

A11103 090127

NAT'L INST OF STANDARDS & TECH R.I.C.



A11103090127

Mamrak, Sandra A/Sizing distributed syst
QC100 .U57 NO.500-, 60, 1980 C.1 NBS-PUB

COMMERCE & TECHNOLOGY:



SIZING DISTRIBUTED SYSTEMS: OVERVIEW AND RECOMMENDATIONS



NBS Special Publication 500-60

**U.S. DEPARTMENT OF COMMERCE
National Bureau of Standards**

500-60

0

NATIONAL BUREAU OF STANDARDS

The National Bureau of Standards¹ was established by an act of Congress on 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 Quantities² — Radiation Research — Thermodynamics and Molecular Science — Analytical Chemistry — Materials Science.

THE NATIONAL ENGINEERING LABORATORY provides technology and technical services to the public and private sectors to address national needs and to solve national problems; conducts research in engineering and applied science in support of 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 ultimate user. The Laboratory consists of the following centers:

Applied Mathematics — Electronics and Electrical Engineering² — Mechanical Engineering and Process Technology² — 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 centers:

Programming Science and Technology — Computer Systems Engineering.

¹Headquarters and Laboratories at Gaithersburg, MD, unless otherwise noted; mailing address Washington, DC 20234.

²Some divisions within the center are located at Boulder, CO 80303.

MAY 13 1980


Not Recd - 01
QC100
-U57
170 500-115
1980
C.2-

COMPUTER SCIENCE & TECHNOLOGY:

Sizing Distributed Systems: Overview and Recommendations

Sandra A. Mamrak

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



Special publication

U.S. DEPARTMENT OF COMMERCE, Philip M. Klutznick, Secretary

Luther H. Hodges, Jr., Deputy Secretary

Jordan J. Baruch, Assistant Secretary for Productivity, Technology, and Innovation

U.S. NATIONAL BUREAU OF STANDARDS, Ernest Ambler, Director

Issued May 1980

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.

Library of Congress Catalog Card Number: 80-600061

National Bureau of Standards Special Publication 500-60

Nat. Bur. Stand. (U.S.), Spec. Publ. 500-60, 21 pages (May 1980)

CODEN: XNBSAV

U.S. GOVERNMENT PRINTING OFFICE

WASHINGTON: 1980

For sale by the Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402

Price \$1.75

(Add 25 percent additional for other than U.S. mailing)

Sizing Distributed Systems:
Overview and Recommendations

CONTENTS

	<u>Page</u>
1.0 INTRODUCTION	1
2.0 SYSTEM SIZING TECHNIQUES	2
3.0 ADDITIONAL SIZING ISSUES	10
3.1 Sizing Problem: Scope and Frequency of Occurrence	10
3.2 Analysts: Availability, Expertise and Credibility .	11
3.3 Availability of Analysis Tools	11
3.4 Availability of Measurement Data	12
4.0 RECOMMENDATIONS FOR SIZING DISTRIBUTED SYSTEMS	12
4.1 Establishing In-House Expertise	12
4.2 Developing a Measurement Center	13
5.0 CONCLUSIONS	14
6.0 REFERENCES	14

FIGURES

Figure	<u>Page</u>
1. Queueing Model	5
2. Computer System Phases Assumed for Hybrid Models . . .	7

TABLES

Table	
1. System Sizing Techniques in Order of Increasing Credibility, Accuracy and Cost	3

Sizing Distributed Systems: Overview and Recommendations

Sandra A. Mamrak

ABSTRACT

Computer system sizing is a complicated process for which a variety of tools have been developed. The choice of tools for a particular sizing exercise is guided by many considerations such as cost, available data and the expertise of the analyst. This report presents an overview of sizing techniques, a brief discussion of the factors that affect choosing one or a combination of techniques, and a set of recommendations for choosing tools for sizing distributed systems. The report is aimed at managerial-level personnel who have developed technical competence with regard to single-processor computer systems and are faced with procurement decisions regarding distributed computer systems or services.

Keywords: benchmarking; distributed systems; hybrid models; queueing analysis; system sizing.

1.0 INTRODUCTION

System sizing is the process of configuring a set of computer hardware and software components so that they can adequately meet the functional and capacity demands of a given workload. It is a complicated process because its success depends not only on specifying the individual components of both a computer system and a workload, but also on capturing the myriad and as yet not understood relationships among all the components. Sizing in distributed computer systems(1) is considerably more complicated than sizing single computer systems because of the number and variety of hardware, software and workload components likely to be present.

Many techniques have been developed for performing system sizing, ranging in sophistication from the application of some rules-of-thumb to the execution of extensive benchmark experiments. These techniques trade off

(1) A distributed computer system is defined to be any system in which a set of host computers or end processors, terminals, and other peripheral devices is interconnected by way of a communications subnetwork.

expected accuracy against expected cost, with the most accurate techniques generally being the most costly. Thus, a decision to use one or the other method for system sizing is largely a cost-benefit analysis, balancing the expected accuracy of a given tool with the budget available for a sizing study.

Other factors besides cost and accuracy also affect the choice of sizing tools. These factors include knowledge of workload, the number of available analysts and their level of expertise, the scope of the sizing problem, the availability of computer-aided design tools, the availability of measurement data, the time-frame for the sizing study and the credibility of the technique to those responsible for decision-making. These factors often dominate any "scientific" considerations in choosing sizing tools. Thus, cost and accuracy considerations must be judiciously balanced with a consideration of all other relevant factors.

This report is aimed at management personnel who are or will be faced with decisions about the procurement of distributed computer systems or services. The purpose of the report is to present a condensed, elementary overview of system sizing options, emphasizing their relative merits with regard to sizing distributed computer systems.

The next section summarizes system sizing techniques as they are generally viewed by computer analysts. The summary is presented primarily to establish a conceptual framework which will facilitate discussion in the rest of the report. Some techniques are discussed in more detail than others to set the stage for recommendations for sizing distributed systems. Section 3.0 presents a more detailed discussion of factors other than cost and accuracy which affect sizing of distributed systems. Recommendations for sizing distributed systems are presented in Section 4.0, based on the issues discussed in the previous two sections.

2.0 SYSTEM SIZING TECHNIQUES

In capacity planning and acquisition of computer systems, the most fundamental question that must be answered is whether or not a proposed computer system configuration will be able to process adequately a current or projected workload. If a system is under-sized, the workload simply will not be able to be processed: a disastrous outcome. If a system is over-sized, computer users are very likely to be paying for extra capacity which they neither desire nor can use.

Adequate processing requires that a computer system be able to meet the functional and capacity demands of a given workload [GSA79]. These demands may represent current or projected processing needs. Functional demands, such as a requirement to support ANS FORTRAN or to provide a hierarchical database system, can usually be evaluated in a straight-forward manner. Most often they can be clearly specified in a vendor-independent language and can be easily assessed to everyone's satisfaction with a yes/no decision. In contrast, a capacity demand, or a requirement to process a workload in a given period of time, is much more difficult to evaluate.

The difficulty in evaluating capacity demands stems from two sources: 1) the need to represent complex interactions among various computer hardware and software components, and 2) the need to accurately represent a projected workload. Since the performance of a computer system can only be evaluated with respect to a given workload [CAL79], an accurate model of a projected workload is essential. The ultimate success of system sizing depends on how precisely the models of both a computer configuration and a test workload represent their real counterparts.

Table 1 lists the range of system sizing techniques. They are ranked in order of increasing credibility, accuracy and cost. Credibility is a subjective criterion, varying from analyst to analyst. Accuracy and cost are difficult to quantify for a given class of techniques. But the ordering presented in Table 1 reflects the relative ranking likely to be assigned by practicing systems analysts. A description of each sizing technique follows.

Table 1. System Sizing Techniques in Order of Increasing Credibility, Accuracy and Cost

1. Subjective Analysis: Rules of Thumb
2. Queueing Models
3. Hybrid Models: Queueing, Simulation
and Other Numerical Components
4. Simulation Models
5. Benchmarking: Real System Running
Synthetic Jobs
6. Benchmarking: Real System Running
Real Jobs

Subjective Analysis: Rules-of-Thumb

This technique involves the application of reasonable "rules-of-thumb" to the system sizing problem. No formal models are employed. Typically, analysts will make subjective, but informed judgments about what the workload requirements are and what hardware/software configurations will support them. These judgments are based on personal experience and on the experience of other analysts who share their expertise through various publications or in informal discussions. An example of a rule-of-thumb that may be applied when sizing a database system is the so-called "80-20 rule" [IDC76]. In a decision about whether to distribute or centralize a database, this rule recommends centralization if 80% or more of database queries are from a local site and 20% or less are from remote sites.

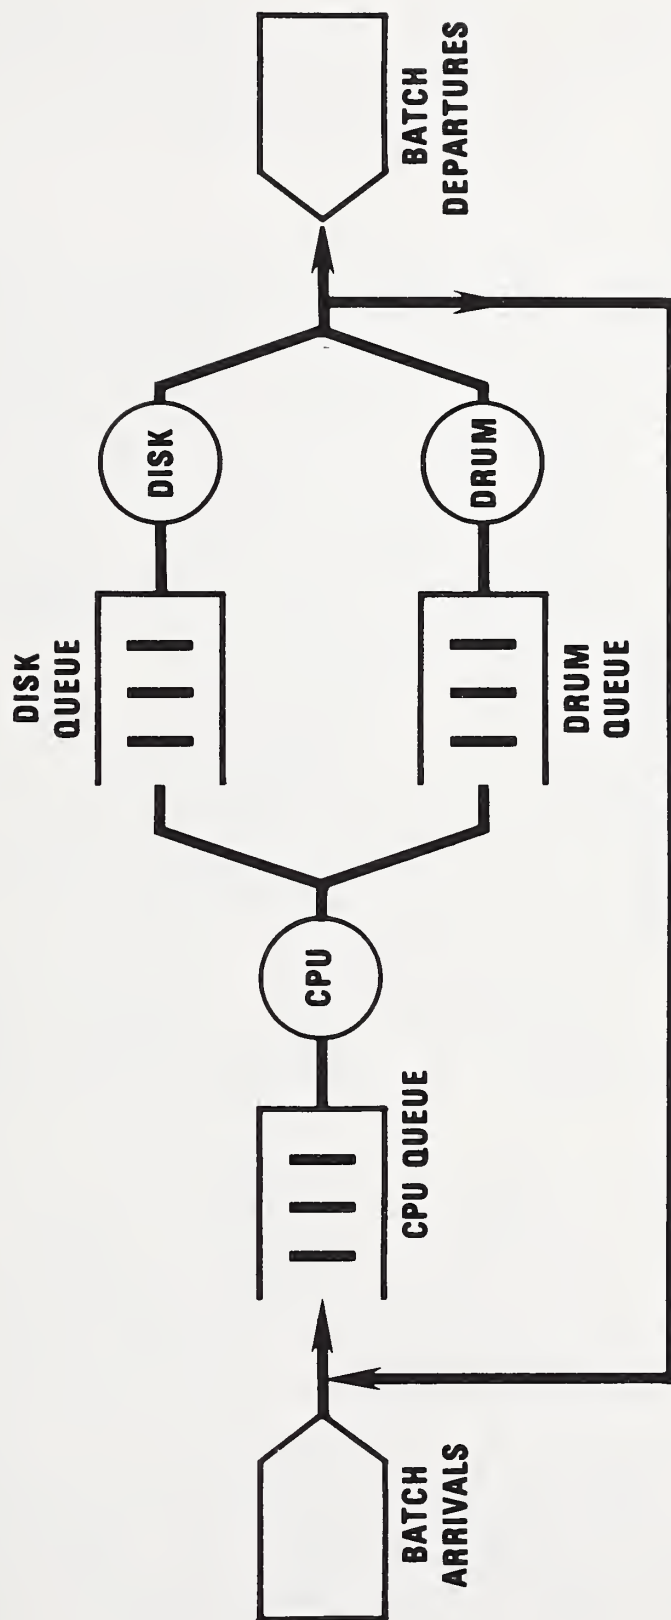
Queueing Models

This technique is characterized by the use of formulas derived from queueing theory analyses of computer systems. The systems are generally viewed as queueing network models [HUG73] with arrivals for service being queued according to various disciplines such as first-in, first-out or processor sharing. Equations are set up to describe system behavior. Solutions to these system equations provide performance quantities such as throughput rates and resource utilizations which are useful for system sizing. Figure 1 shows a typical queueing network model for a batch load.

Many queueing theory models of computer systems exist. (A comprehensive survey of the analysis of queueing network models of computing systems is presented in the September 1978 issue of ACM Computing Surveys.) Although the modeling emphasis has been primarily on single computer systems, computer communication networks have been extensively studied [KLE76] and some models exist for distributed systems [BAB77, LAB77, MCG78, WON78]. A queueing model developed for use in sizing distributed systems, which is unique in that it incorporates economic factors, has been developed by Bucci and Streeter [BUC79].

Queueing network models that represent reality faithfully are often not tractable. That is, if phenomena such as multiple resource holding, blocking, parallel processing and load balancing are incorporated in a queueing model, then the model cannot be analyzed to give exact solutions in a reasonably short time. There has been considerable study devoted to developing approximation methods such as decomposition or diffusion for analyzing queueing network models of computing systems [CHA78]. Although in most cases it is difficult to quantify the error introduced by these approximation techniques, they have proved to be very useful for recognizing and discarding poor design choices (rather than obtaining a high degree of

FIGURE 1
QUEUEING MODEL



precision in predicting performance quantities). Several programming packages exist which incorporate queueing approximation techniques [INF75, REI78, SAU77]. An example which demonstrates the successful use of an approximation technique is the modeling of IBM's Multiple Virtual Storage operating system by Buzen [BUZ78].

Hybrid Models

Hybrid models combine analytic queueing components, simulation components and possible other numerical techniques. The purpose of these hybrid models is to bridge a rather wide gap between pure queueing models and pure simulation models. Simulation models (described in more detail below) have the advantage of accurately representing complex interactions, but they can easily require orders of magnitude more execution time than a simpler analytic queueing model. Queueing models execute quickly, but cannot easily accommodate complex interactions. The basic approach in hybrid modeling is to decompose a complex system into several subsystems and independently choose an appropriate modeling tool for each subsystem in an effort to balance accuracy and speed tradeoffs.

The majority of hybrid models combine analytic and queueing components. A noteworthy exception to this rule is Bard's hybrid model of IBM VM/370, an interactive, multiprogrammed, virtual storage operating system [BAR78]. In the VM/370 Performance Predictor standard methods of queueing network analysis are supplemented by the use of an algebraic transaction flow model and asymptotic formulas for bottleneck analysis.

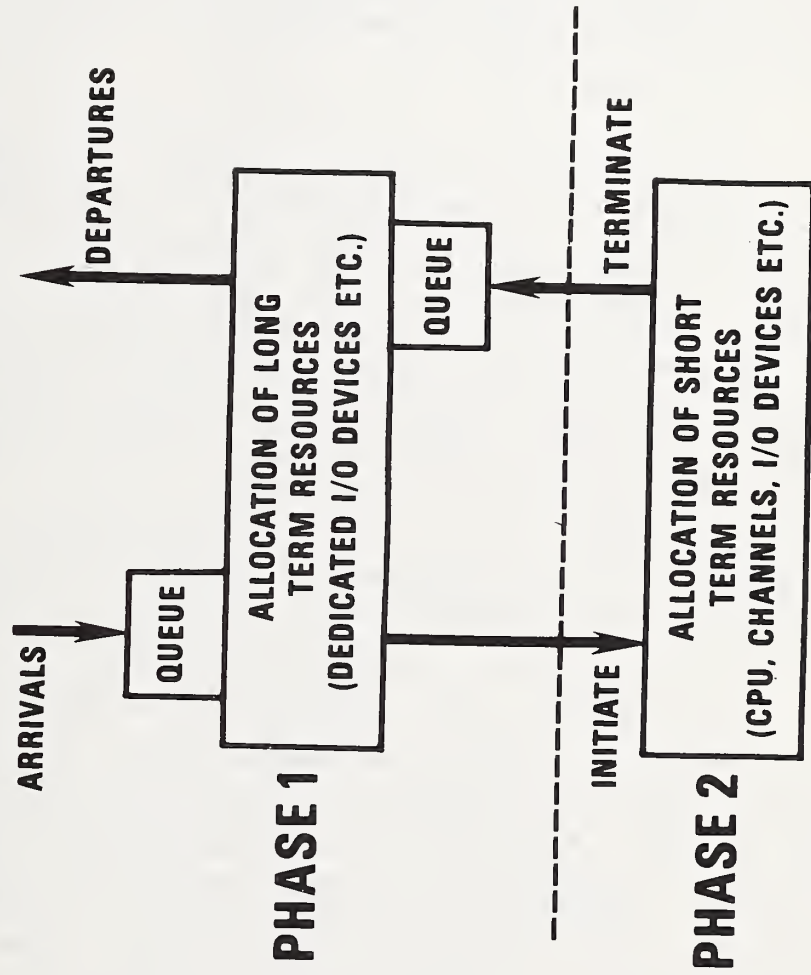
Hybrid simulation models are useful when the computer system to be modeled can be easily decomposed into two phases (see Figure 2): a phase of long-term resource usage (system arrival and departure activity) and a phase of short-term resource usage (CPU, memory and I/O activity).

The first phase is implemented as a simulator, allowing complex job-arrival patterns, arbitrary scheduling rules and allocation policies, and multiple classes of jobs. For example, an arrival rate which is dependent upon the time of day can be easily incorporated as a simulation feature. The time units associated with this phase are typically on the order of seconds or minutes. Implementation of phase 1 as a simulation greatly enhances the accuracy of a hybrid model.

The second phase is implemented as a set of queueing theory equations. In this phase time segments are defined as the time between two successive job initiation or job termination events (arrivals or departures to phase 2). Since the composition of the job mix in this phase is constant over a time segment, analytic prediction of time segment performance is appropriate. The time units

FIGURE 2

COMPUTER SYSTEM PHASES ASSUMED FOR HYBRID MODELS



associated with this phase can be as short as microseconds. It is at this level of detail that simulators tend to execute relatively slowly. Implementation of phase 2 as an analytic model significantly reduces execution time for the hybrid model.

Several languages exist for hybrid simulations and comparison investigations show them to perform nearly as accurately, but in much less time, than simulation models alone [KIM75, SCH78]. An example of the application of hybrid simulation modeling to a complex computer system can be found in Browne's description of a project to model the Advanced Logistics System developed by the United States Air Force Logistics Command [BRO75].

Simulation Models

Discrete-event simulation models represent computer system activity as a series of "events" and simulate running of a computer system by scheduling, executing and collecting data describing a pre-defined sequence of events over some period of time. The level of abstraction of the events depends on the capabilities of the simulation language being used and can vary from "an arrival at a single server queue" to "a retrieval request against a hierarchical database". As mentioned above, simulation models can be used to represent complex interactions at any desired level of detail. Execution time for a simulation is roughly proportional to its level of detail.

Simulations can be written in high-level programming languages or in languages designed for general queueing systems or even specific computer systems. A good survey of simulation languages is presented in SHA75, Chapter 3.

Benchmarking: Real System Running Synthetic Jobs

Synthetic benchmarking is a technique in which the system to be sized is not represented by a model, but by an actual hardware/software configuration. A workload to drive the system, however, is represented at a fairly abstract level by a set of synthetic tasks which are either resource oriented [SRE74], or functionally oriented [CON79]. Resource oriented tasks are designed to consume CPU, memory, channel and I/O device time rather than to perform functionally (e.g. do FORTRAN compiles or text-editing). Functionally oriented tasks are those which perform some pre-defined automatic data processing function like a database query or update. This sizing technique has an advantage over previous methods in that the difficult task of modeling interactions of system components is eliminated. However, the ultimate success of a sizing effort also depends on the accuracy of the workload model (synthetic benchmarks in this case).

A resource-oriented description of a workload is an appropriate one in an environment where alternative systems are essentially homogeneous with respect to hardware and software. It is not appropriate for sizing systems with heterogeneous components, as is very often the case for distributed computer systems. Functionally oriented synthetic benchmarks are valid for use in heterogeneous system selection. They have been used with apparent success in some computer procurements (see MCN77 for example), but there is not yet sufficient data to establish their feasibility in general. The National Bureau of Standards is currently exploring the possibility of establishing a central distribution facility for a highly developed set of synthetic benchmark programs developed by the Department of Agriculture [CON79]. If this is done, use of the benchmark materials could be monitored, thus providing a broader database for accessing feasibility.

Benchmarking: Real System Running Real Jobs

Benchmarking is the most complex and costly technique for system sizing, but it is generally believed to be the most accurate method available for sizing single computer systems. It is the only existing system sizing technique that, when executed properly, is universally accepted by both vendors and procurement agencies as being "fair". In this approach, as in the case of synthetic benchmarking, the proposed computer configuration is used rather than a model thereof. In addition, a complex model of a test workload is constructed, incorporating functional, resource-usage and performance characteristics of the real workload [AGR76, WRI76]. Thus, the technique eliminates as much abstract modeling of the system or workload as is practically feasible, and incorporates all the complex interactions among hardware, software and workload components.

Benchmarking is very expensive. Its cost can easily reach millions of dollars, depending upon system specifications and the number of vendors involved in a procurement bid [PRP78]. A large part of the cost is due to the fact that benchmarking is labor-intensive and can easily occupy a highly skilled team of analysts for months. The cost of benchmarking to size a total distributed system is expected to be too high to justify the benefit.

There are also technological problems that arise when moving from benchmarking large single computer systems to benchmarking distributed systems. These stem from the manner in which benchmark experiments are run. First, there is the need to construct a workload which accurately models a real workload. Although a considerable amount of work has been done to guide test workload construction for single systems [FER79, FIP79], no work has even begun for characterizing workloads on distributed systems. Second, in order to run a benchmark, the proposed hardware/software

configuration must be assembled prior to actual purchase. Vendors have developed elaborate benchmarking centers which allow assembly of components in various combinations for single system sizing. Such assembly would be extremely difficult for distributed systems, especially since most are likely to be composed of multi-vendor components.

3.0 ADDITIONAL SIZING ISSUES

The choice of techniques for system sizing is influenced by a variety of factors other than the accuracy and cost of sizing tools. These other factors are discussed in this section, emphasizing their potential impact on decisions for sizing distributed computer systems(2).

3.1 Sizing Problem: Scope And Frequency Of Occurrence

The anticipated complexity of a computing system as judged by 1) the number and kinds of possible alternatives, 2) the expected frequency with which sizing decisions are to be made and 3) the time available in which to do a system sizing study all impact a choice of sizing techniques.

When sizing distributed systems the number and kinds of possible alternatives will be large. Consideration of hardware components alone presents choices among telecommunication carriers, subnetwork interface components, host computers, user terminals and other peripheral devices. Various combinations of these components provide possibly thousands of alternatives that have to be compared for a given design. Several fast, less accurate tools such as network queueing analysis are needed to eliminate a large portion of unacceptable alternatives. After the field is narrowed, more sophisticated, but relatively slower tools such as the hybrid modeling tools described above are needed. Simulation, and even limited benchmarking if possible, may be employed for decision making among a final small set of options.

The acquisition of a distributed computer system is likely to be an infrequent occurrence for most installations. Such major purchases are likely to be made once in every five to ten years. Under such circumstances it is generally not cost-effective for a purchasing agency to spend large amounts of money building up complex sizing models and developing in-house expertise in the use of the models.

(2) Chandy presents a similar discussion of factors influencing a choice of analysis techniques in CHA78.

The time allowed for any particular sizing study is usually proportional to the anticipated size and cost of the proposed computer system. Even for a large and complex system, however, time for sizing studies is often limited to a few months [BR075, PRP78]. This implies there is little time to build up complex system models from scratch or to build up analyst expertise in using such models.

3.2 Analysts: Availability, Expertise And Credibility

The number of analysts available for a sizing study and their level of expertise relative to various sizing tools are critical factors in a choice of sizing techniques. Even more important, and perhaps the most critical factor of all, is the faith that management has in the analysts' ability to correctly use a set of tools to solve its specific sizing problems.

The complexity of sizing distributed systems dictates that several different kinds of tools be used. This in turn dictates that several analysts be available for the sizing study. The availability of a pool of analysts not only provides for diverse areas of expertise, but also allows for partitioning a sizing study into components that can be studied in parallel, thus shortening the total time required for the study.

The attitude of management toward various sizing tools, and the confidence that management has in a set of analysts and their ability to use those tools, strongly influences a choice of sizing techniques. Two-way communication lines must be kept open between management and an analysis team so that each understands the priorities and constraints under which the other is working. This communication is an absolutely essential requirement if managers are expected to accept the recommendations of an analysis team.

3.3 Availability Of Analysis Tools

The availability of computer-aided tools is a key factor in the choice of a sizing technique. They relieve an analyst of the burden of developing such packages as a part of the sizing study and they often provide friendly interfaces which expedite an understanding of the underlying tools themselves. Several programming packages exist for various sizing approaches and have been referenced in Section 2.0.

A comprehensive interactive program called General Utility for Estimating System Size (GUESS) has been developed by the Network Analysis Corporation [MCG78].

GUESS has been used for determining the relative merits of specific architectural alternatives (i.e. local access, connection, switch and host processor combinations) given a requirements specification. GUESS, along with a tutorial on its use, is available to government agencies for system sizing studies, but is otherwise proprietary.

3.4 Availability Of Measurement Data

The absence of measurement data may preclude the use of certain models which require detailed descriptions of a workload that can only be obtained through measurement. In general, the more sophisticated the model, the more sensitive it is to the accuracy of input data. Thus, in sizing a proposed distributed system the spectrum of sizing tools is likely to range from simpler models which require little input information but allow eliminating bad choices, to more sophisticated models which may even be fed with limited measurement data flowing from a prototype or early, perhaps reduced, system implementation.

4.0 RECOMMENDATIONS FOR SIZING DISTRIBUTED SYSTEMS

The discussion of factors influencing sizing of distributed systems leads to two fundamental recommendations:

1. Establish long-term, in-house expertise in sizing, or hire appropriate outside experts.
2. Develop a measurement center as an integral part of a distributed computing system.

4.1 Establishing In-House Expertise

No "best" methodology or cookbook approach is appropriate for the general problem of sizing distributed systems. It is a complex art, relying on a set of scientific tools that must be used carefully and intelligently if they are to yield valid results. Experience in the use of the available tools is an essential ingredient for sizing success. Therefore, only a knowledgeable, experienced analysis team will be able to properly size distributed systems.

Large government agencies and corporations which have sufficient resources available may find it cost-effective to invest in building up in-house analysis groups with skills

in developing and using all of the available tools for system sizing. They will all be needed as the various stages of sizing progress from original "pencil and paper" analyses through full implementation and support of an operational system.

For those groups without the resources to build up and maintain an in-house analysis team, a system sizing problem will be best handled by contracting out to an appropriate consulting firm, bringing in consultants to work with on-site staff, or hiring sizing experts for the term of a procurement. In the government, for example, FEDSIM, the Federal Computer Performance Evaluation and Simulation Center, could be called upon to provide some form of expert assistance, as could other service selection agencies. Within the normal time frame of a typical sizing study it will not be possible to begin to gather the required personnel and tools, to build up adequate expertise and to actually do the sizing. The problem is simply too complex and the price of error too high.

4.2 Developing A Measurement Center

Distributed systems, after their initial interconnection, are likely to expand in a modular fashion, on a component-by-component basis. System sizing questions will be most often directed primarily at small to medium size host computers and a variety of intelligent peripheral devices. Sizing considerations in this environment can be viewed as one aspect of a comprehensive "capacity planning" process, where capacity planning is defined as the forecasting of future hardware and software requirements.

Capacity planning is best done by using historical performance data [ART78]. This approach precludes the need for using abstract models of either a computer system or a workload. Thus, performance evaluation based on historical performance data has the potential for very accurately reflecting the true behavior of the system.

To be most effective, however, data collection and analysis must be carefully done. It is not sufficient to amass a roomful of tapes which contain a hodgepodge of system performance measurements. A long term commitment is required to properly instrument distributed systems, in their design phase if possible [NBS78], and to establish and maintain an extensive performance measurement database spanning several years of measurement. Questions of where, when, what, why and how to measure must be studied in the context of how the data can be used to feed the modeling [ROS78] and benchmarking [ART78] activities that will be used to maintain and improve levels of performance.

5.0 CONCLUSIONS

The choice of techniques used to size distributed computer systems depends on many factors. Some of these factors introduce scientific issues while others are a function of the particular exigencies of a given sizing problem. When all relevant factors are considered, it is evident that distributed systems will only be successfully sized by an experienced analysis team that is knowledgeable in the development and use of a diverse set of sizing tools. Further, on-going sizing of distributed systems will be greatly enhanced by the development and support of a measurement center as an integral component of a distributed computer system.

6.0 REFERENCES

- AGR76 Agrawala, A. K., J. M. Mohr and R. M. Bryant, "An Approach to the Workload Characterization Problem," Computer, Vol. 9, No. 6, June 1976, pp. 18-32.
- ART78 Artis, H. P., "Capacity Planning for MVS Computer Systems," Performance of Computer Installations, ed. D. Ferrari, North-Holland Publishing Company, 1978.
- BAB77 Babic, G. A., M. T. Liu and R. Pardo, "A Performance Study of the Distributed Loop Computer Network (DLCN)," Proceedings Computer Networking Symposium, 1977, pp. 66-75.
- BAR78 Bard, Y., "The VM/370 Performance Predictor," ACM Computing Surveys, Vol. 10, No. 3, September 1978, pp. 333-342.
- BRO75 Browne, J. C. et al, "Hierarchical Techniques for the Development of Realistic Models of Complex Computer Systems," Proceedings of the IEEE, Vol. 63, No. 6, June 1975, pp. 965-975.
- BUC79 Bucci, G. and D. N. Streeter, "A Methodology for the Design of Distributed Information Systems," Communications of the ACM, Vol. 22, No. 4, April 1979, pp. 233-245.
- BUZ78 Buzen, J. P., "A Queueing Network Model of MVS," ACM Computing Surveys, Vol. 10, No. 3, September 1978, pp. 319-331.
- CAL79 Cale, E. G., L. L. Gremillion and J. L. McKenny, "Price/Performance Patterns of U.S. Computer Systems," Communications of the ACM, Vol. 22, No. 4, April 1979, pp. 225-233.

- CHA78 Chandy, K. M. and C. H. Sauer, "Approximate Methods for Analyzing Queueing Network Models of Computing Systems," ACM Computing Surveys, Vol. 10, No. 3, September 1978, pp. 281-317.
- CON79 Conti, D. M., Findings of the Standard Benchmark Library Study Group, National Bureau of Standards Special Publication 500-38, Washington, DC, January 1979.
- FER79 Ferrari, D., "Characterizing a Workload for the Comparison of Interactive Services," Proceedings AFIPS 1979 NCC, Vol. 49, Montvale, NJ, 1979, pp. 789-796.
- FIP79 Guideline on Benchmark Construction, National Bureau of Standards, Federal Information Processing Standards Publication, in press.
- GSA79 "Use of Remote Terminal Emulation in Federal ADP System Procurements (DRAFT)," Office of Agency Services and Procurement, Automated Data and Telecommunications Service, General Services Administration, March, 1979.
- HUG73 Huges, P. and G. Moe, "A Structural Approach to Computer Performance Analysis," Proceedings AFIPS 1973 NCC, Vol. 42, AFIPS Press, Montvale, NJ, 1973, pp. 109-120.
- IDC76 Distributed Processing, International Data Corporation, Report 1763, December 1976.
- INF75 Users Manual for the ASQ System, Information Research Associates, Austin, TX, 1975.
- KIM75 Kimbleton, S., "A Heuristic Approach to Computer Systems Performance Improvement, I -- A Fast Performance Prediction Tool," Proceedings AFIPS 1975 NCC, Vol. 44, AFIPS Press, Montvale, NJ, 1975, pp. 839-846.
- KLE76 Kleinrock. L., Queueing Systems - Vol. 2: Computer Applications, John Wiley and Sons, New York, NY, 1976.
- LAB77 Labetoulle, J., E. G. Manning and R. Peebles, "A Homogeneous Computer Network: Analysis and Simulation," Computer Network, Vol. 1, 1977, pp. 225-240.
- MCG78 McGregor, P. V. and R. Kaczmarek, "Modeling Network Architectures," Proceedings COMPCON78, IEEE Service Center, Piscataway, NJ, 1978, pp.419-426.
- MCN77 McNeece, J. and R. Sobecki, "Functional Workload Characterization," Proceedings of the 13th Meeting of the Computer Performance Evaluation Users Group, National Bureau of Standards Special Publication 500-18, September 1977, pp. 13-21.

- NBS78 Local Area Networking, Workshop Report, National Bureau of Standards Special Publication 500-31, ed. I. W. Cotton, Washington, DC, April 1978.
- PRP78 Acquisition Team Report, Federal Data Processing Reorganization Study, President's Reorganization Project, June 20, 1978.
- REI78 Reiser, M. and C. H. Sauer, "Queueing Network Models: Methods of Solution and Their Program Implementation," in Current Trends in Programming Methodology, Vol. III: Software Modeling and Its Impact on Performance, edited by K. M. Chandy and R. T. Yeh, Prentice-Hall Inc., Englewood Cliffs, NJ, 1978, pp. 115-167.
- ROS78 ROSE, C. A., "A Measurement Procedure for Queueing Network Models of Computer Systems," ACM Computer Surveys, Vol. 10, No. 3, September 1978, pp. 263-280.
- SAU77 Sauer, C. H. and E. A. MacNair, Computer/Communication System Modeling with Extended Queueing Networks, RC-6654, IBM Research, Yorktown Heights, NY, July 1977.
- SCH78 Schwetman, H. D., "Hybrid Simulation Models of Computer Systems," Communications of the ACM, Vol. 21, No. 9, pp. 718-723.
- SRE74 Sreenivasan, K. and A. Kleinman, "On the Construction of a Representative Synthetic Workload," Communications of the ACM, Vol. 17, No. 3, March 1974, pp. 127-133.
- SHA75 Shannon, R. E., Systems Simulation: The Art and Science, Prentice-Hall, Englewood Cliffs, NJ, 1975.
- WON78 Wong, J. Research in Queueing Models for Computer Communications, Waterloo Research Institute Report, Waterloo, Ontario, Canada, May 1978.
- WRI76 Wright, L. S. and W. A. Burnette, "An Approach to Evaluating Time Sharing Systems: MH-TSS A Case Study," SIGMETRICS, Vol. 5, No. 1, January 1976, pp. 8-28.

U.S. DEPT. OF COMM. BIBLIOGRAPHIC DATA SHEET		1. PUBLICATION OR REPORT NO. NBS SP 500-60		2. Gov't. Accession No.		3. Recipient's Accession No.	
4. TITLE AND SUBTITLE COMPUTER SCIENCE & TECHNOLOGY: Sizing Distributed Systems: Overview and Recommendations				5. Publication Date May 1980			
				6. Performing Organization Code			
7. AUTHOR(S) Sandra A. Mamrak				8. Performing Organ. Report No.			
9. PERFORMING ORGANIZATION NAME AND ADDRESS NATIONAL BUREAU OF STANDARDS DEPARTMENT OF COMMERCE WASHINGTON, DC 20234				10. Project/Task/Work Unit No.			
				11. Contract/Grant No.			
12. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS (Street, City, State, ZIP) same as Item 9				13. Type of Report & Period Covered Final			
				14. Sponsoring Agency Code			
15. SUPPLEMENTARY NOTES Library of Congress Catalog Card Number: 80-600061 <input type="checkbox"/> Document describes a computer program; SF-185, FIPS Software Summary, is attached.							
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.) Computer system sizing is a complicated process for which a variety of tools have been developed. The choice of tools for a particular sizing exercise is guided by many considerations such as cost, available data and the expertise of the analyst. This report presents an overview of sizing techniques, a brief discussion of the factors that affect choosing one or a combination of techniques, and a set of recommendations for choosing tools for sizing distributed systems. The report is aimed at managerial-level personnel who have developed technical competence with regard to single-processor computer systems and are faced with procurement decisions regarding distributed computer systems or services.							
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) Benchmarking; distributed systems; hybrid models; queueing analysis; system sizing.							
18. AVAILABILITY <input checked="" type="checkbox"/> Unlimited <input type="checkbox"/> For Official Distribution. Do Not Release to NTIS <input checked="" type="checkbox"/> Order From Sup. of Doc., U.S. Government Printing Office, Washington, DC 20402. <input type="checkbox"/> Order From National Technical Information Service (NTIS), Springfield, VA. 22161				19. SECURITY CLASS (THIS REPORT) UNCLASSIFIED		21. NO. OF PRINTED PAGES 21	
				20. SECURITY CLASS (THIS PAGE) UNCLASSIFIED		22. Price \$1.75	

**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)

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 Bureau publications in both NBS and non-NBS media. Issued six times a year. Annual subscription: domestic \$17; foreign \$21.25. Single copy, \$3 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, business and industry leaders, teachers, students, and consumers of the latest advances in science and technology, with primary emphasis on 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; 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 worldwide program coordinated by NBS under the authority of the National Standard Data Act (Public Law 90-396).

NOTE: The principal publication outlet for the foregoing 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., NW, Washington, DC 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 standards establish nationally recognized requirements for products, and 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 the above NBS publications from: Superintendent of Documents, Government Printing Office, Washington, DC 20402.

Order the following NBS publications—FIPS and NBSIR's—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. The 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; foreign \$30.

Liquefied Natural Gas. A literature survey issued quarterly. Annual subscription: \$20.

Superconducting Devices and Materials. A literature survey issued quarterly. Annual subscription: \$30. Please send subscription orders and remittances for the preceding bibliographic services to the National Bureau of Standards, Cryogenic Data Center (736) Boulder, CO 80303.

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



SPECIAL FOURTH-CLASS RATE
BOOK
