employed on the Control Data 6000 architecture by individuals,[49,56] as well as by the vendor.[15,25,41,65]

The Control Data 6000 series machines include one or two central processing units (CP's) and ten peripheral processing units (PP's). A CP is used for 60-bit general purpose computation while each PP can be used as an independent 12-bit processor to perform I/O operations, implement the heart of the executive and control the system display. CP-PP communication takes place through the central memory of the machine; the machine state is also maintained in the central memory.

The orthodox mode of operation for the overall machine intends for user (application) programs to execute on a CP, where general purpose computations can be implemented under high level language program control. Whenever the operating system needs to intervene with the user processing, or whenever the application program needs to perform an I/O operation, then that activity is executed on a PP. Thus, the PP's are used for well-defined operations such that PP programs are written, assembled, and tested well before they are eligible for use by a CP program. PP programs extend the hardware, providing a virtual machine environment for user processes that execute on a CP. (PP programs can only be loaded from the system library of PP programs; thus, a user cannot redefine his own PP routines to perform I/O, etc.)

The architecture lends itself well to internally distributed monitoring. A special PP program to gather data, filter, combine, analyze, and record monitor observations can be written and added to the system library; thus, a PP simulates a hardware monitor with probes into the machine state tables in the central memory. For this architecture, the remainder of the host software and hardware need not be altered, since a PP loaded with the specially-written monitor program needs no other facility to monitor the host.

Internally distributed monitoring techniques can be extremely cost-effective in many circumstances. The cost in additional hardware is non-existent in these machines, since parts of a distributed system are used to implement the monitor itself. In some studies that are classified as internally distributed monitors, special purpose hardware has been permanently added to a production machine, see references 20 and 34. As a performance monitor, the approach is ideal except in the case of a system running under processor saturation (hence, a needed resource is removed from the system in order to observe that system). Performance data on program counter distributions can be easily obtained in the Control Data environment (but is much more difficult in the PDP 10 or Univac 1100 environment). Since system tables are stored in a central memory where all processors may inspect them, then machine state is easy to record and analyze; this allows the pseudo hardware monitor to identify causal relationships.

As a diagnostic testing device, the approach also has considerable merit. The pseudo hardware monitor can be used to run diagnostic programs and observe the results.

If the application of a monitor is either performance monitoring or diagnostic testing, internally distributed monitoring techniques lend themselves well to remote monitoring. Control Data currently employs this technique for "Remote Technical Assistance" for customer engineers at the site of the host.[25] The monitor is stimulated by interactive terminals through the conventional time sharing facilities of the host machine.

As a performance assurance or system security remote monitor, the application of internally distributed monitors is less attractive, since the overall operation of the device is ultimately controlled by software. Again, if one could ensure that complex software systems can be proven correct and that they cannot be bypassed or violated, then one might be less skeptical of these applications of internally distributed monitors. The current state of software engineering cannot guarantee that either of these conditions are satisfied.

Programmable Hardware Monitors

Programmable hardware monitors were introduced in the background section of this report; they are characterized as hardware monitors in which the operation of the monitor is programmable and, thus, dynamically reconfigurable as a function of the monitoring environment. In order for a programmable hardware monitor to be classified as a remote monitor, then either the programmable portion of the monitor must be located at the central site, or else the programmable monitor must be able to be controlled from the central site. This form of hardware monitor is, by far, the most popular type at the current time (e.g. see references 27, 60, and 61). The separation of hybrid monitors from programmable hardware monitors is also somewhat arbitrary.

Sperry Univac has been heavily involved in the use of programmable hardware monitors for at least ten years. In 1969, a paper was published describing an 1108 monitoring system in which a major component of the monitor was a second 1108 system;[38] currently, the Univac Eagan Benchmark Facility uses a comprehensive online monitoring system which heavily relies on a Univac 1616 minicomputer.[7,61] The earlier system was composed of three components: a hardware monitor to gather data from the host, data collection software used by the controller system to record data from the hardware monitor, and data reduction software to analyze the collected data. The initial goal for this system was modest, namely to provide a profile of the program counter contents; however, the tools that were developed are clearly of a much wider range of applicability than mentioned in the paper. The hardware monitor component (i.e. the remote monitor) was simple circuitry to detect unconditional branch instruction executions and to record the program counter contents. A critical observation is that the remote monitor ran at the same clock cycle as the host

CPU; and, in fact, was driven off of the same power supply and clock as the host system. (The remote monitor was mounted on a single card rack inside the host.)

Data collection was accomplished by a distinct 1108 system, with remote monitor inputs arriving via a channel, subsequently being recorded on drums private to the controlling system. Monitor data were subsequently moved onto magnetic tapes before analysis.

At first impression, the use of an 1108 as a central site controller seems to be a case of overkill. However, one should note that the combination of central and remote monitor is potentially operating at the cycle time of the host; and, hence, the 1108 controller may be required in order to record all data that were measured.

The Univac Eagan Benchmark Facility[7,61] monitor is composed of a data collection module, a BMD-1100 system, online displays, and an operator interface. A similar facility has been built for the Univac London Benchmark Facility.* The purpose of these facilities is to provide online monitoring displays of a wide variety so that potential customers can see the effect of their workload on the benchmark machine (i.e. an 1100/22 system).

The data collection module is made up of twenty 24-bit comparators, 196 count/time registers with probe plugboard logic to support up to 400 standard Comten high impedance probes. Thus, this portion of the monitor corresponds to a conventional hardware monitor's combination and filter unit. Monitor data can be transferred into the BMD-1100 system via an I/O channel about once every second.

The BMD-1100 (Telecon) component is capable of building real time color displays, "replaying" sketches of a benchmark run with different displays, interacting with the user, etc. It incorporates a Univac 1616 minicomputer (750 ns, 16-bit memory), a tape drive and a cartridge disk. The present uses of the 1616 minicomputer requires only about 10 percent of the CPU cycles to accomplish data collection and display.

The online displays at the Eagan facility include a large color CRT display with a wide variety of display options, a digital display of certain critical elements of the host system's status, and a U-100 CRT terminal to provide an operator's interface.

The current implementations at Eagan and London use the facility as a local hardware monitor; however, either facility is well-suited as a remote monitoring system with no hardware modification. Inasmuch as the

---

* The BMD-1100 system is replaced by a Univac Telecon system. The Telecon system is a standard front end processor for telecommunication applications, where the principal processing is carried out by a Univac 1616 minicomputer; it is similar to the special-purpose BMD-1100 system.

processing element is actually an interactive front end processor, it is obvious that the instrumentation hardware and Telecon front end can be employed as a remote monitor, while the displays and operator's interface are implemented at the central site. Notice that the Telecon processor already includes appropriate hardware and software for interacting between the operator's interface and the remote processor. Although such an installation relies on a partial software solution (at the remote site), a secure implementation might employ ROM for storing Telecon programs, with RAM used as buffer space. The data collection and forwarding codes ought to be simple enough to be made convincingly correct.

Texas Instruments employed an in-house programmable hardware monitor system during the development of their ASC system.[43,44,45,62,63] The reason for building this monitor was performance evaluation. The Texas Instruments System Activity Monitor (TISAM) design was motivated in part by the need for a flexible hardware monitor which could provide a good monitor-analyst interface. The main component of the system is a TI 960A minicomputer system with a card reader, magnetic tape, and interactive (hard copy) terminal. Although the system is capable of performing a wide variety of measurement tasks, it has been used primarily for program counter tracing.

There are a number of other programmable hardware monitors reported on in the open literature, e.g. see reference 3. However, the above discussion adequately discusses the approach that one might use in applying the technique to remote monitoring.

Programmable hardware monitors appear to offer significant promise as remote monitors for most application areas considered in this report. As remote performance monitors, the approach has already been successfully used for as long as ten years. The intelligent console applications (discussed later in the report) are a variant of programmable hardware monitors, and also have been successfully employed as remote diagnostic testing tools. As performance assurance monitors, the picture is not quite so clear; this area has not yet been successfully accomplished (to this author's knowledge). However, for performance assurance applications, the use of the programmable hardware monitor at either the local or remote site seems to have a better chance at remaining secure than either of the previous two approaches. Since the remote portion of the monitor can be simple, its programs can be simple, and physically protected by storing the code in ROM and/or physically sealing the monitor. In this latter case, RAM memory might be loaded only with some special apparatus, which itself may be geographically distinct from the remote monitor. As a system security monitor, it is apparent that this general approach is being used in the WASSO terminal described in a previous subsection.

Hybrid Monitors

The distinction between hybrid monitors and programmable hardware monitors is principally one of application; the components for many programmable

hardware monitors are sufficient to perform hybrid monitoring. A hybrid monitor is composed of a software portion that executes on the host hardware, and a hardware portion that executes a (usually programmable) hardware monitor. The two portions are cognizant of one another and exchange signals and data. Thus, a hybrid monitor reacts to the state of the host system by reconfiguring itself dynamically. It is easy to argue that the proposed WASSO terminal and the BMD-1100 systems, as well as others, are really hybrid monitors.

One early application of the hybrid monitor is a remote tool was given by Grochow in his graduate work at MIT.[17] The Graphic Display Monitoring System (GDM) was designed to observe activity in the GE 645- Multics system developed at Project MAC. The monitoring functions were distributed across a remotely controlled software monitor and a central site programmable facility. The host software portion of the monitor executed as a multics procedure, gathering data from static operating system tables. This procedure was also capable of preliminary filtering of the data and transmission of the partially filtered data to the central site programmable device. The software portion was driven by commands from the central site machine over one data path, and a separate data path was used for monitor data transmission. The host machine central site facility interface was implemented via two data channels on the GE 645 to a 2400 baud modem, through a telecommunications link to a similar modem at the central site. The central site mechanism was composed of a DEC PDP 8 system including a disk, magnetic tape and a display processor. The central site portion of the system used much of its computing capability to format real time displays, such as the usage of Multics core memory pages, the multi-programming state of each process in the system, etc. These analysis and display programs existed in the basic library of functions of GDM. However, the GDM was also user programmable so that monitor functions and displays could be generated whenever new, or unique, requirements were encountered.

The GDM was found to be useful in its flexibility and extensibility. The software portion of the monitor was constructed as a sampling monitor, with the sampling rate determined by requests from the PDP 8 (thus, the sample rate was taken as a function of the display being used). The 2400 baud transmission rate for host-PDP 8 interconnection allowed for a maximum of about 20 samples per second, although many displays were changed only once a second.

In 1971, Aschenbrenner, Amiot, and Natarajan published a paper briefly describing their Neurotron monitor system (implemented at Argonne National Laboratory).[4] The monitor incorporates a minicomputer to control the combination and filer unit, as well as data filtering and recording. The external monitor and the host software monitor interact via conventional I/O ports, as well as through measurement probes. The development goals of Neurotron offer several interesting facets to the use of the system for remote monitoring. Among other goals, the monitor was designed to:

- contain program logic for selecting or filtering events as a function of the current experiment

- allow micro and macro level monitoring of the host

- trigger monitoring activity as a function of host event sequences or at predefined time intervals.

It is not clear exactly how successful the Neurotron monitor was after it had been used for an extended period of time. The Neurotron was used to gather statistics on instruction distribution and correlate them with I/O activity or other events, to investigate memory reference streams, to analyze buffer usage, etc. Perhaps the most interesting claim made about the Neurotron is the portability of the hardware to different systems, each containing a host software monitor. Apparently, a single hardware monitor can be constructed that will successfully operate with a variety of host systems.

In the late 1960's and early 1970's Xerox Data Systems was deeply involved in measurement projects for operating systems and the Sigma series hardware. Hughes and Cronshaw describe another hybrid monitor that was used within XDS.[21] The ADAM hybrid monitor differs little from the Neurotron monitor in its organization; the facility incorporated a minicomputer in the hardware monitor. Interactions between the host and the hardware monitor took place over an I/O channel, while the hardware monitor could also collect data from probes installed on the host. The most significant differences between the ADAM and the Neurotron were in the event recognition algorithms and their versatility. ADAM used an associative memory to quickly recognize event combinations; secondly, the Xerox system was intended only to monitor the Sigma 7, while the Neurotron was aimed at a variety of different host systems.

The technology of hybrid monitors corresponds closely to that of programmable hardware monitors. The strong points of each are similar; e.g. careful design of the remote portion of a hybrid or programmable hardware monitor is likely to be more "trustworthy" than for pure software techniques that execute solely on the host hardware.

Network Monitors

Network monitors are defined to be any monitoring device used in the context of a distributed computer network. The primary motivation for monitoring in this environment has been for performance evaluation. By the nature of computer networks, nearly all measurement studies must employ a remote monitor.

Perhaps the most prominent computer network in the world is the ARPANET which has been operational since about 1971. The ARPANET contains in excess of 50 significant (i.e. medium to large scale) computer nodes, most intercommunicating at a rate of 50 KBPS. There are two network centers devoted to measurements: the Network Control Center at Bolt, Beranek, and Newman, and the Network Measurement Center at UCLA. A few papers have appeared in the open literature that describe measurement activity on the ARPANET. Kleinrock and Naylor provide a report on investigations dealing with message traffic in the network.[23] Each node in the ARPANET

contains an interface processor called an IMP. The IMP is a fully programmable 16-bit minicomputer whose primary function is to interface the host processor to the ARPANET in a standard manner. The IMP must also participate in message packet switching. Each IMP contains software to aid in network measurements; for example the IMP can collect data about its host and forward the monitor data to the Network Measurement Center. The IMP monitor software also may "time stamp" packets which they forward to other IMP's, thus allowing an analysis program to investigate message packet flow rates and routes. The standard ARPANET remote monitor is really a software monitor implemented on an IMP, but under the control of the Network Control Center. Another view of the facility is that it is functionally equivalent to the Univac Telecon monitoring facility discussed earlier; i.e. the IMP corresponds to a Univac Telecon front end processor (Univac 1616 minicomputer).

Another monitor has also been developed for investigating an alternate packet switching mechanism ("packet radio systems").[51] In this discussion a packet radio repeater, rather than an IMP, is used to broadcast packets with a radio transceiver. The packet radio repeater also includes a microprocessor to aid in the broadcast and to act as a monitoring facility. In order to avoid the halo effect of broadcasting the measurement data over the normal data channel, monitor data is saved at the station until the measurement experiment is completed. It may then be transmitted to a central facility for further processing. Another interesting aspect of the work is that the packet radio repeater microprocessor does not maintain a local library of monitoring routines; instead, whenever a measurement test is to be initiated, the monitoring code is transmitted to the packet radio repeater. In the case of performance assurance and system security monitors, this idea has obvious interesting possibilities.

A final note on ARPANET remote monitoring is concerned with measurements of the London node of the network.[50] A PDP 9 is used to collect measurement data on the IBM 360/195 and then analyzed as batch data after the data collection phase.

Considerable thought has gone into the topic of remote monitors for networks at the University of Waterloo.[31] At the time that the paper was published, the group had built a prototype Computer Network Monitoring System (CNMS) designed to:

- observe network performance
- detect malfunctions in the network
- diagnose failures

A basic premise of the measurement group was that software monitors, alone, often produced too much artifact (cf. ARPANET studies), and that pure hardware monitors were not sufficiently flexible to make the required network measurements. Therefore, a hybrid monitor design would be used. The general approach dictated that each host computer should include a remote monitor to be software controlled from the network monitoring center.

The important components of the CNMS are a remote controlled hybrid monitor (RCHM) and the central site network measurement center (NMC as in the ARPANET). RCHM's are potentially allowed to communicate with the NMC either through the normal network data paths or through dedicated links. The philosophy also admits to a hierarchy of measurement centers; e.g. each RCHM could be controlled by a regional NMC (RNMC) which is, in turn, controlled by the NMC.

The prototype RCHM is based on a PDP 11 (probably an LSI-11) processor with extra hardware to implement a combination and filter unit and an event recognition unit. RCHM is attached to the host computer with probes and through the conventional I/O communication lines of the host. The PDP 11 system includes a small disk for buffering raw data, which can be forwarded to a (R)NMC in a condensed form.

The software portion of an RCHM is distributed across the PDP 11 CPU and the CPU of the host computer. Host resident monitor code is system-dependent with system-independent interfaces (between the RCHM-resident software and the host-resident software). The RCHM-resident software is organized as a set of six classes of processes, functions of the classes being:

- experiment manager to schedule and support RCHM experiment control programs
- monitor manager to control the special-purpose monitoring hardware (probes, filters, etc.)
- resource manager to allocate RCHM system resources
- communications manager to receive input from, and forward monitor data to, its (R)NMC
- results manager procedures to record, reduce, and analyze raw monitor data, producing condensed reports
- maintenance manager to provide diagnostic testing codes for the network components (i.e. the host, the RCHM, etc.)

The (R)NMC system is not a special-purpose computer system (although it must have a hardware facility to communicate with the set of RCHM's that it controls). The (R)NMC software contains seven classes of processes that roughly complement the six classes of RCHM processes; the seventh class establishes a user interface with the analyst who controls the experiment.

An instance of a measurement experiment in CNMS might proceed as follows: The purpose of the experiment is carefully analyzed, and a set of measurements that will provide the necessary results is determined. The RCHM's are then configured to take the required measurements by attaching hardware probes, designing and implementing host software probes, and designing and implementing RCHM-resident data collection and reduction software. Since each RCHM hardware system is identical, one set of software is written for all remote monitors. The codes should be able to collect bursts of measurement data from the hardware and software probes, record the data, condense the data (usually to a histogram) and then transmit the condensed monitor data to the controlling (R)NMC.

Once the monitors have been defined, the experiment can be initiated from the NMC. Online control can be exercised by the analyst via the user interface manager that executes on the NMC. The interaction is possible because of real-time analysis and display at the NMC. After an experiment is initiated in the RCHM's, there is no requirement that the analyst continue to monitor the experiment from the NMC; the communication between RCHM's and the (R)NMC need not be continuous. Upon conditions determined by the NMC or the analyst, the experiment can be terminated.

Xerox Palo Alto Research Center has also developed monitoring tools for investigating their internal network.[28] The network connects a number of minicomputers and at least one medium-scale computer system. Each of the minicomputer systems can be loaded with a remotely controlled software monitor. The monitoring facility provides standard software for recording data and for transmitting it to other nodes on the network. The work is similar to the facilities provided in the ARPANET, and not nearly as well-developed as those described in the CNMS.

Tesdata Systems Corporation has two particular monitor systems that can be classified as network monitors. The MSB facility is a product, aimed at small-to-medium sized data centers, for sharing more sophisticated and expensive monitoring equipment with other centers of similar size.[18] The approach taken in the MSB is to install a programmable hardware monitor at each remote site. The hardware monitor is based on a 16-bit microprocessor with RAM for buffer storage and a fixed set of (not more than 64) standard hardware probes. At intervals of time determined by the measurement sampling rate, a central site controller initiates communication with the remote monitor to dump the remote buffers. Data collected by the central controller may or may not have been filtered by the remote microprocessor. The remote monitor is designed with ROM and RAM memory. The ROM contains several built-in functions for the microprocessor, while the RAM is used for buffering and dynamic portions of the monitor code. The RAM portion of the microprocessor is down loaded from the central site controlling processor. Central site computing, performed at a regional control center, produces summary reports on a conventional computer system.

The distributed measurement network is a production facility of Tesdata.[60] The network consists of several secondary systems made up of conventional Tesdata monitors, e.g. MS 38 or MS 58 monitor systems, and a single primary system implemented as an MS 88 monitor system. The secondary systems are connected to the primary system via standard voice grade phone lines, and are controlled by the MS 88 system. The differences between the distributed measurement network and the MSB are quite significant. The secondary systems can operate as autonomous monitors or be a node in a monitoring network. Each secondary system contains complete data collection and archiving equipment, analysis programs, and reporting mechanisms. An MSB remote monitor is not capable of operating totally independently. The distributed measurement network is aimed at collections of large data centers.

National Bureau of Standards has its own Network Measurement Machine (NMM)

and Network Measurement System (NMS). The PDP 11 based NMM is a hybrid remote monitor controlled by central site facilities of the NMS. Consult references 1, 2, and 39 for further information on this work.

Finally, Lombardi has proposed a network monitoring facility similar to the MSB system described above, although implementation had not been completed in September, 1977.[26]

Network monitors are the archetype for remote monitors; the nature of networks implies that remotely located computer systems and network traffic can best be monitored by the combination of a remote monitor and a central control element. Remote monitoring of a single computer system is merely a degenerate case of network monitoring. As remote monitors for performance evaluation, network monitors are already in production, e.g. Tesdata facilities. The CNMS system additionally uses network monitoring techniques for providing a diagnostic testing facility. Applications in the areas of performance assurance and system security are not tested (even for local systems, performance assurance and system security have not been successfully implemented). Of all the remote monitoring classifications discussed up to this point, it is easiest to argue for the probability of a successful implementation using these techniques. In performance assurance and system security applications, critical questions of the monitoring activity are the independence of the remote monitor from the host and the validity of the remote monitor. The possibility of building logically compact hybrid remote monitors which can be proven correct offers some hope for the latter question; the use of ROM for storing remote software also aids in ensuring validity of the monitor. Distributing the remote monitor across hardware and software aids in creating independence between the host and the remote monitor. These factors are an obvious point for further research.

Fault Diagnosis Monitors

Much work has been done in the area of fault-tolerant computing, especially in the area of circuit verification. A fault-tolerant computer system is one in which a single error will not cause the system to fail; the overall system can detect and correct the error. Some obvious examples of fault detection monitoring occur in fault-tolerant computers such as in-flight or on-board computer systems.

Unfortunatly, most work in the area of fault diagnosis is at a rather detailed level. Hardware faults usually occur at the gate level; hence, monitoring equipment must gather data at that level. Thus, fault diagnosis monitors tend to gather such detailed information that the data is only useful for circuit analysis or redundant processing (e.g. see reference 10) and is not generally useful for system level monitoring.

One special-purpose diagnostic facility that has been employed in industry is the Control Data STAR Maintenance Station.[37] Although the maintenance station appears to be at least as applicable of a diagnosis tool as the PP in the 6000 series machines, Control Data personel indicate that it has not been a useful addition to the STAR facilities.[15]

Related work takes place in other areas of hardware-diagnosis using extended system consoles, and that work will be discussed in the next subsection.

Intelligent and Extended Consoles

The final category of remote monitors uses the notion of the intelligent terminal for an operator's console. It is frequently useful to provide more capability in the operator's console than exists in a passive terminal; e.g., the console can be used to drive CRT displays in parallel with the operation of all other facilities in the host system.

The oldest application of the technique must be that of the Control Data 6600 console.[57] One of the ten peripheral processors is dedicated to the task of driving the dual CRT system console. The PP reads input from the operator's keyboard and from the machine's central memory, and processes the data to produce CRT displays describing the state of the operating system.

The Honeywell Remote Maintenance System/62 (RMS-62) is a small system tool to aid field engineers in monitoring, controlling, diagnosing, and patching Level 62 installations.[16,58] The remote portion of the device is a Remote Console Interface Adaptor and a software diagnostic package. The adaptor is a single board component which provides a 300 baud modem to send/receive information to/from the host. The modem is attached to the bus connecting the host system and its console. When the host is experiencing problems, the adaptor is connected to a central site system console via voice grade phone lines. Diagnostic software is then exercised from the central site.

Digital Equipment Corporation has announced a facility, similar to the RMS-62, for use with their PDP11/70.[47] Few details of the work are available since the project is still under proprietory development. However, the adaptor contains, not only a modem for communicating with the central controller, but also a microprocessor with a small amount of RAM. The unit is to be used only for diagnostic testing. One difference between the RMS-62 and the DEC system is that the DEC facility allows communication with either the host's local operator console or the central site; in the latter case, the microprocessor is required to do some console processing at the site of the host. Digital Equipment personnel speculate that the facility is too slow to be used as a performance monitor and will only be used as a diagnostic facility.

The Cray system also employs an intelligent operator's console in its standard configuration.[32] The console is driven by a Data General Eclipse minicomputer (a 400 nanosecond, 16-bit minicomputer). The console is also used as an initial program loading device and for diagnostic testing. Whereas the DEC system extended console saturates its processing facility, the Cray system tends to under utilize its console processor.

The Amdahl 470 computer system also uses a Data General minicomputer as a processor to drive its console.[5] The console hardware includes a floppy disk unit and a modem for telecommunication, as well as facilities to support the operator's console.

The circuit components that are used to build the 470 include the capability for an external medium to read out the value of each latch in the hardware. The first extended capability of the console is to be able to inspect and record each of these latches. These console facilities are used whenever the hardware fails, i.e. the console scans all latches in the machine and records them on a floppy disk. The floppy disk is then mailed to the Sunnyvale site where it can be analyzed by central site experts.

The second application of the extended console uses the (1200 baud) modem. If a machine is experiencing a problem that the field engineer cannot diagnose, he initiates a phone connection of the console/host machine to the Sunnyvale (central) facility. At the central site, another console with a minicomputer is used to remotely control the host computer system. Thus, there is no real-time monitoring at 1200 baud, but incremental stepping of the host can be done from the central site. (Diagnostic software is executed on the remote console processor under the control of the similar central console processor.)

The third application of the console extends the one just described. After a connection is established between the host console and the central console, information routed to the central console is processed locally by a second 470 system. Full trace data can be analyzed by the central facility in order to isolate circuit failures.

The final example of an extended console is the Total Remote Assistance Center (TRACE) facility used at Univac.[30,59] The central facility is centered around a Univac 1108 in Roseville, Minnesota, which is dedicated to diagnostic assistance for field engineers. The central facility can be used to remotely monitor series 90 systems and 1100 series systems. (The remote monitor for each machine, in each series, is unique to that machine model.)

The series 90 remote monitors are portable facilities used only during diagnostic testing. The 1110 and 1100/40 are provided with a Maintenance Controller to serve as the remote (programmable hardware) monitor. The Maintenance Controller provides a mechanism for extending the console to the TRACE facility and to run diagnostic tests. The 1100/80 remote monitor contains slightly more capability than the Maintenance Controller, allowing more of the TRACE processing to be performed at the site of the host.

The TRACE facility bears many similarities to the network monitors discussed in the previous subsection. Although Control Data did not feel that dedicated remote monitors were useful in the context of STAR, Univac has made extensive use of the same technique (so much so that future

Univac systems will also include similar facilities[52]). The most remarkable observation about the TRACE facility is its versatility across a wide variety of machines; the remote monitors have been designed for individual models so that uniform central site processing can be applied to the monitor data.

The remarks at the end of the subsection on network monitors are appropriate for most extended console applications discussed here. The extended consoles are often equivalent to the two-node network (consisting of the host system and the central system). As performance evaluation monitors, the extended console may have inadequate local processing power or be limited in its data transmission bandwidth (e.g. the Honeywell and Digital Equipment facilities). Almost all extended consoles have been specifically designed to perform remote diagnostic capabilities; thus, they are nearly all successful at this aspect of remote monitoring. As performance assurance and system security remote monitors, the question of appropriateness is still open (as mentioned before).

SUMMARY AND CONCLUSIONS

This report has discussed a wide variety of monitors in the context of remote monitoring. The areas of application have been broadly categorized as performance evaluation, diagnostic testing, performance assurance, and system security monitoring. Monitors have been classified into seven groups, based primarily on their architecture. Several individual monitors could easily have been placed in more than one category (the categorization really served to compare similar monitors at one time in a single subsection).

There is no single best monitoring approach for any given application. The choice of a technique is a function, not only of the broad application area, but also of the particular environment in which the monitor is to operate and the type of data that the monitor will be required to collect. Nevertheless, the following paragraphs address the issue in very general terms.

Remote monitoring for performance evaluation is the most popular application area. The monitoring technology has grown complex during the ten or fifteen years that local performance monitors have been used, propagating a large set of principles that could be employed in remote monitoring work. Activity has been the highest in the study of computer networks, although that technology is clearly derived from software, programmable hardware, and hybrid monitoring work. It is also clear that current remote monitor architecture studies have far outstripped the engineering application of the attendant tools. Existing hybrid and network monitor facilities appear to be capable of gathering measurement data in manners that have not really been effectively used. The tool maker has built a tool that is so complex that workers do not know how to use it effectively. A few of the toolmakers seem to be aware of this problem and are attempting to remedy the situation by providing better human-engineered monitors. The

raw power of the existing monitors, as performance monitors, is not likely to be a limiting factor for some time to come.

The area of diagnostic testing using remote monitors is especially well balanced. Manufacturers have extended their console capability, or provided other special facilities, to accomplish precise remote diagnostic monitors. The Amdahl and Univac extended console facilities are not only well-designed (at least at a high level), the features of the monitors are also well utilized. Although other categories of monitors can be used to implement diagnostic testing, the extended console appears to be the best approach to this application area. Conversely, extended consoles can be extended to serve purposes other than diagnostic testing, although remote monitor (host console) processor saturation may be a problem. If the host console processor is embellished to include facilities for filtering and data recording, the extra memory may be left unused for large periods of time, i.e. the hardware investment may not be cost-effective.

Performance assurance monitoring is a difficult application area in which to work. A fundamental axiom of all monitoring studies is to know what data is needed before designing a monitor to collect the data. The question of whether or not performance assurance is even possible is currently unanswered. If, indeed, it is a solvable problem, then one can proceed with an appropriate design. For the sake of discussion, it is assumed that one may be able to find performance assurance metrics, and that a monitor is needed to obtain the corresponding measures. By the nature of the requirement for performance assurance, one cannot assume that the environment of the host computer is totally friendly (for if it were, performance assurance would be unnecessary). The degree of unfriendliness can vary from situations involving, say, mischievous university students, to the clientele of an Eastern European computing center. Hence, implementation issues must be concerned with the correctness, the logical completeness, and the security of the monitor (especially its remote portion).

Correctness of any design can best be ensured by keeping the monitor design as simple as possible (a lesson hard-learned by current operating system designers). When the design is simple, or at least highly structured, then its logical consistency can be checked or proven.

Logical completeness has to do with the absence of "loopholes" in the design and/or bypassing the monitor again. One would like to have proof of logical completeness, but this is most likely impossible. (For example, the IRS has been unable to establish a set of tax laws which effectively plug all tax loopholes.) The best approach toward achieving logical completeness is to, again, keep the monitor design as simple and well-structured as possible; only then can a single human being begin to <u>understand</u> the implications of the design in terms of completeness.

Issues of security are at least as difficult to address as those of logical completeness (by a <u>secure</u> monitor, it is meant that the monitor cannot be violated to change its ability to collect information). Secure remote monitors will almost necessarily be at least partially implemented

in hardware, since software is extremely susceptible to security loopholes. Hybrid monitoring techniques show promise in terms of maintaining secure operation. To minimize the chance of programmable hardware being tampered with, the hardware monitor code should be stored in ROM and/or down loaded from the central facility. Physical security might be enhanced by embedding the monitor into the other facilities of the host system rather than "locking it up" in a box.

System security is the main issue of the WASSO terminal prototype implementation. As in performance assurance, questions of monitor consistency, completeness, and security are critical. Basically, the same observations made about performance assurance monitors also hold for system security monitors.

The overall future of remote monitoring is filled with several unanswered questions. How can remote performance monitors be used to their maximum effectiveness? Is performance assurance decidable in a practical situation involving heuristics? If the question is decidable, what are techniques for ensuring some reasonable level of consistency, completenesss, and security? Hardware solutions will probably tend to surface, since such solutions are becoming less expensive and they can be made to be less volatile than pure software solutions.

ACKNOWLEDGEMENT

The material in this paper has been obtained from the open literature and through interviews with many especially cooperative personnel at various organizations. I would like to mention the following individuals for exceptional cooperation and help: Paul Roth and George Lindamood of NBS, Larry Sheets and coworkers of Univac, Jeff Booth of Amdahl, and Lynn Gallup of Control Data.

References

1. Abrams, M.D., I.W. Cotton, S.W. Watkins, R. Rosenthal, and D.E. Rippy, "The Network Measurement System", unpublished.

2. Abrams, M.D. and S. Treu, "A Methodology for Interactive Computer Service Measurement", Communications of the ACM, Vol 20, No. 12, pages 936-944, December, 1977.

3. Agajanian, A.H., "A Bibliography on System Performance Evaluation", IEEE Computer, Vol. 8, No. 11, pages 63-74, November, 1975.

4. Aschenbrenner, R.A., L. Amiot, and N.K. Natarajan, "The Neurotron Monitor System", AFIPS Proceedings of the FJCC, Vol. 39, pages 31-37, 1971.

5. Booth, J. and C. Neroth, Amdahl Corporation, personal communication, November 22 and November 28, 1977.

6. Browne, P.S., "Computer Security--A Survey", AFIPS Proceedings of the NCC, Vol. 45, pages 53-63, June, 1976.

7. Buchanan, A.T. and L.E. Sheets, Sperry Univac Corporation, personal communication, January 4, 1978.

8. Campbell, D.J. and W.J. Heffner, "Measurement and Analysis of Large Operating Systems during System Development", AFIPS Proceedings of the FJCC, Vol. 33, pages 903-914, 1968.

9. Cantrell, H.N. and A.L. Ellison, "Multiprogramming System Performance Measurement and Analysis, AFIPS Proceedings of the SJCC, Vol. 32, pages 213-221, 1968.

10. Cook, R.W., W.H. Sisson, T.F. Storey, and W.N. Toy, "Design of a Self-Checking Microprogram Control", IEEE Transactions on Computers, Vol. C-22, No. 3, pages 255-262, March, 1973.

11. Drummond, M.E., Jr., Evaluation and Measurement Techniques for Digital Computer Systems, Prentice-Hall, Inc., Englewood Cliffs, N.J., 1973.

12. Estrin, G., D. Hopkins, B. Coggan, and S.D. Crocker, "SNUPER COMPUTER--A Computer in Instrumentation Automation", AFIPS Proceedings of the SJCC, Vol. 30, pages 645-656, 1967.

13. Ferrari, D., "Architecture and Instrumentation in a Modular Interactive System", IEEE Computer, Vol. 6, No. 11, pages 25-29, November, 1973.

14. Fryer, R.E., "The Memory Bus Monitor--A New Device for Developing Real-Time Systems", AFIPS Proceedings of the NCC, Vol. 43, pages 75-79, 1972.

15. Gallup, L., Control Data Corporation, personal communication, January 6, 1978 and January 27, 1978.

16. Gray, W.H., Honeywell Information Systems, personal communication, December 14, 1977.

17. Grochow, J.M., The Graphic Display as an Aid in the Monitoring of a Time-Shared Computer System, Project MAC Technical Report MAC-TR-54 (THESIS), September, 1968.

18. Hart, L., Tesdata Systems Corporation, personal communication, December 20, 1977.

19. Hopcroft, J.E. and J.D. Ullman, Formal Languages and their Relation to Automata, Addison-Wesley, Reading, Massachusetts, 1969.

20. Hornbuckle, G.D., "A Multiprogramming Monitor for Small Machines", Communications of the ACM, Vol. 10, No. 5, pages 273-278, May, 1967.

21. Hughes, J. and D. Cronshaw, "On Using a Hardware Monitor as an Intelligent Peripheral", ACM Sigmetrics Performance Evaluation Review, Vol. 2, No. 4, pages 3-19, December, 1973.

22. Keefe, D.D., "Hierarchical Control Programs for System Evaluation", IBM Systems Journal, Vol. 7, No. 2, pages 123-133, 1968.

23. Kleinrock, L. and W.E. Naylor, "On Measured Behavior of the ARPA Network", AFIPS Proceedings of the NCC, Vol. 43, pages 767-780, 1974.

24. Kolence, K.W., "A Software View of Measurement Tools", Datamation, Vol. 17, No. 1, pages 32-38, January, 1971.

25. Kroesch, J. and J. Merlo, Control Data Corporation, personal communication, January 6, 1978.

26. Lombardi, J.C., "A Programmable Hardware Monitoring and Simulation (PROMOS) System for Computer-Based Process-Control Systems", Proceedings of the IEEE Fall CompCon 77, pages 423-428, September, 1977.

27. Marter, T.M., "Capturing Terminal Traffic Using a Hardware Monitor", Proceedings of Computer Performance Users Group, 13th Meeting, NBS Special Publication 500-18, pages 95-105, September, 1977.

28. McDaniels, G., "METRIC: A Kernel Instrumentation System for Distributed Environments", Proceedings of the Sixth Symposium on Operating Systems Principles, pages 93-99, November, 1977.

29. McDonnell, A.P., "Prototype WASSO Station Functionality (Revised)", DCA Report No. TM-WD-7825/000/02, September, 1977.

30. McManus, W., N. Adams, and D. Rice, Sperry Univac Corporation, personal communication, January 5, 1978.

31. Morgan, D.E., W. Banks, D.P. Goodspeed, and R. Kolanko, "A Computer Network Monitoring System", IEEE Transactions on Software Engineering, Vol. SE-1, No. 3, pages 299-311, September, 1975.

32. Morris, R., Cray Corporation, personal communication, December 13, 1977.

33. Murphy, R.W., "The System Logic and Usage Recorder", AFIPS Proceedings of the FJCC, Vol. 35, pages 219-229, 1969.

34. Nemeth, A.G. and P.D. Rovner, "User Program Measurement in a Time-Shared Environment", Communications of the ACM, Vol. 14, No. 10, pages 661-666, October, 1971.

35. Noe, J.D. and G.J. Nutt, "Validation of a Trace-driven CDC 6400 Simulation", AFIPS Proceedings of the SJCC, Vol. 40, pages 749-757, 1972.

36. Nutt. G.J., "Tutorial: Computer System Monitors", IEEE Computer, Vol. 8, No. 11, pages 51-61, November, 1975.

37. Purcell, C.J., "The Control Data STAR-100--Performance Measurements", AFIPS Proceedings of the NCC, Vol. 43, pages 385-387, 1974.

38. Roek, D.J. and W.C. Emerson, "A Hardware Instrumentation Approach to Evaluation of Large Scale Systems", Proceedings of the ACM National Conference, pages 351-367, 1969.

39. Rosenthal, R., D.E. Rippy, and H.M. Wood, "The Network Measurement Machine--A Data Collection Device for Measuring the Performance and Utilization of Computer Networks", NBS Technical Note 912, April, 1976.

40. Saltzer, J.H., "Protection and the Control of Information Sharing in Multics", Communications of the ACM, Vol. 17, No. 7, pages 388-402, July, 1974.

41. Sebastian, P.R., "HEMI (Hybrid Events Monitoring Instrument)", ACM Sigmetrics Symposium 74, pages 127-139, December, 1974.

42. Shabe, J., Defense Communications Agency, Command and Control Tactical Center, personal communication, November 29, 1977.

43. Shaw, R.W., TISAM Hardware Users Manual (Preliminary), Texas Instruments unpublished report, June, 1974.

44. Shaw, R.W., "An Advanced Hardware Monitor", Proceedings of the Third Texas Conference on Computer Systems, pages 7-2-1 to 7-2-5, November, 1975.

45. Shaw, R.W., Texas Instruments, personal communication, December 2, 1977

46. Sheets, L.E., Sperry Univac Corporation, personal communication, January 4-5, 1978.

47. Skelton, S., Digital Equipment Corporation, personal communication, December 14, 1977.

48. Stang, H. and P. Southgate, "Performance Evaluation of Third Generation Computing Systems", Datamation, Vol. 15, pages 181-190, November, 1969.

49. Stevens, D.F., "System Evaluation of the Control Data 6600", Proceedings of the IFIP Congress, pages C34-C38, 1968.

50. Stokes, A.V., D.L. Bates, and P.T. Kirstein, "Monitoring and Access Control of the London Node of ARPANET", AFIPS Proceedings of the NCC, Vol. 45, pages 597-603, 1976.

51. Tobagi, F.A., S.E. Lieberson, and L. Kleinrock, "On Measurement Facilities in Packet Radio Systems", AFIPS Proceedings of the NCC, Vol. 45, pages 589-596, 1976.

52. Torgerson, J.F., Sperry Univac Corporation, personal communication, January 5, 1978.

53. Turn, R., "Remarks on the Instrumentation of Databank Systems for Data Security", Rand Publication No. P-5151, January, 1974.

54. Waite, W.M., "A Sampling Monitor for Applications Programs", Software--Practice and Experience, Vol. 3, pages 75-79, 1973.

55. Watson, R.A., "Computer Performance Analysis: Applications of Accounting Data", Rand Report No. R-573-NASA-PR, May, 1971.

56. Wisneski, D. and R. Paddock, "SPYUII Description", University of Colorado Computing Center Publication, November, 1976.

57. ---, Control Data 6000 Series Computer Systems Hardware Reference Manual, Control Data Publication N. 60100000, 1975.

58. ---, Honeywell RMS-62 brochure and data sheet, Honeywell Information Systems, 1976.

59. ---, SP 2085 TRACE Presentation Guide, Sperry Univac Corporation, publication No. UDI-572, Rev. 3-68, November, 1976.

60. ---, Tesdata MS brochure, Tesdata Systems Corporation, McLean, Virginia.

61. ---, The Sperry Univac Benchmark Monitor Display System (EMD-1100), Sperry Univac Corporation, publication No. U5830, October, 1975.

62. ---, TIMES-1 Operator's User Guide (Preliminary), unpublished, Texas Instruments report.

63. ---, TISAM Operator's User Guide (Preliminary), unpublished, Texas Instruments report.

64. ---, 1100/80 Systems, Sperry Univac Corporation, publication No. U6004, 1977.

65. ---, 6400/6500/6600 PARTNER Installation Manual and Operating Guide, Control Data Corporation.

BS-114A (REV. 7-73)

U.S. DEPT. OF COMM. BIBLIOGRAPHIC DATA SHEET	1. PUBLICATION OR REPORT NO. NBS SP 500-42	2. Gov't Accession No.	3. Recipient's Accession No.

4. TITLE AND SUBTITLE Computer Science and Technology: A Survey of Remote Monitoring	5. Publication Date January 1979 6. Performing Organization Code
7. AUTHOR(S) Gary J. Nutt, Ph.D	8. Performing Organ. Report No. NBS SP 500-
9. PERFORMING ORGANIZATION NAME AND ADDRESS NATIONAL BUREAU OF STANDARDS DEPARTMENT OF COMMERCE WASHINGTON, D.C. 20234	10. Project/Task/Work Unit No. 11. Contract/Grant No. 7-35839
12. Sponsoring Organization Name and Complete Address (Street, City, State, ZIP) Same as No. 9	13. Type of Report & Period Covered Final NBS SP 500- 14. Sponsoring Agency Code

15. SUPPLEMENTARY NOTES

Library of Congress Catalog Card Number 78-26313

16. ABSTRACT (*A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here.*)

This report describes remote monitoring in the application areas of performance evaluation, diagnostic testing, performance assurance and system security testing. The evolution of remote monitoring is briefly reviewed and, then, remote monitors are categorized into seven classes. Several example systems are discussed for each classification, along with their capabilities in each application area. The views presented in this report represent only those of the author, an independent consultant, and should not be construed as a policy statement of NBS or any other organization.

17. KEY WORDS (*six to twelve entries; alphabetical order; capitalize only the first letter of the first key word unless a proper name; separated by semicolons*)

Diagnostic testing; performance assurance; performance evaluation; remote monitoring; system security testing.

18. AVAILABILITY [X] Unlimited [ ] For Official Distribution. Do Not Release to NTIS [X] Order From Sup. of Doc., U.S. Government Printing Office Washington, D.C. 20402, SD Stock No. SN003-003-02013-3 [ ] Order From National Technical Information Service (NTIS) Springfield, Virginia 22151	19. SECURITY CLASS (THIS REPORT) UNCLASSIFIED 20. SECURITY CLASS (THIS PAGE) UNCLASSIFIED	21. NO. OF PAGES 34 22. Price $1.50

USCOMM-DC 66035-P78

ANNOUNCEMENT OF NEW PUBLICATIONS ON COMPUTER SCIENCE & TECHNOLOGY

Superintendent of Documents,
Government Printing Office,
Washington, D. C. 20402

Dear Sir:

Please add my name to the announcement list of new publications to be issued in the series: National Bureau of Standards Special Publication 500-.

Name _____

Company _____

Address _____

City _____ State _____ Zip Code _____

(Notification key N-503)

☆ U. S. GOVERNMENT PRINTING OFFICE : 1978 281-067/359

NBS TECHNICAL PUBLICATIONS

PERIODICALS

JOURNAL OF RESEARCH—The Journal of Research of the National Bureau of Standards reports NBS research and development in those disciplines of the physical and engineering sciences in which the Bureau is active. These include physics, chemistry, engineering, mathematics, and computer sciences. Papers cover a broad range of subjects, with major emphasis on measurement methodology, and the basic technology underlying standardization. Also included from time to time are survey articles on topics closely related to the Bureau's technical and scientific programs. As a special service to subscribers each issue contains complete citations to all recent NBS publications in NBS and non-NBS media. Issued six times a year. Annual subscription: domestic $17.00; foreign $21.25. Single copy, $3.00 domestic; $3.75 foreign.

Note: The Journal was formerly published in two sections: Section A "Physics and Chemistry" and Section B "Mathematical Sciences."

DIMENSIONS/NBS

This monthly magazine is published to inform scientists, engineers, businessmen, industry, teachers, students, and consumers of the latest advances in science and technology, with primary emphasis on the work at NBS. The magazine highlights and reviews such issues as energy research, fire protection, building technology, metric conversion, pollution abatement, health and safety, and consumer product performance. In addition, it reports the results of Bureau programs in measurement standards and techniques, properties of matter and materials, engineering standards and services, instrumentation, and automatic data processing.

Annual subscription: Domestic, $11.00; Foreign $13.75

NONPERIODICALS

Monographs—Major contributions to the technical literature on various subjects related to the Bureau's scientific and technical activities.

Handbooks—Recommended codes of engineering and industrial practice (including safety codes) developed in cooperation with interested industries, professional organizations, and regulatory bodies.

Special Publications—Include proceedings of conferences sponsored by NBS, NBS annual reports, and other special publications appropriate to this grouping such as wall charts, pocket cards, and bibliographies.

Applied Mathematics Series—Mathematical tables, manuals, and studies of special interest to physicists, engineers, chemists, biologists, mathematicians, computer programmers, and others engaged in scientific and technical work.

National Standard Reference Data Series—Provides quantitative data on the physical and chemical properties of materials, compiled from the world's literature and critically evaluated. Developed under a world-wide program coordinated by NBS. Program under authority of National Standard Data Act (Public Law 90-396).

NOTE: At present the principal publication outlet for these data is the Journal of Physical and Chemical Reference Data (JPCRD) published quarterly for NBS by the American Chemical Society (ACS) and the American Institute of Physics (AIP). Subscriptions, reprints, and supplements available from ACS, 1155 Sixteenth St. N.W., Wash., D.C. 20056.

Building Science Series—Disseminates technical information developed at the Bureau on building materials, components, systems, and whole structures. The series presents research results, test methods, and performance criteria related to the structural and environmental functions and the durability and safety characteristics of building elements and systems.

Technical Notes—Studies or reports which are complete in themselves but restrictive in their treatment of a subject. Analogous to monographs but not so comprehensive in scope or definitive in treatment of the subject area. Often serve as a vehicle for final reports of work performed at NBS under the sponsorship of other government agencies.

Voluntary Product Standards—Developed under procedures published by the Department of Commerce in Part 10, Title 15, of the Code of Federal Regulations. The purpose of the standards is to establish nationally recognized requirements for products, and to provide all concerned interests with a basis for common understanding of the characteristics of the products. NBS administers this program as a supplement to the activities of the private sector standardizing organizations.

Consumer Information Series—Practical information, based on NBS research and experience, covering areas of interest to the consumer. Easily understandable language and illustrations provide useful background knowledge for shopping in today's technological marketplace.

Order **above** *NBS publications from: Superintendent of Documents, Government Printing Office, Washington, D.C. 20402.*

Order **following** *NBS publications—NBSIR's and FIPS from the National Technical Information Services, Springfield, Va. 22161.*

Federal Information Processing Standards Publications (FIPS PUB)—Publications in this series collectively constitute the Federal Information Processing Standards Register. Register serves as the official source of information in the Federal Government regarding standards issued by NBS pursuant to the Federal Property and Administrative Services Act of 1949 as amended, Public Law 89-306 (79 Stat. 1127), and as implemented by Executive Order 11717 (38 FR 12315, dated May 11, 1973) and Part 6 of Title 15 CFR (Code of Federal Regulations).

NBS Interagency Reports (NBSIR)—A special series of interim or final reports on work performed by NBS for outside sponsors (both government and non-government). In general, initial distribution is handled by the sponsor; public distribution is by the National Technical Information Services (Springfield, Va. 22161) in paper copy or microfiche form.

BIBLIOGRAPHIC SUBSCRIPTION SERVICES

The following current-awareness and literature-survey bibliographies are issued periodically by the Bureau:

Cryogenic Data Center Current Awareness Service. A literature survey issued biweekly. Annual subscription: Domestic, $25.00; Foreign, $30.00.

Liquified Natural Gas. A literature survey issued quarterly. Annual subscription: $20.00.

Superconducting Devices and Materials. A literature survey issued quarterly. Annual subscription: $30.00. Send subscription orders and remittances for the preceding bibliographic services to National Bureau of Standards, Cryogenic Data Center (275.02) Boulder, Colorado 80302.

U.S. DEPARTMENT OF COMMERCE
National Bureau of Standards
Washington, D.C. 20234

OFFICIAL BUSINESS

Penalty for Private Use, $300

POSTAGE AND FEES PAID
U.S. DEPARTMENT OF COMMERCE
COM-215

U.S.MA

SPECIAL FOURTH-CLASS RATE
BOOK