SEPTEMBER 1983
VOL. 62, NO. 7, PART 3

THE BELL SYSTEM TECHNICAL JOURNAL

TOTAL NETWORK DATA SYSTEM

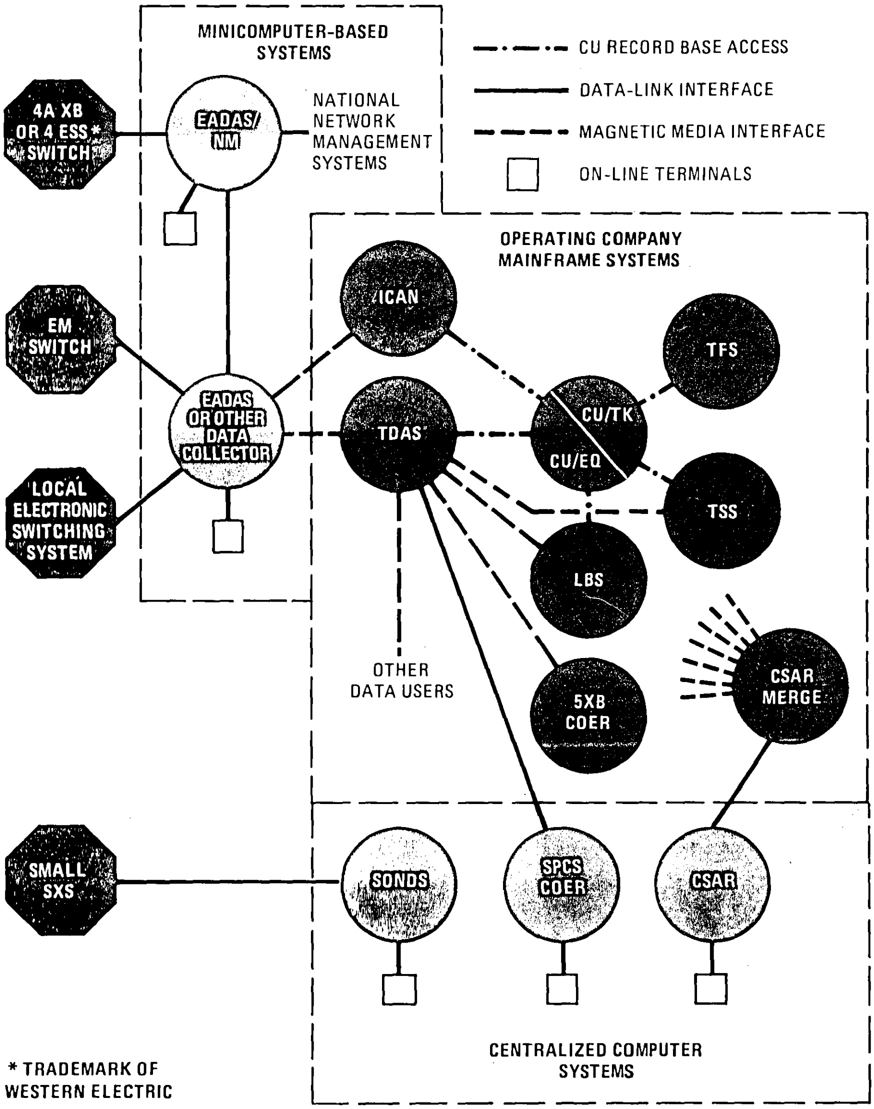

* TRADEMARK OF WESTERN ELECTRIC

THE BELL SYSTEM TECHNICAL JOURNAL

THE BELL SYSTEM TECHNICAL JOURNAL (ISSN0005-8580) is published by the American Telephone and Telegraph Company, 195 Broadway, N. Y., N. Y. 10007; C. L. Brown, Chairman and Chief Executive Officer; W. M. Ellinghaus, President; V. A. Dwyer, Vice President and Treasurer; T. O. Davis, Secretary.

The Journal is published in three parts. Part 1, general subjects, is published ten times each year. Part 2, Computing Science and Systems, and Part 3, single-subject issues, are published with Part 1 as the papers become available.

The subscription price includes all three parts. Subscriptions: United States—1 year $35; 2 years $63; 3 years $84; foreign—1 year $45; 2 years $73; 3 years $94. Subscriptions to Part 2 only are $10 ($12 foreign). Single copies of the Journal are available at $5 ($6 foreign). Payment for foreign subscriptions or single copies must be made in United States funds, or by check drawn on a United States bank and made payable to The Bell System Technical Journal and sent to Bell Laboratories, Circulation Dept., Room 1E-335, 101 J. F. Kennedy Parkway, Short Hills, N. J. 07078.

Comments on the technical content of any article or brief are welcome. These and other editorial inquiries should be addressed to the Editor, The Bell System Technical Journal, Bell Laboratories, Room 1J-319, 101 J. F. Kennedy Parkway, Short Hills, N. J. 07078. Comments and inquiries, whether or not published, shall not be regarded as confidential or otherwise restricted in use and will become the property of the American Telephone and Telegraph Company. Comments selected for publication may be edited for brevity, subject to author approval.

Printed in U.S.A. Second-class postage paid at Short Hills, N. J. 07078 and additional mailing offices. Postmaster: Send address changes to The Bell System Technical Journal, Room 1E-335, 101 J. F. Kennedy Parkway, Short Hills, N. J. 07078.

THE BELL SYSTEM TECHNICAL JOURNAL

DEVOTED TO THE SCIENTIFIC AND ENGINEERING
ASPECTS OF ELECTRICAL COMMUNICATION

Volume 62 | September 1983 | Number 7, Part 3

TOTAL NETWORK DATA SYSTEM

Y. R. Osajima, *Guest Editor*

Introduction 2123
L. Schenker

Environment and Objectives 2127
A. L. Barrese, D. E. Parker, T. E. Robbins, and L. M. Steele

System Plan 2147
M. S. Hall, Jr., J. A. Kohut, G. W. Riesz, and J. W. Steifle

Theoretical and Engineering Foundations 2183
W. S. Hayward and J. P. Moreland

Data Acquisition and Near-Real-Time Surveillance by EADAS 2209
C. J. Byrne, D. J. Gagne, J. A. Grandle, Jr., and G. H. Wedemeyer

Network Management 2239
G. C. Ebner and D. G. Haenschke

National Network Management 2261
W. S. Bartz and R. W. Patterson

Equipment Systems 2281
N. D. Fulton, J. J. Galiardi, E. J. Pasternak, S. A. Schulman, and H. E. Voigt

Trunking Systems 2345
P. V. Bezdek and J. P. Collins

Central Office Equipment Reports for Stored Program Control Systems 2365
R. F. Grantges, V. L. Fahrmann, T. A. Gibson, and L. M. Brown

Small Office Network Data System 2397
D. H. Barnes and J. J. O'Connor

Performance Measurement/Trouble Location 2433
D. R. Anderson and M. J. Evans

Operating Company Perspective 2459
J. Pfeiffer, Jr.

ACRONYMS AND ABBREVIATIONS **2473**

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Introduction

By L. SCHENKER*

(Manuscript received April 18, 1983)

Since the earliest days of telephony telephone traffic measurements have been needed to determine the proper quantities of circuits, operators, and switching systems. In this context "proper" is defined as the efficient and effective utilization of operating factors that provide good service to customers at the lowest possible cost.

The Total Network Data System (TNDS), comprising thirteen different operations systems, is a very large, complex, coordinated set of computer systems developed by Bell Laboratories and now used throughout the Bell Operating Companies. The systems that comprise the Total Network Data System collect, validate, process, archive, and retrieve the traffic data required to fill the entire spectrum of user needs, from near-real-time network management and performance monitoring, through weekly/monthly provisioning and administration, to long-range engineering and planning. The TNDS maintains interfaces with the many telecommunications systems, as well as with operations systems such as the Central Office Equipment Engineering System (COEES), which depend on traffic data.

As the size and complexity of the telephone network expanded, a need evolved for larger quantities of more accurate and timely traffic

* Bell Laboratories.

data. Early measurements of busy equipment were made manually (switch counts, peg counts, etc.) and processed with the aid of desk calculators. The mechanical registers of the 1940s and 1950s evolved into camera register recorders and traffic usage recorders, which partially automated the process. In the mid-1960s, data were being keypunched and then summarized and processed on the large billing computers already in use by comptrollers' organizations. In the early 1970s the advent of the minicomputer facilitated on-line data collection and real-time processing to support such functions as network management. By the mid-1970s, a variety of computer systems collected and processed traffic data to support trunk and switching engineering, administration, and network management. These systems have evolved into a tightly coupled family of operations systems, TNDS, which have been implemented and updated in all the Bell Operating Companies. The systems operate on dedicated minicomputers and large general-purpose computers. The TNDS collects, processes, and utilizes traffic data to provide excellent, efficient, economical telephone service.

The TNDS was developed, used, and enhanced during the thirteen years from 1969 to the present. It has become an indispensable element of Bell operating activity. Without the extensive mechanization provided by the TNDS, it is inconceivable that we could have accommodated the explosive network growth, reconfiguration, and churning. Without the TNDS, operating and investment costs would have risen and service would have been degraded.

This special issue of the *Journal* addresses the TNDS from many points of view. The first two articles describe the TNDS environment and objectives and outline the system plan. The third article describes the conceptual framework and theory upon which the TNDS is built. The eight succeeding, more detailed, articles describe the functions performed by the TNDS and the component operations systems that have been developed to support these functions. The final article, prepared by an employee of Southern Bell, describes the TNDS from an operating telephone company perspective.

As we contemplate the future, we envision continued evolution in two areas: Modifications will be required to keep pace with new services and new technology, to provide interfaces with new telecommunications equipment, and to facilitate new network requirements (such as Dynamic Non-Hierarchical Network Routing); and enhancements will be needed to improve efficiency and make the systems user-friendly. These enhancements will include simplified architecture and new computing facilities, enhanced data communications among the TNDS elements, interactive update features, consolidation of record bases, and on-line user documentation.

This special issue of *The Bell System Technical Journal* appears at an appropriate time—the end of an era. The individual articles reflect the way in which business was conducted prior to 1982. As a result of divestiture, responsibility for much of the TNDS will be transferred to the Central Services Organization, which will be owned and operated by the seven regional Bell Operating Companies. However, the TNDS will continue to provide the independent regional companies with primary traffic data in the future, as it has to the Bell System in the past, without, we hope, missing a beat during the transition.

AUTHOR

Leo Schenker, B.S. (Civil Engineering), 1942, University of London; M.S., 1950, University of Toronto; Ph.D., 1954, University of Michigan; Bell Laboratories, 1954—. After joining Bell Laboratories, Mr. Schenker was involved in the development of telephone station equipment, including *TOUCH-TONE*® dialing and the *TRIMLINE*® phone. In 1968, he was promoted to Director of the Military Electronic Technology Laboratory in North Carolina and, in 1971, became Director of the Loop Maintenance Systems Laboratory, which was closely involved with the development of new software and hardware systems aimed at reducing the expense of maintaining the customer's telephone service. In 1978, he became Director of the Loop Systems Engineering and Methods Laboratory and, in 1979, he was made Director of the Customer Network Operations Systems Engineering Center. In 1980 Mr. Schenker was appointed Executive Director of the Central Office Operations Division. Currently, he is Executive Director, Network Systems Planning Division. Mr. Schenker has been awarded seven patents in connection with *TOUCH-TONE* dialing, *PICTUREPHONE*® meeting service, and *TOUCH-A-MATIC*® repertory dialer. Fellow, IEEE; member, Sigma Xi, Phi Kappa Phi.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Environment and Objectives

By A. L. BARRESE,* D. E. PARKER,† T. E. ROBBINS,* and L. M. STEELE*

(Manuscript received January 19, 1983)

Planning for computer-based tools to help in the measurement and analysis of network traffic data began in the late 1960s. During this same period, the rapid introduction of computer technology was changing the character of the network. The network was becoming more efficient and economical but, at the same time, more sophisticated and complex. This intensified the need to provide timely and complete information to those responsible for the management, administration, engineering, and planning of the network. The computer-based systems that were developed to meet this need are collectively called the Total Network Data System (TNDS). This paper discusses the operating environments of the network and telephone company, and provides a framework for the remainder of the papers in this issue.

I. INTRODUCTION

Managing, engineering, and planning the telecommunications network are essential tasks for the future health and vitality of the network. These complex tasks cannot be performed effectively without detailed data about network traffic. The Bell System has mechanized the collection and processing of such data with a family of computer-based systems collectively called the Total Network Data System

* Bell Laboratories. † AT&T.

(TNDS). This paper provides an overview of the telecommunications network environment, introduces the need for network traffic data in the operating environment of the telephone company, and briefly discusses the primary mechanization objectives for collecting and processing the network traffic data through the TNDS.

II. THE TELECOMMUNICATIONS ENVIRONMENT

To appreciate the need for network traffic data, it is necessary to have a general understanding of the telecommunications network.

2.1 Network description

In the most general sense, a network can be defined as a set of nodes interconnected with links. We can further define the telecommunications network (hereafter referred to as "the network") both on the basis of its function and its physical characteristics. From a functional standpoint, the network carries a variety of telecommunications traffic (e.g., voice, data) between a number of stations that can be connected on demand, or that are permanently connected. From the physical standpoint, the network consists of switching systems (nodes), transmission facilities (links), and stations (sources and receptors of traffic) that are connected together in an organized and controlled manner. These two views are complementary and together provide a framework for understanding aspects of telephone engineering and operations.[1]

Let us examine the physical elements of the network more closely. To construct a telecommunications network that would allow complete, direct interconnection of all stations would be both impractical and cost prohibitive. Therefore, every large communications network is based on the principle of shared facilities, whereby facilities that can be shared among different elements of traffic are utilized extensively when designing and building the network. Central offices are entities built to provide switchable interconnection of customer stations. A central office allows each of its customers, with a single pair of wires connected to the office (i.e., the customer's loop facilities), to talk to any other customer served by that office. Such a central office, having customer station equipment directly connected to it, is said to serve "local" traffic from those stations.

In general, however, customer stations being connected are served by different central offices. Therefore, central offices are connected to one another by transmission paths called trunks, so that customers in one office can reach customers in another. The network of trunks interconnecting central offices is referred to as the Interoffice Facility Network. To enable all stations in the network to be interconnectable while making efficient use of network equipment and facilities, switching and trunking arrangements employ a hierarchical network configuration and the principle of automatic alternate routing.[1] [A new traffic routing technique called Dynamic Non-Hierarchical Routing (DNHR), planned for deployment in the intercity network beginning in 1984, promises significant cost savings in the network.]

The hierarchical network configuration provides for the collection and distribution of traffic and permits switching systems to be completely interconnectable. The hierarchical aspect prevents "loop arounds" that might otherwise occur when calls are automatically alternate routed through the network. Each switching system is given one of five classifications based on the highest switching function performed within the hierarchy, its interrelationship with other switching systems, and its transmission requirements.

With the automatic alternate routing principle, a call that encounters an "all trunks busy" condition in the interoffice facility network on the first route tested is automatically offered sequentially to one or more alternate routes for completion, if such alternate routes are possible for that item of traffic. Arrangements, defined by the Network Routing Plan, that dictate the final routing of traffic are called "homing arrangements." Central offices whose only function is to route calls through the hierarchy are called toll offices. These offices do not directly serve customers (i.e., do not have any connecting loop facilities). Regional Centers, the switching systems at the top of the switching hierarchy, are examples of toll offices. Many central offices perform a tandem switching function (i.e., route traffic by simply interconnecting interoffice trunk groups). These offices may perform a pure tandem switching function or they may also function as local or toll switching systems.

2.2 The concept of service

Because only a small percentage of telephone customers originate calls at any time, the amount of equipment needed to carry the actual simultaneous traffic is only a fraction of that needed to carry simultaneous traffic from all customers. If we install a very limited amount of equipment, there is no assurance that we can meet service demands satisfactorily at all times. Therefore, we can anticipate that a percentage of calls will not be completed (i.e., will be "blocked") during peak traffic periods [e.g., during the busiest hour(s) of the day, during the busy season (busiest three months) of the year]. When calls are "blocked" too frequently, calling customers perceive service as unsatisfactory. Conversely, if we have too much equipment in the network, too much of it will remain unused even during peak traffic periods, which is not cost effective.

The proportion of calls blocked during the busy hour of a switching system is an index of the grade of service rendered. We define the grade of service as the probability that a certain percentage of calls originated during the busy hour will be blocked from utilizing the installed equipment and network facilities. There are strong justifications for providing a high grade of service (i.e., low probability of calls being blocked). For example, when a customer must dial a number more than once because of failure to properly connect, it generates additional traffic, and during peak periods this further aggravates poor service conditions. Acceptable service levels are determined by analyzing both network traffic data and customer reactions.[2] Objective grades of service, or service objectives, are established to balance customer service and network cost. It is the effort to strike a cost-effective balance between providing service and utilizing network equipment that imposes the need for network traffic data.

2.3 Requirements for network data

There are four key network functions that require network traffic data: network management, network administration, network design, and long-range network planning. The function of network management is to control network traffic overloads by distributing loads among circuits and equipment to meet customer service demands in a way that is best from a total network viewpoint. To do this, network management requires a near-real-time view of the traffic in the network. This is made possible through analysis of network traffic data.

The function of network administration is to control the assignment of lines and trunks to take maximum advantage of the installed equipment in the central office for serving the offered call traffic. This involves implementing a plan to spread the office load efficiently over all equipment, as well as to monitor the current load, service levels, and office capacities. A major task within network administration is to use traffic data to monitor the flow of traffic through the central office and to detect changes in office performance (e.g., service degradation) or in offered load. Network administration includes the task of providing sufficient, accurate network traffic data for all functions.

The function of network design is to estimate where, when, and how much equipment will be required within a five-year period so that the necessary additional equipment (i.e., relief equipment) can be ordered and installed in time to satisfy the service objectives.[4] This activity requires data that reflect traffic volumes being carried by existing equipment, as well as knowledge of equipment capacities.

The function of long-range network planning is to determine the most economic growth and replacement strategies for the network to meet estimates of future demand. This future demand is estimated using current traffic load trends and marketing information. The basic output of this function is a broad view of network topology 20 years into the future. Section 3.3 further discusses these key network functions relative to telephone company operations.

In addition to the network functions discussed above, the Bell Operating Company (BOC) marketing organizations request network traffic data to aid in the administration and provisioning of each individual customer's telecommunications service. Due to such factors as increased availability of Stored Program Control (SPC) features and vertical services, increased data sophistication of subscribers and increasing competitive pressures, the number and frequency of these traffic data requests have been increasing consistently. Regulatory needs for these data have also appeared. For example, to support requests for rate increases before regulatory bodies, the BOCs conduct very detailed traffic studies. Regulatory study data requests tend to be complex and difficult to anticipate. These marketing and regulatory needs are only examples of a growing family of special applications that require network traffic data. Though they don't constitute a basic network function, they do represent a significant environmental element that must be considered when planning for the collection and processing of traffic data.

Another function requiring traffic data is operator services force administration. This function involves forecasting the load that will be offered in each half-hour and determining the operator force necessary to carry that load at an objective grade of service. This requires data on traffic volumes, service, and force levels. Because this function has limited interaction with the above-defined key network functions, its data tend to follow a separate but somewhat parallel flow that will not be discussed in this article.

2.4 Network data

The fundamental types of traffic measurements include:

1. Event counts—These measurements simply represent the number of occurrences of an event that occurred in a specific time interval (e.g., hour ending 11:00). "Offered" calls are termed *peg* counts, while "blocked" calls are termed *overflow* counts.

2. Usage—These measurements represent the estimated amount of time an equipment component was busy during a specific time interval, generally an hour. In electromechanical (EM) switching systems, usage is typically obtained through a separate, special measuring device, the Traffic Usage Recorder (TUR). In SPC switching systems, usage measurements (as well as the others) are generated internally by the switch.

3. Delay—This type of measurement usually indicates, on the average, how long a particular type of event would have been delayed, say by an all-servers-busy condition, if it had chosen to wait. Some

different measures of delay include: average delay, average delay of events delayed, probability of delay, and probability of delay exceeding some specified time interval.

4. Status—This measurement, used for network management and administration, indicates the presence or absence of a condition (e.g., equipment outage, severe overload) and is usually in the form of a lamp indication.

With few exceptions, measurements are made at the location of the switching system. All switches make some traffic measurements internally. The interfaces for data collection from EM systems are traffic register leads for peg count, overflow, grouped usage and delay, as well as status-indicating leads. In the case of EM systems, it is usually necessary to provide special traffic measurement equipment (e.g., the TUR) at the switch location. This equipment, in conjunction with special software in the data collection machine, makes possible the gathering of more detailed individual circuit usage data rather than circuit group data. This concept, used for selected EM offices to help increase data accuracy, is called Individual Circuit Usage Recording (ICUR). For SPC switching systems, traffic measurements and status indications are collected by the basic internal programs. The information, equivalent to that stored in EM registers, is retained in memory until the accumulated results are read out on schedules established by the users.[4]

In the Bell System, the volume of network traffic data processed is overwhelming. It was estimated that in the average BOC, over 50 million pieces of traffic data per week were collected and processed in 1980 to support the management, administration, design, and long-range planning of the network.[3] Because of the size, complexity, and importance of the network data job, the management of the data has become a major internal function of the BOCs.

Next we discuss another fundamental element of the network data system environment—the BOC operational environment.

3. THE OPERATIONAL ENVIRONMENT

3.1 Operations planning

Planning for modern computer-based tools to help in the measurement and analysis of network data began in the late 1960s. During this same period, the rapid introduction of computer technology was changing the character of the network. The network was becoming more efficient and economical but, at the same time, more sophisticated and complex. The people responsible for the network's day-to-day operation (i.e., management, administration, engineering and planning) needed more timely and complete information. Bell System management recognized that the largely manual methods in use during the 1960s could not keep pace with the emerging operational needs of the 1970s and 1980s.

Three properties of network data made the planning of computer-based support tools particularly challenging.

1. The data are almost useless until they are processed.

2. A much larger volume of data must be gathered than will be used ultimately because of the way that data are generated and output by the SPC switch and the external measurement devices of the EM switch.

3. The same basic data must be summarized in varying ways to meet the needs of a diverse family of users (see Section 3.3).

Measurement, data processing, and user applications must, therefore, be considered together in planning for the mechanization of an "end-to-end" flow of data.

The formulation of such an end-to-end view of the data flow requires an understanding of the functions performed by various work groups and how these groups relate to one another. This allows the formulation of requirements that describe how the work groups can best be supported by computer-based tools. Finally, the tools must be related to one another and to the elements of the telecommunications network. This concept of a process view, encompassing functions to be performed by people and computers, has been one of the principles used in the design of the individual components, termed *Operations Systems* (*OSs*), that comprise the Total Network Data System (TNDS).[4]

This view of the overall data flow not only facilitated the allocation of functions to the specific TNDS-related OSs, but also resulted in the identification of functions that must be performed by people. Consequently, work groups, termed *operations centers*, were established to oversee the operation of the elements of TNDS and ensure the integrity of the process. The allocation of functions among people and computers, together with a specification of how they interact to accomplish an overall process, was among the initial efforts in a general discipline known as *operations planning*.

Since then, the concept of operations planning has been expanded to encompass virtually all areas of telephone company operations. It is an organized, disciplined procedure to provide an integrated operations structure. This, in turn, will permit the Bell System to meet customer and market needs while minimizing the operational cost through applied computer technology. These efforts have produced a family of operations plans that guide AT&T, Bell Laboratories, and the telephone companies in the planning and deployment of the over 100 currently available OSs.[5]

Planning for the evolution of the network data collection process is now one element of an overall plan known as the Total Network

Operations Plan (TNOP). The planning process is not static; efforts are now under way to define the directions in which operations centers, systems, and processes should evolve to meet the future needs of the telecommunication network and the people who operate it.

3.2 Modeling the environment

A key factor that influences the plan for mechanized support of the network traffic data acquisition and analysis process is the telephone company operating environment. We can gain insight into the requirements imposed by this environment by analyzing models that represent typical situations and by conducting detailed field studies and experiments. These types of activities ensure that both the operations strategy and the operations system developments meet the needs imposed by both the evolving telecommunications and operations environments.

Many existing operating areas in the Bell System have similar characteristics and face similar network operations requirements. Thus, four basic types of model areas can be constructed to represent the various geographies, populations, and network equipment configurations that presently exist in the Bell System.

One, termed Model Area II, represents a large, densely populated metropolitan region, such as Detroit or Cleveland, with about two million main stations. Another, Model Area III, characterizes a state or portion of a large state containing over one million main stations and having an urban center with over two hundred thousand main stations. Indiana and Georgia resemble this model. Figure 1 shows maps of the geographic distribution of local switching equipment and key line operations centers in these two model environments. Table I shows a summary of physical plant statistics in each model environment. Of the two remaining models, Model Area IV characterizes large, primarily rural states, while Model Area I describes New York City and its immediate vicinity.

While line operations functions can be reasonably studied on an area basis, the span of other operations functions, such as trunking network design, often covers more than one operating area. Hence, the model areas were combined to form model BOCs. Figure 2 illustrates the typical deployment of selected operation functions for one such model company. It consists of one Model Area II and one Model Area III and resembles a single-state company like the Illinois Bell Telephone Company. Similarly, model companies can be combined to form a model territory to study toll operations.

These models provide a basis for quantitatively evaluating the effect(s) of the environment on alternative mechanization strategies. Among the key elements that influence the design and evolution of network data mechanization are:

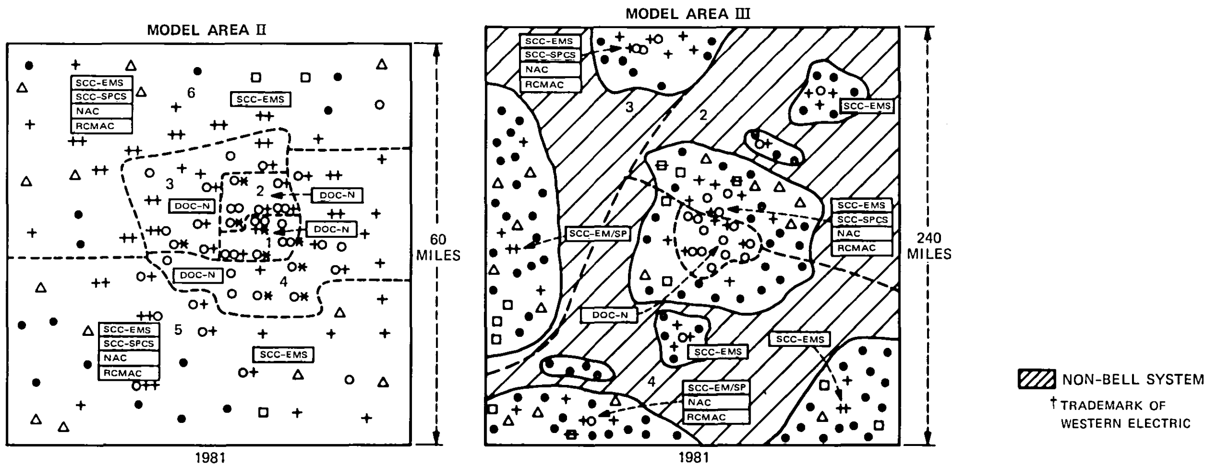

LOCAL SWITCHING PROFILE (Model Area II)

DISTRICT	1	2	3	4	5	6	TOTAL
SWITCH TYPE							
✱ 1XB	3	2	1	3			9
+ 5XB	2	3	12	12	17	19	65
○ 1 *ESS*†	8	8	8	10	7	1	42
△ 2 *ESS*					7	6	13
□ 3 *ESS*					1	2	3
● SXS					8	6	14

CHARACTERISTICS

3600 MI^2 AREA
URBAN AND SUBURBAN AREAS
ONE NPA
NO INDEPENDENTS

LOCAL SWITCHING PROFILE (Model Area III)

DISTRICT	1	2	3	4	TOTAL
SWITCH TYPE					
+ 5XB	4	10	11	12	37
○ 1 *ESS*	11	5	3	2	21
△ 2 *ESS*		5	5	5	15
□ 3 *ESS*		1	3	5	9
⊟ RSS		2		1	3
● SXS		18	26	45	89

CHARACTERISTICS

58,000 MI^2 AREA
TYPICALLY A STATE
TWO NPAs
MEDIUM-SIZE URBAN AREAS
LARGE AREAS SERVED BY INDEPENDENTS

SWITCH AND OPERATIONS CENTERS

DOC-N – DISTRICT OPERATIONS CENTER-NETWORK (LOCATED SCCs, NAC, RCMAC)
EMS – ELECTROMECHANICAL SYSTEMS
EM/SP – COMBINED FUNCTIONS
NAC – NETWORK ADMINISTRATION CENTER
NPA – NUMBERING PLAN AREA
RCMAC – RECENT CHANGE MEMORY ADMINISTRATION CENTER
RSS – REMOTE SWITCHING SYSTEMS
SCC – SWITCHING CONTROL CENTER (MAINTENANCE)
SPCS – STORED PROGRAM CONTROL SYSTEM
SXS – STEP BY STEP
XB – CROSSBAR

Fig. 1—Two representative model operating areas.

Table I—Physical plant statistics in two model areas

	Model Area II (1.93M Main Stations)		Model Area III (1.39M Main Stations)	
	No. of Entities	Main Stations (1000's)	No. of Entities	Main Stations (1000's)
Local switching				
No. 1 Crossbar	9	189	0	0
No. 5 Crossbar	65	738	37	450
Step-by-Step	14	28	89	281
1 *ESS**	42	904	21	542
2 *ESS*	13	65	15	97
3 *ESS*	3	6	9	18
RSS[1]	0	0	3	2
DCO[2]	0	0	0	0
DRSS[3]	0	0	0	0
Total local entities:	146		174	
Wire centers:	98		159	
Toll and local tandem switching				
4 *ESS*	1		1	
1 *ESS*[4]	—		6	
No. 4A Crossbar	3		2	
Crossbar Tandem	5		1	
No. 5 Crossbar with tandem features[4]	3		6	
SXS with tandem[4] features	—		7	
Trunks (1000's)[5,6]	162		118	
Special Service circuits (1000's)[6]	114		80	
T-Carrier channels (1000's)[6]	165		101	
N-Carrier channels (1000's)[6]	—		16	
High-Capacity Carrier channels (1000's)[6,7]	21		26	

1. Remote Switching System.
2. Digital Central Office.
3. Digital Remote Switching System.
4. Local switches with toll or tandem features are included in local entity count.
5. Interoffice circuits. Intraoffice circuits are not included.
6. Count includes intra-area plus one-half interarea circuits.
7. Includes low-capacity carrier systems multiplexed directly to broadband carrier.

* Trademark of Western Electric.

1. Number, type, and geographic distribution of data sources (i.e., switching systems)

2. The relative scope, number, and geographic distribution of operations centers that require mechanized support.

These factors, together with the nature of the functions performed by the various operations centers (see Section 3.3) and insights gained through detailed field studies, form the foundation for a network traffic data mechanization plan for the 1980s. Section 3.3 discusses the key network traffic functions in the BOCs.

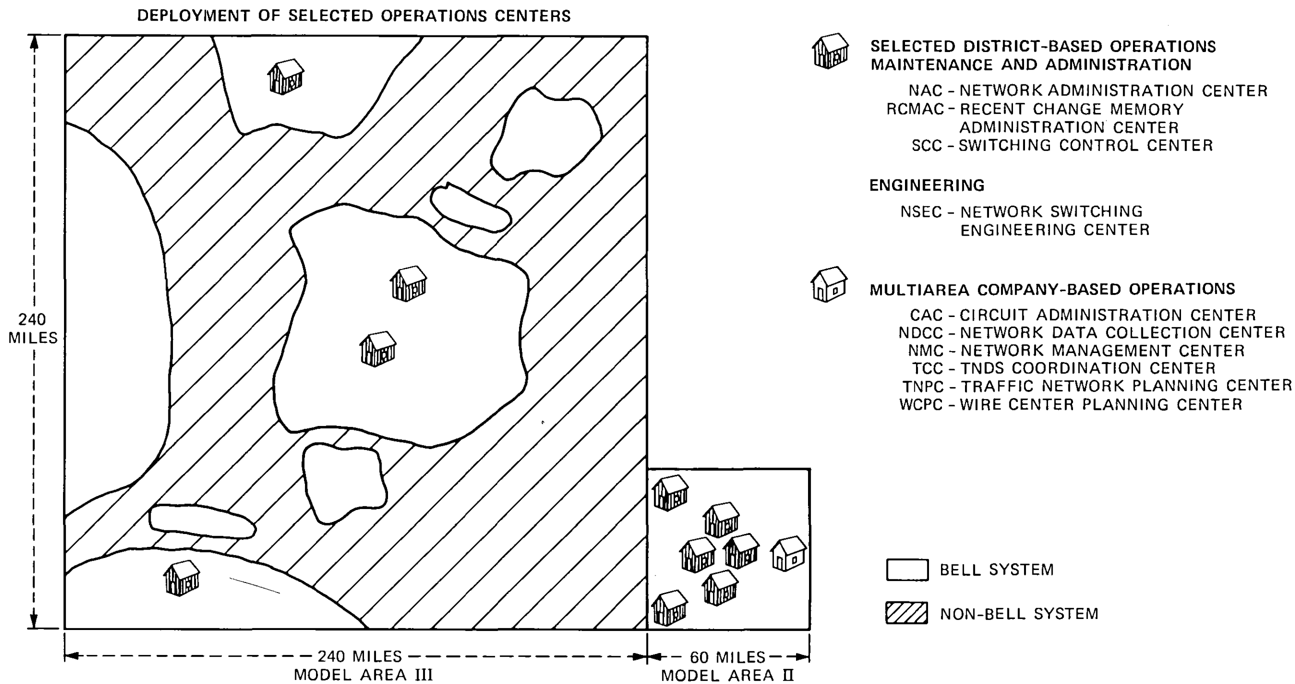

Fig. 2—A model operating company.

3.3 Managing, engineering, and planning the network

The BOC operations centers are responsible for network management, network administration, network design, and long-range planning. This section describes how the centers interact and how the information flow among centers is managed.

3.3.1 BOC Operations Centers

3.3.1.1 Network management. The Network Management Center (NMC) keeps the network operating efficiently when unusual traffic patterns or equipment failures would otherwise result in congestion. The NMC analyzes network performance and prepares contingency plans for situations such as peak days, telethons, and major switching system failures. The NMC also routinely monitors near-real-time network performance data to identify abnormal situations. Once an abnormal situation is detected, appropriate network management controls are implemented. When the critical situation has passed, the network management controls are removed. For further details, see the papers entitled "Network Management" and "National Network Management" later in this issue of the *Journal.*

3.3.1.2 Network administration. The Circuit Administration Center (CAC) and the Network Administration Center (NAC) perform network administration. The NAC is responsible for optimum loading, balancing, and utilization of installed central office equipment. It performs daily surveillance of central office and connecting trunk groups to ensure that service objectives are being met. In addition, the NAC reviews profiles of office load relative to profiles of anticipated capacity growth. It initiates corrective actions to deal with current as well as potential service problems by working with the Network Switching Engineering Center (NSEC) to initiate work orders to increase equipment in service (Section 3.3.1.3).

The CAC ensures that in-service trunks meet current as well as anticipated customer demands at acceptable levels of service. There are two activities, planned and demand servicing, that the CAC performs to discharge this responsibility. In planned servicing, the CAC compares current traffic loads with the forecasted loads for the upcoming busy season. If the loads are consistent, the CAC issues the orders to provide the forecasted trunks. When inconsistencies are found, the CAC examines the variation, develops a modified forecast for the busy season, and issues orders (if appropriate) based on the new forecast. In demand servicing, the CAC reviews weekly traffic data to identify trunk groups that imminently need augmenting beyond what the forecast states, and issues the appropriate trunk orders.

3.3.1.3 Network design. The network design function is performed by the CAC and NSEC work centers. The CAC projects current traffic loads for a one- to five-year period based on estimates of main station and call rate growth trends, and also develops corresponding forecasts of network trunk requirements. The NSEC is responsible for developing an analogous forecast of loads for traffic-sensitive switching equipment, setting office capacities, and determining relief size and timing.

The forecasts developed specify the amount of equipment (switching and trunking) that will be needed for the busy seasons of each of the next few years. Based on these forecasts, construction budget dollars are committed. The CAC and NSEC ensure that these forecasts are consistent with long-range planning.

3.3.1.4 Long-range planning. Long-range network planning is performed by work centers such as the Traffic Network Planning Center (TNPC) and the Wire Center Planning Center (WCPC).

The TNPC conducts studies to determine the most economic growth and replacement strategies for the network to meet its estimates of future demand. The estimates of future demand are based on current traffic load trends and marketing information. The plans developed start with the present environment and provide corporate guidance for future network configurations over a 20-year period. This planning process includes the tandem switching systems, operator services networks, trunks interconnecting all switching systems, and switching terminations to accommodate the trunks.

The WCPC conducts similar studies for the local wire center areas. The WCPC planning process includes the local switch and its interaction with other network elements (such as the subscriber network and interoffice facility network).

The basic output of the long-range planning function is a broad view of the future network topology. This includes the numbers, types, and locations of switching systems and the homing arrangements. The resulting long-range plan embodies various routing rules, such as whether local and toll traffic will flow through the same tandems in a metropolitan network and what sequences of alternate routes will be used. Such planning does not involve commitments to spend money but rather to ensure that the long-term consequences of current decisions are foreseen and that the evolution of the network proceeds smoothly and economically.

3.3.2 Work center interrelationships

Figure 3 illustrates how these work centers interact to perform network management, network administration, network design, and long-range planning.[6]

The TNPC and WCPC provide the CAC and NSEC with the fundamental office and network evolution plans (20-year horizon).

The CAC and NSEC review current trends in traffic offered to in-service network and switching equipment. They develop forecasts of specific equipment needs for the next one to five years, consistent with the fundamental plans. The forecasts result in construction budget allocations and equipment orders.

The CAC and NAC review the current level of traffic offered to in-service equipment. The CAC and NAC identify current (within a year) equipment rearrangement and growth needs. Accordingly, equipment is rearranged, ordered, and installed.

The NMC, with the help of the CAC and NAC, prepares contingency plans for situations such as peak days, telethons, and major switching system outages. The NMC routinely monitors the load, in near-real time, to identify unusual traffic patterns and equipment failures that have resulted or will result in congestion. The NMC will implement the network management plans, as required, to deal with the problems.

As an example of the interrelationships of these centers, Table II outlines the role of the NMC, NAC, and CAC relative to the surveillance of network service. The primary differences that can be noted are: the time frame of interest and action, the geographic domain (network), and the network components of interest.

Table II—Summary of relative service surveillance roles in BOC Operations Centers

Center	Surveillance Objective	Network Traffic Data Used
Network Management Center	Monitors traffic congestion in a portion of the network (e.g., a numbering plan area) and initiates controls to maximize call completion during times of overload or equipment failure	5- to 20-minute traffic data and 30-second network status data for selected central office equipment and trunk groups in selected central offices
Network Administration Center	Monitors load and service status for each switching system and its trunk groups in a central office district, determines if service objectives are being met, detects or is informed of potential or actual service-affecting problems, initiates corrective action when necessary, and verifies that problems are being resolved	Hourly and weekly central office equipment and trunk group traffic data and selected status indicators
Circuit Administration Center	Monitors traffic loads on the message trunk network in an operating area or company, determines the need for near-term trunking rearrangements and additions to resolve conditions that are network service-affecting and that are expected to persist	Weekly and longer-term summaries of trunk group traffic data

3.3.3 Management of information flow

Coordinating the flow of information to and among the work centers is a complex task for which work centers have been established. The Network Data Collection Center (NDCC) coordinates schedules and ensures that the requested data are collected and distributed appropriately (i.e., in the most timely and cost-effective manner). Operations systems have been developed for internal operations and information exchange. The NDCC coordinates the execution of these systems on a day-to-day basis.

Up to this point, we have discussed the network and the BOC operational environment. Now let us turn our attention to the functional objectives of a data system to mechanize the network data process.

IV. MECHANIZATION OBJECTIVES

Until the early 1970s, the collection and processing of network traffic data by the BOCs was a combination of manual and semi-mechanized processes. A number of environmental changes made this type of network data environment inadequate to meet the needs of the business. These influences included:

1. Growing sophistication and complexity in an increasingly SPC network
2. Need for immediate information for network management and for more complete and timely information for all network-related functions
3. Increasing regulatory and competitive pressures.

Accordingly, the trend was toward mechanizing the network data process. The overall objective of this effort has been to automate the process of collecting and summarizing network traffic data so that BOC decision makers have adequate quantities of timely and accurate information to administer and engineer the network. To be effective, the system must be robust to varying BOC environments (e.g., organizational responsibilities, mechanization deployment) and must be expandable to meet the growing processing needs of the BOCs. The system must offer a manageable and cost-effective implementation for the BOCs in the context of their operations.

Before describing the implementation of the Bell System's computer-based network data system, let us look at the basic data function objectives that such a system must meet to satisfy the needs of users performing key BOC traffic functions.

4.1 Data collection

After the traffic measurements have been taken by the switching system or by special measurement equipment located near the switch, it is desirable to collect the measurements at a centralized location where large-scale data processing can be brought to bear. The system must be capable of collecting network traffic data from all types of switching entities—EM and SPC systems, local, tandem, and toll switches.

Because the equipment collecting the data is centralized and has access to the switching systems it serves, it is practical to delegate measurement control (e.g., turning on a TUR to begin taking measurements on the basis of a collection schedule) and network management control (e.g., signaling to a switching system to cancel alternate routing) to the same collection equipment.[4]

4.2 Data administration

The system must also support a number of "data administration" activities. Measurements to be collected from switching systems must be scheduled so they are not collected during periods when they are not required (e.g., during low traffic periods). In addition, the system must allow users to control the data processing schedules (e.g., hours of the day for which the data collected are to be summarized into the various reports).

The measurements must also be linked to records that associate each measurement with the equipment being measured. In addition, records of the individual switching system characteristics (e.g., inventories of installed central office equipment) must be maintained. No less important than the network traffic data, complete and accurate records are absolutely required by the system to transform raw traffic measurement data into useful information. The generation and maintenance of these records are major tasks that must be supported by the system.

At times, more data are collected than are actually needed due to the way measurements are accumulated and "packaged" by the switch or by the special measurement equipment. The system should identify and remove these unneeded data as early as possible to avoid wasteful processing.

It is necessary for the system to validate data (i.e., inspect them to determine if they truly reflect network traffic conditions that they are intended to measure) soon after collection to avoid processing incorrect data and allowing them to contaminate the associated good data.

As a final data administration task, the system must ensure that measurement data are distributed as appropriate between the different processing elements of the system as well as to any other Operations Systems that require those data.

4.3 Report generation

To present BOC users with useful information, the system must perform a number of processing and reporting tasks. As Section 3.3 showed, there are four key traffic-related functions that depend on network traffic data. A primary factor in designing a successful network data system is satisfying the needs of the BOC people performing these functions—both in the type of information reported and the timeliness with which it is made available.[4]

Reported information should be in formats that allow users to quickly identify and interpret the information. The system should provide a set of predefined, standard reports as well as capabilities for selective, demand query of information to allow users to obtain information in formats customized to their special needs. In addition, the system must allow the users to control the routing of processed outputs and must be able to deliver processed results to users in the form to be used in their local operations (i.e., paper reports, microfiche reports, hard copy or cathode ray tube terminals, status board).

Automatic, more detailed, validation should be performed to augment that carried out shortly after data collection. Validation reports should be generated to help users detect suspect data. In addition, the system must allow users to highlight and exclude information that is judged suspect so that it is appropriately treated in affected reports.

The system should provide access to network information in time frames consistent with the spectrum of user functions (see Fig. 3). For instance, the system must provide five-minute network load information in an immediate time frame as the basis for real-time management of the trunk network. However, the system must summarize and maintain information for a year or more to support central office and facility engineering and long-range planning. In general, the system should store data and results for periods of time based on anticipated information usage patterns and economic considerations.

The system must collect and report information on its performance (e.g., data collection availability) to allow the BOCs to promptly identify and resolve system performance problems and to otherwise manage the system.

4.4 Robustness

A mechanized network data system must continue to evolve as new switching machines are introduced, as new network services and features demand the processing of new measurements, as new provisioning or network design algorithms are formulated, and as new functions arise for the network data system. In addition to simply evolving to meet new processing needs, the network data system must do so at or near the time of introduction of new technology to help minimize the impact on Bell System operations.

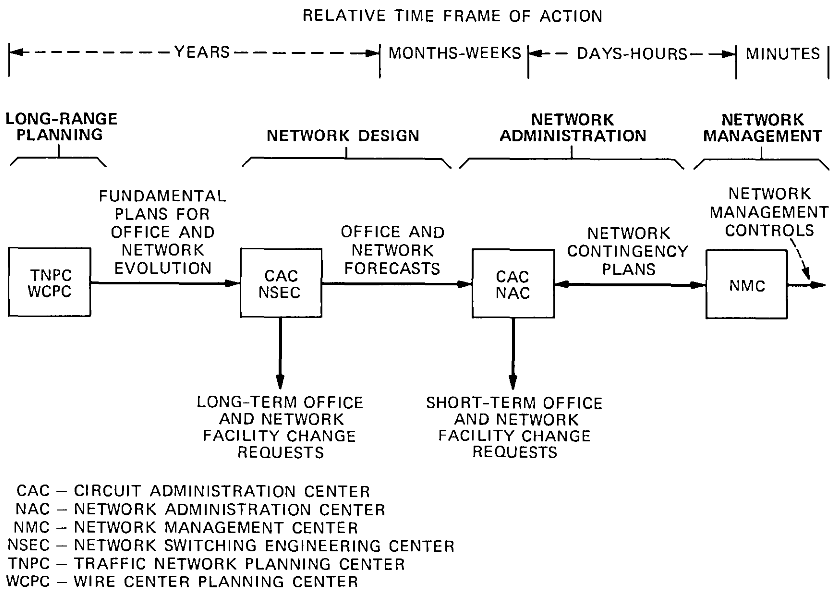

Fig. 3—Work center interrelationship.

V. CONCLUSION

The Bell System has implemented a mechanized process known as the Total Network Data System to meet the needs we have outlined. This TNDS consists of a set of subsystems, each performing part of the overall process. The formulation of and adherence to a system plan has helped assure that all parts of TNDS have been available when required to support the network management, administration, design, and long-range planning needs of the BOCs. The paper that follows discusses this "TNDS System Plan" and the component TNDS subsystems that have evolved.

REFERENCES

1. *Engineering and Operations in the Bell System,* Murray Hill, N.J.: Bell Laboratories, Indiana Publication Center, 1977.
2. David Talley, *Basic Telephone Switching Systems,* Rochelle Park, N.J.: Hayden Book Company, 1969.
3. M. M. Irvine, "An Electronic Watchdog for the Network," Bell Lab. Rec., *58,* No. 8 (September 1980), pp. 267–73.
4. M. M. Buchner and W. S. Hayward, "Total Network Data System," Proceedings of the Eighth International Teletraffic Congress, November 1976, Melbourne, Australia.
5. M. M. Buchner, Jr., "Planning for the Evolving Family of Operations Systems," Bell Laboratories Record, *57,* No. 5 (May 1979), pp. 118–24.
6. R. F. Grantges, G. W. Riesz, and B. E. Snyder, "Evolution of the Total Network Data System," Proceedings of the Ninth International Teletraffic Congress, October 1979, Porremolinof, Spain.

AUTHORS

Anthony L. Barrese, B.E., 1970, M.E.E., 1971, Ph.D., 1978, Stevens Institute of Technology; Bell Laboratories, 1969—. Mr. Barrese was first involved in the development and testing of the Safeguard antiballistic missile system. He then worked as a systems engineer, developer, and tester on the Total Network Data System project. He supervised a group with methods development responsibilities for the Circuit Administration Center. He is currently Supervisor, System Architecture Group, with long-range planning responsibilities for the Operations Systems Network. Member, IEEE, Tau Beta Pi.

Donald E. Parker, Southern Bell Telephone and Telegraph Company, 1967–1968; South Central Bell Telephone Company, 1968–1976; Bell Laboratories, 1976–1979; American Telephone and Telegraph Company, 1979—. Mr. Parker's involvement with TNDS began in 1974 when he was assigned responsibility for training and implementation of Total Network Data System/Equipment (TNDS/EQ) in the Alabama Area of South Central Bell. It has continued through a rotational assignment in the Personnel Subsystems Group at Bell Laboratories, where he was responsible for user documentation and consultation for Traffic Data Administration System (TDAS), Common Update (CU), and Individual Circuit Analysis Program (ICAN). He is currently assigned to the Network Services Department of AT&T as Project Manager for TNDS/EQ and Stored Program Control Systems Central Office Equipment Reports (SPCS COER).

Thomas E. Robbins, B.S., 1968 (Mathematics), Westminster College; M.S., 1971 (Operations Research), New York University; Bell Laboratories, 1968—. Mr. Robbins has been involved in various aspects of the Total Network Data System since 1972, including both development (SPCS COER) and systems engineering (SPCS COER, TNDS/EQ) assignments. He is presently Supervisor of the TNDS Planning Group responsible for planning the feature evolution of the TNDS elements to support new network switches and services. Member, IEEE.

Larry M. Steele, B.S.E.E., 1969, University of Kentucky; M.S.E., 1971, Princeton University; AT&T, 1975–1977; Bell Laboratories, 1969–1975, 1977—. Mr. Steele joined Bell Laboratories in 1969 and employed modeling and simulation techniques to evaluate the accuracy and vulnerability of proposed military missile detection systems. In 1971 he developed the requirements and the software for a system to analyze the performance of No. 1 Crossbar switching entities. He was subsequently given responsibility for the determination of downstream analysis requirements and the formulation of an administrative plan for the individual circuit usage recording feature of a standard Bell System traffic data acquisition system. In 1975 Mr. Steele went on a rotational assignment with the General Departments of the American Telephone and Telegraph Company, where his responsibilities included project management and methods development for a family of traffic data acquisition and analysis systems. Upon his return from AT&T in 1977, he conducted a comparative analysis of system architectures for providing a standard methodology for gathering detailed data relating to point-to-point traffic loads. Between 1978 and early 1981, Mr. Steele supervised a group responsible for studying and recommending improved methodologies and mechanized capabilities in the area of telephone network administration. The group was also responsible for synthesizing administrative plans for incorporation into the Total Network Operations Plan (TNOP). Currently, Mr. Steele is the Head of the Network Operations Planning Department, which is responsible for planning for the integration and use of computer-based systems to support network-related operations functions.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

System Plan

By M. S. HALL, JR.,* J. A. KOHUT,* G. W. RIESZ,* and J. W. STEIFLE†

(Manuscript received July 23, 1982)

The Total Network Data System (TNDS) is a coordinated family of operations systems that work together to mechanize telephone traffic data gathering and reporting. This paper describes the component parts of the TNDS, their functions, their interactions, and their inputs and outputs. The paper also presents the managerial view of TNDS: its product, its size, the number of people involved, how the overall project is organized and coordinated, and how the system evolves to meet the needs of a continually changing environment.

I. INTRODUCTION

The planning effort for traffic data collection and processing for what was to become the Total Network Data System (TNDS) began in the mid-1960s. This plan gradually evolved as the initial steps were taken to mechanize portions of the traffic data collection and data processing functions for the Bell Operating Companies and AT&T Long Lines.

This paper is divided into two parts. The first part, Sections II and III, presents an architectural view of the Total Network Data System.

---

* Bell Laboratories. † American Bell.

The second part of the paper (Section IV) is written from a managerial point of view.

The first part begins (in Section II) with an overview of the current configuration of the TNDS, and the decisions that led to the current architecture; it concludes with brief descriptions of the functions of the individual subsystems and summaries of their interactions. (Subsequent papers in this volume describe the subsystems and their internal operations in detail.) Section III describes TNDS operations from a different point of view, i.e., by following the flow of a representative data item through the system and describing the database update process.

The second part, which begins with Section IV, describes the product delivered to the end users, including software, documentation, and user support, as well as the organization of the development and maintenance effort, and the planning, requirements, development, test, and delivery cycle. Section IV also examines the need to enhance and update TNDS subsystems to meet changing needs, using the integration of the 5 *ESS** switching system as an example. This article ends with a brief look at possible future directions.

II. CURRENT CONFIGURATION

This section introduces the architecture of the Total Network Data System and briefly describes the individual component systems and their interactions. Subsequent articles in this issue of the *Journal* describe the elements of TNDS in greater detail.

TNDS is now widely deployed throughout the Bell System. All or a portion of TNDS is now in use in every Bell operating company. Of the almost 10,000 switching entities, TNDS collects and processes data for more than 7000. Since the switching entities not now connected are primarily the small electromechanical Community Dial Offices (CDOs), TNDS actually handles more than 90 percent of all Bell System traffic information.

TNDS is a coordinated family of operations systems, which work together to mechanize the data gathering and reporting process. It consists of manual procedures and computer systems that enable operating company managers, network administrators, and network designers to analyze traffic data in a variety of ways.

The development of TNDS has been under way since the mid-1960s. It was the Bell System's first set of integrated operations systems, and has required the planning, development, testing, modification, and introduction of a number of new features and generic programs. This

* Trademark of Western Electric.

systems planning served as the forerunner of the main operations planning activities that have since characterized several plans, including the Total Network Operations Plan (TNOP).

TNDS has now evolved into a system that encompasses 13 major component systems. They are:

- EADAS—Engineering and Administrative Data Acquisition System
- EADAS/NM—Engineering and Administrative Data Acquisition System/Network Management
- TDAS—Traffic Data Administration System
- LBS—Load Balance System
- 5XB COER—No. 5 Crossbar Central Office Equipment Reports
- SPCS COER—Stored Program Control Systems Central Office Equipment Reports
- ICAN—Individual Circuit Analysis
- SONDS—Small Office Network Data System
- CU/EQ—Common Update/Equipment
- TSS—Trunk Servicing System
- TFS—Trunk Forecasting System
- CU/TK—Common Update/Trunking
- CSAR—Centralized Systems for Analysis and Reporting.

Figure 1 shows the overall TNDS architecture.

The elements of TNDS are operated in distributed locations. The switching systems, each serving thousands of customers, are generally based in a community's central office. The data collection systems typically employ dedicated minicomputers and gather traffic information from anywhere between 20 and 80 switching systems. Most of the remaining engineering and administrative reporting systems run on either general-purpose operating company computers, or at the AT&T computer center.

Once collected and stored, a number of different systems of the TNDS may process the same items of traffic data to produce a variety of reports.

Despite these distributed operations, the TNDS for any operating company can be thought of as a single, integrated entity designed to provide managers with comprehensive, timely, and accurate information about the network.

2.1 Architectural overview

The Total Network System did not, like Athena from Zeus, spring full blown from the head of a grand planner. Instead it evolved, like most things, from a combination of several individual system developments. R. L. Martin has observed[1] that the architecture of a system

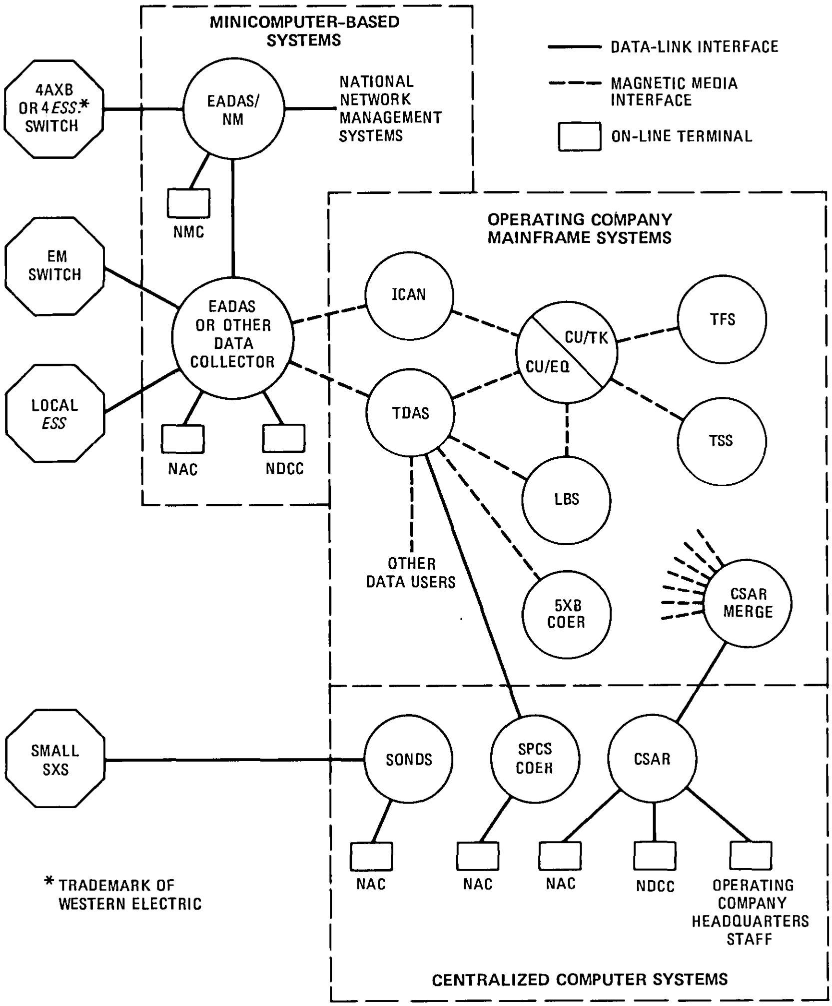

5XB COER – NO. 5 CROSSBAR CENTRAL OFFICE EQUIPMENT REPORTS
CSAR – CENTRALIZED SYSTEMS FOR ANALYSIS AND REPORTING
CU/EQ – COMMON UPDATE/EQUIPMENT
CU/TK – COMMON UPDATE/TRUNKING
EADAS/NM – ENGINEERING AND ADMINISTRATIVE DATA AQUISITION SYSTEM/NETWORK MANAGEMENT
ICAN – INDIVIDUAL CIRCUIT ANALYSIS
LBS – LOAD BALANCE SYSTEM
NAC – NETWORK ADMINISTRATION CENTER
NDCC – NETWORK DATA COLLECTION CENTER
NMC – NETWORK MANAGEMENT CENTER
SONDS – SMALL OFFICE NETWORK DATA SYSTEM
SPCS COER – STORED PROGRAM CONTROL SYSTEMS CENTRAL OFFICE EQUIPMENT REPORTS
SXS – STEP BY STEP
TDAS – TRAFFIC DATA ADMINISTRATION SYSTEM
TFS – TRUNK FORECASTING SYSTEM
TSS – TRUNK SERVICING SYSTEM

Fig. 1—TNDS architecture.

is the product of the history of the organization that builds it, the present and near-present technology, and the intended application. This observation is certainly true for TNDS. Its roots reach back into the 1960s. Its components originated in a number of organizations at

Bell Laboratories and the operating companies. The uses of its output within the operating companies are widespread and diverse. It is, therefore, not surprising that the TNDS architecture (shown in Fig. 1) is not uniform or highly integrated. It is perhaps amazing that it fits together as well as it does.

2.1.1 The primal components

One of the earliest threads leading to TNDS began in 1966. A group of planning department representatives from eight associated companies, along with Long Lines and AT&T, met in New York City and recommended a Bell System effort to centralize the development of a computerized trunk facilities system. Among the components recommended for development were a Traffic Trunk Estimating Subsystem and a Traffic Trunk Servicing System. Development of these two components began shortly thereafter in the Business Information Systems (BIS) area of Bell Laboratories as the Trunk Forecasting System (TFS) and the Trunk Servicing System (TSS). A third program, designed to acquire data from a variety of manual and mechanized sources and to distribute it to client programs, was also included in the planning. This program, called Identify and Edit, later became TDAS.

Identify and Edit served as a data "warehouse." It labeled traffic measurements, stored them, and distributed batches of data to downstream users as they were needed. It and the downstream systems it fed were designed to serve entire operating companies, and in some cases the entire Bell System. These downstream systems were placed on large mainframe computers to take maximum advantage of economies of scale in processing, and to reduce the administrative problems of operating multiple systems. Another factor that led to these systems being company-wide in scope was the need for TFS to design company-wide trunk networks. This capability required a centralized database for trunk traffic measurements. It led to the centralization of the downstream systems with that database. The database, to which was added data for central office equipment, later became the CU/EQ and CU/TK components.

During this early period, the place later taken by EADAS (see Fig. 1) was occupied in planning by the Traffic Data Recording System (TDRS). TDRS was a special-purpose system that collected traffic data from electromechanical switching offices. Also during this period, AT&T planners expected to meet central office equipment data needs by standardizing programs developed by various telephone companies for individual switch types. The programs would receive data from Identify and Edit. Chief among these "Central Office Equipment" programs was one developed by New Jersey Bell for No. 5 Crossbar offices (it was running in 1968). Later standardized by AT&T and still later revised by Bell Laboratories, this program became the present-day 5XB COER component of TNDS.

The basic allocation of processing functions among these early systems has prevailed until now. TDRS was responsible for the real-time collection of traffic measurements for about 40 electromechanical switching offices. This number of offices was selected to balance the cost of switch-to-TDRS data links with the economics of scale available with centralized processing functions. This basic configuration and size have continued with little change even though the role of TDRS has been taken over by EADAS.

The early downstream systems (TSS, TFS, Identify and Edit, and 5XB COER) were designed for batch processing on mainframe computers. This approach was the prevalent technology for large software systems in that era. It provided an economical way to meet user needs with minimal technical risks. Quick turnaround was not a requirement on the reports produced by these systems. These early systems have evolved and been enhanced over the last fifteen years. LBS and ICAN have been added. However, the fundamental architecture established in the late 1960s for this set of operating company, mainframe-based systems has changed little. Only now are on-line, real-time features being added.

The arrangement of Central Office Equipment Reports (COER) systems within the TNDS architecture was driven by different factors than the other downstream systems. Each type of switching equipment has unique traffic data processing requirements. The switch architectures and service features strongly influence these requirements. This characteristic led to separate COER developments for each switch type. Only recently has the technology needed for a generic COER been pursued.[2] Several approaches were taken by the various organizations that developed the COER modules. As discussed above, 5XB COER is implemented on a mainframe system, and SPCS COER (which is a collection of COER modules) is implemented on a centralized, time-shared system (as is SONDS). The selection of time sharing for these systems is discussed below.

2.1.2 Evolution and integration

In the early 1970s economical new minicomputer systems with greater flexibility and versatility replaced the specialized TDRS hardware being used for data collection. Some locally developed and vendor-supplied systems were installed for this purpose. The predominant system in use, however, is EADAS, which is a product designed by Bell Laboratories. Being real-time systems, these minicomputer systems could easily provide the short-term traffic reporting functions that were also required.

Another major evolutionary step was the addition of network management functions in 1975. Network management requires near-real-time data processing and requires an overview of a larger number of switching offices than can be served by one EADAS. However, the EADAS system provided a convenient and economic concentration point for the real-time, local office, traffic measurements needed for network management. Therefore, a decision was made to create a hierarchical architecture with up to six EADAS data collection minicomputers concentrated on an EADAS/NM minicomputer. For 4A Crossbar and 4 *ESS* toll switching offices, however, a direct switch-to-EADAS/NM connection was chosen because the number of these toll offices is relatively small and the data volumes interchanged with them are relatively high. It also provided a higher level of reliability for the communications with these critical switching entities.

Three systems—SPCS COER, SONDS, and CSAR—were implemented on Bell System national-based time sharing systems. All three systems began as experiments conducted at Bell Laboratories to investigate improved traffic data management methods. For SPCS COER and CSAR, centralized time sharing was a means to deploy required on-line processing functions without the need for operating companies to buy additional hardware, or to support software in a number of geographically dispersed sites. These are attractive attributes for an experimental prototype.

SPCS COER and CSAR were standardized on centralized time sharing because it was easier and more economical to evolve the prototype software than it would have been to rewrite it for another computer system and perhaps purchase new hardware. The decision to deploy SONDS on centralized time sharing was more difficult, however.

The SONDS prototype was implemented on a minicomputer system at Bell Laboratories. It was determined that a single computer installation could economically support all the Bell System's small step-by-step offices expected to utilize SONDS. These electromechanical switching offices were expected to be gradually replaced by electronic offices (served by EADAS) over a fifteen- to twenty-year period. The decision therefore was between: (1) standardizing the minicomputer prototype and supporting a centralized, dedicated minicomputer for 15 to 20 years of declining switch population; or (2) rewriting all the prototype software for a general-purpose, time-sharing system. It was decided that users could get better support if the system was on a general-purpose computer. So SONDS joined CSAR and SPCS COER on the AT&T time-shared computer facility.

Integration of TNDS occurred as a series of steps. An early step was linking data collection to downstream processing. Minicomputer-based data collection systems like EADAS wrote traffic measurements on magnetic tapes that could be used as direct input to TDAS. The tape formats used for this link have been changed to improve processing efficiency, but the basic arrangement is still used. Integration of SPCS COER was another major step. Initially SPCS COER received data from punched paper tape via dial-up connections. In 1975, EADAS was upgraded to collect 1 *ESS* traffic data and pass it to TDAS via magnetic tapes. Data transmission from TDAS to SPCS COER was then used to complete the link. Other key integration steps were the automatic input to TFS of trunk group base load data from TSS, and linking the standard CSAR system by means of the downstream CSAR merge program. As part of their initial development, EADAS/NM, ICAN, and LBS were all integrated. TNDS was officially recognized as an integrated system for planning purposes in 1974. This was accomplished by an internal Bell Laboratories memorandum which, in fact, declared it an integrated system and documented its architecture.

The slow pace of evolution from the early architectural choices attests to their workability, but also is evidence of how difficult it is to change established architectures. As discussed at the end of this article, more changes are occurring. They are driven by the needs of users for new processing functions and technological advances that can be effectively incorporated into TNDS. Recently, developers have explored: (1) the use of on-line functions in mainframe systems for database management, report distribution and documentation delivery; and (2) the migration of some functions on centralized time sharing to on-line mainframe systems.

The remainder of this section describes, in more detail, the existing component systems and their current roles within the TNDS architecture.

2.2 Component system functions

To understand the interrelationship of TNDS systems and how they collect and use traffic data, let's look at the four primary functions of TNDS in their general order of occurrence—data acquisition and management, central office equipment reporting, trunk network reporting, and system performance measurement—and at the systems responsible for each function.

2.2.1 Data acquisition collection and management

Managers need data on network performance and traffic loads carried by trunk groups and switching systems to assess the quality of service and to plan for network growth.

This network information begins as bits of data collected in switching offices. For example, in electromechanical offices a specialized data collection device, called a Traffic Usage Recorder, scans the trunks and other components periodically (typically every 100 seconds) and counts how many are busy. Other traffic data such as peg count and overflow are also collected. In *ESS* offices, a similar process is used, but there is no need for specialized equipment, since traffic data are collected by the switching system's central processor.

2.2.1.1 EADAS. Traffic data are transmitted to the first of the 13 systems of TNDS—the Engineering and Administrative Data Acquisition System (EADAS). EADAS runs on dedicated minicomputers located at an operating company's Network Data Collection Centers. As its name suggests, it is the major data-collecting TNDS system. However, a few operating companies use locally developed or non-Bell vendor-supplied systems instead of or in addition to EADAS for this first step in the TNDS process. In addition, the large toll machines, 4 *ESS* and No. 4 Crossbar [with its associated computer, the Peripheral Bus Computer (PBC)] collect their own data and do not connect to EADAS.

Each EADAS serves up to 80 switching offices and, upon receiving traffic data, performs three basic functions.

1. It processes some data in near-real time (shortly after they are received) to provide hourly and half-hourly reports and a short-term database for network administrators.
2. It collects and summarizes data for processing by remaining "downstream" TNDS systems.
3. It performs on-line surveillance and reporting functions.

EADAS also links other TNDS systems by forwarding traffic data to them via a data link or magnetic tape. Three systems receive these data directly from EADAS.

2.2.1.2 EADAS/NM. Two of three direct recipients of traffic data from EADAS are the Individual Circuit Analysis (ICAN) Program and the Engineering and Administrative Data Acquisition System/Network Management (EADAS/NM). They are the only systems that use data collected and processed by EADAS without first having it formatted by TDAS. ICAN is not a data acquisition system, but rather one of the central office engineering and administrative reporting systems, and is described in that section. TDAS is the third direct recipient.

EADAS/NM is located at operating company Network Management Centers, and uses data received from EADAS or directly from some types of switching systems (e.g., 4A and 4E). It watches switching systems and trunk groups designated by network managers, and reports existing or anticipated congestion on a display board at the center. EADAS/NM also provides information to AT&T Long Lines'

Network Operations Center at Bedminster, N.J. This center is supported by an operations system, the Network Operations Center System (NOCS), which performs functions similar to EADAS/NM, but on a national scale.

2.2.1.3 TDAS. The Traffic Data Administration System (TDAS) formats traffic data for use, in turn, by a number of the other "downstream" systems.

But unlike EADAS and EADAS/NM—both of which employ dedicated minicomputers—TDAS runs in an operating company computer center supporting the Network Data Collection Centers. If TNDS were to be pictured as a series of distinct steps, then TDAS can be thought of as the second data acquisition step. Each week it accepts millions of pieces of data from EADAS (or the equivalent of locally developed or vendor-supplied systems) or directly from the large toll machines. In response to requests from downstream systems users, TDAS sorts, labels, stores, and, at the appropriate time, provides the data in the proper format to the engineering and administrative reporting systems. In effect, TDAS acts primarily as a warehouse and distribution facility for traffic data.

TDAS treats its data acquisition job as a basic order/inventory problem. Orders, or Traffic Measurement Requests, for data are manually prepared and sent to TDAS by operating company personnel. These orders are stored in a master database.

As traffic data are received by TDAS, they are matched against the orders held in Common Update. When the data necessary to fill an order have been received, a weekly data summary—either printed or on magnetic tape—is sent to the system that requests it to use in preparing an engineering or administrative report.

Once TDAS has received, formatted, and sent data to the appropriate downstream reporting system, the data acquisition function is complete. At this point, traffic data have moved from the switching system to EADAS and then directly to TDAS and the requesting engineering and administrative reporting system. The next step is to look at how managers employ the TNDS systems to analyze the data that have been gathered.

2.2.1.4 CU/EQ. In addition to the above systems, a common record base is used. TDAS and several of the downstream engineering and administrative systems need much of the same record-base reference information. Those data are maintained in a common system, called Common Update/Equipment (CU/EQ), rather than duplicated in each system. The information for each central office includes the configuration of switching equipment as well as specifications of traffic registers. This database is updated as changes occur in the physical arrangements of central offices.

2.2.2 *Central office equipment reporting*

There are five downstream TNDS engineering and administrative systems that send reports about central office switching equipment to operating company personnel. These systems, which run on either operating company computers or at the AT&T computer center, are the:

1. Load Balance System (LBS)
2. No. 5 Crossbar Central Office Equipment Reports (5XB COER)
3. Stored Program Control Systems Central Office Equipment Reports (SPCS COER)
4. Individual Circuit Analysis (ICAN) Program
5. Small Office Network Data System (SONDS).

The first three receive traffic data from TDAS. ICAN receives data directly from EADAS, but also uses Common Update for some of its reference information. SONDS collects its own data directly from small step-by-step offices.

2.2.2.1 LBS. The Load Balance System, LBS, which is run on the operating company computer, helps assure that the customer traffic load is uniformly distributed over each switching system. Customer lines are connected to the switching system "concentrators," which allow customers to share switching equipment. LBS analyzes traffic data coming to it from TDAS to establish the traffic load on each line group of each switching system. Then, personnel in the Network Administration Center use LBS reports to determine "lightly loaded" line groups to which new subscriber lines can be assigned. This minimizes congestion on a given concentrator. In addition, LBS calculates "load balance indices" for the entire operating company, indicating how effectively each central office has avoided congestion by efficiently distributing traffic.

2.2.2.2 5XB COER, SPCS COER. Two other central office equipment reporting systems—5XB COER for No. 5 Crossbar offices, and SPCS COER for 1, 2, 3, and 5 *ESS* offices—also use traffic data collected by EADAS and formatted by TDAS to support the Network Administration Center. While 5XB COER runs on operating company computers and SPCS COER at the AT&T computer center, the two systems perform similar functions. Both analyze traffic data to indicate the overall load carried and the service provided by the switching systems, and to determine how much of the switching system's capacity is being used. This information helps planners decide when new equipment is needed and answers many other important administrative questions.

2.2.2.3 ICAN. The two remaining central office equipment reporting systems, ICAN and SONDS, do not use TDAS to format their data. ICAN, which receives data directly from EADAS for certain electromechanical offices equipped with an EADAS option called Individual Circuit Usage Recording (ICUR), detects switching system equipment faults by identifying abnormal load patterns on individual circuits. ICAN then produces a series of reports used at the Network Administration Center to analyze individual circuit usage and verify circuit grouping.

*2.2.2.4 **SONDS.*** SONDS—the fifth downstream central office equipment reporting system—is the only TNDS system that performs a full range of data manipulation functions. It economically provides a number of TNDS features for the smaller electromechanical step-by-step offices.

To do that, SONDS—which runs at the AT&T computer center—collects traffic data directly from offices measured, processes them, and automatically distributes weekly, monthly, exception, and on-demand reports to managers at the Network Administration Center via dial-up terminals.

2.2.3 Trunk network reporting

Three TNDS systems—all run on operating company computers—support trunk servicing and forecasting at the Circuit Administration Center:

1. Trunk Servicing System (TSS)
2. Trunk Forecasting System (TFS)
3. Common Update/Trunking (CU/TK).

Together these three systems are sometimes referred to as Total Network Data System/Trunking (TNDS/TK). TSS helps trunk administrators develop short-term plans to relieve unacceptable blocking on final trunk groups. It processes traffic data supplied by TDAS, and computes the "offered load" for each trunk group. TSS calculates the number of trunks theoretically required to handle that traffic load at a designated grade of service—an objective expressed in terms of percent of blocking. TSS produces weekly reports showing which trunk groups are underutilized and which trunk groups are performing below the grade-of-service objective.

The traffic load data calculated by TSS also help support the trunk forecasting function performed by TFS. Those data, in conjunction with information on network configuration and forecasting parameters stored in Common Update, are used for long-term construction planning. The Trunk Forecasting System uses that information to forecast message trunk requirements for the next five years. These forecasts are a fundamental input to the planning process, which leads to the provisioning of additional facilities.

CU/TK provides a common record base for TSS/TFS. It contains information describing the configuration of the trunk network. The traffic data are provided from TDAS.

2.2.4 CSAR system performance measurement

Centralized System for Analysis and Reporting (CSAR) is the newest TNDS system. It was designed to monitor and measure how well data are being processed through TNDS from beginning to end. CSAR quantitatively measures the accuracy, timeliness, and completeness of the data flow for operating company personnel at the Network Data Collection Center, Network Administration Center, and the Circuit Administration Center. It also supplies sufficient information to locate trouble so corrective action can be taken. It does not presently analyze data from EADAS/NM, SONDS, or TFS.

CSAR is an on-line, interactive system that puts TNDS performance information at the user's fingertips. At the conclusion of each run of a system in TNDS, data required by CSAR are placed in a special file. Then, at an appropriate time, they are transferred to the AT&T computer that contains the CSAR program. CSAR then performs the proper association and analysis of data.

Operating company managers can access the report information from terminals at their own work locations. The reports can be arranged in a number of formats, and can provide details on overall TNDS performance or individual system effectiveness. Reports can be broken down by traffic unit (trunk group or switching system), district, division, or area to help identify and resolve problems at various operating company organizational levels.

III. HOW IT WORKS

This section will describe the end-to-end data flow of TNDS using two examples. One example will explain the steps necessary to pass *Touch-Tone** dialing originating register data from a No. 5XB office to the 5XB COER system, and the other will focus on the delivery of trunking data from a 1 *ESS* office to TSS. In both cases considerable preparation is necessary in the switching office and in the databases of the affected TNDS subsystems before data flow can begin. This preparatory work will be described first.

3.1 Addition of Touch-Tone *originating registers in a No. 5XB*

Generally, the marketing department would decide that the level of customer demand was sufficient to justify the cost of purchasing and establishing a *Touch-Tone* originating register (OR) group in a No. 5XB office. Following this decision, the Traffic Engineer would engineer the addition using forecasts of usage and would write a traffic order for the new equipment. A copy of the order would be received

* Registered service mark of AT&T.

by Network Administration personnel whose job it is to assign the *Touch-Tone* equipment to measurement devices. After the assignments are made, central office personnel will connect the leads of the Electronic Traffic Data Collection (ETDC) unit to the *Touch-Tone* equipment.

Traffic counts are maintained on a data collection device (DCD). Before data collection was computerized, traffic measurements were accumulated in mechanical registers. The modern equivalent to these registers are the DCDs, which are the software storage locations of the accumulated counts. For our example these storage locations are in the EADAS memory, although many of the subsystems of TNDS use DCD locations to refer to traffic measurements. The DCD is the common thread that links the traffic measurements from the switching office through EADAS and on to other systems.

A typical No. 5XB may require 1500 DCDs. The important measurements collected in DCDs on the *Touch-Tone* group are identified and defined in Table I below. As illustrated in Fig. 2, the next step in the sequence of preparing individual systems is loading the EADAS database (by defining the new DCDs) to accept the new data and to print exception and hourly reports.

The mapping between the ETDC at the switching office and the EADAS is accomplished via the DCDs. Likewise, the DCD is the link between EADAS and TDAS. The user prepares TDAS to receive this new data from EADAS by modifying the CU/EQ database using input transactions. Each TNDS equipment (TNDS/EQ) system except 5XB COER shares some portion of the CU/EQ database; hence, each system is assigned a series of numerically coded transaction commands to allow the database to be changed. As examples, a 750 transaction is used to associate a DCD with a particular measurement type, such as *Touch-Tone* OR peg count, and the days of the week and hours of the day that measurements need to be processed by TDAS are specified on the Traffic Measurement Request (790) transaction.

The CU/EQ database can be modified to enable 5XB COER to begin processing *Touch-Tone* OR data and printing reports. Once this is accomplished, all pieces are in place to permit traffic data to flow

Table I—DCD measurements on the *Touch-Tone* group

Measurement Type	Definition of Measurement
Usage	An estimate of the total amount of time the *Touch-Tone* ORs were busy for any reason.
Peg Count	The total number of attempts made to access the *Touch-Tone* ORs.
All *Touch-Tone* Registers Busy	The number of attempts to seize *Touch-Tone* ORs when the entire group was busy.
Maintenance Usage	The total amount of time the *Touch-Tone* ORs were busy for maintenance reasons.

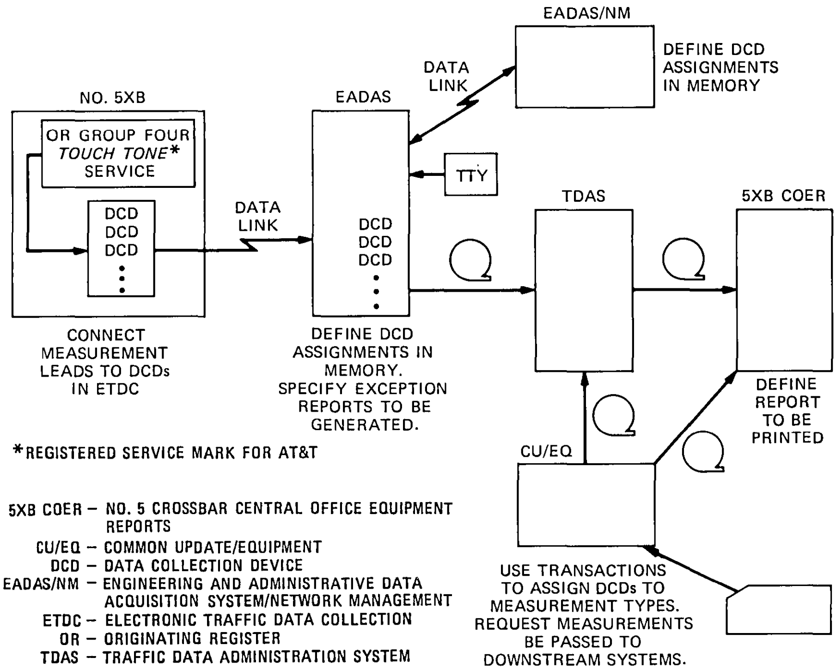

Fig. 2—Database synchronization for *Touch-Tone* originating register data.

from the physical equipment to the Network Administrator and Traffic Engineer. However, two potential users of this data have not yet been included. Those users receive data by an indirect path rather than the main-line path described thus far. Personnel at the network management center receive traffic measurements on critical switching office components in near-real time. To prime EADAS/NM to receive certain of the *Touch-Tone* OR measurements directly from EADAS, the EADAS/NM database is modified by identifying the new measurements (DCDs) that will be available at the EADAS.

Selected data from each of the 27 EADAS/NM systems are sent to the NOCS computer every five minutes. These data describe the status of toll switching offices and trunk groups, but not the status of the local network. Thus, the new No. 5XB data would not be forwarded to NOCS. Data from toll switching offices are sent to NOCS by identifying it in the EADAS/NM database as being of NOCS interest. An EADAS/NM program automatically notifies NOCS that new data will be sent to it by forwarding a database map to NOCS. This program synchronizes the NOCS database to the EADAS/NM database.

The final system to receive these new measurements is CSAR. However, since CSAR measures overall performance rather than specific items, no database changes are required.

Now that all record base components are ready, the flow of new data can be monitored through each system. Figure 3 graphically illustrates this flow.

Data collection starts with the customers; their *Touch-Tone* calls cause attempts (peg counts), usage, and, during high load periods, an occasional all-busy condition on the *Touch-Tone* ORs. (Maintenance usage is generated by the central office work force when they remove an OR from service to maintain or repair it.) The ETDC is signaled each time an attempt is made to seize a *Touch-Tone* OR. The ETDC, in turn, codes a message, which is sent to EADAS when a seizure is attempted. This coded message is simply the DCD for that measurement. The ETDC, located at the switching office, is connected to the EADAS, which generally is physically remote from the ETDC.

For each switching office served by an ETDC, a block of DCDs in the memory of the EADAS records all the traffic measurements being

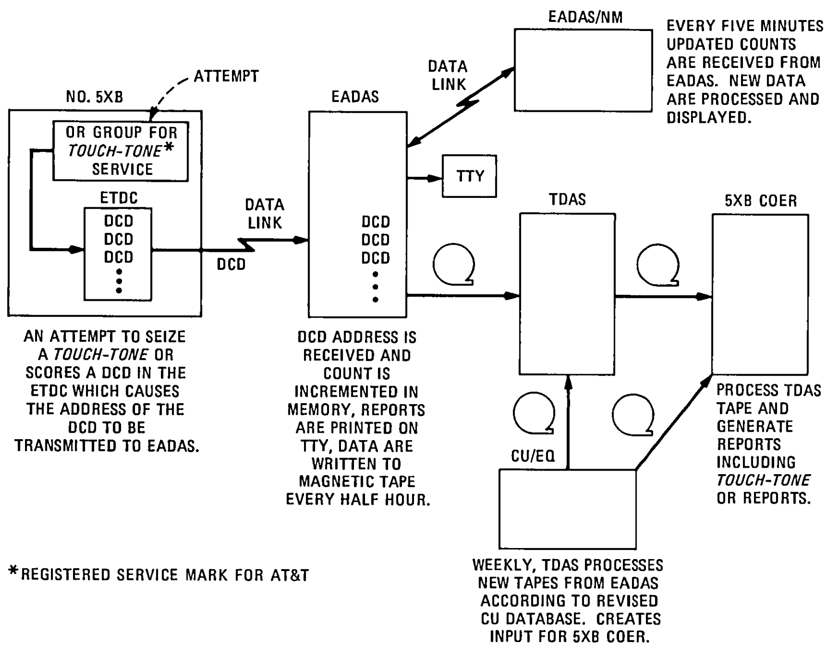

5XB COER – NO. 5 CROSSBAR CENTRAL OFFICE EQUIPMENT REPORTS
CU/EQ – COMMON UPDATE/EQUIPMENT
DCD – DATA COLLECTION DEVICE
EADAS/NM – ENGINEERING AND ADMINISTRATIVE DATA ACQUISITION SYSTEM/NETWORK MANAGEMENT
ETDC – ELECTRONIC TRAFFIC DATA COLLECTION
OR – ORIGINATING REGISTER
TDAS – TRAFFIC DATA ADMINISTRATION SYSTEM

Fig. 3—Flow of *Touch-Tone* originating register data through TNDS.

collected. These banks of DCDs are forwarded to EADAS/NM every five minutes when a command is sent from EADAS/NM over a data link. EADAS transmits data from all of the offices requested by EADAS/NM, including the No. 5XB with the new *Touch-Tone* ORs. After receiving data, EADAS/NM performs calculations on that data and displays the results nearly instantaneously to the network manager. If data from the No. 5XB were of national interest, every five minutes they would be forwarded to the NOCS where calculations would be performed and data displayed to national network managers.

Network management requires 5-minute data, whereas traffic engineering requires hourly data. Starting either on the hour or at half past the hour, EADAS writes the current DCD counts to magnetic tape and resets these registers to zero.

One of the major benefits of EADAS is its exception-reporting capability. The network administrator can define an exception condition such as some number of occurrences of all ORs busy in a 30-minute period. If this frequency occurs or is exceeded, a report would be printed at the EADAS teletypewriter (TTY) and the network administrator would take the appropriate action.

Periodically, the magnetic tape will be forwarded to a data processing center. These tapes will be run through TDAS, which will perform some validations and "warehouse" the data. TDAS will then create new tapes, one for each TNDS downstream system. One tape produced by TDAS would be processed by 5XB COER, which would then perform validation tests and calculations on the data and generate output reports to be used by the network administrator and traffic engineer. These reports would be used to determine if the new equipment was being utilized effectively while providing an acceptable grade of service to the customer. In most companies, the 5XB COER program is run weekly with a TDAS tape containing each day of the week used as input. The key report generated by 5XB COER is the Machine Load Service Summary (MLSS) report, which displays the ten high-day loads and the average busy season load on most switching office components, including the new *Touch-Tone* ORs. The MLSS report is also used by traffic engineers to determine equipment quantities for the next busy season.

At key checkpoints, traffic measurements are monitored by CSAR to ensure that they are received, processed, and are reasonably accurate. If, for example, the DCD associated with *Touch-Tone* OR usage were improperly wired, the occupancy of this equipment could erroneously appear to be greater than 100 percent. CSAR would be notified of these invalid measurements and would include this problem in the calculation of the index. Besides producing the index, CSAR identifies problem areas, thus facilitating their correction.

3.2 Addition of a new trunk group to TSS

The previous example gave an overview of how TNDS collects and processes equipment data from an electromechanical office. In contrast, the next example describes how trunking data is collected from a 1 *ESS* office and processed by TSS. As before, we will start with a description of the record base process needed to support this data flow. In this case, we will describe the effort needed to add a new trunk group between two 1 *ESS* offices. For this example, measurements will be collected at only one end of the group, hence, data will be needed from only one 1 *ESS* office.

Generally, the recommendation to add a new trunk group would come from TFS as part of the forecasting process. At the time the group need is forecasted, it is added to the CU/TK database. The trunk servicer would write an order to have this new group installed and specify the schedule for data collection. The order would pass through organizations responsible for assigning and installing equipment to make the trunk group operational. One step of this operation would be assigning the group a trunk group serial number (TGSN). The TGSN is similar to the Common Language Location Identifier (CLLI); it is the unique identification for a trunk group in the electronic switching system office. These offices do not require ETDCs; they store counts in their memory and later forward them to the data collection computer.

Figure 4 illustrates the steps necessary to prepare each subsystem to collect and process data from this new trunk group. These steps are similar to those in Fig. 2; the major difference occurs at the switching offices. Traffic measurements are specified in software of the 1 *ESS* switching equipment. This includes the EADAS/NM requirements for five-minute data on the new trunk group and requirements for half-hourly data on the TSS and the other downstream systems. The new trunk group is added to the EADAS and EADAS/NM databases with the TGSN being used to identify the group and the DCDs used to identify individual measurements. If the trunk group is of national interest, it will be marked in the EADAS/NM database to be sent to NOCS.

Record base information for this new group is entered for TDAS and TSS via CU/EQ and CU/TK. As before, transactions are used to define the new measurements, and the time period during which these measurements will be processed, and to specify the new trunk group. Once the record bases have been modified, all subsystems are prepared to process this new group. The next section will trace the data flow from the office to the end users.

Figure 5 shows each step of this data flow. The major difference between this example and the previous one occurs at the front end—

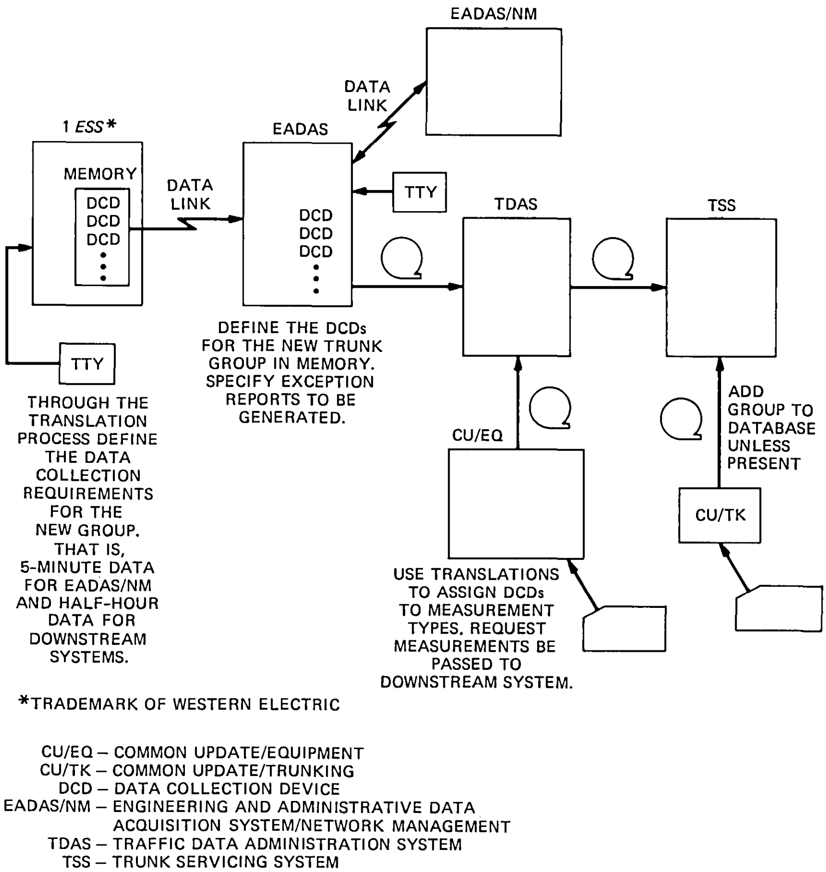

Fig. 4—Database synchronization for trunk group data.

between the switching office and data collection computer. When an attempt is made to seize the new group, a peg count is stored in a memory location at the 1 *ESS* office. The system passes these data forward upon request from the EADAS. Every five minutes EADAS/NM requests data from the EADAS. In the previous example, data were already in the data collection computer. However, for this example the EADAS must request data from the *ESS* offices that it serves. Raw counts are forwarded to the EADAS, which then subtracts the previous readings and forwards them to EADAS/NM. As before, EADAS/NM processes these data and displays them to the network manager. If this 1*ESS* is a key part of the toll network, data from the new group could be forwarded to NOCS for display.

To satisfy the needs of engineering, much more data are passed from the ESS office to the EADAS on the hour and half hour. These

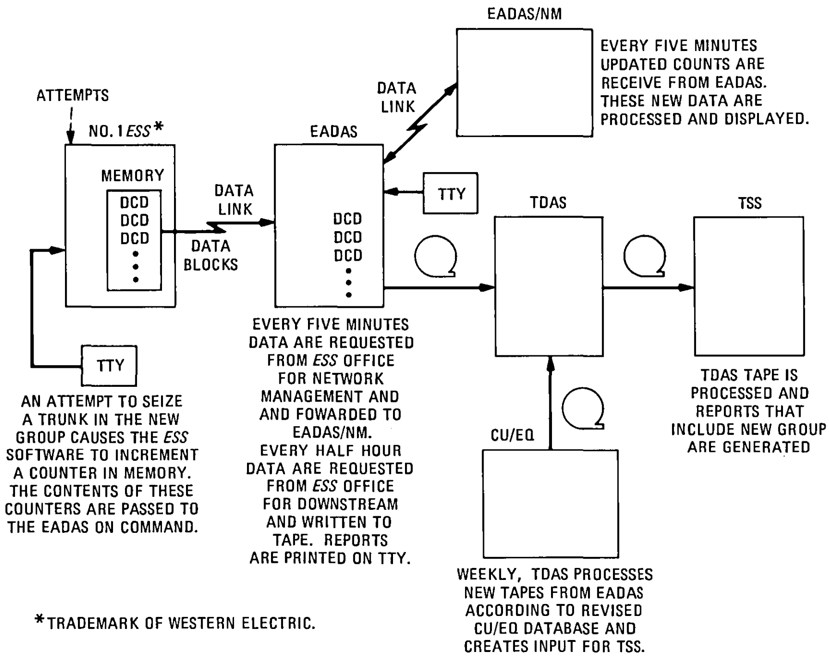

CU/EQ – COMMON UPDATE/EQUIPMENT
DCD – DATA COLLECTION DEVICE
EADAS/NM – ENGINEERING AND ADMINISTRATIVE DATA ACQUISITION SYSTEM/NETWORK MANAGEMENT
TDAS – TRAFFIC DATA ADMINISTRATION SYSTEM
TSS – TRUNK SERVICING SYSTEM

Fig. 5—Flow of trunk group data through TNDS.

data are written to tape and handled identically to the previous example. TDAS generates a separate tape to be processed by TSS and, as frequently as weekly, the TSS subsystem is run to process the data and generate reports. In the case of the new group, the servicer would closely examine its performance to determine if the number of circuits in the group was a good match to the load being offered to the group. If a mismatch occurred, an adjustment in the circuit quantity would be made after a pattern of performance for that group was established by running the TSS program several times with new traffic data.

As was the case for equipment data, CSAR monitors the flow of trunking data, calculates an index on this portion of TNDS, and produces reports that highlight problem areas.

IV. MANAGING SYSTEMS ENGINEERING AND CONTINUING DEVELOPMENT

From the descriptions above, it is clear that from the user viewpoint, TNDS is large and complex. It is even more complex from a software engineering viewpoint. This section describes how the software engineering for this large and complex system is managed.

4.1 End products and services

4.1.1 What is provided

The TNDS organization provides various end products and services to the users. Although the specific details vary among the systems in TNDS, the end products and services can be categorized as follows:

1. Software
2. Documentation
   a. End user
   b. System Operation
3. Training
4. User Support.

The software is in the form of load modules for the specific machine on which the system is run. The load modules for systems run in each operating company are sent by tape or high-speed data transmission and then loaded. System operations personnel install the load modules using the installation guide provided. For centralized systems (e.g., SPCS COER, CSAR, and SONDS), development personnel install the new load modules.

The documentation provided falls into two categories: end user and system operation. End user documentation provides network personnel with information on how to interface with the system. This documentation is usually in paper form, but in the case of SPCS COER and CSAR, it is on-line at the user terminal with the option of printing. Each operating company also receives documentation for system operation from an Electronic Data Processing (EDP) viewpoint for the systems run at that site. This documentation includes information on backup and recovery procedures, resource requirements, etc.

The TNDS/EQ-TNDS/TK training philosophy is to train the trainers—TNDS coordinators and trainers in each operating company—by giving them the information they need to hold similar training courses in their companies. Both user and EDP training is provided. In addition, special release-oriented training usually supplements the basic system training.

User support for each operating company is usually in the form of an assist line (or "hot line") number, which the user can call when questions or problems arise.

4.1.2 Size of the product

Table II shows the relative size of each end product and service for each system in the TNDS. The lines of source code are in various high-level languages (PL/I, COBOL, FORTRAN, and C).

Table II—TNDS software and documentation

	Lines of Code (Thousands)	Pages of Documentation	
		User	EDP
EADAS/NM	900,000	700	800
EADAS	200,000	3500	500
TDAS	75,000	1725	1334
CU (CU/EQ and CU/TK)	119,000	586	2527
ICAN	43,000	1445	318
LBS	42,000	1764	560
5XB COER	128,000	2215	340
SPCS COER	450,000	1000*	20
CSAR	125,000	200*	48
TSS	166,000	1805	641
TFS	80,000	2086	2788
SONDS	178,000	380*	30†
Total (Rounded)	2,328,000	17,000	10,000

* On-line lessons.
† Binders.

4.1.3 TNDS deployment

Another important factor that affects managing the TNDS is its extensive deployment. Table III illustrates the extent of deployment of each component system in TNDS. Looked at another way, some or all of these systems are deployed in all of the operating companies and several are deployed in Bell of Canada.

4.2 Managing change

Although TNDS is a mature system that has been widely deployed and operating in Bell operating companies for some time, the systems are still undergoing extensive development to keep up with the changing operating environment and to furnish improved capabilities. A goal of this continuing development process is to furnish users periodically (usually yearly) with new versions of the systems that incorporate needed changes in a timely manner. To achieve this goal, several factors must be considered.

1. The total number of possible improvements is generally much greater than the resources available for a particular release.

2. The calendar time required to develop a given release is much longer than one year.

3. A change frequently has impact in more than one of the systems of the TNDS.

These factors result in a need for careful planning to choose among the capabilities to be incorporated into a release, to begin planning well before the targeted release date, and to carefully coordinate planned changes among the systems. A formal approach called the TNDS Release Cycle is used to satisfy these needs.

Table III—TNDS subsystem development and utilization

Function	Subsystem	Parameter	Percent Coverage (End 1982)
Data Acquisition	EADAS (or equivalent)	Main and equivalent main telephones	93
	TDAS/CU	Number of companies	100
Network Management	EADAS/NM	Class 1,2,3,4 and sector tandem offices	84
		End offices with over 5000 telephones	63
Central Office Equipment	No. 5 Crossbar COER	Marker groups	97
	SPCS COER	Control groups	100
	SONDS	Step offices over 2000 lines	37
	ICAN	Eligible traffic units	60
	LBS	Eligible traffic units	98
Trunking	TFS	Message and auxiliary trunk groups	84
	TSS	Message and auxiliary trunk groups	93
TNDS Performance	CSAR	TNDS systems monitored	100

4.2.1 The nature of change

TNDS software and hardware are modified and enhanced for many reasons. First of all, designers must run to stay even with the rapidly evolving telephone network equipment, services, and operating procedures. Next additional features are needed by TNDS users and operators. And to prevent the system from becoming technically obsolete, it is necessary to make changes that improve the efficiency, maintainability, and expandability of software and hardware components. Examples of changes to the telephone network that have been accommodated in TNDS include the addition of new switching system equipment such as 5 *ESS* (the management of this change is discussed further below) and the use of new extreme value load engineering procedures. Recent additions include mechanization of *ESS* busy-hour determination studies, and new reports for No. 5 Crossbar administrators. There have been technical improvements to TNDS. Some examples are run time and disk utilization improvements, the complete rewrite of the No. 5XB COER module to improve its maintainability, and the conversion of EADAS to the *UNIX** operating system, which simplified the addition of new features.

New releases of software modules usually contain a mixture of enhancement types. A recent TNDS/EQ release is typical. It contains features that support 5 *ESS* data collection, a new 2B *ESS* traffic data collection interface, and improvements to allow automatic transfer of *ESS* capacity statistics from EADAS to SPCS COER. This release also eliminates software that supported now obsolete data collection equipment. Several software functions were simplified, including the interface with one non-TNDS system, the distribution of traffic register definition listings, and company parameter tables. Also, thirteen separate changes to improve operational efficiency were made.

4.2.2 The release cycle

The TNDS release cycle covers the overall planning, development, and project control processes to be employed to produce the major releases of systems in the TNDS. For release purposes, these systems are grouped into seven sets as outlined in Fig. 6.

Each of these groupings has a separate release cycle, all of which are coupled to varying degrees. In particular, the cycles of EADAS, TNDS/EQ, SPCS COER, and CSAR are tightly coupled because of feature interactions. These systems all support the Network Administration function. On the other hand, release of EADAS/NM and TNDS/TK (supporting network management and trunk administration, respectively) are less dependent on the other systems and on

* Trademark of Bell Laboratories.

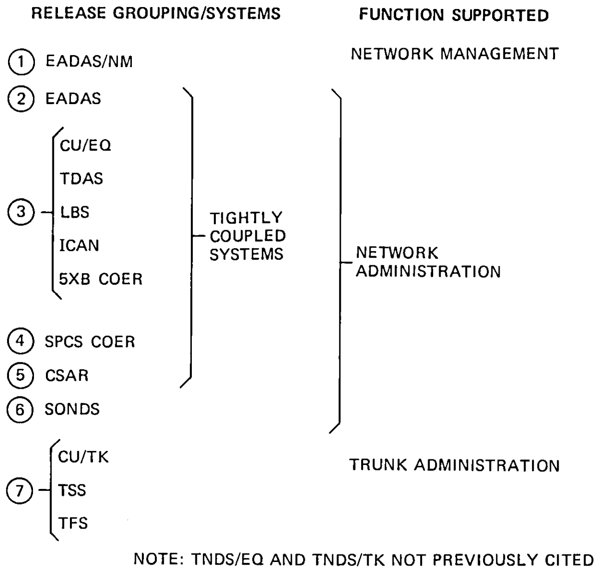

5XB COER – NO. 5 CROSSBAR CENTRAL OFFICE EQUIPMENT REPORTS
CSAR – CENTRALIZED SYSTEMS FOR ANALYSIS AND REPORTING
CU/EQ – COMMON UPDATE/EQUIPMENT
CU/TK – COMMON UPDATE/TRUNKING
EADAS/NM – ENGINEERING AND ADMINISTRATIVE DATA ACQUISITION SYSTEM/NETWORK MANAGEMENT
ICAN – INDIVIDUAL CIRCUIT ANALYSIS
LBS – LOAD BALANCE SYSTEM
SONDS – SMALL OFFICE NETWORK DATA SYSTEM
SPCS COER – STORED PROGRAM CONTROL SYSTEMS CENTRAL OFFICE EQUIPMENT REPORTS
TDAS – TRAFFIC DATA ADMINISTRATION SYSTEM
TFS – TRUNK FORECASTING SYSTEM
TSS – TRUNK SERVICING SYSTEM

Fig. 6—TNDS release coupling.

each other. Finally, SONDS is a complete mini-TNDS for a specific class of offices (SXS CDOs), and consequently is more independent of intersystem interactions than other systems in the TNDS.

The following discussion applies in general to all the TNDS systems, but in detailed examples, the release cycle for TNDS/EQ will be used.

4.2.3 Project phases

The project phases in the release cycle define a continuing development process. The phases are divided into major categories and subcategories as shown in Fig. 7 and as listed below:

1. Planning
2. Implementation
   a. Requirements
   b. Design, code, and test
   c. Soak

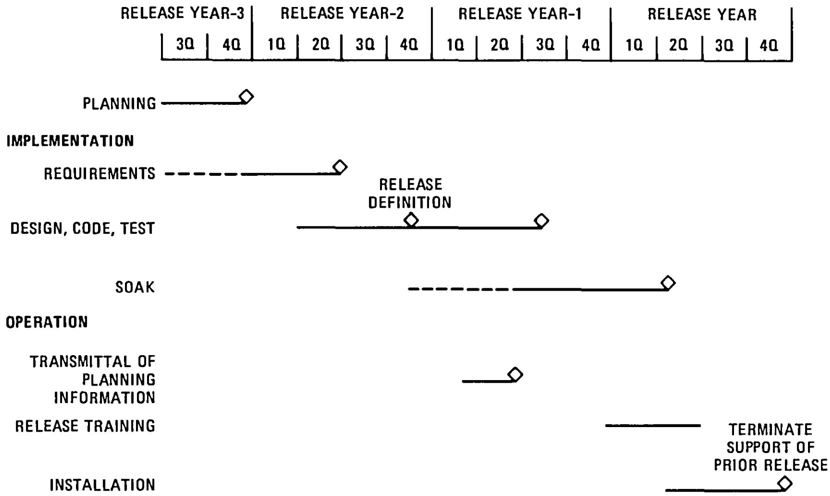

Fig. 7—TNDS release cycle.

3. Operation
   a. Transmittal of planning information
   b. Release training
   c. Installation.

4.2.3.1 Planning. For the continuing development process, inputs come from many sources—operating companies, AT&T, and Bell Laboratories. Operating companies submit change requests to the AT&T project managers. These managers as well as Bell Laboratories designers and systems engineers also initiate proposals. These inputs are diverse. They range from requests to fix bugs or design errors to major development programs. The AT&T project managers and Bell Laboratories systems engineers, often with the advice of telephone company user representatives, perform an initial screening and classification of the inputs. In this step the inputs are classified into maintenance items, desirable improvements, and rejected suggestions. Maintenance items are transferred to the maintenance organization for the component system involved. Rejects are returned to the originator with an explanation of the reason for rejection. Desirable improvements are included in one or more "feature" packages.

Feature packages are specific development items that can be assigned individual priorities. They usually involve only one component system, but may be dependent on feature packages for other component systems for implementation. The assignment of individual items to feature packages has several advantages. Related items can be developed more efficiently as a group rather than separately. Items with similar objectives can be combined or modified to create improvements with more global and more universally useful objectives. And since improvements often are suggested in terms of implementation, the same objectives sometimes can be accomplished as part of an already scheduled or more easily accomplished feature.

Certain time-critical features may be worked into ongoing development of forthcoming releases. These "expedited enhancements" generally must be small and very important because they displace other planned work. The level of management approval required for expedited enhancements depends upon the nature of the work being displaced. The remaining features are subjected to a much longer obstacle course. They must compete for the limited resources available for TNDS ongoing development.

Feature packages are carefully evaluated by Bell Laboratories systems engineering and development, by AT&T project managers, and by telephone company user representatives to establish their priority. Metrics for comparative evaluation include economic benefits, estimated development cost, and estimated length of time to develop. Other intangible factors that are considered in setting priorities are potential quality-of-service improvements, needs to adhere to AT&T corporate policies or operating practices, and the need to respond to FCC or state regulatory commission requests for data.

Feature lists in priority order are maintained on a permanent basis and are updated periodically. Separate lists are maintained for TNDS/TK, TNDS/EQ, SPCS COER, EADAS, and EADAS/NM. Any dependencies on other TNDS or non-TNDS features are noted. These priority lists are used as input for planning specific new TNDS software releases. Priorities may be time-dependent, that is, a certain feature, such as support of a new switching entity may not be needed until some future date. After that date, the importance of the support increases with the expected deployment of the switch.

Release planning is the responsibility of Bell Laboratories systems engineers working with each of the TNDS component systems. They consider the engineering resources available, the priority of proposed features, and length of time for development to formulate the content of proposed releases. These proposals are reviewed by AT&T project managers, telephone company user representatives and, when appropriate, Western Electric product line managers. Changes, if any, are made, and the priority of each of the features in the release package is established. This release priority is used to eliminate work when the inevitable resource constraints and expedited enhancements arrive. While approvals for the proposed release packages are being obtained, work will usually start on requirements for the features in the proposed package. Sign work, however, starts after final approval.

4.2.3.2 Implementation. The first step toward implementing a release is to complete detailed requirements for each feature. These are reviewed and approved, and they are one of the key project control documents.

The design step then translates the requirements into the code changes to be made. Design is divided into two parts: (1) system design, which identifies the modules, files, and interfaces (e.g., transaction and report layouts, intersystem interface definitions) to be changed; and (2) detail design, which identifies the coding changes to be made in each module. After system design is complete, the development organizations formally commit themselves to the contents and schedule for the release. After detail design, the changes are then coded and tested.

Finally, a soak test of the product is performed in one operating company site. The soak process determines if the end product performs in accordance with the requirements, is capable of satisfying user needs, and is worthy of systemwide release. The soak site is selected based on environmental requirements needed to exercise the features in the release. Soak periods vary by system but usually last from three to six months. After a soak is successfully completed, the release is made available for installation in other sites.

4.2.3.3 Operation. The first step an operating company takes in installing a release is receipt of planning information in the form of a Product Release Description, which is usually made available approximately four months prior to the general availability date.

Also prior to general availability, release training information is made available and training classes are held. Documentation for the release is also made available in this time frame.

When the release is available it is installed in other sites. Each site is expected to install the release within a reasonable period (usually six months), at which time continued support of the previous release is terminated.

General availability dates are usually chosen to allow installation during a time period that is not in the traffic busy season when data collection is critical.

4.2.4 Release cycle timing

The phases of the release process are integrated into an overall release cycle as shown in Fig. 7. This figure depicts the major activities within the cycle and the time intervals typically required to conduct each of these activities. Roughly three years are required to complete all of the phases for a particular release. With releases planned for annual delivery, this implies that at any point in time, development activities will be in progress on three different releases. The overlapping of releases is shown in Fig. 8.

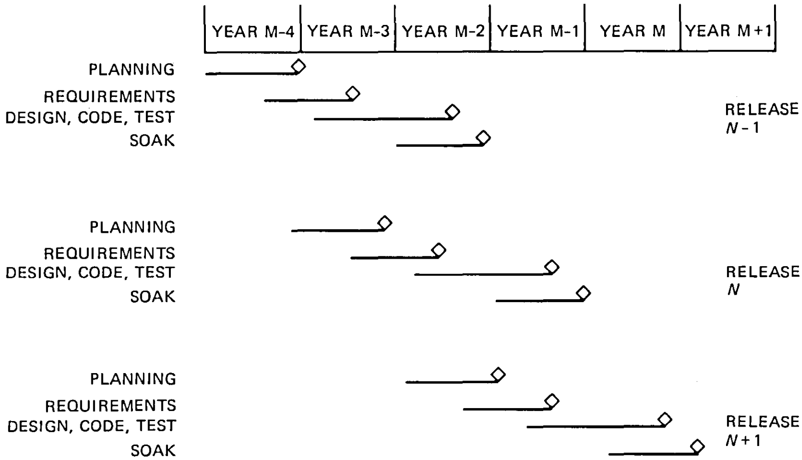

Fig. 8—Parallelism between releases.

4.2.5 Interproject coordination

For the subset of TNDS features that involve development in multiple component systems or in non-TNDS systems, an interproject coordination activity must be overlaid on the internal TNDS control process. Features in this category are few but often large, important development efforts. An example of such a development is supporting a new switching system like the 5 *ESS* system. Coordination involves establishing overall requirements for TNDS support, identifying specific "features" to be developed in each component system, and negotiating a coordinated development program for these releases. The development program may involve work in both TNDS and non-TNDS systems. Often, project committees are formed to coordinate requirements, design, testing, and soak of both the hardware and software components. Always there is an exchange of documentation at each step in the process. Interfaces are defined and documented as early as possible.

To make the continuing development process work, it is essential that continual communication among all the people involved be maintained. A permanent AT&T, Bell Laboratories, Western Electric management committee has the responsibility to see that communications are maintained and that problems in the continuing development process are resolved. The process has enabled TNDS to evolve with changes in the telephone business, become more efficient, grow in services to its users, and introduce new technologies, where appropriate.

4.3 The organization structure

TNDS software engineering takes a large number of people working effectively together with well-defined individual responsibilities. Because of the size of TNDS, the organization is structured functionally with each organizational group having responsibility for one or more of the project phases defined above and for one or more of the systems in the TNDS.

4.3.1 Functional structure

Three types of organizations divide up the work in the release cycle—Systems Engineering, Development, and Western Electric Product Engineering Control Center (PECC).

4.3.1.1 Systems engineering. Systems Engineering organizations have responsibility for the planning and requirements phases, and they participate in the soak evaluation.

4.3.1.2 Development. The development organizations are responsible for design, coding, and test, and they also participate in the soak evaluation.

4.3.1.3 Role of Western Electric. Western Electric PECC organizations are responsible for (1) developing documentation, (2) developing training and giving classes, (3) coordinating and performing soak evaluation, (4) maintaining installation support, and (5) providing customer consultation.

4.3.2 Project structures

Figure 9 shows the structure of the TNDS organization in terms of responsibility for systems engineering and development. Each circle represents a separate department in Bell Laboratories. Those organizations on the left are responsible for systems engineering, and those on the right for development. The systems for which the organization is responsible are indicated in the circle. Also indicated above the circle is the physical location of the organization (Holmdel, Columbus, Piscataway, West Long Branch, and AT&T Data Systems in Piscataway). Note that many Western Electric organizations involved in TNDS are not shown on the chart.

4.4 A continuing development example—the marriage of TNDS and 5 ESS

The 5 *ESS* switching system is a new, Bell Laboratories-designed system that recently went into service with full TNDS support in place. Planning and implementation of this support has occurred in parallel with the development of the 5 *ESS* switching equipment. To illustrate the TNDS continuing development process, we will use the history of TNDS support for 5 *ESS* switching equipment.

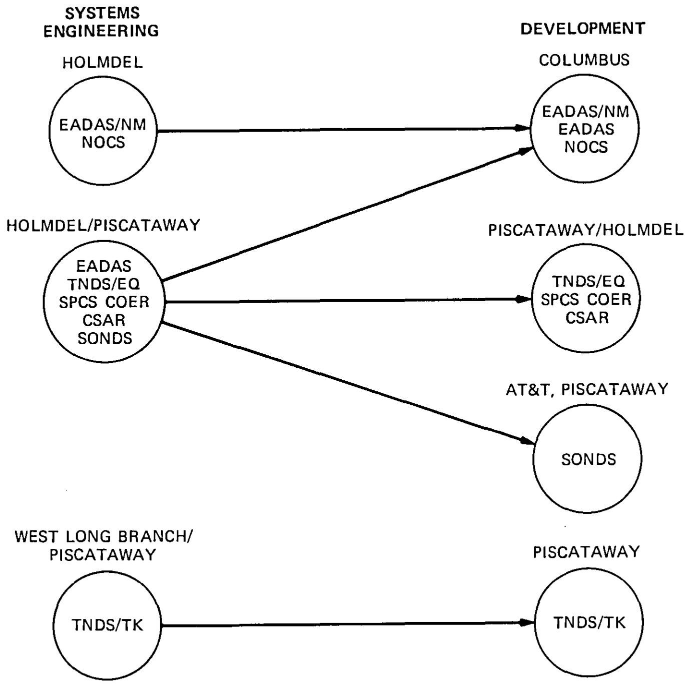

CSAR – CENTRALIZED SYSTEMS FOR ANALYSIS AND REPORTING
EADAS/NM – ENGINEERING AND ADMINISTRATIVE DATA ACQUISITION SYSTEM/NETWORK MANAGEMENT
NOCS – NETWORK OPERATIONS CENTER SYSTEM
SONDS – SMALL OFFICE NETWORK DATA SYSTEM
SPCS COER – STORED PROGRAM CONTROL SYSTEMS CENTRAL OFFICE EQUIPMENT REPORTS
TNDS/EQ – TOTAL NETWORK DATA SYSTEM/EQUIPMENT
TNDS/TK – TOTAL NETWORK DATA SYSTEM/TRUNKING

Fig. 9—TNDS organizational interactions.

Almost four years before the scheduled cutover of the first 5 *ESS* switching equipment, its development organization began to consider alternatives for handling the traffic data, and the needs for specific types of measurements. A joint AT&T, Bell Laboratories task group was formed to consider 5 *ESS* traffic data-handling needs. Experts on 5 *ESS* development, traffic engineering methods, TNDS, and telephone company data needs were assigned to the group. This task force recommended an overall strategy for handling 5 *ESS* traffic data and identified a number of specific work efforts needed to implement the strategy. The recommendations of this group established the objectives for the several work activities that followed. Key recommendations of this group included:

1. Allocating all traffic data-handling functions except basic measurement and measurement distribution to TNDS.

2. Using extreme value engineering techniques for timing and sizing of 5 *ESS* additions.

3. Collecting traffic data for trunk and switching equipment engineering, division of revenue, service provisioning, switching administration, and network management through EADAS.

4. Selecting peak values in EADAS for extreme value engineering.

Work activities included:

1. Developing extreme value engineering methods and measurement requirements for 5 *ESS* switching equipment.

2. Formulating a TNDS plan to support 5 *ESS* switching equipment.

3. Formulating requirements for 5 *ESS* features needed to deliver required measurements to EADAS.

The task group's recommendations were accepted by AT&T and Bell Laboratories managers, and work started on the above tasks. The work proceeded in parallel.

The first step in the TNDS planning effort was identifying TNDS outputs required for 5 *ESS* switching equipment and the TNDS developments needed to provide these outputs. Descriptions were developed, by individual component system, of the feature packages needed to support 5 *ESS* switching equipment. Table IV identifies the changes needed to TNDS. Following identification of the necessary features a tentative schedule was established for requirements formulation and component system design. This schedule was coordinated with the estimated schedules for the engineering methods work,

Table IV—TNDS developments for initial 5 ESS service

System	Development Required
EADAS	X.25 interface with 5 *ESS* switching equipment Peak measurement selection Complete Network Operation Report Generation (NORGEN) report package
EADAS/NM	No development for first service. (Full EADAS/NM support will be provided in stages with development of advanced versions of 5 *ESS* switching equipment.)
TDAS CU/EQ	5 *ESS* measurement identification Peak measurement handling 5 *ESS* statistics for CSAR
SPCS COER	Complete package of engineering and administrative reports for processing extreme value statistics 5 *ESS* statistics for CSAR
LBS	No development needed for first service. (Development to support a new 5 *ESS* load balance index will be required later.)
TSS, TFS, CSAR	No development required as current software is sufficiently general to handle 5 *ESS* switching equipment.
5XB COER ICAN SONDS	These systems support electromechanical switching entities. No changes for 5 *ESS* switching equipment required.

and the 5 *ESS* development work. The resulting TNDS "Integration Plan" for 5 *ESS* switching equipment was reviewed by Bell Laboratories and AT&T management, and then was distributed to all organizations having responsibility for the TNDS component systems. The TNDS integration plan was completed two years before the scheduled first 5 *ESS* service.

At this point the features needed to support 5 *ESS* switching equipment were added to the feature priority lists for each of the TNDS component systems. These 5 *ESS* features thus became part of the overall continuing development process for TNDS. Because of the expected widespread deployment of 5 *ESS* switching equipment in the operating companies and the large economic penalties that would occur without TNDS support for 5 *ESS*, the 5 *ESS* packages generally received the highest priority for development. Approval for development and for requirements formulation was obtained for each of the component systems in time to meet the initial 5 *ESS* service date. Work then started on requirements and design under the normal TNDS release process.

Three interface definitions (5 *ESS* switching equipment and EADAS, EADAS and TNDS/EQ, TNDS/EQ and SPCS COER) were formulated. These were required because of the decisions to utilize an X.25 protocol between EADAS and 5 *ESS* switching equipment, and the need to specify new schedules for extreme value data being selected by EADAS. Other requirements for new EADAS, SPCS COER, and TNDS/EQ reporting and processing were written. Actual design and coding of the various TNDS features began one to two years before the scheduled first 5 *ESS* service.

A committee of development representatives was formed to coordinate development activities for first service. Experts from 5 *ESS* switching equipment, EADAS, and TNDS/EQ served on it. This group worked out problems in design details and coordinated laboratory-to-laboratory testing activity, which generally proceeded in stages from basic hardware communications to application process communications. In addition, load simulators were used in the development organizations to do further interface testing and debugging. A large number of measurement definition changes occurred during the 5 *ESS* development. These changes had to be accommodated in the TNDS during development.

The first 5 *ESS* is owned by the Illinois Bell Telephone Company. To test the TNDS features designed to support the 5 *ESS* switching equipment, it was necessary to arrange soak testing of EADAS, TNDS/EQ, and SPCS COER releases in Illinois Bell. Planning for this soak testing began about one year before the scheduled first office service.

Final end-to-end laboratory tests were conducted three months before first service. Initial soak testing with the first 5 *ESS* began about six weeks prior to service. A simulated call load on the 5 *ESS* switching equipment was used for these tests. When the 5 *ESS* switching equipment started handling live traffic, TNDS was in place to measure the traffic. Coordinated testing of TNDS components under live traffic conditions continued for several months after first service. Then the new release packages were made available to other TNDS users.

TNDS was ready to serve the first 5 *ESS* switching equipment because the need for TNDS support and planning was recognized by the 5 *ESS* developers early in their development cycle. Planning, requirements formulations, design, code, test, and soak proceeded in an orderly manner with careful coordination at each step. Problems in scheduling, interface details, and detailed measurement definitions were all solved through close coordination of the development efforts.

4.5 Future directions

TNDS will continue to evolve gradually. Setting the direction of this continuing evolution is the purpose of TNDS long-range planning. Long-range planning for the TNDS is a continuous process whose output is a series of TNDS feature packages designed to avoid obsolescence and to take advantage of new technology. The results of recent long-range planning studies have indicated that further major improvements are desirable and appear to offer significant economic benefits. Several improvements relating to additional support for network administration functions are being studied. These studies are aimed at reducing the volume of paper reports produced by the TNDS, and providing users more direct control of TNDS operations, and flexible user-defined output reports. These objectives are planned to be achieved through on-line updating capabilities, expanded user programmability features, and provision of a simple standardized human interface to TNDS component systems. The next steps are to expand the scope of the planning work to include support for other telephone company organizations and to define the specific features that will be implemented. Through this process of continuous renewal guided by a periodically revised long-range plan, it is expected that TNDS will continue to serve the traffic-data handling needs of operating companies into the next decade and beyond.

REFERENCES

1. R. L. Martin, "Automated Repair Service Bureau: System Architecture," B.S.T.J., *61*, No. 6, Part 2 (July–August 1982), pp. 1115–29.
2. R. F. Grantges, et al, "Central Office Equipment Reports for Stored Program Control Systems", B.S.T.J., this issue.

AUTHORS

George W. Riesz, B.S. (Mechanical Engineering), 1950, M.S., (Electrical Engineering), 1952, Princeton University; Bell Laboratories, 1952—. Mr. Riesz joined Bell Laboratories in 1952 as a member of the Analog Computer Design Group. During 1955 and 1956 he served in the U.S. Army at Redstone Arsenal, where he participated in the initial ballistic missile defense studies. In 1956, he joined Nike-Zeus Systems Engineering, and from 1958 to 1960, was responsible for the design of Nike-Zeus battery control equipment. In 1962, he became Head of the Advanced Design Department. He then worked as Head of the Discrimination Measurements Department in Kwajalein from 1963 to 1965, when he became Head of the Nike-X Systems Analysis Department. In 1968, he was named Head of the CAMAR Data Systems Department, working on data systems design for advanced measurements radar. In 1970, he became Head of the Resource Allocation Studies Department, and in 1971, Head of the Traffic Network Studies Department. In 1974, he was named Head of the Network Administration Department. He assumed his present position as Head of the Operations Systems Implementation Planning Department in 1979. Mr. Riesz holds one patent and is the author of several technical articles. Member, Phi Beta Kappa, Sigma Xi.

J. Wayne Steifle, B.S. (Electrical Engineering), 1968, Clemson University; M.S. (Electrical Engineering), 1969, Stanford University; Bell Laboratories, 1968–1983; American Bell, 1983—. After joining Bell Laboratories, Mr. Steifle worked on radar and ABM system studies, and then spent three years in the Air Force working at NASA on guidance systems for manned space programs. After returning from military leave, he worked on guidance design for Safeguard, and later on systems engineering and development of various Business Information System projects. From 1977 to 1980 Mr. Steifle was Supervisor of the TNDS/EQ Engineering Group, where he was responsible for systems engineering for the TNDS/EQ systems. He is currently Supervisor of the Interconnect Software Group in the AIS/NET 1000 Node Hardware Department at American Bell. Member, IEEE.

John A. Kohut, B.A. (Mathematics), 1964, Gannon College; M.S. (Systems Engineering), 1967, Stevens Institute of Technology; Bell Laboratories, 1964–1977, 1980—; AT&T, 1977–1980. Early in his career at Bell Laboratories, Mr. Kohut evaluated the performance of various antiballistic missile systems using computer simulation techniques. From 1970 to 1977, he analyzed the behavior of the message telephone networks and various control systems during failure and overload conditions. From 1977 to 1980, he worked at AT&T, where he was the project manager of an operational support system that aids in the real-time management of the network. Presently, Mr. Kohut is Supervisor of the Trunk Servicing Engineering Group, which plans new capabilities for the Trunk Servicing System (TSS).

Milton S. Hall, Jr., B.E.E. 1959, Georgia Institute of Technology; M.E.E. 1962, New York University; Bell Laboratories, 1959–1968, 1970—; AT&T, 1968–1970. Initially Mr. Hall worked in the data systems engineering organization. His responsibilities there included computer performance analysis, requirements for store-and-forward switching, and plans for wideband data services. He also spent six months on active duty for training in the U.S. Army Signal Corps. In 1968, he transferred to AT&T, where he was involved in long-haul transmission facilities planning. Then in 1970, he returned to

Bell Laboratories to supervise traffic engineering studies. In addition, his group designed a large, on-line system for processing telephone traffic data. From 1977 to October 1981, when he assumed his current assignment, Mr. Hall supervised a group responsible for network administration systems planning. In this work, his group investigated on-line access to information and records as a means to eliminate paper reports and to provide direct user updating of a distributed datebase. Since October 1981, he has supervised a group responsible for operations systems evolution planning. Member, IEEE, ACM, Tau Beta Pi, Phi Kappa Phi, Phi Eta Sigma.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Theoretical and Engineering Foundations

By W. S. HAYWARD* and J. P. MORELAND*

(Manuscript received July 14, 1982)

The Total Network Data System (TNDS) is a coordinated family of computer-based systems that collect and process network measurements to aid the engineers, administrators, and managers of the Bell System network in efficiently meeting service objectives. This paper describes these service objectives, the nature of telephone traffic and traffic measurements, and the theories and engineering assumptions underlying the use of these measurements in the design and administration of the trunk network and switching systems.

I. INTRODUCTION

The Total Network Data System (TNDS) is a coordinated family of computer-based systems that collect and process network measurements to aid the engineers, administrators, and managers of the Bell System network in efficiently meeting service objectives. In this paper we describe these service objectives, the traffic measurements used to monitor and design the network, and the theories underlying the use of these measurements in the various TNDS systems.

It is important to realize that, as a result of the continuing changes in the switching systems and methods of routing traffic in the Bell

* Bell Laboratories.

System network, new mathematical models and new service criteria are likely to be used as time goes by. Therefore what we are reporting here is a snapshot of what is going on today. We fully expect that there will be significant changes in the years to come.

The TNDS is described in the other articles of this issue. Our presentation of the theoretical foundation of the activities that the TNDS supports is divided into two parts, trunking and central office. We start with trunking because its simplest model is widely known and serves as an introduction.

II. TRUNKING

2.1 Service objectives

Consider the simple two-node network shown in Fig. 1. The trunk group joining end offices A and B provides the only route for calls between A and B and, as such, is called an only-route trunk group. Call attempts which arrive when all trunks are busy are said to be blocked and the customer is requested (by a reorder tone) to hang up and try the call at a later time.

The grade-of-service for an only-route group is defined to be the unweighted average blocking observed in the time-consistent busy hour of the busy season (defined below). For a given hour, the blocking is defined to be the ratio of the total number of blocked calls (overflows) to the total number of call attempts (peg count). When the hourly blockings are averaged over the time-consistent busy hour of the busy season, the resulting number is the observed grade-of-service for the group. The service objective is 0.01 average blocking.

The measure of telephone traffic used to define the time-consistent busy hour of the busy season is called the trunk-group offered load. Offered load is measured in units called erlangs, and is equal to the average number of busy trunks for a situation in which there is no blocking. Of course, in practice, offered load cannot be measured directly since calls are blocked. It can, however, be estimated from measurements of the carried load (average number of busy trunks) and blocking, as explained in Section 2.3.

When the hourly offered loads are averaged over the same hour for 20 consecutive business days, the maximum of these averages defines the time-consistent busy-hour load for the 20-day period.* The busy season is then defined to be the 20-day period for which the time-consistent busy-hour load is a maximum.

Formally, if PC_j and O_j denote, respectively, the number of call

* A different definition of busy hour is used for trunking than for central office equipment. The difference has grown up over the years because of fundamental differences between trunks and central offices, particularly with regard to forecasting capability, available measurements, and ability to respond to unforeseen shifts of traffic.

Fig. 1—Only-route trunk group.

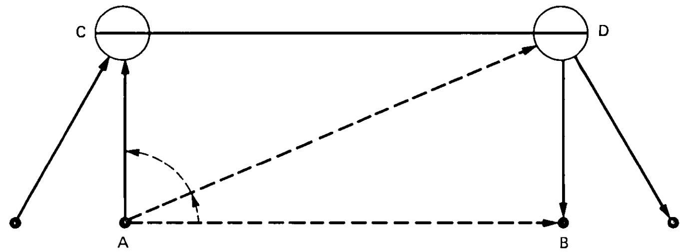

Fig. 2—Alternate-routing network.

attempts and blocked attempts during the jth time-consistent busy hour of the busy season, the observed blocking, b_j, during hour j, is defined to be

$$b_j = \begin{matrix} O_j/PC_j, & \text{if} \quad PC_j > 0 \\ 0 \quad \text{otherwise}, & \end{matrix} \tag{1}$$

and the observed grade-of-service is defined to be

$$b = \frac{1}{n} \sum_{j=1}^{n} b_j, \tag{2}$$

where n is typically 20 consecutive business days. As explained in Section 6.3, the statistic b is used to detect service problems (blocking significantly greater than the 0.01 objective) and, if necessary, trigger corrective actions (e.g., trunk augments).

Figure 2 shows a more complex traffic network consisting of high-usage trunk groups (dashed lines) and final trunk groups (solid lines). Calls originating at end ofice A and destined for end office B are first offered to the primary high-usage group AB. If, however, all AB trunks are busy, the call is alternate routed to the intermediate high-usage group between A and the tandem switch D or, if all AD trunks are busy, to the final group AC. If the call arrives at D it is offered to the final group DB; if the call arrives at C, it is offered to the final group CD. Whenever a call is offered to a final group with all trunks busy, the customer receives a reorder tone or recorded announcement.

The grade-of-service for final trunk groups is defined in the same way as described above for only-route groups; the service objective is again 0.01 average blocking.

Service objectives are not, however, specified for high-usage groups. As explained in Section 2.4.1, they are sized so as to minimize the cost of the trunk network.

2.2 Traffic models

Modern trunking theory shows that three traffic parameters are required to predict busy-hour busy-season average blocking[1]: (1) average offered load during the time-consistent busy hour of the busy season, (2) peakedness (the ratio of the variance to the mean of the number of busy trunks for a trunk group so large that there is no blocking), and (3) the day-to-day variation of the busy-hour loads (variance of the daily offered loads in the time-consistent busy hour of the busy season).

2.2.1 Poisson/Erlang formulas

Prior to 1970, the Poisson formula,

$$P(c, a) = e^{-a} \sum_{n=c}^{\infty} \frac{a^n}{n!}, \tag{3}$$

where c is the number of trunks and a is the offered load in erlangs, was the standard formula for sizing trunk groups in the Bell System. It gives the blocking probability when Poisson traffic (i.e., traffic with random call arrivals) is offered to a group of c trunks and blocked calls are held. That is, calls are assumed to remain in the system—either waiting or holding a trunk if one is available—for intervals that are independent of whether they were initially blocked.

The Erlang loss formula,

$$B(c, a) = \frac{a^c}{c!} \Big/ \sum_{n=0}^{c} \frac{a^n}{n!}, \tag{4}$$

published in 1917 by A. K. Erlang of the Copenhagen Telephone Company, also assumes Poisson arrivals but assumes that blocked calls are cleared; i.e., blocked calls are assumed to leave immediately and to have no further effect on the telephone system. Note that both the Poisson and Erlang formulas are independent of the distribution of holding times.

While the Erlang formula was derived from an apparently more realistic assumption of blocked call behavior, it, in practice, underestimated the measured blocking corresponding to an average busy season load a. Accordingly, after considerable discussion during the 1920s, AT&T decided to use the Poisson formula, which predicts a higher probability of blocking than the Erlang formula. It was generally thought that blocked call behavior, while not the same as assumed for the Poisson formula, was the prime reason for the Poisson formula's superiority. Much later it was found that day-to-day load variation was the major underlying cause.

2.2.2 *Wilkinson's Equivalent Random method*

The assumption of Poisson (random) call arrivals accurately models most first-offered traffic, but does not adequately describe overflow traffic from high-usage trunk groups. That is, overflow traffic is more variable than Poisson traffic because it arrives in bunches; consequently, more trunks are required than the Poisson formula predicts when final groups receive overflow traffic. Accordingly, in 1956, shortly after the introduction of automatic alternate routing of customer-dialed calls, R. I. Wilkinson developed the Equivalent Random method to size final trunk groups in an alternate routing network.[2]

The basis of Wilkinson's method, as illustrated in Fig. 3, is the assumption that the superposition of the individual overflows offered to a final trunk group with c trunks can be represented by a single overflow from a (fictitious) group of s trunks with Poisson offered load a^*. The parameters a^* and s^* are chosen so that the resulting overflow has the same mean, α, and variance, v, as the actual traffic offered to the final group. With this approximation, the Erlang loss formula, with $c + s^*$ trunks and offered load a^*, can be used to estimate the blocking on the final group. That is, the Equivalent Random approximation to blocking probability, which simulation studies have shown to be remarkably accurate, is given by

$$B(c,\ \alpha,\ z) = \frac{a^*}{\alpha}\, B(c + s^*,\ a^*), \tag{5}$$

where $z = v/\alpha$ is called the peakedness of the traffic offered to the final. Formulas for computing z, a^*, and s^* are given in Ref. 1.

2.2.3 *Day-to-day variation*

Wilkinson showed that day-to-day load variation causes trunk group average blocking to be significantly higher than that predicted under the assumption of a constant offered load.[3] Furthermore, he showed that the distribution of measured daily offered loads is well approximated by a gamma distribution, $\Gamma(\alpha \mid \bar{a},\ v_d)$, with mean $\bar{a}$ and variance v related to the mean $\bar{a}$ by

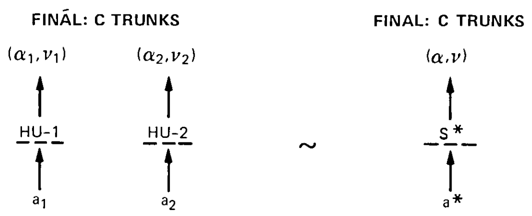

Fig. 3—Wilkinson's Equivalent Random method.

$$v_d = 0.13\bar{a}^{\phi}, \tag{6}$$

where ϕ is a parameter that describes the level of day-to-day variation. Wilkinson's studies—which showed that v_d is relatively larger for overflow traffic than for Poisson traffic—led to the use of three values of ϕ for engineering applications: $\phi = 1.5$, $\phi = 1.7$, and $\phi = 1.84$, which are referred to respectively as low, medium, and high day-to-day variation.

Using this model for day-to-day variation, Wilkinson's formula for predicting trunk group average blocking is given by

$$\bar{B}(c, \bar{a}, z) = \int_0^\infty B(c, \alpha, z) d\Gamma(\alpha \,|\, \bar{a}, v_d). \tag{7}$$

Since it was not practicable to apply Wilkinson's method without the use of a computer, his now famous $\bar{B}$ capacity tables were not available in the Bell System until about 1970.

2.2.4 Neal-Wilkinson tables

In 1976, Hill and Neal refined Wilkinson's $\bar{B}$ tables by developing mathematical models that account for the effects of the finite (one-hour) measurement interval.[1] That is, the service objective is defined in terms of the expected single-hour blocking, $E(O/PC)$. Thus, since eq. (7) provides an estimate of the probability of blocking, $E(O)/E(PC)$, it must be modified to remove the assumption that the measurement interval is infinite. Furthermore, Hill and Neal modified the formula for day-to-day load variation to account for the fact that part of the measured variation in daily offered loads is due to the finite measurement interval and must, therefore, be subtracted from the observed variation. Their work led to the new Neal-Wilkinson trunk group capacity tables, which are now the Bell System standard for sizing final and only-route trunk groups.

2.3 Traffic measurements and data transformations

This section describes the traffic measurements used to monitor network service, and the conversion of these measurements into estimates of the traffic parameters used to design the network.

2.3.1 Measurements

The Trunk Servicing System (TSS) derives estimates of average blocking on final groups. These estimates are used to detect service problems that trigger the demand servicing action described in Section 2.5. In addition, TSS derives estimates of trunk-group offered load, peakedness, and day-to-day variation. These estimates are used in demand servicing to correct existing service problems and in planned servicing to forecast future trunk requirements. These estimates are derived from three traffic measurements: peg count, overflow, and usage (i.e., carried load).

In electromechanical offices, usage is measured in units called CCS (hundred call seconds) by a traffic usage recorder (TUR) which scans the trunk group every 100 seconds and increments a register for each busy trunk. Thus, the usage or carried load in erlangs (average number of busy trunks) is obtained by dividing hourly register count by 36; i.e., one erlang equals 36 CCS. In electronic switching system offices, a similar process is used to estimate usage, but there is no need for specialized equipment since traffic data are collected by the switching system's central processor.

Because of the discrete scanning (i.e., once every 100 seconds) there is, of course, some statistical sampling error associated with the usage estimate. However, our studies have shown that this error is negligible compared with the unavoidable statistical errors associated with the finite measurement interval.

2.3.2 Data transformations

The process used to derive estimates of blocking, offered load, peakedness, and day-to-day variation consists of three basic steps: data validation, computation of traffic parameters, and data substitution. Data validation is a mechanized process used to determine whether there are unusual traffic conditions (indicated by unusually high blocking due, for example, to a snowstorm) or problems in the data collection process (indicated by inconsistent measurements such as overflow exceeding peg count or usage exceeding 36 CCS per trunk). Such data must be detected and deleted; otherwise they would propagate through the trunk provisioning process and possibly cause unnecessary trunk augments to be made.

Under normal conditions, when peg count (PC_i), overflow (O_i), and usage (U_i) are available for each day in the study period (normally twenty consecutive business days), the required traffic parameters are estimated as follows:

1. The study period average blocking, b, is estimated as given by eqs. (1) and (2).

2. The study period offered load a (in erlangs) is estimated by

$$a = \frac{1}{n} \sum_{i=1}^{n} a_i, \tag{8}$$

where

$$a_i = \frac{u_i/36}{i - b_i} \tag{9}$$

is an estimate of the daily offered load.

3. The observed variance, v, of the daily offered loads is estimated by

$$v = \frac{1}{n-1} \sum_{i=1}^{n} (a_i - a)^2. \tag{10}$$

Peakedness is not directly measured. Instead, to reduce data collection costs, a value of z is inferred from the relation

$$b = B(c, a, z), \tag{11}$$

where c is the number of trunks in the group and $B(c, a, z)$ is the equivalent random blocking formula.

Since there are many cases when complete UPCO (usage, peg count, and overflow) data are not recorded or are invalidated, TSS includes procedures to estimate traffic parameters with less than complete UPCO data. For example, if PCO is available but U is not, and if $b_i \neq 0$, an estimate of the daily offered load a_i is obtained by solving the equation

$$b_i = B(c, a_i, z), \tag{12}$$

where a typical value for z is assumed. Alternatively, if U is available, but PCO is not, the daily blocking is estimated by solving the equation

$$u_i/36 = a_i[1 - B(c, a_i, z)] \tag{13}$$

for a_i; then

$$b_i = 1 - \frac{u_i/36}{a_i}. \tag{14}$$

To minimize the impact of sampling error, TSS requires a minimum of three days per week of both peg count and overflow or three days per week of usage measurements. If less than this amount of data is available, the traffic parameters are estimated by using the most recent historical values; this process is called data substitution.

2.4 Trunk forecasting

The Trunk Forecasting System (TFS) is used by most Bell operating telephone companies to provide forecasts of interoffice message-trunk requirements for each of the next five years. The process consists of two basic steps: the estimation of future busy-season traffic loads and the design of a traffic network that minimizes the cost of the trunks required to satisfy these anticipated demands.

Since the network design process determines the traffic loads that must be forecast, we first discuss the concepts underlying the design of minimum cost traffic networks and then describe the load forecasting process.

2.4.1 Network design

The objective of the network design process is to determine the number of trunks in each high-usage and final trunk group that minimizes the cost of the trunks provided subject to the constraint that the average blocking on final groups does not exceed 0.01 in any hour.

To illustrate the basic concepts, we first consider the simple case of engineering the network shown in Fig. 4 for a single hour. The problem is to determine the value of n, the number of trunks in the high-usage group, which minimizes the cost function

$$\text{COST} = nC + N_a C_a, \tag{15}$$

where C is the cost per trunk on the high-usage route, C_a is the cost per circuit on the alternate route (which, in reality, consists of at least two trunk groups and a tandem switch), and N_a is the number of alternate route circuits required to meet the service objective. The load offered to the high-usage group is a; the traffic offered to the alternate route, A, consists of overflow traffic from the high-usage group under consideration plus "background traffic" consisting, in general, of overflow from other high-usage groups and first-routed traffic.

Qualitatively, the trade-off is between the less expensive high-usage trunks (i.e., the direct route has no switching cost and is usually physically shorter) and the more efficient alternate route circuits (i.e., the total number of circuits, $n + Na$, is minimized by providing all circuits in a common group where they can be shared by all traffic).

Since N_a depends upon both the mean A and the peakedness Z of the traffic offered to the alternate route, we should account for the change in both A and Z as n is varied. However, the procedure is considerably simplified without significant loss in accuracy by ignoring the variation of Z with n and by assuming that the rate of change of A with respect to N_a is a constant, γ. Thus, with these approximations, the condition for minimum cost may be written as

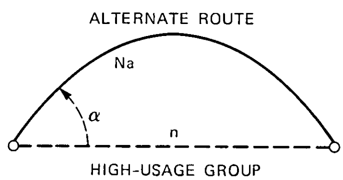

Fig. 4—ECCS engineering.

$$-\frac{dA}{dn}=\frac{\gamma}{C_R},$$

where $C_R = Ca/C$ is called the cost ratio.

Since only the overflow portion of A varies with n, and since the overflow load is $aB(n, a)$, where $B(n, a)$ is the Erlang-B blocking formula, the value of n that minimizes the COST is determined by

$$-\frac{d}{dn}[aB(n, a)] = \gamma/C_R. \tag{16}$$

The quantity of the left is approximately $a[B(n-1, a) - B(n, a)]$ and is called the load carried on the last trunk. The quantity on the right is called the economic load on the last trunk or, when the load is expressed in CCS, the ECCS; in fact, a more careful derivation, which recognizes that n is restricted to integer values, shows that the optimum value of n is the largest value for which the load on the last trunk is greater than or equal to the ECCS.

If the time-consistent busy hour for each group in the network were the same, that hour would be the only one needed in the sizing process. However, since such complete coincidence seldom occurs and since the existing algorithms are designed to use only a single-hour load, the question arises as to which hour to use to size each group.

At first glance, the solution may appear simple: since the service objective must be satisfied in all hours, final-trunk groups must be sized for their busy-hour loads, and high-usage groups should be sized for their offered loads in the alternate route's busy hour. However, there are two problems here. First, the various legs of the alternate route may have different busy hours and second, these busy-hours may depend upon the size of the subtending high-usage groups.[4] Thus, in theory, it is not possible to preselect a single hour to use in designing a minimum-cost network that satisfies the service objectives in all hours.

In practice, however, a heuristic, called the significant-hour method, has been found to produce networks that do not differ significantly from those obtained by using all hourly data. To illustrate this method, consider the two-level network shown in Fig. 2. For group AB, for example, two significant hours are considered: the A-office originating cluster busy hour (the hour for which the total load offered to the set of high-usage groups and the final originating at A is maximum) and the B-office terminating cluster busy hour (the hour for which the total load terminating at B is a maximum). Of these two hours, which are preselected in TSS from current traffic measurements, the one for which the AB load is larger is called the control hour and the group is sized using its control-hour load. A more complete discussion of this method is given in Ref. 4.

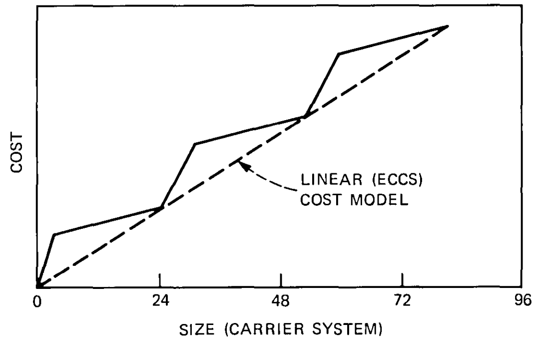

Fig. 5—Modular engineering.

The trunk quantities provided by the ECCS method must be modified to account for several factors. First, the administrative cost for a high-usage group is not included in the ECCS model. Since this cost is independent of group size, its inclusion leads to a prove-in or minimum group size; for larger groups it plays no role in the determination of the optimum size. While the exact prove-in size depends upon the specific costs, in practice we use a three-trunk minimum for local networks and a six-trunk minimum for toll networks.

Second, the ECCS model assumes a linear relationship between the cost and size of a high-usage group. However, for carrier systems, such as the T1 system, which provides the capacity for 24 time-division multiplexed channels on two pairs of wires, the cost is the step-like function illustrated in Fig. 5. Consequently, the ECCS solution is rounded to the nearest module of 24 for two-way groups or 12 for one-way groups. This rounding procedure is called modular engineering.

Finally, the ECCS method does not address the dynamic aspects of the trunk forecasting process, specifically, the time-varying nature of demands (which may call for the removal of trunks in one year only to have them added back in a future year) and forecast uncertainty. Currently, these factors are accounted for by manual adjustments to smooth the trunk requirements and avoid uneconomical disconnects and to introduce some reserve capacity to limit the amount of activity required in demand servicing.

2.4.2 Load forecasting

The objective of the load forecasting process is to estimate future busy season (control hour) first-route loads for each trunk group. First-route load, i.e., total offered load minus overflows from subtending high-usage groups, is projected since future offered load depends, of course, on the future sizes of subtending high-usage groups.

The standard load-forecasting algorithms currently in use in the

Bell System obtain estimates of future busy-season first-route loads by multiplying the most recent measurement of busy-season first-route load by an aggregate growth factor, for example, the average of the growth factors obtained by trending the total office loads at each end of the trunk group. Descriptions and comparisons of the various algorithms currently in use are given in Ref. 5.

As we explained in Ref. 5, the existing algorithms have two significant sources of error: First, on account of the finite amount of data, measured loads can have large statistical errors; standard deviations fall in the range of 5 to 10 percent for trunk-group data. Second, individual trunk-group growth factors can differ from the aggregate growth factor. These errors are significant since they lead to an increase in the reserve capacity required to satisfy anticipated customer demands.[6]

To reduce forecast error, and hence reserve trunking capacity, a new algorithm, called the Sequential Projection Algorithm (SPA), has been developed to forecast busy-season traffic loads within the Bell System.[5,7] SPA is based on a linear Kalman filter model, which establishes a linear trend for individual trunk group loads, together with logic for detecting and responding to outlier measurements. A complete description of SPA is given in Ref. 7.

2.5 *Trunk demand servicing*

To compensate for the effects of forecast error, demand servicing determines which trunk groups should be augmented to satisfy current busy season demands. The process uses current traffic measurements to detect service problems on finals (blocking significantly greater than objective) and to determine which high-usage and/or final groups should be augmented to correct these problems in a cost-effective manner. In theory, the demand-servicing process could also direct the removal of trunks in groups providing significantly better than objective service. In practice, however, the decision to remove trunks is, generally, part of the trunk-forecasting process described in Section 2.4.

The first step in demand servicing is to determine when action should be taken. This decision is based upon the observed average blocking b defined by eq. (2). Specifically, when b exceeds a threshold b_u, the final group is declared to be overloaded and corrective action is taken; otherwise the measured blocking is assumed to be acceptable.

Because of the statistical nature of the demands, the finite number of days in the study period, and the finite (one-hour) measurement interval each day, the measured blocking can be smaller or larger than the underlying true mean blocking.[8] For example, Fig. 6 shows a

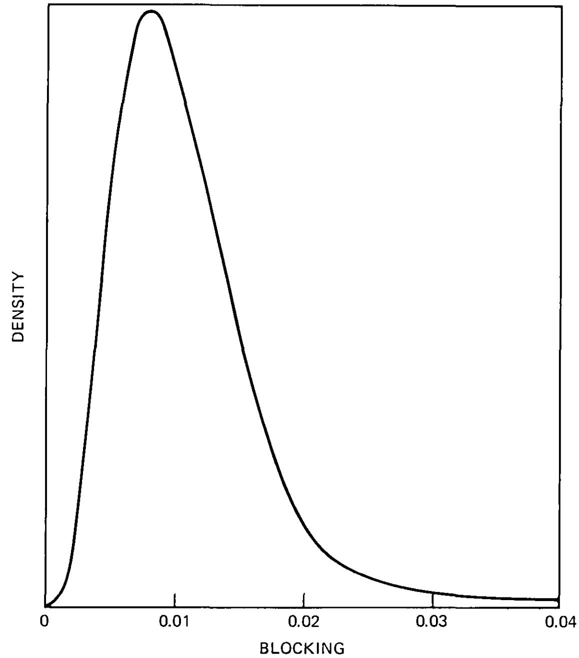

Fig. 6—Distribution (density) of measured blocking.

typical distribution of observed blocking for a correctly sized final group.

Accordingly, the threshold b has been selected to achieve a reasonable balance between two types of servicing errors: false alarms, or Type I errors, occur when b exceeds the threshold but the group is not overloaded; misses, or Type II errors, occur when a service problem is not detected; i.e., when the measured blocking falls below the threshold but the group is underprovided. The threshold u has been selected to give a false alarm probability of less than 2.5 percent; in most cases this threshold corresponds to a miss probability of less than 10 percent when the group is overloaded to at least a 0.05 blocking level.[9] For the standard 20-day study period, the threshold b_u falls in range of 0.07 (small groups) to 0.02 (large groups).

If a service problem has been detected, the next step is to determine which groups to augment. At present, this decision is based on the servicer's judgment, but a new trunk-demand servicing policy, described below, has been developed for possible inclusion in TSS.[9]

Although the simplest procedure would be to augment the final group, this procedure would tend to drive the network away from the optimal balance between high-usage and final trunks. Consequently, the demand servicing policy should attempt to resolve a service problem at its source. That is, the problem should be corrected by augmenting those groups—starting at the lowest level in the hierarchy—that are significantly underprovided (based upon the ECCS design) and are contributing to the service problem on the final group.

Since the determination of when a high-usage group is underprovided is subject to the same sources of statistical error as blocking estimates, thresholds that complement those for final groups have been developed for high-usage groups.[9] Thus, when the difference between the required number of trunks and the number of trunks in service exceeds the threshold, the group is a potential candidate for servicing. However, since only those groups contributing to the blocking problem are considered as candidates for servicing, an overloaded primary high-usage group that subtends an acceptably loaded intermediate high-usage group is not selected as a candidate.

Thus, as explained in Ref. 9, the proposed demand-servicing process is an iterative one that starts at the top of the hierarchy (the final) to identify those undertrunked groups that subtend undertrunked groups at the next higher level; such groups are selected as candidates for servicing. The lowest-level candidate, whose servicing will contribute the maximum reduction in load offered to the final, is then augmented. The blocking on the final is then recomputed. If a problem still exists, the trunk requirements and candidacy of all groups is then reevaluated. The algorithm is terminated when the blocking on the final has been reduced to an acceptable level.

III. SWITCHING SYSTEMS

3.1 Service objectives

Switching systems typically divide the process of switching calls into functional parts. Each part may require specialized equipment, such as customer receivers to receive dial pulses, or dedicated areas of memory in which to store information relative to the function, such as the number dialed. Equipment and memory of this kind are generally provided only as required to handle the traffic of each installation—a traffic engineering job requiring data from TNDS. The central processor which time-shares the control of most functions has a finite capacity, so it too must be traffic engineered. Because of the different kinds of traffic and the different function of a switching system in serving customers, the service criteria and hours of measurement differ, not only from those applied to trunks but even among switching systems according to their place in the network. This and the next section discuss service objectives and measurement timing; the remaining sections discuss application.

In setting service criteria to balance customer service against the cost of providing it, a comprehensive consideration has to be given to what blocking or delaying calls means to a customer. In toll switching and in trunking, any particular subscriber is likely to place only a small proportion of calls through a particular trunk group or switch. If average blocking is achieved here, the customer is not likely to build up a feeling of frustration when a call is blocked; the next call is very likely to go through. In a local office, however, an overloaded switching system can affect all calls placed by the customer, regardless of destination, and blocking may remain high for a considerable time. The feeling of frustration can be much higher therefore, and service criteria must be chosen for closer control of service than can be achieved by considering averages alone. Day-to-day variation is, obviously, most important. Another important factor is balance; during the high-traffic hours, a switching system can give excellent *average* service while giving consistently poor service to a subset of customers.

A further difference applies to administration. As point-to-point loads vary with time, the components from which trunks are constructed can be reconfigured to meet changing demands. Very little of this kind of administration of equipment can be done with a switching system. Additions can normally be made only at the end of an engineering period, which is usually two years.

In addition to criteria of call blocking, which are expressed in terms of fraction of calls not completed because of congestion, delay criteria are used for engineering and administering switching systems. Delay criteria, which were not discussed in the section on trunking, are expressed in terms of the probability of exceeding an objective delay time. They require measurement methods different from those of load or loss. Ideally, delays would be measured separately for each call; however, a very busy period is a poor one for the switching system to divert its activity from call processing to service measurement. Instead, the usual delay measurement is made by placing test calls, at fixed intervals either within the system or "externally" by an attached test-call originator. In this way, a fixed amount of time is devoted to the measurement which is independent of the load on the system. The total number of test calls and the number of test calls that are delayed over an objective time are reported to the data acquisition system of TNDS.

3.2 Measurement timing

As with trunking, most traffic measurements of switching systems are based on a concept of a busy hour, that is, a single clock hour of the day such as 9:00 to 10:00 am, which is identified as the hour of the day in which, on the average, more traffic is carried than in any other. It has long been recognized that the busiest hour of a particular day may not coincide with the "busy hours" as defined above. A newly developed engineering concept to take this into account is described at the end of this section.

As we mentioned previously, definitions involved in traffic measurement differ from those for trunks. Various "engineering hours" used in switching are described in the following sections.

3.2.1 Average busy-season busy hour

Through special studies, often made at the start of the busy season (which is itself usually well known through previous data) by means of special programs available in TNDS, the hour of the day that exhibits the highest average load, when averaged over the selected days of the study period, is identified as the busy hour. The load may be number of calls, usage, or both, depending on the type of switching system component being measured. Different components may have different busy hours. Measurements are taken in the busy hours on all days throughout the year. The traffic measurements of those three months that have the highest average busy-hour traffic are averaged to get the average busy-season busy-hour traffic. (Weekends and holidays are usually excluded from this calculation.) Provision is made within TNDS to measure loads in more than one candidate busy hour, to handle cases where different switch components have different busy hours, or when two hours are so close together in load that a clear decision between them cannot be made ahead of time.

3.2.2 Average 10-high-day busy hour

As the name implies this measurement period still uses the busy hour; however, the average usage or call volumes of the ten highest busy hours of all the months of the year are used rather than just those of the busiest three months.

3.2.3 High-day busy hour

This is the highest of the 10-high-day busy hours for a given year. Provision is made for deleting measurements made on days of a very unusual type that are not expected to recur. Service impairment on days with such unusual traffic extremes is generally accepted by the public because the cause is obvious: blizzard, flood, tornado, local or national disaster. Provision of sufficient equipment to handle such extremes without service impairment is uneconomical; network management serves under such circumstances to help the network complete as many calls as possible.

To assist the traffic engineer in evaluating the 10-high-day and high-day busy-hour values, a comparison is made by the TNDS Central Office Equipment Report (COER) system based on a theoretical model of traffic variation. Studies have shown that the probability distribution of the busy hour traffic carried by central office components can often be approximated by a gamma distribution. The gamma distribution is determined by its first two moments, so that a fitting curve can be generated from the first two moments of the observed data. COER makes such a calculation for the current year of data being collected and prints out in parallel columns the values of the observed highest 15 measurements and the values computed for a gamma distribution with the same mean and standard deviation as all of the busy-hour measurements taken to date. A quick visual check will give the engineer some idea of whether this traffic is typical. In particular, an impression can be given of whether there is more or less volatility in the traffic under study. Also, if high-day engineering is required, the difference between actual high-day load and the gamma projection will help the administrator judge whether the high-day load may have been a data outlier.

3.2.4 Extreme-value measurement period

This type of measurement period differs markedly from the three previous ones. As the costs of collecting data have changed, it has been possible to do a limited amount of processing at the collection point so that long-term storage of many measurements can be eliminated. The hour with the highest measured value for the day can be retained and all others discarded. The resulting measurement can be used to obtain reliable estimates of the extremes of traffic and service to be encountered. A different criterion of service is possible now; instead of average or high-day service, the objective can be set in terms of the expected number of hours or days in which service fails to meet a given criterion. For example, one objective now used is that on the average no more than one day a month will contain an hour in which dial tone delay exceeds 8 percent over 3 seconds. Because the criterion applies to only one day of the month, the probability of delay can be made significantly higher than the average 1.5 percent over 3 seconds used for average busy-season busy hour. In this example the criterion was picked so that in an average office there would be no noticeable change in service on changing from busy-hour engineering to extreme-value engineering.

Extreme-value methods are still being incorporated into the TNDS. Studies of the statistics of extremes indicate that extreme-value methods result in much more accurate projection of peak period loads than do current methods.[10,11] This leads to more accurate forecasting and engineering.

3.3 Types of equipment

Measurement periods and the accompanying criteria of service vary among the three general types of central office components that are described next. Also in this section, we discuss both load (traffic intensity) and call volume. For example, the common control elements of a system are affected primarily by the number of calls processed, while the switching network is affected primarily by the average number of simultaneous conversations in progress.

3.3.1 Switching network

The capacity of the switching network portion of a switching system is almost always a function of the conversation load. Load capacity is specified in units of erlangs or CCS. (See Section 2.3.1.)

Electromechanical switching networks in general react with rather slowly rising blocking to increasing load, so that a criterion of loss averaged over the busy hours is used similar to the trunking criterion. On the other hand, time-division networks generally display very low blocking until high load is reached where blocking rises quickly.[12] In the latter case a high-day criterion is used.

The theoretical reasons for this can be seen by considering the effect of the number of parallel paths on the load-loss characteristics. Blocking in a typical three-stage network tends to follow the simple formula:[13]

$$B = [1 - (1 - a)^3]^n, \tag{17}$$

where:

B = blocking (loss)
a = average load in erlangs carried by a link connecting two stages
n = number of parallel paths.

This formula is plotted in Fig. 7 for 10, 30, and 128 parallel paths to show the effects that different blocking characteristics may have upon the selection of network service criteria. For illustration purposes, the ratios of high-day to 10-high-day to average busy season, busy hour are selected at 1.15:1.10:1.00. If the objectives are $B = 0.02$ for the average busy-season busy hour, $B = 0.05$ for the 10-high day average, and $B = 0.10$ for the high-day, three different networks might require three different engineering criteria.

The curves show the result of having selected the appropriate criterion. For the network with ten parallel paths the *ABS* criterion is appropriate; the 10-high-day and high-day service will be better than required. For a network with 30 parallel paths, however, the 10-high-day criterion is picked; the high-day and *ABS* service will be better than required. Finally, for the network with 128 parallel paths only the high-day is of concern; essentially no calls will be blocked on any other day. In fact, for this network it is likely that a service criterion will not be used directly, but that a maximum occupancy for the high-

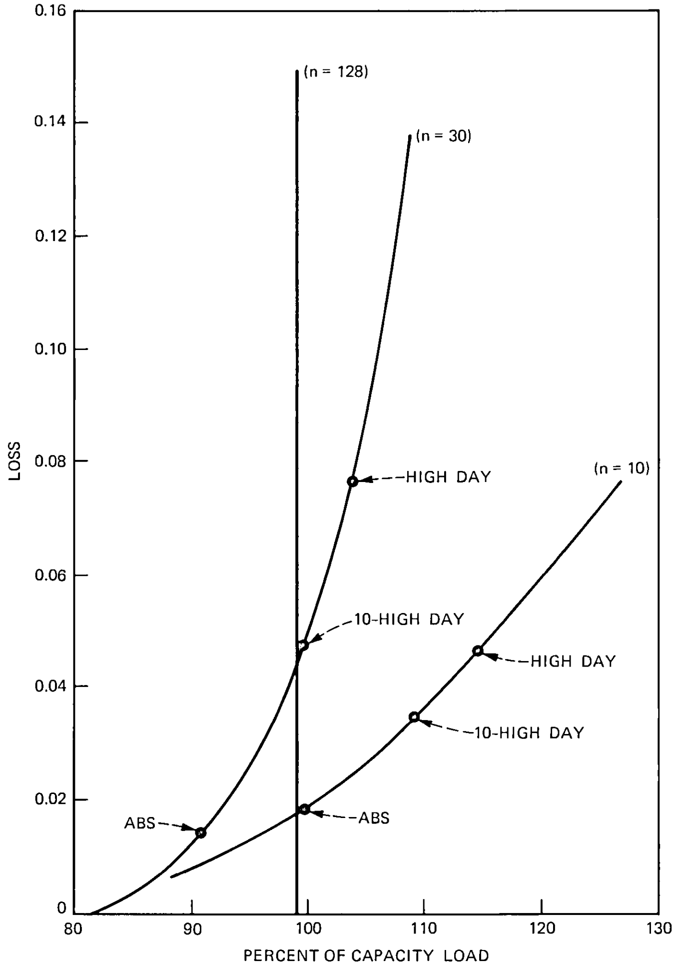

Fig. 7—Effect of parallel paths on service criterion—based on Loss = $[1 - (1 - a)^3]^n$.

day load will be adopted in order to leave room for error in predicting the high-day load.

3.3.2 Service circuits

Service circuits are the circuits that are provided in switching systems to give tones or to record customer dialing, to accept in-band signals from other offices, or to transmit such signals to other offices. Depending on the size of the office, these are provided usually in small groups ranging from five to ten, but they may occur in large groups of well over 100. Because of the way in which service circuits are used by the switching systems, service circuits delay rather than block customers' calls during periods of congestion. The delay criterion of service used here is expressed in percent of calls delayed over a specified number of seconds. Depending on the group's size and holding time, the average, 10-high-day or high-day criterion may be used. The choice of criteria depends to a large extent on the judgment of the traffic engineer. Customer reaction, cost of equipment, stage in the progress of a call, traffic volatility, and interaction with other parts of the switching system all play a part.

The capacities of many service circuit groups are determined from delay calculations using the erlang delay formula when service times are highly variable or using the Crommelin-Pollaczek delay curves when holding times are nearly constant.

3.3.3 Central processors and controls

The control components of switching systems that are mainly used in setting up a call generally have very short holding times and are provided in small quantities, often only one per system. The probability of delay is very high, but, because of the very low holding time per call, delays are usually unnoticeable. Criteria are based on delays on the highest day, because when delays start to become noticeable the processor is usually close to saturation, where no further increase in load can be handled. In such cases if additional load is offered, customer calls will be blocked. Customers try again when blocked, so the number of attempts will rise faster than the carried load. Eventually, so much processing time is used for uncompleted attempts that the number of good calls completed by the switching system is reduced.

3.4 Traffic engineering models

The engineering of most central office equipment is based on assumptions of random call originations with constant source loads during the busy hour but with day-to-day variation. In general, peakedness is assumed at unity, although it is known that there are phenomena, such as customer retrials, that can produce some of the effects of peakedness. Also, toll offices high in the hierarchy will handle a large quantity of overflow traffic, and thus peaked traffic. The effect of this peakedness may seriously affect capacity in the periods of overload. In spite of the differences from the model underlying the Poisson formula, that is, constant traffic intensity and random call arrivals, the formula has been the work horse for much of central office engineering because it was the best choice in predicting actual system performance.

The Poisson formula is not applicable to internal switching networks. Formulae of the general form of eq. (17) give remarkably close approximation to the performance of this type of network. Observations show that a particular network will display a load-loss relationship that parallels such a theoretical curve. Differences may be attributed not only to the simplified assumptions of the formula but also to imbalance effects, day-to-day variation, systematic load variation within the hour, or minor variations in the network structure from that assumed in the model. The TNDS supplies data from which comparisons of actual with theoretical load-loss relationships can readily be made.

Common control equipment, including call processors, present a different traffic modeling problem from what has been considered so far. Loss or delays are insignificant most of the time. The event of call processor overload is so infrequent that measurements are usually too few for making useful load-loss curves to compare with theory.

With electromechanical switching systems, the proportion of time that a processor, such as a marker, is busy can be measured with little complication. A plot of occupancy against calls gives a clear picture of the number of calls that will fully occupy the processor. Capacity is usually picked at a lower number, say 95 percent of full occupancy, and applied to high-day busy-hour traffic.

A stored program controlled processor, however, is busy all of the time; when it is not processing calls, it is performing maintenance or audit functions or looking for work. The control program is designed to reduce the time spent in non-call processing activities as the call processing load increases. In many systems, this is done by organizing the processing into a fixed order of tasks. Processing jobs that must be done with little delay are given priority over those with less stringent time requirements. So that some maintenance and audit work is done even in periods of heavy load, the search for work is made in a cyclical fashion over the various priorities with many repeat searches for high-priority work. Thus the time taken between visits to the lowest priority level becomes longer with increased processing load. The capacity of the processor, in terms of high-day busy-hour calls, is then found by analyzing observations of offered originating-plus-incoming calls with the corresponding observations of basic processor cycles, i.e., visits to the lowest priority level. A first-order least-squares fit to these points is made and a 90-percent confidence line is computed and drawn below the fitting line (90 percent of the observations are expected to fall above this line). Next a horizontal line is drawn at the minimum number of visits that allow acceptable service as determined by simulations and field studies of many systems. The intercept of the 90-percent line and the minimum-number-of-visits line gives the estimated call capacity for the high-day busy hour. This is illustrated in Fig. 8.

The choice of 90-percent confidence is, of course, a judgment choice.

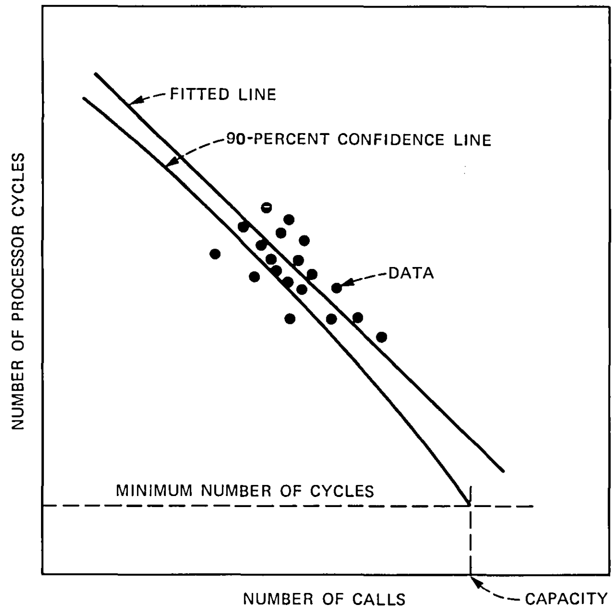

Fig. 8—Use of measurements to determine central processor capacity.

It is based on experience gained from many switching systems. A load greater than capacity will often cause no problem; on the other hand, a load less than capacity may rarely overload the processor.

3.5 Load balance

The switching networks of most switching systems require balancing of offered traffic over the network switches and frames. Unbalance arises from the nature of traffic sources—some lines or trunks generate more traffic than others. Correction of unbalance is particularly needed in the concentrating stages of the network, where the number of input lines exceeds the number of output links. While the first objective of load balancing is to equalize service among customers, it also achieves a second objective of improving the load-carrying capacity of the network. The latter effect of load balancing arises from the nonlinear load-loss curve that exists for most networks. As Fig. 9 illustrates, the higher loss in a section of network loaded above average is not completely compensated for by the lower loss in a section loaded, by the same amount, below average. Therefore, reducing the spread of the distributions of loads around the mean will reduce average blocking at a given load or enable increased loading at a given service objective.

It is an objective of line-load balancing to control the assignment of

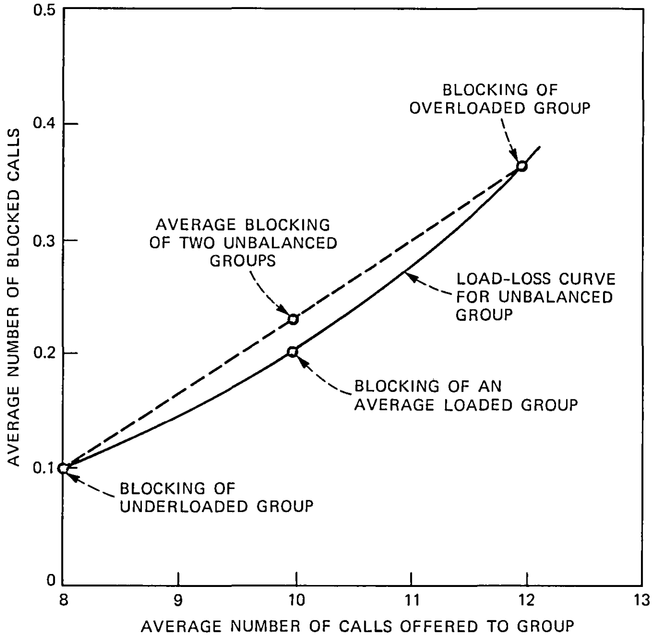

Fig. 9—Effect of unbalance on blocking.

new lines in such a fashion that all units stay in traffic balance. In a 55,000-line 1A *ESS** switching equipment for example, this implies keeping track of a number of separately tracked "loading modules" that may approach a thousand. Even in smaller offices the problem is far from trivial. The measurement and processing systems of TNDS provide the means of controlling load balance—without excessive manual effort—by accumulating line-group usage for all of the designated hours of the week. Also, measurements of individual line usage are taken for guidance in selecting lines when extreme unbalance requires shifting lines from one particular group to another.

The first step in balancing concentrator loads is to assign lines in such a fashion that every concentrator has the same number of lines of each identifiable traffic or class of service. Unfortunately, this is not enough to ensure balance.

The first systematic application of traffic measurements to load balance was reported in a study by D. H. Barnes in 1958.[14] Barnes proposed a method of statistical control that would separate effects of chance load variations from those of systematic load balances. It was implemented on the computers of that day as well as in a manual

* Trademark of Western Electric.

procedure (the "score method") that would permit the calculations to be made with simple mathematical operations.

The type of solution proposed by Barnes has been refined over the years as more computing power has become available and more measurement refinement has become possible. The operational problem to be solved remains the same however: to separate the random variations from the systematic ones out of the large variations that occur in the hourly measured traffic on a concentrated group of lines. This is the same problem faced in the design of a statistical quality control procedure. It involves estimating load variance and measurement error, and displaying the results in a form that will enable managers to take appropriate action without moving lines needlessly or letting service deteriorate. The TNDS Load Balance system takes over the computation of balance indices and processing of the balance data.

IV. NETWORK MANAGEMENT

The background of network management is covered in a companion article to this one and in Ref. 15. The main difference in the requirements placed by network management on TNDS from those of other traffic data applications is the very fast response required between traffic events and the reporting of them. Because of the long-established need to keep network managers informed within seconds of a major control loss or within minutes of serious call completion problems, the traffic-gathering parts of TNDS must have—and do have—the capability of passing data quickly and accurately to the network management computer and relaying control orders quickly back to the measured switching systems.

V. CONCLUSION

Throughout the evolution of TNDS there has been a continuous increase in the accuracy and availability of traffic data. This in turn has made possible more accurate traffic models and has spurred activity for better traffic theory and better network operating practices. What has been reported here is the state of the art at this time. The coming years may be expected to bring more changes and improvements.

REFERENCES

1. D. W. Hill and S. R. Neal, "Traffic Capacity of a Probability-Engineered Trunk Group," B.S.T.J., *55,* No. 7 (September 1976), pp. 831–2.
2. R. I. Wilkinson, "Theories for Toll Traffic Engineering in the U.S.A.," B.S.T.J., *35,* No. 2 (March 1956), pp. 421–514.
3. R. I. Wilkinson, "Some Comparisons of Load and Loss Data with Current Teletraffic Theory," B.S.T.J., *50,* No. 9 (October 1971), pp. 2807–34.

4. M. Eisenberg, "Engineering Traffic Networks for More Than One Busy Hour," B.S.T.J., *56,* No. 1 (January 1977), pp. 1–15.
5. A. J. David and C. D. Pack, "The Sequential Projection Algorithm: A New and Improved Traffic Forecasting Procedure," 9th International Teletraffic Congress, Torremolinos, Spain, 1979.
6. R. L. Franks et al., "A Model Relating Measurement and Forecast Errors to the Provisioning of Direct Final Trunk Groups," B.S.T.J., *58,* No. 2 (February 1979), pp. 351–78.
7. J. P. Moreland, private communication.
8. S. R. Neal, "Blocking Distributions for Trunk Network Administration," B.S.T.J., *59,* No. 6 (July 1980), pp. 829–44.
9. C. R. Szelag, "Trunk Demand Servicing in the Presence of Measurement Uncertainty," B.S.T.J., *59,* No. 6 (July 1980), pp. 845–60.
10. D. H. Barnes, "Extreme Value Engineering of Small Switching Offices," 8th International Teletraffic Congress, Melbourne, Australia, 1976.
11. K. A. Friedman, "Extreme Value Analysis Techniques," 9th International Teletraffic Congress, Torremolinos, Spain, 1979.
12. J. F. Huttenhoff et al., "No. 4 ESS: Peripheral System," B.S.T.J., *56,* No. 7 (September 1977), pp. 1029–55.
13. C. Y. Lee, "Analysis of Switching Networks," B.S.T.J., *34,* No. 6 (November 1955), pp. 1287–1315.
14. D. H. Barnes, "Statistical Methods for the Administration of Dial Offices," 2nd International Teletraffic Congress, The Hague, 1958.
15. D. G. Haenschke, D. A. Kettler, and E. Oberer, "Network Management and Congestion in the U. S. Telecommunications Network," IEEE Trans. Commun., *COM-29,* No. 4 (April 1981), pp. 376–85.

AUTHORS

Walter S. Hayward, Jr., A.B., 1943, S.M. (Electrical Engineering), 1947, Harvard University; Bell Laboratories, 1947—. Mr. Hayward has worked in the field of telephone traffic and switching systems engineering. In 1961, he was appointed Head, Electronic Switching Studies Department. In 1964, he was appointed Director of the Traffic Studies Center. He is now Consultant, SPC Studies Center. Member, IEEE, ORSA, and ACM.

James P. Moreland, B.S.E.E., 1964; M.S.E.E., 1964; Ph.D. (E.E.), 1967, Ohio State University; Research Associate, Electroscience Laboratory, 1964–1968, Instructor, Electrical Engineering, 1967–1968, Ohio State University; Bell Laboratories, 1968—. At Ohio State, Mr. Moreland worked on studies of scattering theory and optical heterodyne detection. At Bell Laboratories, he has been concerned with clock-synchronization schemes for digital communications networks, optical-fiber transmission studies, and traffic and facility network planning. He is presently a Supervisor in the Trunk Traffic Engineering Department. Member, IEEE, Eta Kappa Nu, Tau Beta Pi, Sigma Xi.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Data Acquisition and Near-Real-Time Surveillance by EADAS

By C. J. BYRNE,* D. J. GAGNE,* J. A. GRANDLE, Jr.,* and G. H. WEDEMEYER†

(Manuscript received May 3, 1982)

The Total Network Data System (TNDS) requires a facility to collect the traffic data generated by electromechanical and electronic offices. The Engineering and Administrative Data Acquisition System (EADAS) fulfills this function. EADAS systems that serve electromechanical offices employ unique self-testing circuitry to interface to central office signals. A novel buffering scheme also improves system efficiency. For electronic offices, a specialized file system has been developed, and the input data are specified in a high-level language. All of these features permit EADAS to be a cost-effective, flexible, and reliable data collection system. EADAS also provides real-time surveillance and reports for the switching offices it serves. A set of reports is predefined; however, the user may also specify reports using the Network Operations Report Generator feature (NORGEN). System capacity is specified in simple formulas which are derived and presented. With these features, EADAS provides a comprehensive facility for presenting real-time traffic data to the telephone companies.

I. INTRODUCTION

High-quality telephone service requires facilities that collect and process data on the traffic being handled by a telephone system. In

* Bell Laboratories. † AT&T.

particular, it needs data on the number of calls being handled by the system, the duration of each call, and the number of times a call encounters difficulties as it proceeds through the system to its destination.

Such data are needed for surveillance of the quality of service being provided. This requires, for example, measuring in real time the interval required to obtain dial tone, and informing responsible personnel when such intervals exceed acceptable limits.

Traffic data are also used to determine the proper quality and distribution of central office equipment and trunk facilities. Other uses are to provide telephone customers with special reports on the performance of their own particular facilities and grade of service, and to assist telephone company personnel in proper maintenance of equipment.

In the Bell System, traffic data are collected and processed by several different systems. Among these are the Telesciences, Inc. "Automatic Traffic Recording and Analysis Complex" (AUTRAX) System, the Conrac Corp. "Alston Traffic, Engineering and Management Information System" (ATEMIS), and the Western Electric "Engineering and Administrative Data Acquisition System" (EADAS). The operation of EADAS will be described in more detail in succeeding paragraphs to illustrate how traffic data are typically collected and processed.

II. SYSTEM CONCEPT

Figure 1 shows the overall system concept of EADAS. A centrally located minicomputer collects data from the central offices. Within certain limitations described later, these offices can be electromechanical or electronic and can be large or small, with and without toll features.

Traffic data are collected via a data-link facility between each central office served and the EADAS central unit (CU). As mentioned above, EADAS provides information to telephone company network administration personnel for near-real-time surveillance on the quality of telephone service being provided. EADAS also makes available information to assist maintenance personnel in remedying equipment problems before they can affect service. This information is in the form of reports produced by the Network Operations Report Generator (NORGEN) feature of EADAS. These reports can be in printed form or displayed on a Cathode Ray Tube (CRT). NORGEN produces a wide variety of such reports, which are made available on demand, on a previously established schedule, or automatically when an abnormal service condition (exception) develops.

NORGEN was introduced in 1977 to replace the demand for re-

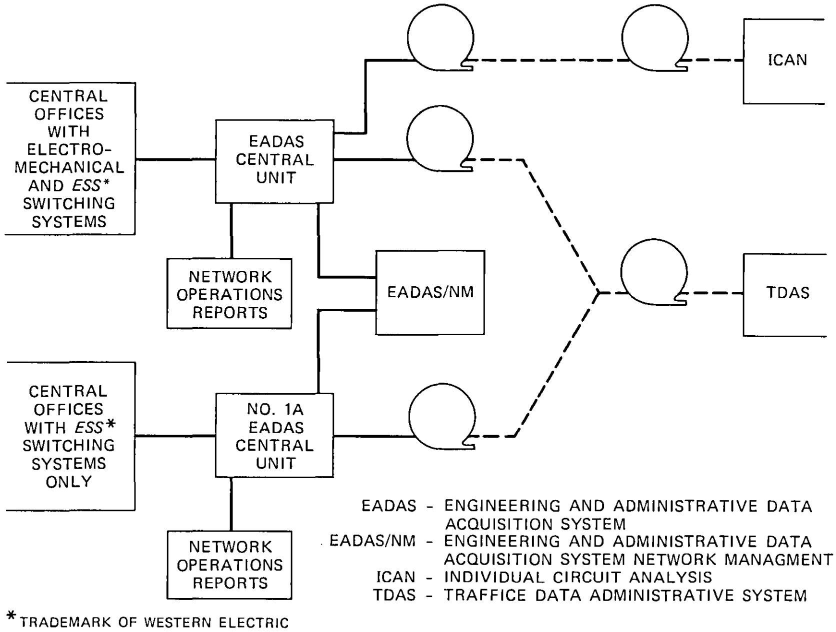

Fig. 1—The Engineering and Administrative Data Acquisition System concept.

porting features that had exceeded the calculation capacity and flexibility of the original reporting features of EADAS. NORGEN provides a calculation capacity six times larger and with greatly increased flexibility, as well as a set of standard applications programs designed to meet those near-real-time report needs required on a systemwide basis. A programmability feature allows the user to modify the standard reporting features and to provide new ones to meet local conditions.

Figure 1 shows that traffic data collected by EADAS can also be recorded at regular intervals on magnetic tapes. These tapes are physically transported to the Traffic Data Administration System (TDAS) and to the Individual Circuit Analysis (ICAN) program described in the article on TNDS equipment systems[1] in this issue. TDAS formats the data for use by other downstream software systems. These systems report on longer-term engineering functions, such as determining how much of a switching system's capacity is being used and forecasting future trunking requirements. ICAN analyzes subsets of the data collected by the Individual Circuit Usage Recording (ICUR) feature of EADAS to detect record base errors, equipment faults on individual circuits in electromechanical offices, and malfunctions in the data collection and processing functions performed by the central office and EADAS facilities.

Figure 1 also shows the data link between EADAS and the EADAS/

Network Management (EADAS/NM) System. Over this link, traffic data are sent at five-minute intervals to alert the local Network Management Center or Regional Operations Center to network problems. Signals from EADAS/NM, which activate various network controls such as trunk reroute, and signals to EADAS/NM, which indicate the status of these controls, are also transmitted over this link.

III. FIELD OF APPLICATION

There are two versions of EADAS. The original system, No. 1 EADAS,* uses a 16-bit minicomputer with 128K words (word equals two bytes) of memory. This system was designed to serve No. 1 and No. 5 Crossbar, Crossbar Tandem, and Step-by-Step type electromechanical offices, as well as 1/1A and 2/2B *ESS*† switching equipment, and Remote Switching System (RSS) offices. The No. 1 EADAS was first installed in Kansas City, Missouri, in 1973. In 1974, the ICUR feature was first installed in Miami, Florida.

A later version—the No. 1A EADAS—was designed to serve *ESS* switching equipment offices only. The No. 1A EADAS employs a faster minicomputer (three times the processor speed) with expanded memory capability (larger address range and use of cache memory). The No. 1A EADAS has approximately twice the capacity for serving *ESS* switching equipment offices, with a lower installed cost than No. 1 EADAS and the flexibility to grow with the continued expansion of Stored Program Control electronic switching in the Bell System. All No. 1A EADAS installations have the NORGEN feature. The No. 1A EADAS was first installed in New York City in 1977. The No. 1A EADAS is presently arranged to serve 1/1A, 2/2B, 3, and 5 *ESS* switching equipment offices, as well as RSS offices. In addition, coverage is provided for Northern Telecom's DMS-10‡ offices.

The No. 1 EADAS met the Network Administration needs of electromechanical offices. However, the Bell System trend toward adoption of Stored Program Control (SPC) for new and replacement offices is reducing the number of electromechanical offices in service. This, together with the advantages of No. 1A EADAS for SPC entities, has caused a leveling off in the number of No. 1 EADAS installations. On the other hand, the number of No. 1A EADAS continues to increase in proportion to the installation of Stored Program Control offices.

IV. DATA COLLECTION

The data collection process differs between the electromechanical offices and the electronic offices. Electromechanical data are collected

---

* The offical designation of the original system is simply "EADAS."

† Trademark of Western Electric.

‡ Trademark of Northern Telecom.

from several terminal types. The predominant terminal type provides real-time data, while data are also accumulated and held by other types of terminals in end offices. The accumulated data sources operate over either a dialed-up or dedicated link, resulting in three different interfaces for the electromechanical offices. On the other hand, electronic offices send only accumulated data to the EADAS, and each electronic system has a unique output format. The data collection process in electromechanical offices will be discussed first.

A number of requirements were considered in developing the electromechanical design:

1. It was advantageous to use the data collection equipment of the existing Traffic Data Recording System (TDRS) already installed in a number of central offices. The main concern here was the high cost of installing new traffic measurement equipment in electromechanical offices. This is caused by the large number of points to be monitored, which are scattered throughout the office and require many individual wiring runs.

2. Data were to be collected in real time to provide information for exception reports and the network management feature (see other articles on Network Management).

3. The system was to provide new features such as Individual Circuit Usage Recording (ICUR), described later, and office control functions for the network management feature mentioned above.

The traffic data equipment used by the then existing TDRSs, known as the Traffic Data Converter, operated as follows. The unit, located in the central office, is connected to the data collection center via a dedicated link. The occurrence of an event in the central office (i.e., generation of a call or detection of a busy on a usage scan) causes the converter to generate a data word representing the input address or connection point for the event. Thus, the function of the data collection center is to sort and sum the occurrences of these address words by address for the desired measurement period. While this function is simple in concept, it can quickly become complex to provide a capacity of 100 data channels, each capable of generating 1000 input addresses. Since all of the real-time data are available at the collection center at any instant, centralized accumulation is ideal for real-time data collection. Thus, the network management and real-time reporting functions can be readily accomplished.

4.1 Input network

The desire to provide new features and a capacity greater than that of the TDRS led to the design of a new converter for collecting data that could be used in those offices not already containing TDRS equipment. In addition to collecting peg counts, this design provided remote office control, fast scan usage data, individual input scaling (to reduce the data channel load for rapidly occurring events), concentration, and ICUR. This converter, called the EADAS Traffic Data Converter or ETDC, had a number of interesting features. Perhaps the most interesting is the input circuitry used to interface the system to the central office. This input circuit is shown in Fig. 2. The circuit has to provide a high impedance input which could monitor ground closures provided to relays (without changing operating characteristics), while absorbing inductive spikes and contact bounce associated with the relays. In addition, a periodic self-test should be provided. The input circuit provides these features as well as an input verification feature.

Meeting the high impedance requirement was a problem. Transistor-transistor logic (TTL) of the early 1970s employed 2K pull-up resistors, and a quick analysis will show that to achieve a low on the input, the maximum resistance to −48V could be no greater than 20K. This was not anywhere near the 100K to 300K desired. By power switching the TTL gate, the high input impedance can be obtained since the TTL gate is effectively removed until the input is scanned. A capacitor (C) added to the circuit to remove noise pulses serves as a charge storage point and holds the input state during the short interval when the power is switched onto the TTL gate. With this arrangement, the levels fed to the TTL gate are determined by the external voltage divider R_P and R_I.

This network is also testable. If the power switch is turned on permanently, all of the input capacitors will charge through the TTL pull-up resistor, and all of the input points should read a high state if

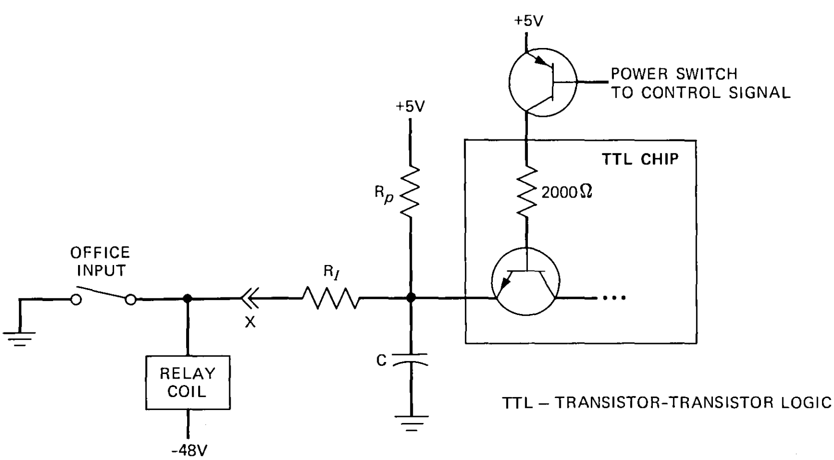

Fig. 2—Input circuitry used to interface EADAS to the central office.

a failure has not occurred. To test the low state, the +5V supply is removed from the resistor R_P. This results in all inputs being low since the voltage feeding the capacitor is either negative or zero.

Another test feature is available. Note that if a positive voltage is applied to the input at point X in Fig. 2, it is possible to introduce a high input on the TTL gate. Thus, if one takes a positive supply and pulses an input, it should register a count. This feature has two uses. First, input connections can be verified in a central office by going to the input source and pulsing it with positive voltage. Even though other inputs are active, only the positive input will count. Second, using this same technique, any crosses to other inputs will appear as additional counts or generation of additional addresses.

4.2 Electromechanical CU operation

Having described some features of the central office circuitry, let us turn to the central data collection operation. As stated earlier, this requires the ability to collect data from up to 100 links with 1000 addresses each. The simplest approach would be to use 100,000 words of memory and use the incoming address along with its link number as the address of a location to be incremented by one. Needless to say, this operation is memory intensive, and considering that the capacity of the minicomputer is 128K words, not very efficient.

A more appropriate implementation is to store the data on a bulk storage device, such as a fixed head disk, and buffer the incoming data until a sufficient amount is collected to update the disk. Once again complications arise. Disks can only be written one sector or one track at a time. For the disk selected, 2048 words represented a track. This was used to hold two incoming channels with 48 words for header. Since the rotation speed of the disk is 1800 rpm, and two rotations are required to update two channels (one read and one write), the maximum rate for updating 100 channels is

$$1800\,\frac{\text{rotations}}{\text{minute}} \times 1\,\frac{\text{channel}}{\text{rotation}} \times \frac{1\text{ update}}{100\text{ channels}} = 18\,\frac{\text{updates}}{\text{minute}},$$

if no other accesses are allowed. However, in actual operation the data must be read and used for processing so that only half of all disk accesses are allocated to input processing. Therefore, nine channel updates/minute are provided. Thus, the incoming data must be buffered for one ninth of a minute or 6.67 seconds. Since the arrival rate at maximum channel capacity is one address every 12.5 ms (an address is 15 bits long including start, stop, and parity and is transmitted at 1200 baud), then a buffer of size 534 is required for each channel. For 100 channels, the buffer contains 53,400 words of memory.

If one looks at actual data storage needs, one would find that at any given time nearly half of the cells in these buffers are empty. Thus, in reality only 26,700 words are needed if an algorithm can be developed to access the space. This is done using a buffering scheme that employs the triangular geometry discussed below.

First, assume that every time data are gathered into the buffer for all channels, only one channel can be updated from the buffer to the disk. Second, assume a system with five channels. Then, the buffering structure will be as shown in Figs. 3 and 4. The process is as follows. For the last channel updated (initially use channel 1), place the scanned incoming data sequentially in the vacated locations, starting with the first entry in the row assigned to that channel (channel 1 has row 1; channel 2, row 2; etc.). When the right row edge is reached, the remaining data points are placed in the column under the right edge. This operation is shown for five scans in Fig. 4, where once five scans are completed the appropriate cells for the data from the next scan are empty and the process can be repeated. Generation of addresses for using this algorithm on a digital computer is quite simple. The cells are numbered as shown in Fig. 5. Then, given:

S = Scan or channel being processed,

X = Start address of buffer − 1,

N = Number of input sources,

a. To find the first location C_1,

$$C_1 = X + S$$

b. To find the remaining locations 2 through N,

$$C_A = C_{A-1} + N - (A - 1) \quad \text{for} \quad 2 \le A \le S,$$

$$C_A = C_{A-1} + 1 \quad \text{for all other } A, \quad S < A \le N.$$

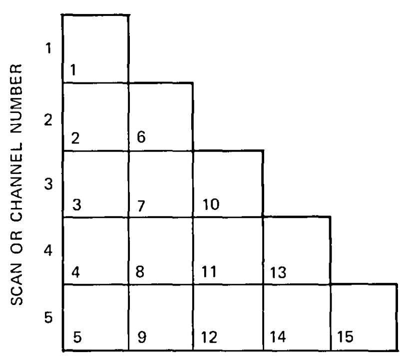

Fig. 3—Buffering scheme used to develop the algorithm that determines actual data storage needs.

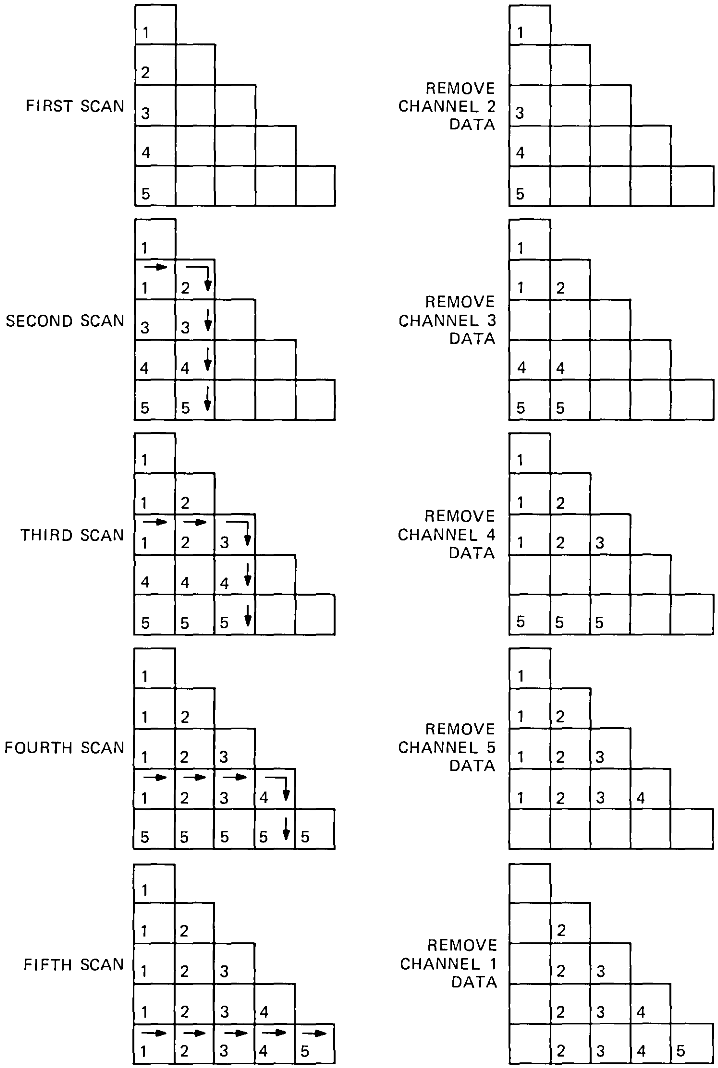

Fig. 4—Remaining data points in column under right edge of buffer structure, for system with five channels.

The above example assumed one update process for each data scan. In EADAS, scanning must occur once every 12.5 ms, while the update process takes four disk rotations for every two channels or two rotations per channel, which is equal to 66.6 ms (i.e., 66.6/6, which is the implementable rate closest to 12.5 ms). If scanning occurs every 11.1 ms, then six buffers of edge size 100 can be employed. This requires 6 $(100^2 + 100)/2$ memory locations, or 30,300. This number is slightly

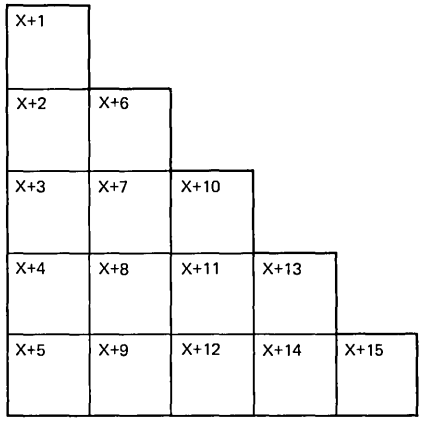

Fig. 5—Scheme for numbering buffer so that it is a contiguous string of cells.

larger than required by the 11.1-ms actual scan rate. Thus, with this algorithm, memory requirements for data collection were reduced, resulting in additional space for other functions.

In addition to regular group usage (data which are collected over all trunks in a group), the EADAS Traffic Data Converter (ETDC) includes circuitry to provide Individual Circuit Usage Recording (ICUR). These data require individual transmission of each of 3600 inputs on up to four 4A Traffic Usage Recorders (TUR), or a maximum of 14,400 inputs, as well as the 1000 inputs already available to the ETDC. Since the ETDC design had a 15-bit word (two start bits, an ICUR bit, ten data bits, a parity bit, and one stop bit) with only ten data bits, a multiword format was necessary. Two words were used to indicate the state of six inputs. The ICUR bit is set in each word, and the remaining 20 data bits are employed as follows. Two bits in one word serve as a first or second word indicator, while the remaining 18 bits indicate the six input states, and the address of those six inputs [two bits for the TUR number (0–3), four bits for the horizontal (0–9), and six bits for the vertical position (0–59)]. These data are intercepted by the scanning process, assembled and buffered, and then used to update individual circuit data, as well as form grouped totals which are merged with the accumulated peg count information. The updating process involves reading the usage and the peg count data from the disk, merging, and then writing the merged peg count data to the disk. The reads and writes are performed using Direct Memory Access (DMA) techniques. Dual-port memory arrangements are needed to accommodate the cumulative load of this processing and the data collection process. The system uses the second port on the ICUR disk as a transfer port, loading data into the added memory port. The individual circuit data are passed downstream and analyzed by the ICAN program described in the article on TNDS equipment systems in this issue.[1]

As memory costs decreased, devices with internal data storage became available. These sources could accumulate data in the central office terminal and forward it to the central unit on command. This method of operation had the potential of a low initial cost for a small number of offices by reducing the initial complexity of the central data collection facility. In addition, there was the opportunity to handle small offices via dial-up data links where dedicated links were not cost effective. For small offices only a few inputs are required, network management data are not needed, and only busy hour information is collected. To serve small offices, the Pollable Data Terminal was designed, and an interface specification published to permit traffic-data-gathering equipment to be connected in this mode. Operation of the Pollable Data Terminal is over a dial-up link, and depending upon busy hours, one EADAS interface port can serve many offices.

4.3 Electronic CU operation

The No. 1A EADAS is designed to collect data from electronic offices only, since plans to replace electromechanical offices were well under way and sufficient No. 1 EADAS Systems would soon be deployed to accommodate these offices until replaced. Initial software was developed by deleting the electromechanical features from the No. 1 EADAS. The initial No. 1A EADAS software was written in assembly language and ran under a very specialized operating system. The No. 1A EADAS software has since been rewritten in C language and now runs under the *UNIX** operating system.

Conversion to the *UNIX* operating system led to several interesting software operations. First is the reception of the accumulated data from the *ESS* switching equipment offices. Each switching machine type has its own unique output format (i.e., 1 *ESS* electronic switching is different from 2 *ESS* switching equipment, from 3 *ESS* switching equipment, etc.). Thus, a unique program was written to receive the data from each type, convert it, and store it in a common form. This task is handled by using LEX, a lexical scanner generation program, and YACC (Yet Another Compiler Compiler) to produce a scanner and parser to process the incoming data stream. The incoming data from the office are treated as a grammar. The description of this grammar is fed to LEX, producing a scanner which identifies the

* Trademark of Bell Laboratories.

tokens. A second description is fed to YACC to generate the parser. The token output from the scanner is passed on to this parser, which then prepares the data for storage by converting the data to binary and removing the header text. This approach significantly reduces the effort necessary to add new office types to the data collection portion of No. 1A EADAS. In addition, if errors are detected in data specification (either the format has changed or the specification was improperly interpreted), modifications are easily made by changing the descriptions of the grammar.

4.4 Overview of the No. 1A EADAS File System

The function of the No. 1A EADAS File System is to provide storage and retrieval capabilities for near-real-time switching network data collected from *ESS* switching equipment offices. These data are needed to produce status and exception reports for network administrators, and for transmission to EADAS/Network Management Systems, as well as downstream data analysis systems. To provide these functions, No. 1A EADAS must store large amounts of data collected from the *ESS* switching equipment in a database. This task is complicated by the high volume of data, the data storage and retrieval speed required of the file system, the large number of offices from which data are collected, and the necessity for storing data on disk over time for use by the report generation programs. The impact these requirements had on the development of the file system for No. 1A EADAS is described in this section.

4.4.1 File system requirements

The No 1A EADAS File System provides the following capabilities:

1. It allows efficient input of near-real-time register data collected from *ESS* switching equipment offices. EADAS is able to collect data simultaneously from at most 48 entities and store it in the database.

2. Many different types of data collected from *ESS* switching equipment offices are stored in the database: discretes, five minute, hourly or half hourly, daily, weekly, etc. The type of data actually stored is a function of the particular office involved, and the number of registers varies by office and type of data. For example, the 1 *ESS* switching equipment may send about 7000 registers every half hour for report generation, but only about 1400 registers for its daily transmissions.

Although register data are by far the largest in volume and require the most rapid input and retrieval, there are several other types of data that are stored in the EADAS database. These include per-office reference data, thresholds for data exception generation, and time schedules for distributing reports to various network terminals connected to the system. The file system accommodates these files as well as the register files. Currently, 36 types of data are stored in the EADAS File System.

3. Most EADAS data are collected on a periodic basis (e.g., every five minutes or every half hour). There are several reasons why successive intervals of data must be stored in the EADAS database. In some cases, calculations are performed on data received during successive collection intervals. In other cases, there is a need to have access to "old" data so that reports may be generated from special studies or for investigating office problems. Finally, a buffering scheme is needed so that reports can be generated simultaneously with data collection.

The number of collection intervals typically varies by type of data. For example, half-hour data are stored for 128 intervals (two and one-half days), whereas daily data are stored for seven days. The file system is able to store different amounts of data for each office and type of data. The total number of files necessary for EADAS to store in its database can be as high as 15,000 to 20,000.

The volume of registers collected varies not only by office and type of data, but also by collection interval. Thus, office A may normally send 2000 registers for data type 1, but at some point may decide to increase the transmission to 2500. Although this event does not occur very often, EADAS is able to accommodate the increase so that no data are lost and does so without increasing file overhead for any other office, data type, or interval. No backup of transient (register) data is necessary. The cost of a rare loss does not justify the cost of regular backups.

4.4.2 The No. 1A EADAS File System

The EADAS File System was implemented using raw (i.e., unbuffered) input/output (I/O) to meet the above requirements. It is a higher-level interface to the Logical File System (LFS), which is used in several other projects, as well as in EADAS. The overall function of the LFS is to subdivide an area of disk into contiguous files and perform all the necessary file management. The following is a list of the most important features of the LFS, as well as some information about how EADAS uses it:

1. Unlike the standard *UNIX** operating system, which provides a hierarchical file system using ASCII path names, LFS provides a "flat" file system. The entire LFS disk area is treated as a single *UNIX* system file that when opened allows access to all logical files. Each file corresponds to a unique number, called a logical file number (lfn),

* Trademark of Bell Laboratories.

that is the sole means of referencing the file. Up to 65,535 lfn's per file system are permitted.

2. The LFS performs all space management and low-level I/O, which is transparent to the user. Supported functions are creating, deleting, reading, writing, and switching logical files. LFS disk areas may be mounted and unmounted much like *UNIX* file systems.

3. One of the major functions of the LFS is to perform the mapping from logical file number to physical file location. The LFS requires no opening of individual files since they are "opened" by simply referencing them. Although efficient, access by file numbers is very cumbersome since they reveal little about file contents. Thus, the EADAS File System uses a specialized "open" function to translate retrieval keys (office, type of data, and time) into lfn's. Compared with the *UNIX* file system, the EADAS File System (using the LFS) allows many more open files and avoids costly directory searches in the opening process.

4. All logical files consist of contiguous disk storage. By computing offsets into contiguous files, the LFS avoids the overhead inherent in indirect blocks. Contiguous file storage also allows the LFS to read or write many blocks with one I/O request. This allows the physical block size to be determined by the access requirements for data in a *particular* file rather than be determined arbitrarily.

5. Since the LFS is implemented using the *UNIX* system's raw I/O, data may be transferred directly to/from disk from/to user memory.

In the present No. 1A EADAS configuration, the No. 1A EADAS File System uses three LFSs, with a maximum of 48,000 logical files. Typical EADAS database sizes range from 45M to 70M bytes, depending on the mix of switches from which data are being collected. Thus, the LFS results in efficient storage to permit the No. 1A EADAS to effectively store data from the electronic offices.

Once data are collected, the remainder of the system programs are executed to provide the reports and outputs necessary to administer the network.

V. NORGEN REPORTS

The intent of the Network Operations Report Generator feature is to provide a time-shared programming environment for the development of reporting functions, and a set of reports designed to meet universal reporting needs including near-real-time exception reports. The reports (referred to as Network Operations Reports) are designed to meet the basic reporting needs for each switching entity type served by No. 1 and No. 1A EADAS. These programs provide users with a package that may be modified or supplemented with locally developed features for the support of local reporting needs.

5.1 Basic objectives

Most of the processes that use traffic data operate on data collected over a long period of time, to assure statistical accuracy. Examples of such long-range functions are equipment and facility provisioning, load balancing, and optimized utilization of present equipment and facilities. The reporting functions associated with the above are performed by the "downstream systems" of TNDS. The report results are generally made available to the user within one to several weeks after the end of the study period.

However, there are a number of processes that must be completed in short time ranges. Network performance levels must be monitored for service-affecting problems on a daily basis. Analysis of the traffic data provides valuable insights to the causes for many types of service-affecting problems. This problem-related information is most useful when provided on a near-real-time basis. Furthermore, validation of the data should be done in near-real time to minimize loss of data when something goes wrong. The EADAS Network Operations Reports are designed to meet the near-real-time information needs, while the long-range data collection needs, as indicated earlier, are handled by the TNDS downstream processing systems.

5.2 Report types

The following report types are provided in the generic program:

1. Service Exception Reports are generated when high-priority service faults occur which may require prompt attention or corrective action.

2. Component Exception Reports provide an early warning indication of component conditions which, if allowed to persist, could affect service.

3. Trunk Exception Reports are printed when the percent overflow on final group(s) exceeds an assigned threshold. This report also provides supporting data that indicate where the pressure on the final group(s) is coming from.

4. Summary Reports provide a temporary record of busy-hour central office and trunking conditions. These reports are retained by network administrators until similar data become available from downstream processing.

5. Special Reports provide information which is useful to customers who use centrex, tie line, and multiline hunt-group facilities. The information in these reports is used by Marketing to advise customers

about the performance and needed changes in their communcations services.

6. Daily Reports are generated once a day during the early morning hours. The reports use data that was accumulated by EADAS or added up over the day by the switching system and transmitted daily to EADAS.

7. Network Switching Performance Measurement Plan reports are used to measure central office efficiency. They are organized for easy interpretation, and data are clearly labeled to simplify communications among network administrators, maintenance personnel, and traffic engineers.

8. Division of Revenue Reports provide data that help separate the cost of a switching entity into toll service, local service, and other rate base categories.

9. Measurement Apparatus Exception Reports are used to indicate failure of the data collection devices and associated data links. The report is also used to track corrective action, especially when support from maintenance organizations is required.

5.3 Report users

The principal user of the Network Operations Reports is the Network Administration Center (NAC). The Switching Control Center (SCC) and the Network Data Collection Center (NDCC) also use the reports, but to a lesser degree.

The NAC is responsible for the quality of the local dial service provided by the local central offices and the connecting trunk groups. A major NAC objective is to detect problems as quickly as possible before they can affect service. It is the NAC's responsibility to take the necessary action or make the necessary recômmendations to the appropriate organization for correction. These problems may be the result of underprovisioning of equipment or trunks, traffic imbalances, equipment outages, etc.

The NAC is also involved in network transition and cutover activities associated with growth. Real-time analysis of traffic data is a way to ensure that additions and rearrangements work properly.

The SCC has primary responsibility for the equipment maintenance-related aspects of local dial service. Maintenance data such as that contained in the Network Switching Performance Measurement Plan, which is generated by EADAS on a daily basis, are used as an aid in detecting and correcting equipment faults, and in setting maintenance priorities.

The NDCC is responsible for the administration of network and special studies data collection, supervising the operation and maintenance of EADAS, the data links, and data collection apparatus, such as the EADAS Traffic Data Converter, Traffic Usage Recorder, and Pollable Data Terminal. The center is also responsible for coordinating the operation of the TNDS Central Office Equipment (TNDS/EQ) downstream systems. This center is responsible for reviewing measurement Apparatus Exception Reports (AER) and acting to correct problems which cause loss of data.

5.4 Users of special reports

The near-real-time needs of various users continue to evolve. For example, special traffic studies on facilities associated with large business customers have increased substantially in the past few years and all indications are that they will increase even more in the near future. Organizations such as Business Services and Marketing are currently increasing the number of requests for special studies data. These requests are based on increased regulatory and competitive pressures as well as the increased availability of new features made possible by the widespread installation of stored program control systems and a higher level of sophistication of telephone customers. These factors have also combined to increase the frequency of these requests. The study results are often required in almost near-real time, making EADAS the logical system of TNDS to generate them. Studies not required in near-real time, or involving large amounts of data, such as those used for long-range planning, may be more efficiently produced downstream by other TNDS systems, such as TDAS.

The Network Operations Reports deliver the requested information and provide the necessary surveillance for all studies, both near-real time and long range.

5.5 Typical report use

Dial tone speed tests results are part of the Network Operations Reports and are used in the real-time service analysis process. An example of a 1 *ESS* dial tone speed service exception report is shown in Fig. 6. The intent of this report is to alert the network administrator that an abnormal dial tone delay condition has occurred so that corrective action may be taken. This report is divided into four sections: dial tone speed results, customer digit receiver (CDR) data, remote switching system, and supporting data. The first section contains the measured percent dial tone speed results by customer digit receiver type, i.e., *Touch-Tone** dialing and Dial Pulse (DP) with the associated test and delay counts. This report is only generated when the percent *Touch-Tone* dialing or DP dial tone delay exceeds a user-

* Registered service mark of AT&T.

```
* * * 1A ESS SWITCHING EQUIPMENT SERVICE EXCEPTION * * *

ENTITY:  XXXXXXXXXXX   DATE:    8/11/82   TIME:  12:30   INTERVAL:  30 MIN

*  DTS SERVICE EXCEPTION  *

                  % DTD             TSTS              DLYS
  * TT             0.3               363                 1
  * DP             0.0                87                 0

  CDR GROUPS            % OFL     OFL      PC   % OCC    # MB     NCI      HT
  * COM                   0.0       0       0     0.0     0.0      16
  *  TT                   0.0       0   13916    55.9     8.0     115     7.3
  *  DP                   0.0       0   10646    47.1    11.0     114    11.1

  ROB Q PC  =             ROB OCC  =                 ROB QTY  =     0

  RSSID  =        XXXXXXXXXXXX
                  % DTD             TSTS              DLYS
  *  TT/DP         0.0               275                 0

                  ORIG PC  =     224          LN-LN USAGE  =   467

                      % OFL      OFL       PC   % OCC    # MB      NCI       HT
    CHANNEL             0.0        0      428    24.0     0.0      108    109.1


  E-E CYCLES    =    8504    LN SCAN COMP  =   8504
  ORIG PC       =   24401    INC PC        =   3594
  A LK CCS      =   27899
  BLK DT Q PC   =       0
  POB %OCC      =    20.2    POB OFL       =
```

*REGISTERED SERVICE MARK OF AT&T

TT – *TOUCH-TONE*® SERVICE

Fig. 6—Example of 1 *ESS* switching system dial tone speed service exception report alerting network administrator of abnormal dial tone delay.

specified threshold. The remaining three portions of the reports provide supporting data for customer digit receivers by type.

When an excessive dial tone delay condition exists, the dial tone delay statistics are printed for both *Touch-Tone* dialing and DP. An asterisk is printed on the line(s) for which the threshold failure(s) have occurred. A threshold test of the CDR group(s) is then performed. If a common CDR group is provided, a threshold test is made for excessive overflow. If the test fails, the results are printed for the common group with an asterisk followed by two lines of information for the DP and *Touch-Tone* CDR groups. Otherwise, the common group results are printed without the asterisk. If a common CDR group is not provided, threshold tests on the DP and *Touch-Tone* dialing CDR groups are made and the results are printed with an asterisk, where required.

After the CDR results are printed, the next section will be provided only if the *ESS* switching equipment is an RSS host office; otherwise this section of the report is omitted.

The last section of the report is always provided and contains overall call processing load data.

Dial tone speed data are analyzed by the Network Administrator following reception of a dial tone speed service exception report. This report is generated when the percent dial tone delay threshold is exceeded for DP and/or *Touch-Tone* dialing. If a common CDR group is provided and it exceeds the percent overflow threshold, then the following analysis would be performed:*

1. The CDR percent occupancy is observed to determine if it exceeds the engineered capacity. If it is too high, then the CDR maintenance busy item (number of circuits made busy) is examined.
2. If the maintenance busy value is too high, then the network administrator requests restoral of the made busy CDRs.
3. If the maintenance busy value is not too high or if there would be a problem even after they are restored, then an analysis of the DP and *Touch-Tone* dialing CDRs is performed. If both the DP and *Touch-Tone* dialing CDR overflows are high and their holding times are unusually long, then permanent signal conditions are analyzed.
4. If holding times are normal and no extraordinary demand exists (such as a severe storm produces), then this condition would be referred to the traffic engineer if it persists.

If *Touch-Tone* dialing percent overflow is high and the DP percent overflow is not, then the DP *Touch-Tone* dialing CDRs should be rebalanced or *Touch-Tone* dialing CDRs added.

If DP and *Touch-Tone* dialing CDR percent overflow is not high, then a check of all related database, circuit quantities, and register assignments should be performed.

The Network Operations dial tone speed service exception report provides the network administrator with the necessary data to support this problem analysis.

In rare cases traffic peaking can be a factor in high CDR overflow. Examination of the 15-minute traffic count report which is generated by the 1 *ESS* switching equipment would provide the necessary 15-minute data required for this analysis to supplement the 30-minute data available to No. 1A EADAS. There are three other areas of investigation related to the dial tone speed analysis. They are Processor Overload, Network Blockage, and Translations. The first two areas would make use of a number of Network Operations Reports in the analysis procedure. The translation analysis would use the dial tone speed service exception report.

VI. NORGEN IMPLEMENTATION

The development of the NORGEN feature was an ambitious under-

* Familiarity with Bell System traffic terms is assumed for this section.

taking. Data coming from 48 or more entities are interpreted and analyzed in real time. Formatted reports are prepared and distributed. The internal architectures of a large number of switches determine the report structure and algorithms in fine detail. The program can be distributed to all operating companies and run without local programmer support. Yet the system is tolerant of user programming modifications.

6.1 The challenge

A major challenge was to implement NORGEN in No. 1 EADAS, with a total Random Access Memory (RAM) capacity of 256K bytes (only 64K bytes available to NORGEN), without compromising NORGEN's ultimate capability in the No. 1A EADAS, which has essentially unlimited RAM capacity (the current memory size is 1M byte).

The solution was to use two entirely different software architectures. Yet the feature appears very similar to the users of the two systems.

6.2 NORGEN in No. 1 EADAS

In the No. 1 EADAS, programs are stored in interpretive language to conserve RAM. Program is brought into core in 512-byte blocks and executed block by block. A very linear coding style is used, with a minimum of subroutine calls to hold down the need for calling in extra blocks of code.

The interpretive language (object) is compiled from source code written in K language, which was designed for the project. K language is similar in style to languages with the same level of generality, but has fewer elaborate commands. Thus, K language is easier to learn than C, although C has more "power" in some applications (that is, certain features of C allow some functions to be performed in fewer lines of source code). A text editor and source-to-interpretive language compiler, both written in assembly language, are supplied with the generic.

In K language, all variables reside logically on disk; there is no distinction between local and global variables. The system keeps a cache of disk blocks in memory to improve the speed of operation. This structure was chosen because of the small amount of memory available, but has been found to ease coding problems because the programmer need not consider whether a variable is of one type or another.

The disk file access has three levels. A directory contains all named files and file offsets for particular data items. The directory is accessed using a backup function. When an item is needed, a fast access directory is checked first, to see if the file is already in RAM; then the backup function is used to access the disk directory; and then the file itself is called from disk. Most requests are found in RAM; nearly all others are filled by two disk accesses (one for the directory and one for the actual data requested).

6.3 NORGEN in No. 1A EADAS

In the No. 1A EADAS, programs are stored in assembly language, compiled from C language source code by the *UNIX* compiler. This permits more subroutine calls and a more conventional programming style. The *UNIX* text editor and C language compiler are provided with the generic.

The *UNIX* file system is used for program files and parameters, but the specialized two-level Logical File System as previously described is provided for accessing the traffic data. Otherwise, multiple file accesses would be required for the multilevel directory structure of *UNIX*.

6.4 Parameter administration

In this article, the term "parameters" refers to all the user-controlled tables which the generic program reads to relate to local conditions. NORGEN parameters include scheduled times for reports to be run, exception thresholds, register definitions, etc. Any parameter administration system must have the functions add, delete, change, display, and backup.

In both the No. 1 EADAS and No. 1A EADAS, these functions are served by using text files, in controlled format, to hold all the data. The text editors, which are quite similar, are used for all of the above functions except backup; this is accomplished by dumping all the parameter files onto tape (approximately weekly) and reading them from tape in an emergency.

Programs are provided to convert each of the text files, by type, into binary object files for efficient access. When any text file has been changed (by the editor), the file is recompiled.

The text editor enables the user to work in a relatively English-like environment, minimizing training and making it easy to spot errors by browsing displays of the files.

6.5 Raw register access

The No. 1 and No. 1A EADAS use very different register access mechanisms. Both use key words to provide mnemonic access to individual or summed groups of registers. However, the No. 1 EADAS systematically translates the packed register data from data collection into rigidly structured tables, with each key word assigned a fixed location (for each switch type). On the other hand, the No. 1A EADAS uses mapping algorithms and parameter tables to extract each key word from the raw registers as it is called for from the program. This newer technique provides a great savings in disk space, allowing a fourfold increase in the number of intervals of data that can be stored, since most entities have only a minority of the total possible key words actually in use. The real time used by the two methods appears to be similar. The older technique is more efficient when key words are used repeatedly, but very few key words are used more than once or twice.

6.6 Scheduler

Again, there are strong differences between the two schedulers. In the No. 1 EADAS, the scheduled reports are run one at a time, to minimize the swapping of files into and out of the severely limited RAM space. In the No. 1A EADAS, two modules are run at one time (level 2 multiprocessing). This achieves an efficiency gain due to interleaving processor and disk I/O functions.

Both executives are time-shared, allowing user access to demand reports, dump key-word values, etc. The No. 1 EADAS executive allows new users high priority. As recent run time increases for a given user, that user's priority (probability of being run at the next opportunity) decreases. This automatically allows users high priority for short, easy tasks like editing or key-word dumps while lowering the priority for compilations and running long reports.

The No. 1A EADAS uses the *UNIX* priority feature, giving manual users priority over scheduled work, independent of their tasks.

6.7 Reports

The report modules are written in C language, using straightforward, linear style. This facilitates user programming as well as Bell Laboratories and Western Electric modifications of the design. Subroutines are used sparingly, so that calculation details are grouped in one place in the source program.

Special functions are defined for pervasive features. In particular, a threshold check function automatically seeks the threshold parameter and applies the appropriate check (less than, equal to, range). A special print function prints a blank in a field when appropriate. (For example, data pertaining to items which are unequipped in a given entity are carried in memory as negative numbers, which are suppressed by the print function.)

Standard page headers and section header routines are called to provide uniform structure of the reports. Automatic pagination is provided (through the special print function).

The layout of the reports emphasizes the grouping of related data for easy reading. White space is used freely to set off the groups.

Mnemonic labels are provided; abbreviations are chosen to be consistent with related Bell System documentation.

6.8 Report distribution

After the ASCII image of a report is delivered to the spooling system, it is distributed to appropriate user terminals. Each report has a message class number. A parameter table lists the destination terminals for each entity and each class. In No. 1 EADAS, this table was centrally administered; in No. 1A EADAS, the user at each terminal designates the entities and message classes it is to receive. A particular report thus may be sent to many terminals.

VII. TAPE WRITING

In the No. 1 EADAS, tape writing is done by an assembly-language-coded program that runs under the data collection executive. Data are written each half hour, under control of schedule parameters. The data are recorded in ASCII characters; each register is represented by a six-digit number.

In the No. 1A EADAS, tape data are prepared in two steps. The first step is to process and format the data; the Tape Data Formatter (TDF) module runs under the NORGEN executive. This provides opportunity for much more powerful processing of the data. In particular, the TDF combines two successive half-hour data messages into a single one-hour message. In most cases, registers from each half hour are added together; but certain registers, such as those containing averages, peaks, or trunk group numbers, are combined by specialized algorithms. Improved validity checking is done to avoid passing bad data downstream.

In the second step, the TDF writes the formatted data in a spooling area, where they are held for up to several hours. This spooling capability improves operational flexibility; for example, it is possible to hold data while a tape drive is being repaired.

A new tape format is provided in No. 1A EADAS. Data are usually written as 16-bit binary words, thus fitting a register into two bytes, where the No. 1 EADAS requires six. The effective reduction is approximately 2 to 1 since there is less compression in headers, and no compression of record gaps.

Industry standard tape labels in the new format facilitate data-link transmission of tapes and inventory control.

Two tape drives may be provided in the No. 1A EADAS; when one tape is filled, the other begins to write. This feature, together with the binary format, combination of half-hour data into one-hour data, and spooling, allows data collection to continue unattended in most installations over long holiday weekends.

VIII. DATA FLOW

Following the course of data through the No. 1A EADAS shows how major program modules relate to one another. Let us follow the data of and *H* schedule* from a polled 1A *ESS* switching equipment office—the schedule for the interval from 10:30 a.m. to 11:00 a.m., in particular.

The 1A *ESS* switching equipment office will have collected the data and made it ready for transmission shortly after 11:00 a.m. on its clock. The No. 1A EADAS will wait several seconds after 11:00 a.m. on its clock to allow for system time differences.† Then it will send a poll asking for the first block of 256 registers of *H* schedule data. When the block has successfully been received, it will ask for the next block, and so on, until the proper number of registers (a channel parameter) has been validly received.

As the blocks of data are received, they are stored in the logical file system. A new entry has been made in the LFS directory, so that the data may be accessed by entity name, end time of data interval, and type of data.

When the schedule is complete in the LFS, a message is sent via a pipe to the NORGEN Data Analyzer (NDA) executive. The executive checks schedule parameters to see which program modules are to run on this particular *H* schedule.

Suppose that the busy hour for this entity is 10:00 a.m. to 11:00 a.m. Then the Load-Service Summary module will be scheduled among others. The executive will make an entry in its job queue, listing entity, module, and end time of data. When this entry works its way to the head of the first-in-first-out queue and a new process can be started, the executive will run it. This particular module (the Load-Service Summary) works on a full hour's worth of data. Therefore, it retrieves two successive intervals of data from the LFS and adds their register values together to form an equivalent one-hour schedule.

Next, the module accesses data to prepare its report. For example, it reads the peg count of incoming calls by use of key word INCPC (incoming calls peg count) as a variable in a program. The compiler has recognized this key word as a register reference and has compiled a function call to the key-word access routine. The key-word access routine references the Data Collection Device (DCD) parameter table to find the register number and return the corresponding value to the Load-Service Summary module.

The module prepares the ASCII text of the report, calling special

* An "*H* schedule" is a set of registers that represent traffic through the common parts of the switch (as opposed to the trunk groups).

† The No. 1A EADAS periodically requests time-of-day from each 1A *ESS* switching equipment office; if the time difference is large, a message is printed.

subroutines to form the header, subsection header, and end of report headers. Pagination is provided automatically; to make this possible, a PRINTN function is defined, which performs special NORGEN report functions before calling the PRINT functions.

The ASCII text is stored in a spooling area, but the spooler directory entry is not completed until all modules for that entity and interval are completed. This allows the NDA executive to mark the sequence of reports in a fixed order, so that they will always appear in the sequence expected by the user. This facilitates search and filing of specific reports.

The report spooling program distributes each report in order. Each report type is assigned a message class. The terminal tables are searched to determine which terminals have designated the entity name and message class for the Load-Service Summary Report. The report is sent to these terminals through the multiplexor driver program.

IX. CAPACITY

In the No. 1 EADAS, NORGEN is usually the limiting capacity consideration. Consequently, the capacity characteristics of the reports have been analyzed carefully.

The capacity criterion of the system was to complete the report generation load of an "equivalent half hour" in one half hour, to prevent excessive delays in the availability of reports. The concept of an "equivalent half hour" was introduced because the report load is irregular. For example, for a polled interface, some reports run continuously (every half hour), some run every hour, and some only after busy hours. After discussions with users, it was decided that a backlog can build up during a busy hour, as long as the system recovers within one and one-half hours.

The load on the system for an equivalent half hour was taken to be:

continuous reports: fully weighted
hourly reports: one-half weighted
busy hour reports: one-third weighted.

The capacity analysis was thus reduced to estimating the run time of each report module. The weighted run times are then added to determine whether the system can keep up with its load.

Report run times are composed of two major components; processor time and disk access time. Each of these resources is subject to priority seizure by data collection activities, which extends the report run times significantly. Extensive modeling and measurement of the No. 1 EADAS and No. 1A EADAS Systems have been performed to determine the quantitative effects of data collection and relate those to the mix of sizes and types of remote data sources (switching entities).

9.1 Modeling

The application of analytic modeling to combination real time and time shared systems has been challenging. It has been successful because a limited accuracy goal of plus or minus five percent was set. As a result, effects which sum to only a few percent were neglected, or assigned a fixed allocation which was neither best nor worst case, but somewhere in between.

Attention was thus directed at the first-order effects. Measurements were carried out in the laboratory, using a field configuration system under approximately 50-percent load. The load was applied by a separate minicomputer acting as a load test generator. Traffic data used was based on data collected from "typical" offices in the field. These model offices were chosen to be reasonably complex in terms of the number and types of registers, but worst-case offices were avoided since they do not dominate a fully loaded system.

By taking measurements on a system loaded in the 50-percent range, the effects of approximations in the model are minimized.

9.1.1 No. 1 EADAS capacity model

In the No. 1 EADAS, the scheduled report modules are run sequentially, rather than multiprocessed, because of the limited availability of RAM. In this system, data collection takes a large amount of processor time at interrupt level (up to 60 percent in a heavily loaded system), especially for the ICUR feature. The effect of this data collection load on a given report module depends on how much processor time it needs; the data collection load has no first-order effect on disk access time. The effect, then, is to extend the processor time component of the report run time:

$$T_{\mathrm{p}} = \frac{T_{\mathrm{po}}}{1 - F},$$

where:

T_{p} = processor component of the report run time in the presence of data collection load,

T_{po} = processor component of the report run time in the absence of data collection load,

F = fraction of processor used by data collection.

In the above model, it is important to note that the data collection work is not time shared with the report modules, but is performed at priority interrupt.

A second-order effect, which is typically large enough to be significant (of the order of ten percent), is caused by base-level data collection load. This load has no effect on processor time because it always takes place during times when report modules are waiting for disk accesses. However, when a disk access is complete, base-level work is usually in progress and neither a further disk access nor a report processor segment can be started until it completes. Effectively then, a fraction (taken to be 1/2) of the base-level data collection time (stretched out by interrupt data collection time) is added to disk latency.

$$L = L_o + \frac{1}{2}\frac{a}{1 - F},$$

where:
L = disk latency in the presence of data collection load,
L_o = disk latency in the absence of data collection load,
a = average base-level time for data collection,
F = fraction of processor used by data collection at interrupt.

The effect on the disk portion of the report run time is given by:

$$T_d = T_{do}\frac{L}{L_o},$$

where:
T_d = disk component of the report run time in the presence of data collection load,
T_{do} = the disk time measured in the absence of data collection load.

The rule for capacity estimation is to add the processor and disk component times for the relevant report modules (for the given mix of office type and report schedules, according to local operation practices) weighted by the appropriate factor according to the report frequency, to apply the model equations to account for the estimated data collection loads, and to compare the resultant time to the 1800 seconds of each half hour. Ten percent of the report capacity is allocated to allow for manually generated load running in time share with the reports.

While the above process may seem complex, it has been reduced to a straightforward application of tables and worksheets and has been well accepted by the operating companies who must estimate their capacity.

9.1.2 No. 1A EADAS capacity

The No. 1A EADAS is quite similar in its modeling except that only one effect of data collection load is significant—priority access to the disk. Since scheduled reports are run in time share, two at a time, and since the reports need about 60-percent of disk I/O time and only 40-percent of processor time, the reports are strongly disk limited. Measurements show that about 85 percent of the disk is utilized, which is quite constant even when 50 percent of the disk is taken by data collection. A change in data collection architecture has resulted in heavy disk loads for collection of ASCII data in the No. 1A EADAS.

The system is found to be so disk limited, that the distinction between processor and disk components can be neglected, and the *effective* run time (which is measured total run time in multiprocessing, divided by two, since two reports are run at the same time) can be considered to be all disk time.

Then the effect of data collection disk load is given by:

$$T = \frac{T_0}{1 - F},$$

where:

T = effective report run time in the presence of data collection load,
T_0 = effective report run time in the absence of data collection load,
F = fraction of the disk used by data collection load.

Estimation of capacity is then carried out in essentially the same manner as for the No. 1 EADAS.

X. SUMMARY

EADAS has met its original operational objectives by providing for the collection of network data for engineering, network managment, and real-time data analysis. The result has been an overall improvement in the monitoring of network operations, thus meeting the increasing need for more efficient network utilization to optimize both capital investment and customer service. Eliminating the manual effort formerly associated with the collection and processing of network data has resulted in substantial economic savings. EADAS will continue to be enhanced to provide interfaces with new switching machines and operations systems.

XI. ACKNOWLEDGMENTS

The No. 1 and No. 1A EADAS represent the work of many people over a full decade. The features of the system are reported here without specific accreditation, which would be difficult at this time. However, the authors wish to acknowledge the special help of J. R. McSkimin and R. L. Hardin in the preparation of this paper.

REFERENCES

1. N. D. Fulton, J. J. Galiardi, E. J. Pasternak, S. A. Schulman, and H. E. Voigt, "Total Network Data System: Equipment Systems" B.S.T.J., this issue.

AUTHORS

J. A. Grandle, Jr., B.E.E., 1967, University of Virginia; M.S.E.E., 1968, Purdue University; Bell Laboratories, 1967—. Mr. Grandle has helped design Traffic Measurement Circuits; hardware and software for No. 1 Engineering and Administrative Data Acquisition System (EADAS), No. 1A EADAS, EADAS Individual Circuit Recording, and EADAS Network Operations Report Generators; CDA; No. 1 Automatic Message Accounting Recording Center central hardware; PDT-2A hardware and firmware; and processor terminals, including the PLURAL multiprocessor terminal. He also has three patents and has published several articles. ("A New Test Line for Coin Phones," and "A Versatile Data Gathering Tool") published in the *Bell Laboratories Record.*

Donald J. Gagne, B.E.E., 1943, Pratt Institute; Bell Laboratories, 1943—. Mr. Gagne has participated in the development of order wire and alarm facilities, mobile and overseas control terminals, and No. 1 and No. 5 Crossbar circuits. He has also conducted systems engineering studies for new recorded announcement systems, new electromechanical and electronic PBXs, *PICTUREPHONE*® meeting service and Automatic Number Identification for PBXs, and new Service Observing and Traffic Measurement facilities. He supervised early Systems Engineering for EADAS and at present is supervising a group responsible for EADAS Systems Engineering.

George H. Wedemeyer, Certificate in Aircraft Design and Technology, 1959, Academy of Aeronautics; Bell Laboratories, 1962–1981; AT&T, 1981—. At Bell Laboratories Mr. Wedemeyer participated in the hardware and software design of data collection devices. In 1981, he became a Member of Technical Staff while working for the Network Administration Systems Department in Holmdel, New Jersey. His responsibilities included formulating requirements for traffic-data collection and reporting systems, economic analysis, and system evaluations. Currently, Mr. Wedemeyer is a staff manager at AT&T responsible for the project management of EADAS.

Charles J. Byrne, B.S.E.E., 1957, Renesselear Polytechnic Institute; M.S.E.E., 1958, California Institute of Technology; Bell Laboratories, 1958–1962, 1968.— Before rejoining Bell Laboratories, Mr. Byrne was employed at Bellcomm, Incorporated, where he worked on systems engineering for the Apollo Project. At Bell Laboratories he has worked in communications systems research and operations systems development. He is currently Supervisor of the International Sales Planning Group of the Switching Operations Systems Laboratory. Member, Sigma Xi, Tau Beta Pi, Eta Kappa Nu.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Network Management

By G. C. EBNER* and D. G. HAENSCHKE*

(Manuscript received June 9, 1982)

This paper describes steps that have been taken to respond to the need for better network management in the evolving telecommunications network. In near-real time, network management functions recognize the onset of an overload and respond with control actions that change normal call routing through expansive or restrictive traffic controls. Emphasis is being placed on improved automatic controls that are built into the network and its switching systems. Advances in manual controls and real-time network performance monitoring have been accomplished through the introduction of computer-based systems that provide network managers with preprocessed network performance data and with the ability to intervene in problems that require human judgment.

I. INTRODUCTION

Network management consists of real-time network performance monitoring and control of the telephone network. It is a technique designed to optimize the call-carrying capacity of the network when the network is under stress due to traffic overload or failures. In recent years significant steps have been taken to respond to the need for better network management for the evolving Message Telecommunications Service (MTS) network. Advances have been made in manual

* Bell Laboratories.

network management controls and real-time network performance monitoring capabilities. These advances were accomplished primarily by the introduction of computer-based systems that support the operation of centralized network management centers. Network management centers employ automatic and manual capabilities that recognize the onset of overloads and respond with traffic control actions within a time span ranging from seconds for automatic controls to minutes for manual actions. Each center is supported by an operations system that collects the network data, monitors the status of the switching and signaling systems, and permits the implementation of appropriate control measures during overload and failure conditions. Advances have also been made in improved automatic network management controls that are built into Stored Program Controlled (SPC) switching and signaling systems. The current approach to network management is to provide an economical balance between automatic and manual network management capabilities, with greater emphasis on improvements in automatic controls.

This paper presents the status of network management and its application to the evolving MTS network. The network performance under overload and the motivations for network management are described in Section II. Section III discusses the centralized network management operations systems and the automatic network management SPC switching system controls. Section IV gives a description of the operations system Engineering and Administrative Data Acquisition System/Network Management (EADAS/NM). Section V discusses future network management needs for the evolving MTS network.

II. MOTIVATION FOR NETWORK MANAGEMENT

2.1 Network performance under overload

The North American Message Telecommunications Service (MTS) network functions as a single, integrated entity to which customers have shared access for voice telephone calls, data calls, and other uses such as facsimile transmission. The network is currently being enhanced by the rapid introduction of Stored Program Controlled (SPC) switching systems interconnected via the Common Channel Interoffice Signaling (CCIS) system.[1] In effect, the CCIS network is a separate high-speed, packet-switched data network that allows SPC switching systems in the message circuit-switched network to communicate with each other. The interconnection of SPC switching systems by a reliable, high-speed, high-capacity signaling system permits more efficient use and control of the network.

Under overload, modern telecommunications networks that employ common control and automatic alternate routing can be forced into a congested, inefficient state, as marked by a significant decline in the number of calls that can be completed. The loss of network capacity under overload was described by P. J. Burke.[2] Figure 1 presents a key study result obtained through simulation of a 16-node hierarchical network. This figure indicates that there is a decline in the amount of traffic that can be completed when the offered load substantially exceeds the design limits of the network (about 1600 erlangs). There are two basic reasons for the loss of capacity of a network in congestion: excessive alternate routing and regenerative switching delays that are compounded by customer reattempts. Excessive alternate routing causes an increase in the average number of links a call will use before it is completed, thereby increasing call blocking and reducing the efficiency of network utilization. This increase in blocking, with increased alternate routing for a network in congestion, can be mitigated by trunk reservation schemes and by constraining the number of links in alternate routes.[3]

Switching delays are the dominant cause for the loss of the call-carrying capacity of the network under overloads. If left uncontrolled, switching delays feed on themselves and can quickly spread throughout the network. They cause time-out conditions during call setup and occur when switching systems become severely overloaded. Switching congestion time-outs result in short-holding-time attempts on trunk groups, replacing normal-holding-time messages. Today, it is recognized that the telecommunications network's response to overload and the magnitude of the decline in call-carrying capacity of an overloaded

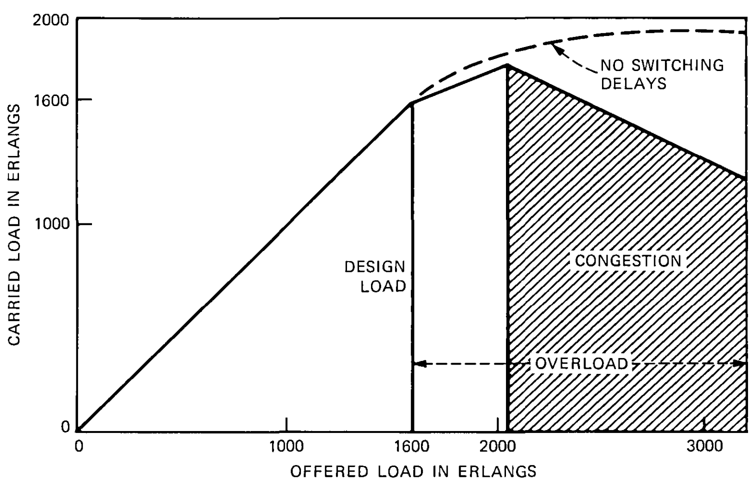

Fig. 1—Network performance under overload.

network depend on network architecture, call setup and routing procedures, signaling techniques, and switching system structure.[2-7]

2.2 Network management benefits

The primary motivations for network management in the Bell System are to maintain the integrity of the network and to ensure optimal performance of the network during overloads and failures. Telecommunications networks play a vital role during emergencies and natural disasters. The benefits of investments in modern network management capabilities are realized during such emergencies.[4] Network management, coupled with redundancy in critical switching components, assures network integrity in the Bell System at less cost than alternatives that do not involve network management. Maximizing call completions during overloads by network management techniques usually produces more revenue than network management capabilities cost.

III. NETWORK MANAGEMENT SYSTEMS

3.1 Automatic control systems

Network managers at the Network Management Centers (NMCs) discussed in Section 3.2 plan the employment of automatic network management controls in the MTS network. The NMC is responsible for monitoring the performance of the automatic controls through data furnished to the network management support system EADAS/NM by the switching systems. Occasionally, the NMC must fine tune the control response actions of the automatic controls that are built into SPC switching systems.

3.1.1 Automatic protective controls

Based on results of network overload characteristic studies, automatic network management controls, such as Dynamic Overload Control (DOC) and Directional Trunk Reservation (DRE), were introduced in the mid-1960s and proved to be effective during peak-day overloads such as Mother's Day and Christmas. However, these controls are not very effective during congestion caused by focused overloads, which are characterized by a surge of traffic originating in many parts of the network to a single destination, because they cannot selectively discriminate and control traffic based on destination codes. Focused overloads, stimulated by mass-media advertising, preplanned call-ins, or natural disasters, are occurring with increasing frequency in telecommunications networks.

As the network evolved to *ESS** switching system and CCIS,

* Trademark of Western Electric.

network management control capabilities also evolved. A major advancement was the introduction of improved automatic controls with the introduction of the 4 *ESS* switching system in 1976.[8] The automatic controls that were introduced and are now in operation are summarized in Fig. 2. They include two automatic protective controls, Selective Dynamic Overload Control (SDOC) and Selective Trunk Reservation (STR), as well as a new automatic expansive control called Automatic-Out-of-Chain routing, discussed in Section 3.1.2.

SDOC responds to switching congestion by dynamically controlling the amount and type of traffic offered to an overloaded or failed switching system. SDOC response actions are taken based on control request messages from the overloaded switching system, typically transmitted via the CCIS data network. STR, on the other hand, does not require the transmission of a control message; it responds to trunk congestion in the outgoing trunking field and is triggered on a particular trunk group when less than a certain number or circuits are idle in that group. The novel feature that makes these two protective controls selective is that they control traffic to hard-to-reach (HTR) points more severely than other traffic. An HTR code is a three-digit or six-digit destination code to which calls have a low probability of completing.

Based on real-time analysis of three-digit and six-digit destination code completion statistics, HTR traffic is automatically detected by the 4 *ESS* switching system. Every five minutes the number of Ineffective Machine Attempts (IMA), defined as attempts which fail within the 4 *ESS* switching system, and the number of Ineffective Network Attempts (INA), defined as the number of calls failing in the network between the node of interest and the call's final destination, are determined by each 4 *ESS* switching system in the network. The INA count is based on the number of calls that abandon after outpulsing without obtaining answer supervision from the called customer line.

OVERLOAD OR FAILURE IN	CONTROL	RESPONSE
SWITCHING SYSTEM	SELECTIVE DYNAMIC OVERLOAD CONTROL	PROTECTIVE
TRUNK GROUP	SELECTIVE TRUNK RESERVATION	PROTECTIVE
HIERARCHICAL ROUTES	AUTOMATIC OUT-OF-CHAIN ROUTING	EXPANSIVE

Fig. 2—4 *ESS* automatic network management controls.

The current 4 *ESS* switching system code detection and control system[8] uses the 4 *ESS* switching system for identification of HTR traffic destination codes and, in the near future, will use the CCIS network to distribute this information to other nodes in the network for control purposes.[3,9] Thresholds established at switching systems by the NMC identify HTR traffic destination codes. Traffic for such codes will be selectively controlled at the switching systems by SDOC, STR, and during congestion in the CCIS network.

Concurrent with the introduction of 4 *ESS* switching system was the initial deployment of the CCIS network. Failures are rare because of the high degree of redundancy built into the CCIS network. Nevertheless, new protective automatic controls had to be provided that sense overloads and failures in the CCIS network and cause SPC switching systems to respond with control actions that can affect the flow of traffic in the telecommunications network. For instance, a focused overload can overload the terminal buffers of CCIS data links. To avoid the loss of CCIS messages and corresponding calls for the circuit-switched telecommunications network, SPC switching systems apply automatic protective controls in response to CCIS terminal buffer overloads and CCIS processor congestion. The controls temporarily restrict the amount of new messages that are offered to trunk groups that signal via the overloaded CCIS link terminals. CCIS link overloads are usually caused by unsuccessful reattempts to HTR points in the network and can, therefore, be controlled most effectively by code-selective controls. Recent studies utilizing single server queueing models and mean value analysis indicate potential benefit in increased call completion (throughput) due to selective control on CCIS messages.

Figure 3 shows the predicted effectiveness of these selective controls for a simulated focused overload in a 24-node toll network. The nonselective automatic controls are contrasted with SDOC and STR, assuming full deployment in all toll switching systems. Performance is measured in Fig. 3 by the average number of completed calls in progress versus time. While the improvements due to the selective nature of these controls are greatest during focused overloads, simulation results and experience indicate improvement also for peak-day overloads. Other automatic processor controls are incorporated in the network, such as line load control which, when enabled, denies customer access to the switching system in congestion.

3.1.2 Automatic expansive controls

In contrast to protective controls, which restrict access to overloaded network resources, expansive controls take advantage of idle capacity on out-of-chain routes, i.e., on routes that are not within the design of

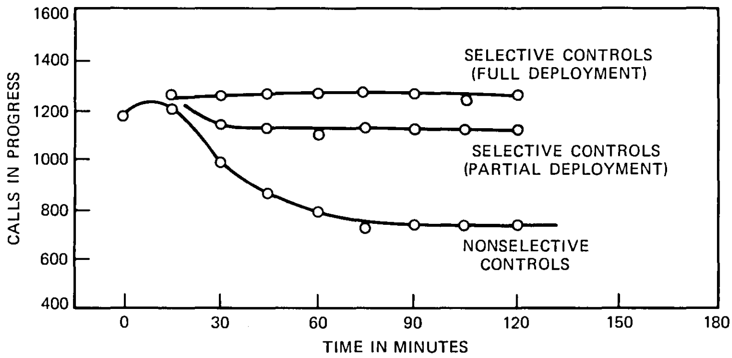

Fig. 3—Automatic control of focused overload.

the hierarchical routing structure. Automatic Out-of-Chain (AOOC) routing is a control that expands the route selection beyond the hierarchical routing constraints by offering calls overflowing the normal final-choice in-chain routes to up to seven out-of-chain routes. AOOC routing is the first step towards improved network utilization by taking advantage of network capacity that is often available through traffic noncoincidence. AOOC routing is made possible by the capabilities of the CCIS network, which permits AOOC routed calls to be classmarked, counted, and restricted to out-of-chain routes with idle capacity.

3.1.3 Controls for 800 Service with CCIS

The evolving SPC network with its inherent CCIS interconnectivity is also the key to new, innovative, communication services. These new services are made possible by computer-based Network Control Points (NCPs) to which switching systems are given direct switching access via the CCIS network. NCPs are specialized network nodes that provide the logic and routing information to handle enriched network services.

An example of new network management opportunities made possible by the SPC-CCIS network is the handling and control of 800 Service, formerly called Inward Wide Area Telephone Service (INWATS). Today, more than ten percent of the total toll calls are 800 Service calls; 800 Service mass calling is a frequent cause of overloads in the telecommunications network. In the SPC-enriched network, such calls first access a centralized NCP via the CCIS network. The NCP translates the 800 number to a new destination number and returns the number to the switching system, and the call is routed as a normal telephone call.

Network management controls are provided that protect the NCP, the CCIS network, and the circuit-switched telecommunications network from overloads. Code-selective controls are initiated automatically (or manually from the Network Operations Center described in Refs. 10 and 11) at the NCP. Call attempt thresholds are established at the NCP to detect excessive calling volumes. The Action Point (ACP) in the network (the switching system at which access to the NCP via the CCIS network is initiated) will recognize the control signals and limit attempts and NCP inquiries for high call volume destinations. The NMC will be informed via EADAS/NM when ACP controls are initiated.

3.2 Hierarchy of network management operations systems

Network managers in NMCs intervene in problems for which automatic system solutions would be excessively expensive and in problems requiring human judgment. In the past, NMCs were associated with major toll switching systems on a one-to-one basis. Each NMC displayed information that pertains to only one switching system and was able to affect call completions through manual controls in only that switching system. Later, real-time network performance monitoring capabilities were expanded by telemetering data to the NMC from a few distant lower-ranking switching systems. The main drawbacks of these older NMCs were decentralization and inadequate real-time network performance monitoring capabilities.

To overcome these drawbacks, the Bell System has developed a minicomputer-based system called EADAS/NM, which supports the operation of a centralized NMC.[9,12,13] Figure 4 shows this system collecting data in near-real time from toll and local switching systems in a large geographical area, such as an entire metropolitan area or a state. This allows such an area, called a "cluster," to be managed as a whole, and breaks away from the older concept of managing individual switching systems. Local and small toll switching systems are connected to the EADAS/NM system through an intermediate traffic data collection system. This intermediate data collection system consists of the Bell System Engineering Administrative Data Acquisition System (No. 1A EADAS) or a data acquisition system supplied by the General Trade. Large toll switching systems and CCIS Signaling Transfer Points (STPs) which have their own traffic data collection systems, interface with EADAS/NM via direct data links. As explained in more detail in Section 4.1.1, the EADAS/NM processor performs exception calculations on five-minute traffic measurements and displays the results in the NMC on a wall panel, printers, and Cathode Ray Tube (CRT) display "pages."

Figure 5 shows manual network management centers organized in

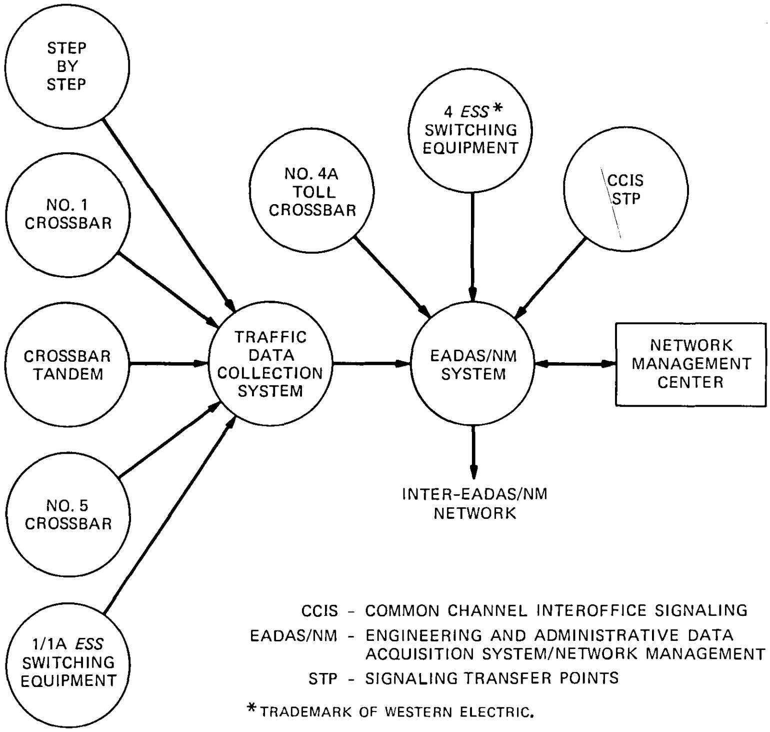

Fig. 4—EADAS/NM interfaces.

a three-level hierarchy. The first or basic level is occupied by the EADAS/NM-supported Network Management Center. Twenty-seven such centers are now in service, spanning the entire United States.[11] The switching systems and CCIS STPs of the telecommunications network are shown as dotted lines. Each NMC is jointly staffed with Bell Operating Company, Long Lines, and in some cases, Independent Company personnel.

The Regional Operations Center (ROC) occupies the next level in this hierarchy. Each ROC is supported by an EADAS/NM computer system and provides a higher level of network performance monitoring and control than the basic Network Management Center. Whereas the basic NMC has the direct responsibility for the control and real-time performance monitoring of the toll and local networks in its cluster, the ROC has the responsibility for the coordination of activities between the two or three NMCs in its region. There are ten switching regions in the United States and therefore ten Regional Operations Centers. The ROCs are also focal points for real-time

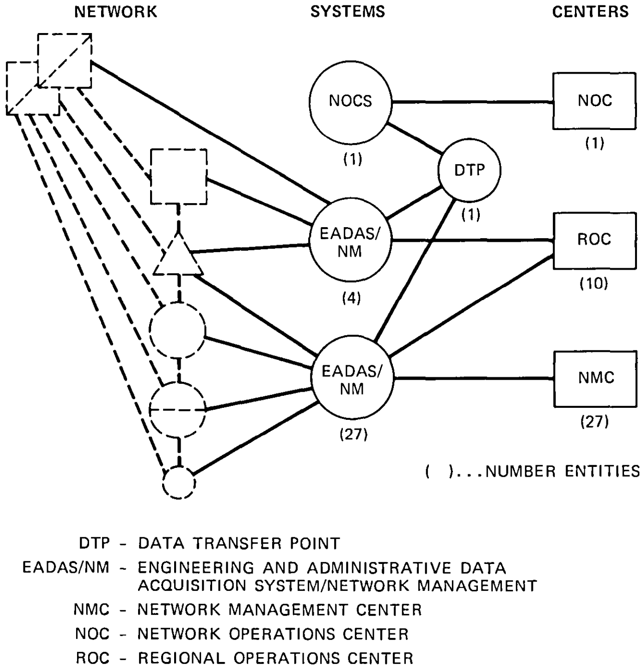

Fig. 5—Network management structure.

performance monitoring of portions of the CCIS network. As shown in Fig. 5, four ROCs are supported by dedicated EADAS/NM computer systems. Each of the remaining six ROCs is supported by an EADAS/NM system which also supports an NMC. There are 31 EADAS/NM network management operations support systems now in use.

A single center called the Network Operations Center (NOC)[10,11] makes up the top level of this hierarchy. As described in more detail in Ref. 11, this center is supported by its own computer system and it is responsible for coordinating interregional network management problems between the 10 ROCs, the 27 NMCs, and 2 regional network management centers in Canada. It also has the main responsibility for international network management and real-time performance monitoring of the entire CCIS network. A large amount of data are transmitted between the EADAS/NMs and the NOC system via a data network. The hub of this network is a data switch called the Data Transfer Point (DTP), which furnishes the NOC with a constant flow of high-speed performance monitoring data for all key toll switching systems, CCIS STPs, and international switching systems in the United States.

3.3 Centralized network monitoring and manual control

Even a highly sophisticated automatic system cannot respond optimally to all network disturbances. Thus, manual controls are required to provide additional adjustments to traffic flows in the network.

Figure 6 shows an example of an EADAS/NM-supported NMC in Boston, with the wall display panel in the background and interactive CRT systems in the foreground. The display system, which is described further in Section 4.1.2, is organized so that the activation of an indicator on the wall display panel directs the manager to CRT display pages for problem investigation and control activation. From one EADAS/NM-supported NMC, network managers can monitor and control several hundred switching systems.

Manual controls are activated from an NMC in many switching systems simultaneously through interactive CRT control pages. Control commands are communicated back to the switching systems by EADAS/NM over the same interface links over which data are collected from these systems. As shown in Fig. 7, manual network management controls are either protective or expansive in nature. Protective trunk group controls include trunk group controls which deny certain traffic access to a trunk group (cancel and skip controls) as well as trunk group directionalization. Protective controls are typically employed to limit traffic and thereby prevent the spread of traffic congestion. Expansive controls include controls that reroute traffic away from overloaded or failed facilities to facilities with idle capacity. Some computer-based tools that permit network managers to identify and solve traffic problems between regional switching systems are available at the NOC. These problems are solved through reroutes of which hundreds are implemented on peak days. Code block controls, which stop all or most of the traffic destined for overloaded points in the network, are employed mostly during focused overload congestion situations.

Because of the high number of reattempts, code block controls are not very useful in a mass-calling situation. With the deployment of CCIS, call-setup times are significantly reduced. Thus, the number of ineffective attempts per trunk can exceed 10,000 calls per hour, thereby defeating any mechanism to "choke" traffic flow based on the provision of a few, temporary, segregated trunk groups for traffic destined to a particular code. As shown in Fig. 8, a code block control based on a percentage restriction will still result in a high traffic volume under extreme conditions. Thus, a new control, "call gap," has been introduced. It consists of an adjustable timer that blocks all calls to a specified destination code for the set interval of time. The next call that arrives after the expiration of the time interval is allowed, after which the time gap begins again. This rate is strictly limited and

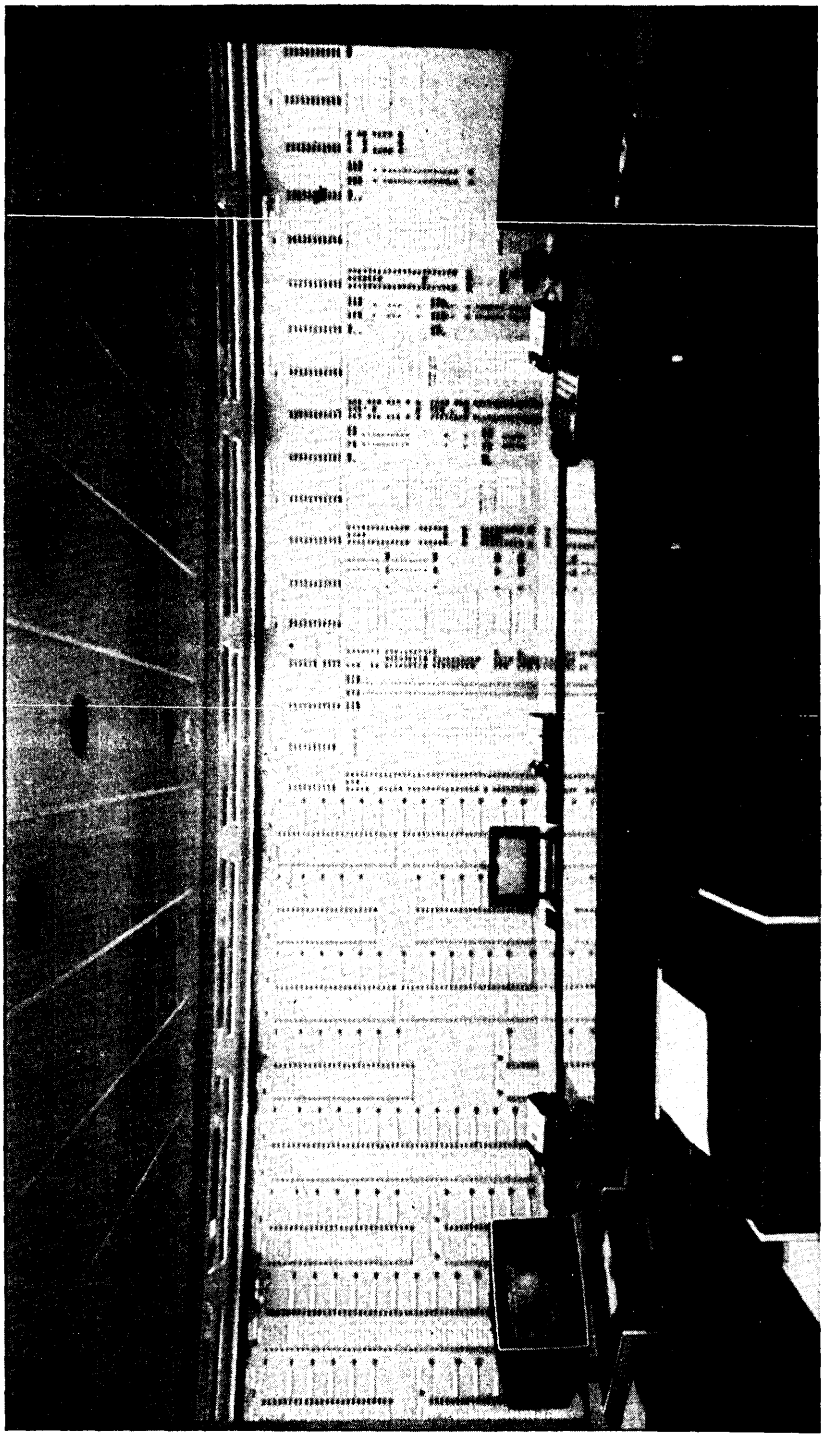

Fig. 6—Network management center.

PROBLEM	CONTROL	RESPONSE
NETWORK OVERLOADS OR FAILURES	CANCEL-TO/FROM SKIP TRUNK DIRECTIONALIZATION	PROTECTIVE
	REROUTE	EXPANSIVE
FOCUS OVERLOADS OR MASS CALLING	CODE BLOCK CALL GAP	PROTECTIVE

Fig. 7—Manual network management controls.

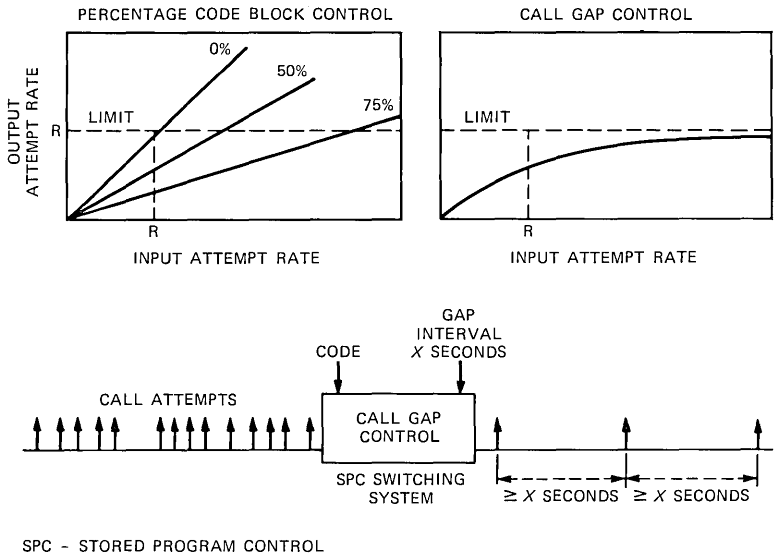

Fig. 8—Comparison of code controls.

approaches the inverse of the gap interval. This control will be useful in situations that require a set rate of traffic flow, particularly in mass calling and focused overloads.

IV. EADAS/NM—THE COMPUTER SYSTEM

4.1 Basic EADAS/NM functions

EADAS/NM is the minicomputer-based operations support system for NMCs and ROCs.[14] Its major functions are: (1) collection of network traffic and performance occurrences on an event basis every 30 seconds and associated traffic counts every five minutes; (2) analysis of traffic data to ascertain exception conditions and output of results to a wall display, printers, and interactive CRT terminals; (3) implementation of network management controls by transmission of control messages to appropriate switching systems; (4) record base maintenance by auditing selected information from SPC switching systems and by manual input of record items; and (5) transmission of selected information to the NOCS for national network performance monitoring.

4.1.1 Data collection and exception calculations

Data collection is done on a five-minute basis for traffic counts and on a 30-second basis for discrete (event) information. At the start of a five-minute interval (or a 30-second interval for discretes), EADAS/NM dispatches poll messages over direct data links to 4 *ESS* switching systems and PBCs (Peripheral Bus Computers), and through EADASs or equivalent General Trade systems to various local and small toll switching systems. The PBCs are administrative adjuncts to No. 4A Toll Crossbar Electronic Translator Switching Systems (No. 4A ETSs) and to STPs. Upon the completion of data acquisition for a given switching system, exception calculations are performed by comparing traffic data to preset thresholds. Once all polling and exception calculations are complete, the wallboard is updated to alert network managers of any changes in network conditions that may require further attention. Critical discrete changes update the wallboard on a 30-second basis. Exceptions of interest to the NOC (usually those involving switching systems and trunk groups in the upper three levels of the MTS network hierarchy) are routinely sent over a dedicated data link through the DTP to the NOCS.

The operation of an ROC EADAS/NM system is similar to the operation of an NMC EADAS/NM system except that the ROC EADAS/NM system obtains data both by direct polling and by passive monitoring. Typically, an ROC EADAS/NM system directly polls the regional switching system and the PBCs associated with STPs. It also receives data from passive connections to selected 4 *ESS* switching systems and PBCs. Figure 9 illustrates this arrangement.

4.1.2 Data display

The data display systems for EADAS/NM include the wallboard and exception printer as alerting devices, a monitor printer, and a repertoire of over sixty CRT display pages. The exception printer prints out a record of exceptions as they occur. The monitor printer allows the network manager to designate particular trunk groups (often those which are being used in a network management reroute control) for continuous observation.

The EADAS/NM CRT subsystem is comprised of a number of fixed-format displays which provide demand access to traffic data for

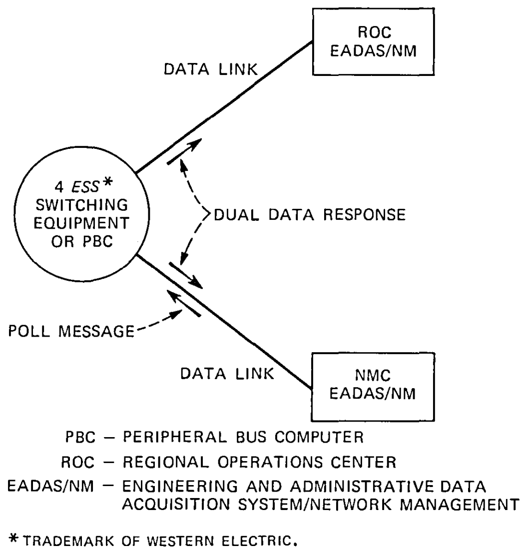

Fig. 9—ROC passive data collection feature.

the last 20 minutes and to current control and discrete status. Displays also may be used by network management personnel to implement or remove network controls, to modify certain trunk group-related reference database information, and to initiate requests for continuous observation of network components through the monitor system.

The CRT displays are divided into the following categories: *trunk group*, for traffic data on trunk groups at a selected switching system or switching systems within a geographical area; *machine*, for equipment and control status for a selected switching system; *control*, to display and implement network controls at one or more switching systems; *exception*, for resolution of data triggering exception indicators on the wallboard; *input and monitor*, for certain database and monitor system requests; and *analysis*, for analyzing traffic data to ascertain, for example, where there is idle capacity in the MTS network for network management reroutes.

The human interface to the CRT subsystem employs standard CLLI (Common Language Location Identification) codes to designate switching systems and trunk groups. In this way, a consistent identification of network components is achieved throughout the NMC/ROC/NOC hierarchy.

4.1.3 Controls

Network Management controls are activated and deactivated from selected CRT pages, after which appropriate messages are transmitted to the switching systems where the controls are actually instituted. A control log is maintained in EADAS/NM so that network managers have a record of what actions were taken. Control changes affecting traffic in the upper levels of the MTS network are sent to the NOCS. If, for some reason, a network management control is applied locally at a switching system, an alerting discrete is picked up by EADAS/NM on the next 30-second discrete poll. EADAS/NM then does an audit to determine and record the new control status for that switching system.

4.1.4 Record base definition and maintenance

EADAS/NM requires a very large record base which describes the "cluster" or portion of the network being managed. A full set of database management tools are provided for network management personnel (system users) to input record information and to validate it. Manual and automatic audit capabilities are provided to keep the record base current. For example, if more trunks are added to a trunk group, EADAS/NM receives an alerting discrete that triggers an automatic audit to update its database.

Reference information is also exchanged between EADAS/NM and the NOCS. Specifically, the NOCS transmits to EADAS/NM a list of switching systems that are of interest to the NOCS, along with a set of switching system and trunk group exception threshold levels. In this way the NOC can control what network information it receives. The NOCS can also audit EADAS/NM for the status of controls in switching systems that are of NOC interest.

4.2 System hardware configuration

Figure 10 shows a block diagram of the EADAS/NM minicomputer system and its relation to the data collecting and display devices. A brief description of the components follows:

EADAS/NM Minicomputer—The data processor for the system; it interfaces to all of the data collection, storage, and display peripherals, and executes the main system programs.

CRT Display Devices—Up to five DATASPEED Model 40 cathode ray tube terminal station arrangements in a KDP (Keyboard, Display, Printer) configuration.

Wall Display Board—Consists of an array of fluorescent, reflective, electromagnetically controlled indicators. A maximum of 4095 indicators can be used with a single wallboard driver. Up to two wallboard drivers may be equipped on an EADAS/NM system.

Disk Drives—Mass storage system using magnetic disk packs. The disk packs have a capacity of 80 million 16-bit words.

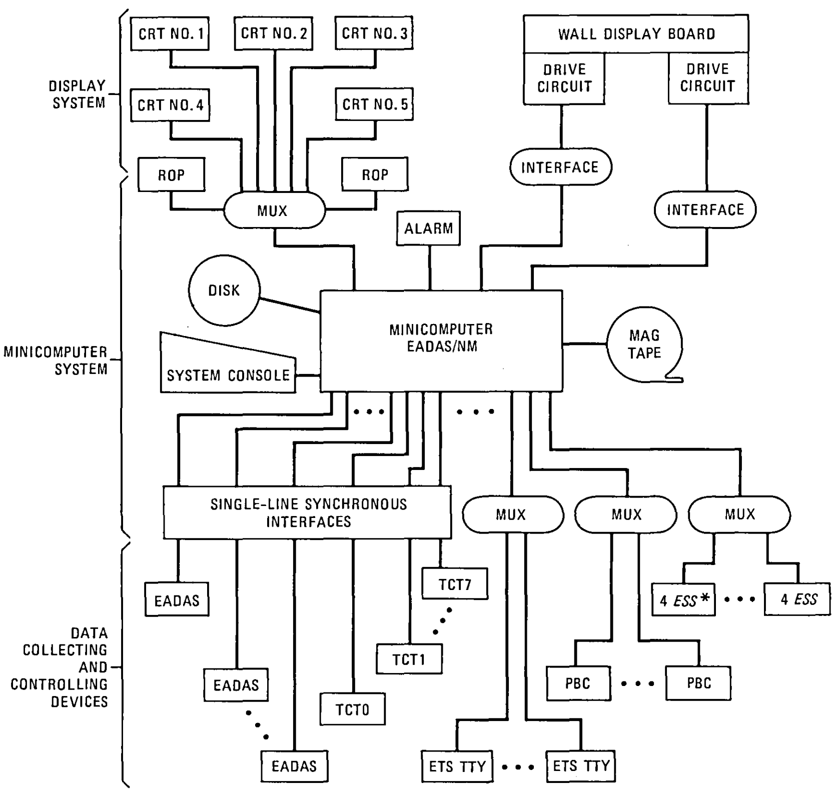

CRT – CATHODE RAY TUBE
EADAS – ENGINEERING AND ADMINISTRATIVE DATA ACQUISITION SYSTEM
EADAS/NM – ENGINEERING AND ADMINISTRATIVE DATA ACQUISITION SYSTEM/NETWORK MANAGEMENT
ETS – ELECTRONIC TRANSFER SYSTEM
MUX – MULTIPLEXOR
PBC – PERIPHERAL BUS COMPUTER
ROP – RECEIVE-ONLY PRINTER
TCT – TELEMETRY COMPUTER TRANSLATOR
TTY – TELETYPEWRITER
* TRADEMARK OF WESTERN ELECTRIC.

Fig. 10—Hardware configuration.

Command Console—Data terminal which may be used for any of the system commands.

Magnetic Tape System—Standard nine-track tapes at either 800 or 1600 bits per inch (bpi) used to initially load the programs, to back up the database, and to record real-time data for subsequent playback. The playback feature helps network managers review and critique specific events of network management interest using the wall display board and CRT displays. Such events can be played back in slow, real, or fast motion.

Telemetry Computer Translator (TCT) Interface—The E2A telemetry computer translators, each of which interfaces with up to four remote terminals, are used to receive discretes and send controls to

No. 4A ETS and Crossbar Tandem switching systems via the E2A telemetry system.

General Purpose Interface—Interface for both the wallboard and the E2A telemetry system.

Single Line Synchronous Interface—Interface capable of transfers at up to 9600 characters per second. These are used to transfer data from an EADAS to EADAS/NM and from EADAS/NM to NOCS.

Multiplexor (MUX)—Connects the processor with up to 16 asynchronous serial communication lines. One multiplexor is required for the EADAS/NM CRT terminals and receive-only printers (ROPs), and one or more additional multiplexors are required for No. 4A ETS, PBC, and 4 *ESS* switching system data links.

ETS TTY—100-baud teletypewriter channels used to transmit certain network controls to No. 4A ETS switching systems.

4.3 Software architecture

The EADAS/NM minicomputer software consists of the *UNIX** operating system and a collection of application programs that perform the EADAS/NM tasks. The application programs can be divided into real-time tasks, which perform data gathering, processing, and display functions; and certain non-real-time tasks. All applications programs are written in the C programming language. There are 12 major subsystems comprised of over 70 independent programs and over 1100 C-language modules. The overall software structure is depicted in Fig. 11 and will be discussed in more detail in this section.

4.3.1 The UNIX *operating system*

The *UNIX* operating system is general purpose, multiuser, and interactive. Utilized by EADAS/NM, its features include the following:

1. A hierarchical file system incorporating demountable files.
2. Compatible file, device, and interprocess Input/Output (I/O).
3. The ability to initiate asynchronous processes.
4. A system command language that can be selected on a per-user basis.

It should also be noted that the *UNIX* operating system is used for program development of the EADAS/NM project, providing a complete environment for the entry, compilation, loading, testing, maintenance, and documentation of the software product.

4.3.2 The EADAS/NM database

The EADAS/NM reference database consists of a number of user-

* Trademark of Bell Laboratories.

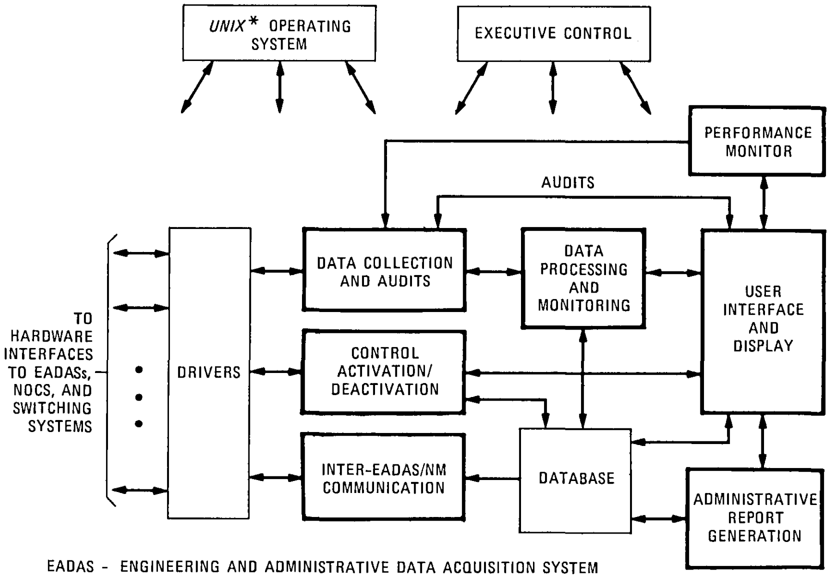

EADAS - ENGINEERING AND ADMINISTRATIVE DATA ACQUISITION SYSTEM
EADAS/NM - ENGINEERING AND ADMINISTRATIVE DATA ACQUISITION SYSTEM/NETWORK MANAGEMENT
NOC - NETWORK OPERATIONS SYSTEM
*TRADEMARK OF BELL LABORATORIES.

Fig. 11—EADAS/NM software structure.

built text files. These files are entered into the system using the *UNIX* operating system text editor.

In general, the information contained in these files falls into two categories. First is the information that pertains to the NMC itself, which includes the configuration of the center, the wallboard layout, and the threshold tables. The second category of information is concerned with the network to be monitored. This includes details about the switching systems and trunk groups in the NMC's cluster.

Once the database files are prepared, an operational database can be built using the database commands.

4.3.3 Executive control program

The executive control program is responsible for system initialization and for controlling the various real-time application processes. It accepts initial system parameters, starts the real-time application processes, monitors those processes for abnormal termination, and in general establishes and maintains the heartbeat of the system.

4.3.4 Interface drivers

A collection of driver programs interface the application software with the interface hardware. These programs provide the basic communications protocols for transmitting and receiving information on the data links that connect to EADAS systems, the NOCS, and the various switching systems.

4.3.5 Application programs

Those blocks that are heavily outlined in Fig. 11 represent the actual EADAS/NM application programs. The major functions of these programs include data collection and audits, control activation and removal, inter-EADAS/NM communication, data processing, and user interface and display. In addition, a performance monitor maintains performance statistics on major system events such as data collection, exception calculations, and delivery of information to the NOCS. Also, a set of administrative report programs automatically generate needed switching system and trunk group summary reports.

The data collection and audit subsystem includes programs for polling both five-minute traffic data and 30-second discrete information. For example, the 4 *ESS* switching equipment traffic data polling programs poll 4 *ESS* switching equipment for switching system, trunk group, and HTR code data and then store these data on disk for subsequent retrieval by the data processing programs. Discrete polling programs perform similar functions and also update the wallboard based on the discrete information. Audits are essentially demand polls that may either be requested manually by the network manager or automatically whenever an alerting discrete (see Sections 4.1.3 and 4.1.4) is received. The audit program collects the new information and updates the system reference or control data accordingly.

Control activation and removal programs receive control requests from the user interface and display subsystem, format the control messages in the appropriate switching system command protocol, and send them to the switching system via the interface driver programs. The control programs perform checks to verify that controls are properly activated and deactivated. They also maintain up-to-date control status records in the database and the control log.

The inter-EADAS/NM subsystem handles all communication with the NOCS. It includes programs that receive appropriate information from the NOCS and then establish in the database a record of what switching systems and trunk groups are of NOC interest. Other programs arbitrate queues that contain exceptions, control status changes, and HTR code data for transmission to the NOCS. Still other programs format the messages and pass them to the NOCS interface driver for transmission. The inter-EADAS/NM subsystem can also handle messages that come from or go to other EADAS/NMs via the DTP.

The set of data processing and data monitoring programs perform calculations on five-minute traffic data to determine whether exception thresholds have been exceeded and to support trunk group data monitoring requests. These programs also update the wallboard, based upon the results of the exception calculations, and feed the inter-EADAS/NM subsystem with NOC-interest exceptions. Discrete monitor programs monitor user-specified discretes for user-specified time intervals and produce a discrete monitor list which is passed to the user interface and display subsystem for display.

The user interface and display subsystem includes all of the various CRT page programs (see Section 4.1.2), the programs that control the printing of information on the exception and monitor printer, and the programs that receive and interpret user commands.

V. FUTURE NETWORK MANAGEMENT NEEDS

Future network management studies are aimed at providing improved methods for controlling the evolving switched network under overload and failure conditions. These efforts are expected to increase the reliance on automatic controls and improve centralized operations procedures and supporting system capabilities.

Simulation and analytical studies are under way to characterize the overload dynamics of a network dominated by electronic switching systems and CCIS. Specific studies focus on control strategies for new NCP-based services such as enhanced 800 Service, and the proposed evolution from the present hierarchical network to nonhierarchical, time-varying routing.[15]

REFERENCES

1. A. E. Ritchie and J. Z. Menard, "Common Channel Interoffice Signaling: An Overview," B.S.T.J., *57,* No. 2 (February 1978), pp. 221–24.
2. P. J. Burke, "Automatic Overload Controls in a Circuit-Switched Communications Network," Proc. Nat. Elec. Conf., December 1968, pp. 667–72.
3. D. G. Haenschke, D. A. Kettler, and E. Oberer, "Network Management and Congestion in the U.S. Telecommunications Network," IEEE Trans. Commun., *COM-29,* No. 4 (April 1981), pp. 376–85.
4. L. A. Gimpelson, "Network Management: Design and Control of Communications Networks," Elec. Commun., ITT Corp., *49,* No. 1 (1973), pp. 4–22.
5. Y. Nakagome and H. Mori, "Flexible Routing in the Global Communication Network," Seventh International Teletraffic Cong., Stockholm, June 1973, Paper 426.
6. J. H. Weber, "Some Traffic Characteristics of Communications Networks and Automatic Alternate Routing," B.S.T.J., *41,* No. 2 (March 1962), pp. 769-96.
7. J. H. Weber, "A Simulation Study of Routing and Control in Communications Networks," B.S.T.J., *43,* No. 6 (November 1964), pp. 2639–76.
8. T. V. Greene, D. G. Haenschke, B. H. Hornbach, and C. E. Johnson, "No. 4 ESS: Network Management and Traffic Administration," B.S.T.J., *56,* No. 6 (September 1977), pp. 1169–202.
9. D. G. Haenschke, "Network Management—An Overview of Controls and Systems," Proc. National Electronics Conf., October 1980, pp. 364–67.
10. E. Oberer, "The Role of the AT&T Network Operations Center," Proc. of the

Comm. Techniques Seminar, Princeton University, March 1980, pp. 1-1 through 1-6.
11. W. S. Bartz, and R. W. Patterson, "Total Network Data System: National Network Management," B.S.T.J., this issue.
12. R. M. Averill, Jr. and R. E. Machol, "A Centralized Network Management System for the Bell System Network," Int. Switch. Symp. Record, Japan, October 1976, pp. 433-3-1 through 433-3-7.
13. D. G. Haenschke, "Network Management Systems in the United States," Proc. 9th Int. Teletraffic Cong., October 1979, Torremolinos, Spain, Paper 632.
14. R. L. Bennett, "EADAS/NM—The Computer System Supporting the Area and Regional Centers," Proc. Nat. Elec. Conf., October 1980, pp. 370-5.
15. G. R. Ash, R. H. Cardwell, and R. P. Murray, "Design and Optimization of Networks With Dynamic Routing," B.S.T.J., *60,* No. 8 (October 1981), pp. 1787-820.

AUTHORS

D. G. Haenschke, Diplom Vorpruefung, 1951, Diplom Ing., 1953, Technical University, Vienna: Institute for Telecommunications, Vienna University, 1953–1955; Bell Laboratories, 1955—. Mr. Haenschke has planned communications systems for air traffic control, performed traffic studies applied to data communications, and engaged in long-range planning studies for the telecommunications network. In 1968 he became Supervisor of a group responsible for performance appraisals of operator services traffic systems, and the evaluation of a *Picturephone*® service trial. Since 1969 he has been responsible for the planning of traffic network management and overload control features for the telecommunications network and its interfaces with network management operations support systems. He is currently supervising a group concerned with network management capabilities for new switching and signaling systems and for new services for the Stored Program Control Telecommunications Network. Senior member, IEEE; member, AAAS, Verband Deutscher Electrotekhniker.

George C. Ebner, B.S.E.E., 1964, M.S.E.E., 1965, Massachusetts Institute of Technology; Ph.D. (Electrical Engineering), Ohio State University; Bell Laboratories, 1965—. Mr. Ebner's initial work at Bell Laboratories was in the exploratory development of switching processors and their associated peripherals. Later he was involved in the development of CCIS equipment. In 1973, Mr. Ebner became Supervisor of the CCIS Engineering and Current Planning Group, a responsibility whose major thrust was systems engineering for the rapidly evolving CCIS network. In 1979, he became Supervisor of the Network Management System Group, which develops generic software for the EADAS/NM. Member, Eta Kappa Nu, Tau Beta Pi, Sigma Xi.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

National Network Management

By W. S. BARTZ* and R. W. PATTERSON*

(Manuscript received June 14, 1982)

During periods of network stress, network management (real-time monitoring and control of the network) assures optimum call-carrying capacity. To assure timely reactions to network overloads and failures, the national Network Operations Center coordinates the activities of a set of network management operations centers. Supporting the Network Operations Center, the Network Operations Center System collects data for major network elements and, within minutes, provides a national overview of exceptional conditions.

I. INTRODUCTION

Real-time monitoring and control of the network through a coordinated set of network management operations centers assures optimum call-carrying capacity during periods of stress caused by overloads and failures. Switching systems are controlled and monitored in detail in specific geographical areas at the 27 Network Management Centers (NMCs) that blanket the country. Regional Operations Centers (ROCs) coordinate NMC activities within each of the ten switching regions, and monitor and control the regional switching and signaling systems. To coordinate and manage the activities of the NMCs, the ROCs, and the two Canadian regional management centers, the na-

* Bell Laboratories.

tional Network Operations Center (NOC) at AT&T Long Lines' headquarters was placed in operation in 1977.[1]

The Network Operations Center System (NOCS), a one-of-a-kind system, supports the NOC. Twenty-four hours a day, seven days a week, NOCS collects data for the major elements of the telecommunications network, thousands of trunk groups, and hundreds of switching systems. It then checks these data for exceptional conditions, and displays information for these conditions within minutes of the occurrence. This national overview allows personnel at the NOC to coordinate appropriate reactions to national network situations, assuring uniformity of control.

The preceding paper in this issue of the *Journal*[2] explains the reasons for network management of the telecommunications network and describes the system that provides area and regional support for this function. In this paper, we describe the functions of the NOC, its support system, NOCS, and network management coordination for the North American telecommunications network.

II. NETWORK MANAGEMENT COORDINATION

2.1 The network management operations structure

As described in the preceding paper in this issue,[2] network management centers are organized in a three-level hierarchy. The national NOC is at the top level of this hierarchy and is supported by NOCS. Personnel at the NOC have ultimate responsibility for coordinating network management activities for situations that transcend regional boundaries. By assuring that actions taken optimize the use of network facilities, they protect areas hit by disasters and extraordinary calling volumes, coordinate the use of idle capacity when failures cut the capacity of normal routes, and assure consistent, compatible control application and removal.

While the NMC/ROC structure covers the entire domestic network, this center structure does not extend overseas. Most of the information about overseas destinations comes from data collected at the international gateway switching systems. Because no more than one gateway switching system is covered by any one NMC, the NOC plays a more direct role for overseas network management. Data from all gateways come together at the NOC. Thus, personnel at the NOC are in a unique position to see how calls from the domestic network as a whole are completing to overseas points. Because of this, NOC personnel take a more direct role in coordinating management activities at the international gateways and in working with overseas partners to assure consistency in methods and procedures.

2.2 *Network management methods*

In addition to monitoring the network and coordinating management activities in real time, the NOC is responsible for establishing a unified network management methodology. NOC personnel work with AT&T and Bell Laboratories to determine future directions and capabilities for network management. Guidelines for the selection of automatic control responses and for manual control strategies are sent to the ROCs and forwarded to the NMCs. The NOC organizes preplanning for anticipated problems and post-event analysis and critiques. Finally, the NOC is responsible for training network managers by arranging on-site rotational assignments, administering a training school, and issuing newsletters covering recent developments affecting network management.

Putting the responsibility for network management methods at the NOC assures a consistent, unified application of network management techniques throughout the domestic telecommunications network.

III. THE TELECOMMUNICATIONS NETWORK

The telecommunications network for which the NOC personnel are responsible comprises three distinct but highly interdependent networks.

First there is the top portion of the North American network hierarchy of toll switching systems and the trunk groups that interconnect them. This includes all ten United States regional (Class 1) switching systems and the two Canadian regional switching systems, about 70 sectional (Class 2) systems, 200 primary (Class 3) systems, and some toll (Class 4) systems. In total there are about 300 to 350 switching systems of interest to the NOC. The NOC does not monitor the local switching systems or their associated trunk groups.

The second network that the NOC personnel monitor is the overseas network, which comprises seven gateway switching systems through which all overseas traffic flows, trunk groups to several hundred overseas destinations, and a few trunk groups in the United States dedicated to overseas traffic between gateway systems. The seven gateway systems also serve as switching systems in the domestic network, where they range from Class 1 to Class 3 switching systems. The several hundred overseas trunk groups terminate in over one hundred foreign countries. There may be one or many trunk groups to any given country. Three different types of trunk groups within the United States are dedicated to overseas traffic. First, there are intergateway trunk groups reserved for overseas calls that overflow a particular gateway's trunk group to the destination country and must be routed to another gateway to use its trunk group to the desired destination. Second, there are intergateway trunk groups that are used exclusively for incoming overseas calls that are merely transiting the U.S. to get to another overseas point. Third, there are trunk groups that connect the gateway systems to the overseas operations at the international operator centers.

The third network monitored by the NOC personnel is the Common Channel Interoffice Signaling (CCIS) network.[3] This high-speed, highly reliable signaling medium provides for fast call setup, efficient use and control of the network, and network features that require access to centralized databases. It is composed of several pairs of Signal Transfer Points (STPs) and thousands of signaling links that are of interest to the NOC. The signaling links connect the STPs to each other, to the switching systems, and to centralized databases called network control points.

Collectively these three networks are composed of many elements. To receive data on all of them constantly at one location would be an enormous task. If it were done, only a small part of the data collected would be used most of the time. Because of these two facts, the NOCS does not receive data on all elements all the time. Rather, it receives data primarily on an exception basis: data are transmitted to the NOCS when a problem develops in an area. While there are a few items that are always reported, such as data on interregional final trunk groups and overseas trunk groups, data for most entities are rarely reported.

IV. OVERVIEW OF THE NETWORK OPERATIONS CENTER

The philosophy behind the Network Operations Center is exception alerting, followed by detailed reporting on the upper levels of the network. Because status information on all parts of the network would be much too voluminous to sort through manually, NOCS analyzes and displays information based on calculation results that exceed normal thresholds. Managers were alerted to these exceptional conditions in the network primarily by a wall display. This wall display, partially shown in Fig. 1, is visible from all operating positions. It is divided into network management and facility management sections. The display is functionally separated into two pieces; the right portion of the network management section displays information on the North American network, while status of overseas trunking is displayed on the left. The facility management portion of the wall display and the facility management function of the NOC are not discussed in this article.

The main part of the North American display is arranged in a matrix wherein each column represents data from one of the 12 switching regions. Conditions in any particular region affecting the

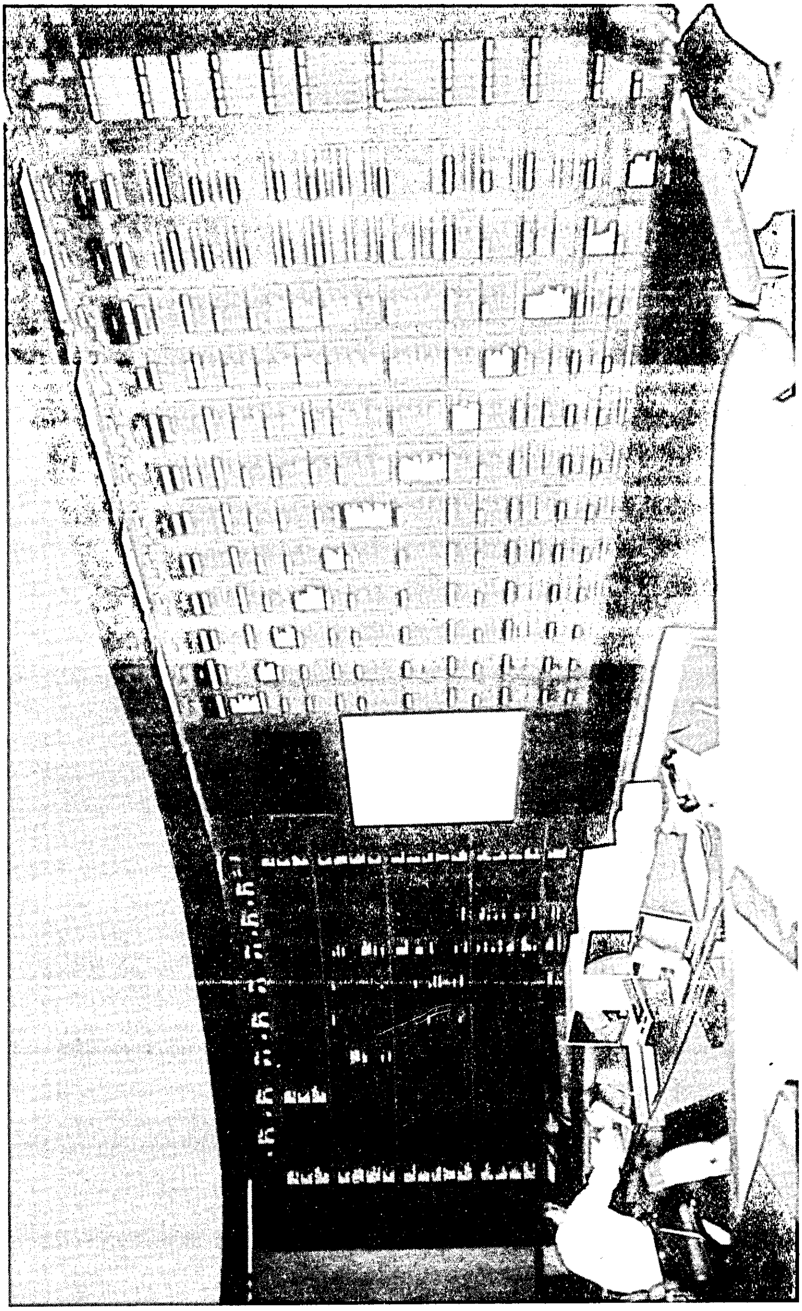

Fig. 1—NOC network management wall display.

rest of the network can be quickly deduced by scanning the appropriate column of the matrix. Similarly, conditions affecting a region can be detected by scaninng the appropriate matrix row. Thus, the network manager can quickly assess the location and scope of a network problem by viewing the domestic part of the wall display.

The main part of the international wall display is also laid out in a matrix fashion. The columns of this matrix represent the seven international gateways. The rows represent individual foreign countries. For backup and record-keeping purposes, wall display information is also printed on exception printers at the NOC.

For focus on individual exception conditions, CRT terminal displays provide resolution of data that triggers wall display exceptions. Other displays allow managers to analyze traffic patterns and aid in selecting and monitoring routes outside the normal routing chain when the normal in-chain routes cannot accommodate the traffic offered to them.

Analysis of troubles in the CCIS network is handled by an interactive, color graphics display, which provides information for both an overview and a detailed analysis of CCIS network problems.

V. NETWORK OPERATIONS CENTER SYSTEM

The primary need of the personnel at the NOC in carrying out their job is information—data telling them about problem conditions in the telecommunications network and helping them decide what, if anything, they can do to alleviate the problems. Making this information available quickly and in the most usable format possible is the task of the NOCS.

5.1 System architecture

Reference 2 describes how area and regional network management is supported by the Engineering and Administrative Data Acquisition System for Network Management (EADAS/NM). One function of these systems is data collection and screening for the NOCS. These systems collect data for all entities in the domestic network that are of interest to the NOC personnel. (The two Canadian regions are not monitored through EADAS/NM; data collection for those switching systems is provided by an E2 telemetry system.) The EADAS/NMs perform calculations on the data they collect and compare the results to thresholds established by the NOC personnel. For those entities that have exceeded one or more thresholds and for those entities that the NOC personnel have scheduled, the calculation results are transmitted to the NOCS.

Working in conjunction with NOCS, the inter-EADAS/NM network collects data from the various EADAS/NMs. This network consists of the Data Transfer Point (DTP) together with the set of links to the EADAS/NMs over which information is sent. The EADAS/NM links to the DTP operate asynchronously at 1800 baud. As we will note later, these links are used in both directions and do more than just transmit results of real-time data calculations to the NOCS.

Figure 2 shows the configuration of the multiprocessor NOCS with the DTP. Because of its important function in the network management environment, it is a duplicated system. It is possible to switch

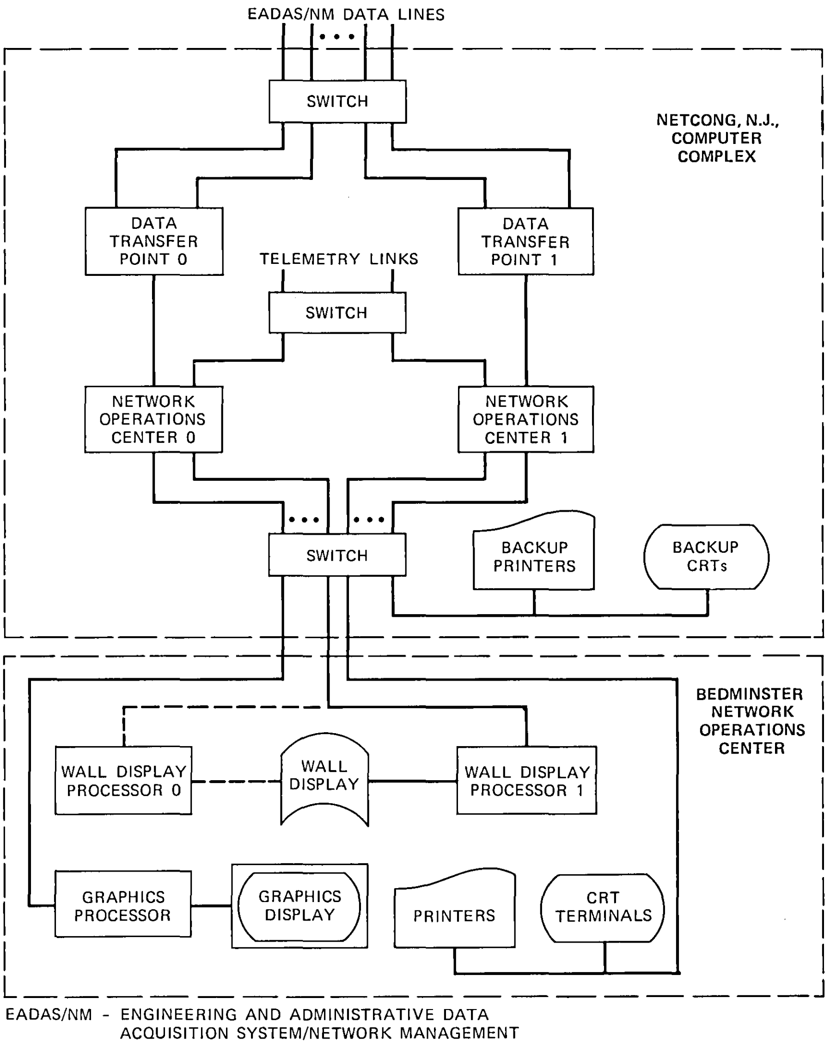

Fig. 2—System configuration for NOCS.

to the backup system within 15 minutes if a failure occurs on the active system. To achieve this, the backup system is always kept operational and is identical to the on-line system. Manual switches allow all EADAS/NM lines and all input/output (I/O) devices to be switched from one system to the other.

The primary computers in the system are not colocated with the primary display devices. As we stated earlier, the NOC is located in the Long Lines headquarters building in Bedminster, New Jersey, where the primary set of display devices and the small peripheral processors are located. The main computers, though, and the terminations of the EADAS/NM links are located about 25 miles away in Netcong, New Jersey. Also located at Netcong are a minimal set of backup display devices that can be used if a failure prevents the use of the facilities in Bedminster.

A simplex configuration of the systems would consist of the following elements:

1. The Data Transfer Point
2. The Network Operations Center System (NOCS) processor (labeled NOC in Fig. 2)
3. The Wall Display Processor (WDP) connected to the wall display
4. The Graphics Processor (GP) connected to the graphics display and the joystick
5. The Receive-Only Printers (ROPs)
6. CRT terminals.

Each DTP-NOC pair has a permanent high-capacity link joining the two processors. The WDPs are connected through the switch arrangement to a NOCS processor via a 4800-baud link. The wall display is 12 feet high, about 85 feet long, and uses luminescent, electromagnetically controlled indicators. These indicators are about one-inch square with black on one side and a reflective, highly visible color on the other side. Most of the wall display surface is actually taken up with labels. These identify the entities associated with the conditions represented by the indicators. The GP is also connected through a switch to a NOCS processor via a 4800-baud link. The ROPs are 4800-baud printers. The system supports both 1200-baud and 4800-baud CRTs. The telemetry links shown in Fig. 2 collect data from the two Canadian regions.

5.2 DTP functions

Although the DTP performs several functions in the network management support system environment, its primary function is transferring data among the links that terminate on it. These links connect with the EADAS/NMs and the NOCS processor. Address information received in each message enables the DTP to determine on which link or links the message is to be transmitted. Sometimes the message cannot be routed because of insufficient address information. In these instances the DTP must retrieve additional address information from its database before it can transmit the message to its proper destination(s). Keeping this additional address information in the DTP database eliminates the need to maintain it at all the EADAS/NMs.

In Section 5.1 we stated that the EADAS/NMs determine what data the NOC personnel want and then transmit that data to the NOCS. This information arrives at the DTP addressed to the NOCS, which is where the DTP routes it. In addition to this type of message the DTP also receives messages from the NOCS processor addressed to one or more EADAS/NMs. These messages supply information such as thresholds to be used to determine what data to send to the NOC. The DTP also receives messages from EADAS/NMs addressed to other EADAS/NMs.

The basic data update interval used by the network management centers and the NOC is five minutes. Although some events are reported as they occur or soon thereafter, the bulk of the network management data consists of counters scored locally by the switching systems or STPs. These are read and zeroed once every five minutes. The timing of these five-minute intervals is another function that the DTP performs. When a period begins, the DTP sends a message to the NOCS processor and to all the EADAS/NMs. As soon after the beginning of the period as possible, the EADAS/NMs start collecting data from the entities for which they are responsible, performing the calculations on these data, and sending the exceptions and scheduled data to the DTP. The DTP receives this five-minute data and passes it on to the NOCS processor until all the EADAS/NMs have reported that they have sent all their data. In the event that a cutoff time near the end of the five-minute interval is reached, the DTP stops the EADAS/NMs from further transmission and forwards the available data to the NOCS processor. When all EADAS/NMs have completed data transmission or the cutoff point is reached, the DTP sends a message to the NOCS and all EADAS/NMs to inform them that this period is over. Five-minute data for a period received by the DTP after the cutoff point is not passed on to the NOCS processor.

The DTP also monitors the inter-EADAS/NM links for problems and keeps its database up to date. If any of the links from the EADAS/NMs are having trouble, the DTP disconnects those, and then tries to reconnect them to clear the condition. The current status of all links is constantly available and can be accessed by the NOC personnel. The DTP receives all its database information automatically from the NOCS processor. Whenever the NOCS processor or the DTP is restarted, or the NOC personnel install new information in the NOCS database that could affect the DTP, the NOCS processor updates the DTP database without any manual intervention.

5.3 Reference data

The NOCS must have a large amount of reference data to display information. This reference data reflects the characteristics of all the switching, trunking, and signaling entities discussed in Section III. The NOCS reference data must be consistent with the reference data used at the EADAS/NMs. To achieve this consistency, and to handle the many database changes that occur, NOCS derives a major portion of its reference data from reference data tramsmitted by the EADAS/NMs to the NOCS processor over the inter-EADAS/NM links. It generates the remainder from data that are input by the NOCS database administrator.

The reference data transmitted from the EADAS/NMs to the NOCS processor are in binary format. When these data are received, they are converted to an ASCII format and stored in the NOCS. The NOCS database administrator can then modify and augment this information as required before using it to create the active NOCS database. Some information that is needed in the NOCS database is not received from the EADAS/NMs and must be manually entered and updated at the NOCS. The NOCS database administrator has a set of commmands to create, modify, examine, and audit additional reference information. The creation and installation of the active database from the combined EADAS/NM and locally supplied reference information are also under the control of the NOCS database administrator.

In addition to the reference data sent to the NOCS from the EADAS/NMs, there is also reference data sent from the NOCS to the EADAS/NMs. These include such items as thresholds to be used by the EADAS/NMs and a common set of location (switching system and STP) identifiers. To ensure that all EADAS/NMs are using the same version of these tables, each version has a unique identifier sent along with it. Whenever an EADAS/NM is reconnected to the DTP, it sends the NOCS a list of its current table identifiers. If they are not the correct set, the NOCS sends the EADAS/NM the proper tables automatically.

5.4 Real-time data

The types of real-time data that the NOCS receives from the EADAS/NMs include:

1. Five-minute data calculation results
2. Switching system and STP status indicators (hereafter referred to as machine status indicators)
3. Control activations and deactivations.

We discussed the five-minute data calculations in Section 5.1. The EADAS/NMs collect data, perform calculations on the data, compare the results to NOCS established thresholds, and transmit exceptions and scheduled data to the NOCS. The EADAS/NMs collect data on trunk groups, signaling links, switching systems, STPs, and destination codes. These raw data are used by the EADAS/NMs to calculate rates and percentages, which can then be compared to thresholds. Whenever it is determined that a calculation result exceeds its threshold, all the data for the entity concerned are sent to the NOCS.

The NOCS stores all the five-minute data calculation results that it receives in an interval for twenty minutes. Therefore, at any time network managers can retrieve the data for the four most recent five-minute time periods. When data for a new five-minute period are available, exceptions are immediately displayed and automatically recorded for the network managers. We discuss this further in Section 5.5.

NOCS uses the machine status indicators, received as often as every 30 seconds, to notify the NOC personnel of any significant switching system or STP problem as quickly as possible. The events that are reported include congestion and failure conditions measured at switching systems and STPs. Whenever new machine status information arrives, NOCS immediately displays it to the network manager. In addition, a history of all machine status changes is kept and is available to the network managers (see Section 5.5).

The third type of real-time data that NOCS receives is control activations and removals. NOCS receives information on code controls and trunk group cancel, skip, trunk reservation, and reroute controls. At the time that any of these controls are applied to or removed from a NOC interest trunk group, the NOCS is informed. Furthermore, the NOCS control audit capability permits it to request an update of all the controls that are active in all the switching systems monitored by an EADAS/NM.

5.5 Data display

Sections IV and 5.1 described briefly the functional characteristics of the NOCS I/O devices. This section describes how these devices provide an interface to the user.

5.5.1 The wall display

The primary objective of the wall display is to alert the network managers to problems in the network. In general it is not used to determine exactly what a problem is or how severe it is; this is done by other I/O devices available on the system. The indicators on the wall display are used to inform the NOC personnel that some calculation result has exceeded its threshold, some system condition has activated a machine status indicator, or some control has been invoked. About 30 percent of the wall display on the left is used for overseas items, about 40 percent in the middle is used for domestic conditions, and the remaining portion on the right is used for a display of transmission facilities. This right portion is not controlled by the NOCS and is not discussed in this article.

The domestic section has four general areas, as shown in Fig. 3. Area 1 is the largest portion. Its 12 × 12 matrix display of intraregional and interregional trunking conditions was described in Section IV. Each column has the regional center name at the top and the names of all the regional and sectional centers below it. To the left of each switching system name are four indicators that reflect conditions from the regional center listed at the top to that particular system. Area 2 also lists all the regional and sectional centers. To the left of the names this time are seven indicators that reflect code controls or trunk group controls activated at the named systems. To the right of these names are six (for sectional centers) or seven (for regional centers) more indicators showing the machine status indicators. Area 3 contains more specific information on which machine status indicators are active, as well as some miscellaneous items about the NOCS, threshold tables in use, and other conditions. Area 4 is similar to Area 1 in that there is a column for each of the twelve regions. Each column shows the status of the STPs in the region and the signaling links that emanate from them.

The overseas portion of the wall display is shown in Fig. 4. The layout in this portion is straightforward; all the countries that have trunk groups are listed in three columns. Each country also has the names of its gateway switching systems listed. To the left of the country names two items are shown: (1) a label listing the International Operator Center(s) that serves the country, and (2) indicators displaying code controls in effect. To the right of the country names

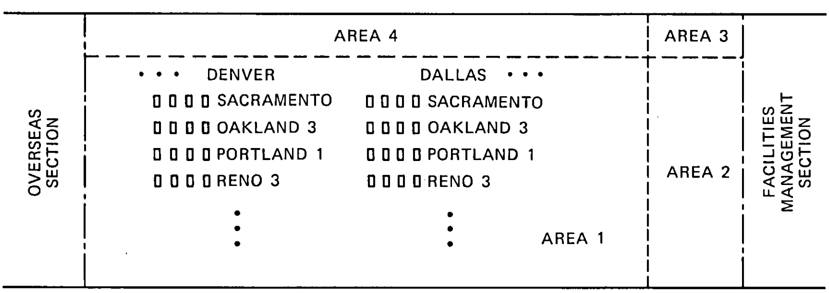

Fig. 3—Domestic section of the wall display.

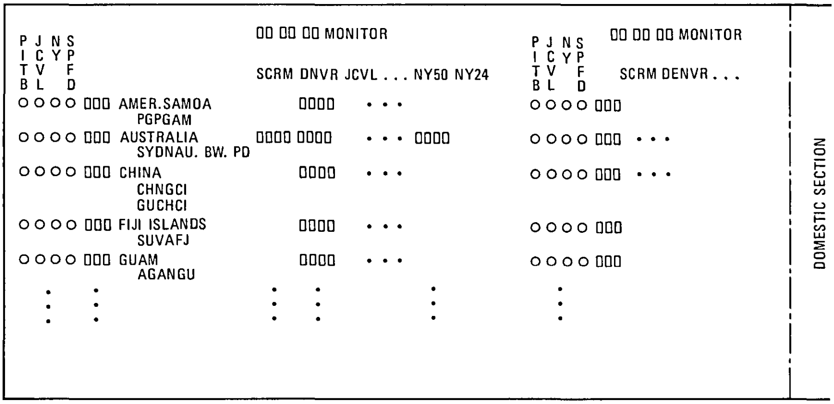

Fig. 4—Overseas section of the wall display.

are seven columns for the seven U.S. gateway switching systems. If a gateway serves a country, there are four indicators that show conditions to that country.

5.5.2 *The CRTs*

The CRTs are the primary tools that enable network managers to retrieve switching and trunking information from NOCS. Once alerted to a problem by the wall display, network managers analyze details of the problem using the CRTs. Since the network managers at the NOC have to obtain information from the system as fast as possible, a series of fixed format displays have been developed for use on the CRTs. They provide quick access to almost all the data in the system with a minimum of input. These displays or "pages" are divided into the following categories:

1. Analysis pages, for analyzing traffic data patterns
2. Exception pages, for resolution of data-triggering exception indicators on the wall display
3. Input pages, for some of the database modifications.

The pages provide access to data for the four most recent five-minute periods, current machine status indicators, and current controls. A page consists of up to 68 lines, 80 characters in length. It contains fixed information, known as "background" and variable input and output areas, referred to as "windows." Input windows enable network managers to select parameters quickly by entering a single character designate (+), or a short string of characters, such as the first few characters of a Common Language Location Identifier. After designating the appropriate input parameters or using the default values where appropriate, the network manager depresses the SEND key. The system then responds by filling in the output windows with the desired information.

Of the eleven analysis pages in the system, seven are for analyzing domestic trunk group conditions, such as reroute controls in effect or interregional trunk group congestion. The other four are for looking at overseas trunking conditions, such as the traffic from a gateway to a country. The 13 exception pages provide data for machine exceptions, domestic and overseas trunk group exceptions, and code-related completion problems.

Figure 5 is an example of an analysis page. This page is for viewing trunk groups from one geographical area to another. In the area at the top and along the left edge, the network manager enters selection parameters in input windows. The remainder of the page is the display area (output windows) for the trunk groups that meet all the input selection criteria. This area identifies the trunk groups by the Common Language Location Identifiers of the switching systems at each end, and data shows calculation results and conditions affecting the trunk group.

In addition to providing access to current data, status, and controls, the CRTs serve other purposes. The NOCS maintains a history of machine status indicators. Using a system command entered on any of the CRTs, network managers can obtain specific subsets of this file tailored to their current interest. A report may cover from 30 minutes to several days of data, and may specify a single machine or all machines, as well as one or more specific status indicators. The data in the report can be ordered by time or by machine.

```
AN85 OUTGOING NETWORKS                                  FROM AREA ( ) [              ]
MF( )  CCIS( )  ALL(s)      FNL( ) HU( ) ALL(s)        TO   AREA ( ) [              ]
-------------------------------------------------------------------------------------
     SC  SN  DN  DL  ST  NR  RC  PI  WA  WH  MT  RE  FR  ALL  R   S  PT  ALL          TIME
FM  ( ) ( ) ( ) (s) ( ) ( ) ( ) ( ) ( ) ( ) ( ) ( ) ( ) ( )  (s) ( ) ( ) ( )  0-5  0-15  0-20  5-20
TO  ( ) ( ) ( ) ( ) ( ) ( ) ( ) ( ) ( ) ( ) ( ) ( ) ( ) (s)  (s) (s) ( ) ( )  (s)  ( )   ( )   ( )
-------------------------------------------------------------------------------------
 OR  ( )    | ACH >  5 and %OCC  <  75                             %              %
 AND (s)    | FROM OFFICE         TO OFFICE                 EQ2W  OFL  ACH  CCH  OCC  TBIT CTRL % M
            | DLLS TX TL 34T( )  JCVL  FL  CL  03T  N2H      60    0    9    9   65  ---A RRT  50*
XCPT  ( )   | DLLS TX TL 34T( )  ALBQ  NM  MA  02T  N2H     120    0    5    5   72  ----
%OF[      ] | DLLS TX TL 34T( )  BLNG  MT  MA  02T  N2H      12    0    5    5   67  ----
ACH[      ] | DLLS TX TL 34T( )  DNVR  CO  MA  01T  N2H     120    0    7    7   69  ---- SKIP 75-
CCH[      ] | DLLS TX TL 34T( )  PHNX  AZ  MA  01T  N2H      72    0    9    9   75  ----
%OC[      ] | DLLS TX TL 34T( )  SLKC  UT  MA  01T  N2H      48    0   11   11   75  ----
            | DLLS TX TL 34T( )  BLTN  IN  01  04T  N2H      23    0    7    7   51  ---- CANT 75-
TBIT  ( )   | DLLS TX TL 34T( )  CHCG  IL  CL  46T  N2H     108    0    6    6   75  ----
   [      ] | DLLS TX TL 34T( )  DESM  IA  DT  04T  N2H      36    3   10   10   75  ----
            | DLLS TX TL 34T( )  FARG  ND  BC  03T  N2H      24    0    6    6   60  ----
CTRL  ( )   | DLLS TX TL 34T( )  GDRP  MI  BL  40T  N2H      36    0   15   15   36  ----
   [      ] | DLLS TX TL 34T( )  MPLS  MN  DT  04T  N2H      84    0    9    9   70  ----
-------------------------------------------------------------------------------------
PART [ 1 ] OF 4  DATA FOR 11:35-11:40                       TRANSFER TO SINGLE-TG ( )
SEPT 1 82  11:47:46 NWT      PRINT( )                  DIRECTORY[    ]     PAGE[    ]
```

Fig. 5—Example of an analysis page.

A large set of system commands is available to administer and monitor the NOCS and its database. These commands are also entered on the CRTs and can be directed to either the NOCS processor or the DTP. The DTP commands permit such functions as the attaching and detaching of EADAS/NMs; the starting and stopping of exception transmission from an EADAS/NM; and the printing of DTP data tables, inter-EADAS/NM network status, and various other network statistics. The NOCS processor commands permit such functions as initializing and testing the wall display and graphics processors, administering the database (as discussed in Section 5.3), auditing EADAS/NMs for current status indicators and controls, and controlling the playback function (to be described in Section 5.6).

5.5.3 The graphics displays

The graphics subsystem of the NOCS provides an interactive, color display that is used as both an alerting device and a detailed data source. It is interactive in that its fast update rate permits the user to receive near instantaneous response to information requests. It is capable of displaying 17 colors in 15 different shades and producing two- and three-dimensional shaded displays. This device permits adjustment of three input coordinates and the selection of items in the displays. Cursors on the monitor show the current values of the three coordinates.

The operating system is designed to support up to 15 functions, or collections of closely related displays. These can show information from the four latest five-minute data periods along with the current controls and machine status indicators. The operating system can also automatically update the current displays with new data when the NOCS processor indicates it is available. Currently, the functions available on the graphics system are:

1. CCIS monitoring
2. Domestic trunk group monitoring and analysis
3. A geographic editor.

Since the graphics system is a relatively recent addition to the system, this list of functions is expected to grow.

The CCIS display pictorially represents all the elements in the CCIS network, and color-coded symbols alert the network manager to exception conditions. The display shows the latest sample of five-minute data and current CCIS-related machine status indicators for the entire CCIS network. This display allows the network manager to determine the general nature of CCIS problems. At that point the network manager can request that particular link sets with exceptions be expanded and that labels and specific data calculation results be shown. The domestic trunk group display graphically represents the switching and trunking entities to be monitored and color codes problems to allow the network managers to focus quickly on the worst problem areas or to determine the type of problem that exists. The geographic editor is used for database administration only.

5.5.4 The printers

The receive-only printers on the NOCS are used as exception printers, as monitor subsystem printers, and for database work. The exception printers exist to maintain a permanent record of all domestic and overseas exceptions that are reported by the EADAS/NMs. Trunk group, machine, and signaling link exceptions are recorded.

The monitor subsystem gives the network manager a printed chronological record of data for up to 150 trunk groups. The trunk groups for which print criteria are selected and set are listed on a separate printer called the monitor printer. Furthermore, the network manager can activate special indicators on the wall display or set off an audible alarm when a monitor system threshold is exceeded. These wall display indicators exist on a per-region basis on the domestic part of the wall display and on a per-gateway-system basis on the overseas part of the wall display.

5.6 Playback

The playback subsystem of NOCS allows the status of the switching network, as reported to the NOCS by EADAS/NM and the E2A telemetry interface of NOCS, to be recorded and later replayed on the system. The recording is done at any time on the active NOCS processor. All the data messages that come over the DTP-NOCS link and all the data that comes from the E2A telemetry are copied to a file on the disk. This file is large enough to hold more than a day's worth of data at normal system load. There are two such files on the disk so when one is filled, the other can be activated. Either file can be copied to magnetic tape for later analysis if desired. The recording of incoming data in the playback system disk file does not affect normal operation of the NOCS. When the record mode is activated, all pertinent system information that will be required to replay the data is also recorded.

Previously recorded network data can be replayed at any time on the off-line NOCS processor. If the data to be replayed are not currently in a playback system disk file, they are read in from magnetic tape. Using the switches, the system administrator can connect any combination of I/O devices, i.e., wall display, CRTs, graphics, or printers, to the system running replay. These devices will then appear as though they are connected to a NOCS processor that is receiving live data. When the playback system is running in the replay mode, it retrieves the data stored on the disk file and feeds it to the NOCS as if it had just come from the EADAS/NMs or the E2A telemetry. Several options are made available to the replay operator. Replay can be made to run the system at a "normal" speed or faster or slower than normal. It can also be made to start at a certain point in the file and to continue from there, or to read in one period of data and not advance.

5.7 NOCS software features

Both the DTP and NOCS processor run *UNIX** operating systems and all the application code in both systems is implemented in the C programming language.[4] This paper will not attempt to go into detail on the software architecture of the NOCS. Rather, a few salient points of the implementation will be discussed briefly because of the part that they play in making the NOCS a successful real-time operations support system.

DTP software takes the data from the EADAS/NM lines and "funnels" it into one data stream for internal processing. Most of the messages are destined for the NOCS processor, so the DTP "blocks" as many as possible of the short EADAS/NM messages together into larger messages for efficient transfer through the DTP and to the NOCS.

Figure 6 shows the primary subsystems in the NOCS processor. All data coming in from the DTP or the E2A and all data going out to the DTP must pass through the Message Processing System (MPS). The MPS sorts incoming data into two basic categories: reference data and real-time data. It passes the reference data on to the Reference Data Update System (RDUS) to be processed as discussed in Section 5.3 while the MPS enters the real-time data into the database as they arrive. These processes also make certain data items immediately available to other processes so that the network managers can be alerted as fast as possible to potential problems. When all the data for a five-minute period are in, the MPS notifies all interested subsystems.

A prominent part of the NOCS software is the Data Access System (DAS), which handles all interactions with the database. The DAS is a large set of subroutines that provide common, standard access to the real-time and active reference database files. This single point of access to the system database has many benefits; the two most prominent ones are:

1. Standardized, centrally maintained retrievals that can be used by a multitude of processes

---

* Trademark of Bell Laboratories.

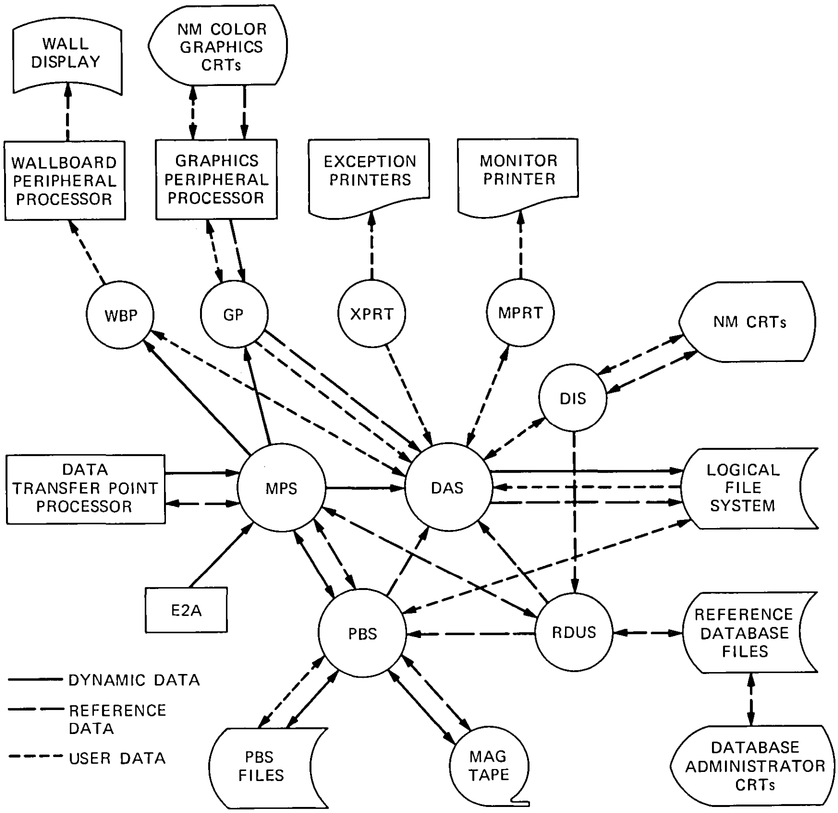

DAS - DATA ACCESS SYSTEM
DIS - DATA INQUIRY SYSTEM
GP - GRAPHICS PROCESS
MPS - MESSAGE PROCESSING SYSTEM
MPRT - MONITOR PRINT PROCESS
NM - NETWORK MANAGER
PBS - PLAYBACK SYSTEM
RDUS - REFERENCE DATA UPDATE SYSTEM
WBP - WALLBOARD PROCESS
XPRT - EXCEPTION PRINT PROCESS

Fig. 6—Major NOCS software systems.

2. Control of the reading and writing of individual files by a multitude of processes.

In addition to these benefits, the DAS also provides relief on system disk I/O through a per-process cache buffering mechanism that allows frequently accessed disk sectors to be kept in core, thereby avoiding additional disk reads.

In the NOCS processor the real-time data and active reference database are not kept in the standard *UNIX* file system but rather are maintained in a special file system called the Logical File System (LFS). The LFS is implemented using the raw I/O facility in the *UNIX* operating system. It provides a fast excess file system that:

1. Isolates the applications programs from having to know the physical disk location of files.

2. Enables application programs to create, delete, read, write, switch, and copy logical files.

3. Contiguously stores LFS files. This results in reduced disk I/O transfer time.

4. Transfers LFS files directly between the disk and the application data area, thereby reducing system overhead on each disk transfer.

The portion of the grahics subsystem that resides in the GP (Graphics Processor) runs under a specially designed operating system. The system provides functions required by the applications processes that drive the graphics displays. Most of the code that resides in the GP is implemented in C language; the remainder, though, is implemented in assembly language or is microcoded. This non-C code exists only in the operating system and is used to achieve the speed required to produce the interactive displays.

VI. CONCLUDING REMARKS

The NOC, supported by NOCS, plays a key role in coordinating and managing the North American telecommunications network. Its role is evolving from coordination of different geographical areas of the MTS network toward coordination and management of different subnetworks and service capabilities of the telecommunications network. We discussed the North American message network, the CCIS network, and overseas trunking. Another function of the NOC, not yet supported by NOCS, is monitoring and control of System 800, as noted in Ref. 2. As the network evolves with topological changes, new approaches to common channel signaling, and new service capabilities, we expect coordination and management of new service capabilities and the interactions of different network components to become an increasingly important part of the NOC's job.

REFERENCES

1. E. Oberer, "Role of the AT&T Network Operations Center," Proc. Comm. Techniques Seminar, Princeton University, March 1980, pp. 1-1 through 1-6.
2. G. C. Ebner and D. G. Haenschke, "Total Network Data System: Network Management," B.S.T.J., this issue.
3. A. E. Ritchie and J. Z. Menard, "Common Channel Interoffice Signaling: An Overview," B.S.T.J., *57,* No. 2 (February 1978), pp. 221–4.
4. H. Cohen and J. C. Kaufeld, Jr., "The Network Operations Center System," B.S.T.J., *57,* No. 6, Part 2 (July-August 1978), pp. 2289–304.
5. K. E. Greisen, "Keeping the NOC Informed—The Computer System That Provides the Data," Proc. Nat. Elec. Conf., *34* (October 1980), pp. 382–6.

AUTHORS

Robert W. Patterson, B.S.E.E., 1976, University of Tennessee; M.S.E.E., 1978, M.I.T.; Bell Laboratories, 1976—. At Bell Laboratories, Mr. Patterson

initially worked on network management surveillance and control capabilities for electronic switching systems and new services. He currently supervises the systems engineering for network management operations centers and support systems. Member, IEEE.

William S. Bartz, B.S.E.E., 1970, Manhattan College; M.S.E.E., 1971, Cornell University; Bell Laboratories, 1970—. From 1970 to 1978 Mr. Bartz worked in the Toll Crossbar Systems Department on the Peripheral Bus Computer system development. Since 1978 he has supervised the Network Operations Center System group in the Switching Operations Systems laboratory.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Equipment Systems

By N. D. FULTON,* J. J. GALIARDI,* E. J. PASTERNAK,*
S. A. SCHULMAN,* and H. E. VOIGT*

(Manuscript received March 10, 1983)

This article describes a set of five systems collectively called the Total Network Data System Equipment Systems, which support a number of network administration and engineering functions in the operating telephone companies. These systems are run by the operating telephone companies on large mainframe computers. In an average telephone company they support the management of the collection and distribution of over 50 million individual items of network data each week. They also provide reports to assist in the engineering and administration of No. 5 Crossbar switching offices, and the load balancing of all local switching systems. Initial development of these systems occurred in the early to mid-1970s. They are installed in all Bell System operating telephone companies and Bell Canada. Currently, on-line record base updates and inquiries and on-line viewing of report outputs are being developed.

I. INTRODUCTION

1.1 Role of TNDS/EQ in support of network functions

Basic to the operation of the telephone business is the collection, processing, distribution, and analysis of traffic data. The usage and availability for service of each of the many millions of components in

* Bell Laboratories.

the telephone system are measured by evaluating these data. In a typical week, the average Operating Telephone Company (OTC) collects over 50 million separate traffic measurements. Aided by the reports formulated from these data, telephone company personnel monitor how well customers are being served, place work orders as needed to optimize usage of currently installed equipment, and formulate traffic trends to forecast future needs. Almost all major decisions in the managing of the telephone network are influenced by the conclusions derived from traffic data.

The group of systems referred to as the TNDS Equipment (TNDS/EQ) systems primarily support data administration and central office switching engineering and administration functions in network organizations. Included in data administration is the collection, verification, and delivery of traffic data, both to the network and to other users. (As an example of the latter, local/toll call separations data are delivered to Division of Revenues.) One function of central office switching engineering and administration is to monitor usage patterns of the many individual traffic-sensitive components in a central office to ensure good customer service along with proper equipment utilization.

1.2 Organization of article

The remainder of Section I provides an overview of TNDS/EQ and discusses some of the history and decisions involved in its evolution. Section II describes the data collection environment, which serves as the input to the TNDS/EQ systems. Section III discusses the computer operational environment under which these systems run, and describes some common controls that were initiated to improve this environment. Sections IV through X describe in detail each of the TNDS/EQ components, and Section XI indicates some of the future directions for TNDS/EQ.

1.3 Overview of TNDS/EQ

TNDS/EQ comprises five component systems (see Fig. 1) and two special facilities. As discussed in Section 1.4, planning for these systems goes back to the mid-1960s with most initial development occurring in the early to mid-1970s. They are currently used throughout all Bell System operating telephone companies and Bell Canada. TNDS/EQ operates systems on local IBM 370 or equivalent mainframe computers, and are run weekly in a batch processing mode. Comprising over one-half million instructions, the programs have been written in a combination of COBOL, FORTRAN, PL/1, and Basic Assembly Language. Products delivered to the operating companies include executable load modules, IBM job control language, installa-

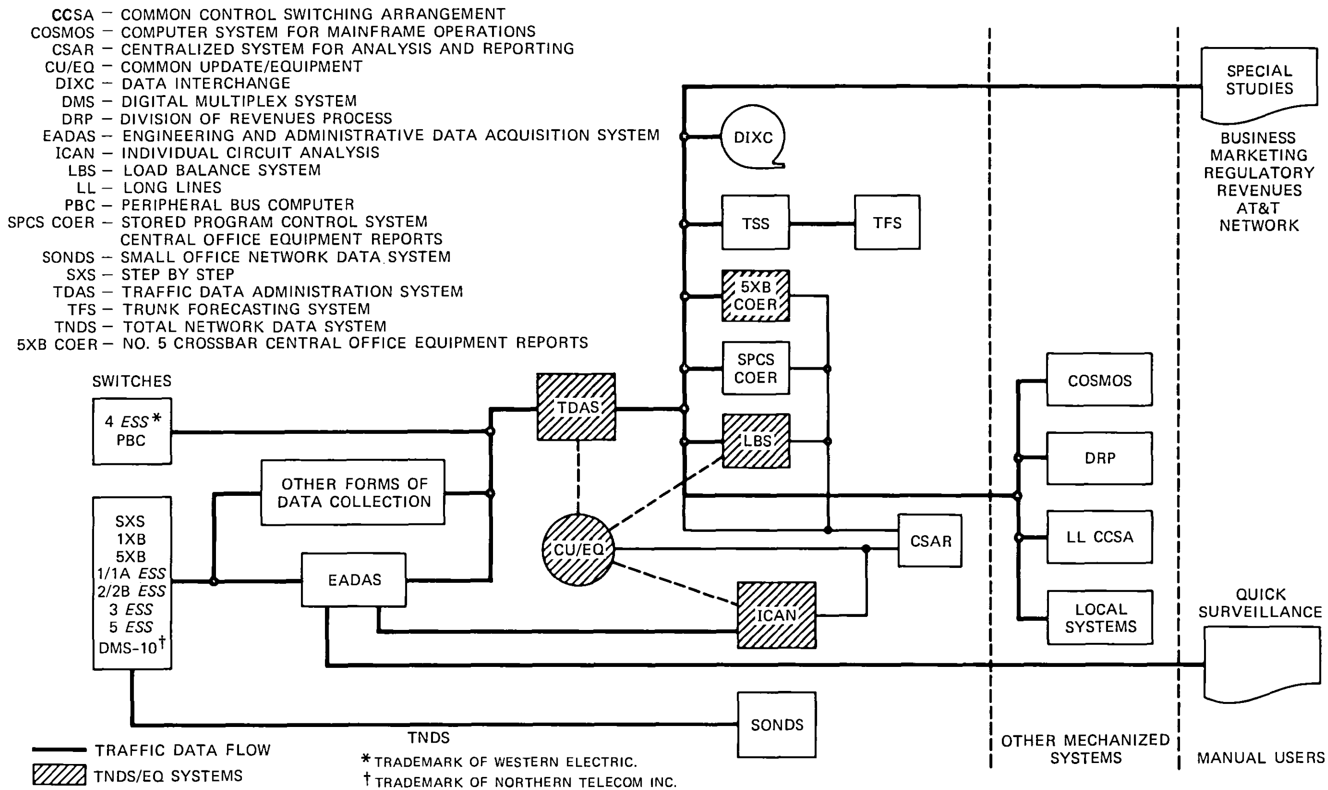

Fig. 1—TNDS data processing.

tion and conversion instructions, and documentation to support both computer center operations personnel and network users.

Three other systems are sometimes considered part of TNDS/EQ—the Stored Program Control System Central Office Equipment Reports System (SPCS COER),[1] the Centralized System for Analysis and Reporting (CSAR),[2] and the Small Office Network Data System (SONDS).[3] Because of their different operating environment—running on a single, time-shared computer for all using companies, rather than being deployed to each operating telephone company—they are presented in other articles in this issue.

One of the TNDS/EQ components, Common Update (CU), maintains master files used by a number of the other TNDS component systems (see Section IV). These files contain the basic records needed to process the data. As an example, the CU master files contain information that defines what data are to be transmitted to each of the downstream systems.

The Traffic Data Administration System (TDAS) is the key TNDS element serving the data administration functions in the operating companies (see Section V). TDAS accepts traffic data from most data collection sources, including paper tape, punched card, and magnetic tape. Among the sources of data are the Western Electric Engineering and Administrative Data Acquisition System (EADAS),[4] several non-Bell System data acquisition systems, and those switching systems, such as the Bell System 4 *ESS** switching equipment, which internally summarize traffic data. Using CU files, TDAS identifies, verifies, and labels each individual incoming data item, saving those items wanted for further processing and discarding those that are not wanted. On a weekly basis, scheduled data are transmitted to the requesting downstream systems. These output (downstream) interfaces are unlimited, as TDAS can effectively fill any order for data provided the data have been collected and delivered to it.

Included among the downstream systems are two component systems of TNDS/EQ: the Load Balance System [(LBS), see Section VIII], and the No. 5 Crossbar Central Office Equipment Reports System [(5XB COER, see Section VI]. LBS reports assist the network administrator in assigning customer lines to the switching machines in a manner that balances the traffic load on the switch. It also provides telephone company management personnel with quantified measures (known as the Load Balance Index) on how well the office balance is being maintained. 5XB COER reports on weekly, monthly, and annual traffic load for each individual component in a No. 5 Crossbar switching office.

---

* Trademark of Western Electric.

Another component of TNDS/EQ, the Individual Circuit Analysis System (ICAN), serves two basic functions (see Section VII). First, it provides reports to help administer the record base for the Individual Circuit Usage Recording (ICUR) option of EADAS; second, it analyzes individual component usage data to identify unusual measurements or patterns that could be indicative of central office or trunk maintenance problems.

In addition to these five component systems, TNDS/EQ has two special facilities, the Report Distributor (Section IX) and the Management Reporting System [(MRS) see Section X]. The former helps to deliver outputs to the network users and produces both paper and microfiche output. The MRS enables an operating company to access selected portions of the record base and output files to produce its own specialized reports.

1.4 History and major decisions

Elements of TNDS/EQ originated in several different development organizations within Bell Laboratories, AT&T and the Operating Telephone Companies. Most of the components were developed prior to the overall concept of a coordinated Total Network Data System. As a result, many activities over the past decade have been directed at integrating the systems. Some of this history has been addressed in an earlier article.[5] Even though there originally was not a complete plan, the TNDS architecture quickly and naturally fell into an "upstream" and "downstream" split. Placed upstream were the functions associated with data collection and real-time exception reporting, activities suitably accomplished with minicomputer technology. On the other hand, the heavy processing functions involved with crunching tens to hundreds of millions of weekly data items into usable engineering and administrative reports readily fell into a downstream processing environment of large mainframe batch computers.

The history of downstream TNDS components goes back to 1966 (see discussion in Ref. 5), with a decision to centralize the development of a computerized Trunk Facilities System. This decision led to a major development program in the Business Information Systems (BIS) area of Bell Laboratories. Included in the program were three systems that became part of TNDS—The Trunk Servicing System (TSS), The Trunk Forecasting System (TFS), and TDAS. The former two are now part of the Trunking System (TNDS/TK).[6] These three systems, together with Common Update as a supporting record base system, became known as Servicing and Estimating, under Bell Laboratories BIS development. Around the time Servicing and Estimating got its start, a stand-alone reporting system for No. 5 Crossbar central office data was developed by New Jersey Bell. It was standardized and integrated into the TNDS downstream structure by AT&T and a mechanized interface to TDAS was implemented in 1972.

One of the earliest and perhaps most important architectural decisions was the placement of TDAS as a single point of entry for all data to the "downstream" world. In this position, TDAS served as a data warehouse—receiving data on its input side from many different data collection processes and distributing data on its output side to the many different users (both individuals as well as other systems). This basic architecture, established when most data collection was manually transposed and keypunched from Traffic Usage Recorder film outputs, has been extremely robust under the stress of rapidly changing data collection methods and data processing needs. On the output side, new downstream processing systems have been able to satisfy their needs for traffic data by simply placing orders, known as Traffic Measurement Requests (TMRs), on the TDAS data warehouse. And on the input side, new collection systems have been able to deliver their data by defining interfaces to TDAS.

In the mid-1970s, ICAN and LBS were developed and added to the TNDS downstream systems. ICAN was developed by the Network Planning area of Bell Laboratories, while LBS was developed by the BTL Servicing and Estimating organization. Shortly after ICAN was developed, all TNDS downstream systems (Equipment and Trunking) migrated to the Bell Laboratories BIS organization. By 1977 5XB COER, ICAN, and CSAR (which had also been developed by Bell Laboratories Network Planning) had been moved. Around this time, because of the growing size of the development organizations, the downstream systems were split into two BIS departments, along the lines of network user communities (trunking and equipment), to form TNDS/TK and TNDS/EQ as separate system groupings.

Combining all the downstream TNDS systems (TNDS/EQ and TNDS/TK) into one organization enabled common standards and features to be established including standard run control procedures (discussed in Section III), and a report distribution facility (discussed in Section IX).

With the expansion in the mid-1970s of both the number of systems as well as the number of operating telephone company installations, it became necessary to provide strong controls over the release of program modifications—both fixes and new features. Coordination was necessary from many standpoints. For developers, the incorporation of a new feature often required simultaneous changes to several systems. For system maintainers, keeping track of each of 20 operating company configurations of approximately 200 program modules required mechanized aids. And for each operating company, where hundreds of individuals depended on TNDS/EQ outputs to carry out their daily jobs, there was a need to plan for new features and to be kept informed of the status of fixes to problems.

To cope with these needs that are common to large systems developments and large user communities, three major changes were instituted:

1. A coordinated annual release was established for all TNDS/EQ systems for packaging of new features. The scheduling of the release was set to occur during periods of slack telephone traffic so that any disruptions resulting from installing the new release would not interfere with processing busy season data.

2. A development environment was established based on the *UNIX** operating system, utilizing the Source Code Control System (SCCS) and Modification Request Control System (MRCS). These provided management control over different program versions and provided status tracking of trouble reports (modification requests).

3. A functional organization was established with separate requirements, development, test, installation, and consultation groups. Along with this organization, standard procedures and schedules were established for internal development activities. Also, greater emphasis was placed on formal documentation, both to support the internal functional organization and to support telephone company personnel in their planning and training for the annual TNDS/EQ releases.

Starting in the late 1970s attention was directed to the need for modernization of TNDS/EQ, in line with technological advances that had occurred over the past decade. Of all the decisions made, this one was the most difficult. Over the years, the major design focus for TNDS/EQ had been to keep up with network changes, to improve the efficiency of its heavy-duty data processing activities, and to streamline the flow of data from data collection systems, through TDAS to the downstream processing programs. For those TNDS/EQ functions associated with this processing of traffic measurement data, batch processing remained appropriate, even with current state-of-the-art computers.

Batch operations were awkward in areas involving direct user interfaces: inputs to update the record base (Common Update), and deliveries of output reports. On the input side, Common Update's weekly processing cycle introduced long delays in making record base changes and verifying their correctness. When errors occurred, several weeks could elapse before they would be corrected. On the output side, since the batch process does not retain outputs, telephone company personnel would often obtain output volumes far exceeding their needs "just in case" some information might be needed at a later date. It is evident

* Trademark of Bell Laboratories.

that an on-line, interactive interface with the user would significantly improve the environment. But in spite of these opportunities for improving human interfaces, net benefits were difficult to quantify. On the one hand, costs for terminals, printers, telecommunications lines, and computer on-line processing are easily quantified and generally are expensive. On the other hand, savings gained in personnel time from improved input error processing and less paper shuffling are very subjective.

Several activities were conducted to obtain a better view of benefits. In 1980, time-motion studies were conducted at two operating companies to obtain a view on how people, particularly clerical personnel, spent their time. The results of these two studies were consistent and showed that there were significant opportunities for savings through modernization. As a result of these studies, an experimental on-line system was installed in the second half of 1981 at a Chesapeake & Potomac Telephone Company network location and remained in place for about one year. Even though limited in scope, the on-line capabilities were received enthusiastically by the C & P Company personnel, and comparative time-motion studies conducted both before and during the experiment again showed that there were opportunities for savings. As a final activity, each Bell Operating Company (BOC) was requested to conduct its own cost-benefit study to determine whether on-line features would be economically beneficial. These studies were conducted during the summer of 1982 and resulted in the decision to proceed with an on-line system for record base inputs and output report viewing. This system has been named On-line Records and Reporting System (ORRS). The current plan is to provide these capabilities over the 1983 to 1986 time period, with first operating company application occurring in 1983.

II. DATA COLLECTION

Because TNDS/EQ (specifically TDAS) needs to interface with multiple sources of data, a short discussion on data collection is presented here. Methods of data collection as well as types of data collected have changed greatly over the years. TNDS/EQ serves the vital function in the overall TNDS design of translating these multiple inputs to a common output format and identifying the data in terms of Bell System Common Language. The full list of measurement types supported by TNDS/EQ is given in Table I and were discussed in an earlier article.[7]

There was a need to collect telephone traffic data well before there were mechanized systems for collecting the data. In the early 1900s operators who manually completed calls recorded data on the frequency of calls to specific locations. This information was used for

Table I—TNDS/EQ measurement types

Measurement Type	Description	Switching Machines Applicable
PC	Peg count	All
USG	Usage	All
OVF	Overflow	All
MTU	Maintenance usage	1E, 2E, 5XB, 1XB, XBT
INU	Incoming usage	1 *ESS*, 4A, 2 *ESS*
ATB	All trunks busy	SXS
LTB	Last trunk busy	SXS
CMP	Completions	SXS, 5XB, XBT, 1XB
IPC	Incoming peg count	4E, DMS
MBC	Maintenance busy count	2E, 3E
DGU	Detector group usage	SXS, 5XB, XBT, 1XB
RTS	Reroute to seizure	4E

establishing new routes between cities and resizing existing trunk groups. With the development of electromechanical switching systems, the collection of counts was mechanized. Counts were accumulated on mechanical registers. Periodically, the register readings were recorded by clerical personnel and compared with earlier readings to obtain a difference count. The manual recording of traffic register values was later replaced by automatically photographing these registers at predefined intervals.

During the 1960s, the rapid introduction of computer technology brought about three major changes in the data collection process. First, general-purpose computers were used to calculate the register differences. Second, the camera-register method began to be replaced by centralized data collection hardware, designed to collect measurements directly from the switching machines and to record the counts on magnetic tape for later processing by general-purpose computers. The first system of this type was the Traffic Data Recording System (TDRS). This system, implemented as a hard-wired computer, was soon found to lack the flexibility and capacity of the rapidly evolving minicomputer technology. The third major change in the 1960s was the development of electronic switching machines, which could perform their own internal data collection.

During the 1970s, the collection of real-time data introduced new capabilities to support office maintenance and network management. Also during this period, various support systems were centrally developed to run on general-purpose computers, providing features and flexibility applicable to all companies. These centrally developed systems standardized the format and labeling of traffic data so operating telephone companies could conveniently interchange common data.

As a result of this evolutionary process, data are collected many different ways in the Bell System. Table II lists the 35 different collection types currently supported. One of the major ongoing activ-

Table II—Traffic data input formats

Data Collection Type	Description	Switching Machines Supported
R	Register readings	Electromechanical
D	Register differences	Electromechanical
S	Minicomputer scanned data	Electromechanical
A	Minicomputer accumulated data	Electromechanical
P	Minicomputer polled data	Electromechanical
I	Minicomputer individual circuit data	Electromechanical
X-A	4A trunking data	4A
X-D	4A central office data	4A
1-C	1/1A *ESS* switch continuous data	1/1A *ESS*
1-H	1/1A *ESS* switch hourly data	1/1A *ESS*
1-W	1/1A *ESS* switch weekly data	1/1A *ESS*
1-D	1/1A *ESS* switch daily data	1/1A *ESS*
1-T	1/1A *ESS* switch traffic separations data	1/1A *ESS*
1-E	1/1A *ESS* switch extreme value data	1/1A *ESS*
2-C	2/2B *ESS* switch continuous data	2/2B
2-H	2/2B *ESS* switch hourly data	2/2B
2-W	2/2B *ESS* switch weekly data	2/2B
2-D	2/2B *ESS* switch daily data	2/2B
2-R	2B *ESS* switch record verification data	2B
2-E	2B *ESS* switch extreme value data	2B
3-C	3 *ESS* switch continuous data	3 *ESS*
3-H	3 *ESS* switch hourly data	3 *ESS*
3-W	3 *ESS* switch weekly data	3 *ESS*
3-D	3 *ESS* switch daily data	3 *ESS*
4-5	4 *ESS* switch trunking data	4 *ESS*
5-C	5 *ESS* switch continuous data	5 *ESS*
5-H	5 *ESS* switch hourly data	5 *ESS*
5-D	5 *ESS* switch daily data	5 *ESS*
5-R	5 *ESS* switch record verification data	5 *ESS*
5-E	5 *ESS* switch extreme value data	5 *ESS*
M-C	DMS-10 continuous data	DMS-10*
M-H	DMS-10 hourly data	DMS-10
M-D	DMS-10 daily data	DMS-10
M-E	DMS-10 extreme value data	DMS-10
J-C	DMS-200 continuous data	DMS-200

* Trademark of Northern Telecom Inc.

ities of TNDS/EQ is to keep current with the collection environment as it continues to evolve.

III. OPERATIONAL ENVIRONMENT AND CONTROLS

3.1 Introduction

When TNDS/EQ and TNDS/TK were introduced in the early 1970s, the OTC computer operating environment was frequently not prepared to handle them. For some companies the TNDS systems were the first centrally developed BIS product to be installed. Interfacing with Bell Laboratories for reporting troubles, transmitting diagnostic data, and receiving corrections represented new concepts and procedures. Of even greater difficulty, each operating company had its own local rules for documentation and computer job operations. Conventions for computer job termination, job completion, and error detection, recovery, and restart varied. In some cases local procedures required extensive manual verifications to ensure proper program job execution. From these early dissimilarities, a number of activities emerged. Working with the central development organizations and the operating telephone companies, the AT&T Information Systems organization established basic rules. And within the TNDS project, a very important activity was the establishment of "run control" standards for the TNDS/EQ and TNDS/TK systems. These controls were designed to detect error conditions as they occur, inform operations personnel of their existence, and inhibit further processing until the problem can be corrected. Once the problem is corrected, these controls permit simple restart with minimum reprocessing. Run control has been implemented in all TNDS/EQ systems and has been effective in minimizing the overhead associated with operating centrally developed software in remote batch production environments. The remainder of Section III provides details on the interface of run control with the operating environment.

3.2 Run control

The attributes of the TNDS run control standard are described in this section.

3.2.1 Load module execution sequence control

To carry out a specific program function, program modules (known as load modules) that are delivered by Bell Laboratories to the processing sites are grouped into larger entities, called jobs. The execution sequence of jobs and of the load modules within jobs usually follows a predefined arrangement. Communications between programs within a job and between separate jobs are carried out by use of intermediate files stored on tape or disk. If through operational error, a defined job sequence is violated, run control detects and inhibits further processing until the proper sequence is resumed. Before run control was used, the jobs would have continued either to completion or until an error was detected. Because of the delay in detecting the trouble, location of its source and subsequent recovery was often difficult.

3.2.2 Load module termination conventions

When a load module termination occurs, either because it has completed its task or because of an abnormal condition, it is necessary to identify the reason for its termination. This information is needed both internally by other load modules executing in the same job and externally by the computer operations staff. The run control standard specifies use of a mechanism called the condition code to indicate normal load module termination and an "abort end" or abend code mechanism to indicate abnormal load module termination. The operations staff are thereby notified of abnormal terminations in a consistent way and job processing is stopped whenever a load module terminates abnormally. Similar to the previous attribute, early detection of a trouble condition is therefore improved.

3.2.3 Restart aids

A third aspect of run control ensures that restarts after a failure are handled as simply as possible. To help achieve this, run control:

1. Automatically selects the restart load module among the set of load modules within a job and alerts operations staff in the event of an error.
2. Provides a standardized interface to the operating system "checkpoint/restart" facility, which under certain operating conditions restarts the job at an intermediate point in the processing.
3. Automatically removes those intermediate files that must be recreated as part of the restart.
4. Automatically restores the contents of update files to their pre-execution contents before beginning the restart.

Figure 2 illustrates these concepts by showing a hypothetical job sequence and listing the items that must be performed to complete a restart.

3.2.4 Logging system processing activities

Run control also specifies standards for logging information about the processing activities of the system. This information is a summary of processing activities and has been very helpful later in diagnosing problems. The following information is logged for this purpose:

1. Indication of every module executed, recorded by date, time, module name, and unique version identification.
2. Indication of when the execution of the load module was completed, and whether such completion was normal or abnormal.
3. Identification of every file processed, recorded by its name and unique version identification. In addition, for sequential files, the number of records processed is recorded.
4. Recording of any anomalous conditions encountered by the program.
5. Recording of the Central Processing Unit (CPU) time and wall clock time needed to execute the load module.
6. Changes in value of any run control information.

The log can be readily transmitted back to the central development facility for detailed analysis.

3.2.5 Automatic verification of file versions and sequential file record counts

The fifth attribute of run control has improved the integrity of the

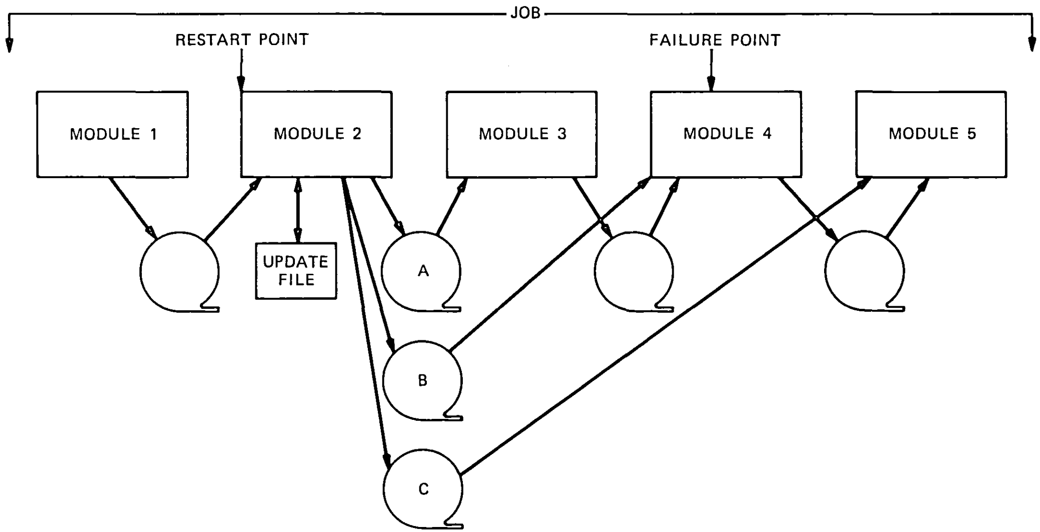

STATUS: SUCCESSFUL PROCESSING THROUGH MODULE 4, WHICH IS ABNORMALLY TERMINATED.

TO RESTART A JOB AFTER FAILURE:
1. DETERMINE RESTART POINT, SAY MODULE 2.
2. RESTORE UPDATE FILE TO PRERUN CONTENTS.
3. REMOVE INTERMEDIATE OUTPUT FILES (A, B, C) BEFORE INITIATING RESTART.
4. SPECIFY JOB RESUMPTION AT THE BEGINNING OF MODULE 2.

Fig. 2—Restart after failure.

file system. During a production cycle of TNDS/EQ, permanent files are modified to reflect current information. In addition to retaining the updated file, the operating system also retains the version of the file as it existed before the cycle so that a rerun can be performed if needed. The operating system provides a mechanism for using a family of files [called Generation Data Group (GDG)] to maintain information across multiple production cycles. The input to a program would be generation zero (0) and the output would be generation plus one (+1). At the completion of execution, the system records the new output as (0) and the original input as (−1). Thus generations are "rolled" from run to run. Figure 3 illustrates the use of GDGs.

Unfortunately, referencing a specific member of a GDG family requires manual intervention and can lead to accessing an incorrect member of the family. Run control detects such situations, alerts the OTC computer operations staff, and halts further processing pending correction of the problem.

Another file problem that can occur, particularly with the processing of sequential files, is loss of some of the records. A common way for this to occur is through inadvertent deletion of one of the middle reels of a multireel tape file. Run control detects this condition by maintaining knowledge of the number of records placed on the file at creation time versus the number actually processed when the file was accessed. Prior to run control, this verification was often done manually by computer operations personnel matching file counts.

3.3 Run control example

The following example illustrates how the parts of run control work together to ensure correct processing. Consider a job consisting of four load modules (LM1, LM2, LM3, and LM4). Each module creates one intermediate output file that is passed to the next load module in

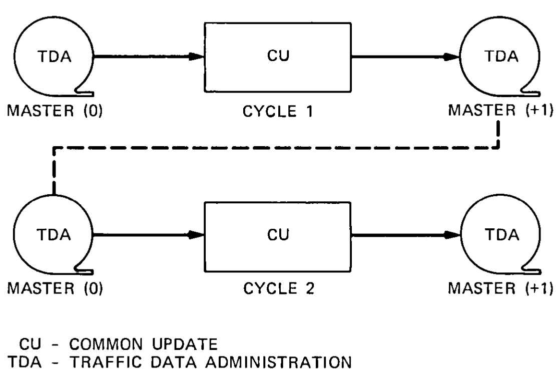

Fig. 3—Use of generation data groups.

sequence. Table III provides the contents of the run control external file. This file indicates that the current production cycle has been stopped, with only LM1 having completed execution. In this state, it is possible that LM2 might have been started, but was stopped because of an error such as a program problem, a hardware difficulty, or an incorrect tape mount.

Given the control information in Table III, when the job is executed again (the problem having been corrected), the following events occur as part of the run control functions:

1. Run control restart module executes.
   (a) A log entry is made showing run control module name, date, time, and version identification (id).
   (b) The name of the program to be executed next (LM2) is obtained from the run control external file.
   (c) The intermediate work file associated with LM2 (F2) is erased. This information is obtained from the table entitled "Association of Load Modules to Files Created" (see Table IV).
   (d) A condition code is issued to cause module LM1 to be skipped by Job Control Language (JCL) processing.
   (e) A log entry is made showing the run control module name, date, time, CPU time, and normal completion.

2. Load Module LM1 is skipped by the operating system JCL processor.

3. Load Module LM2 executes.
   (a) A log entry is made showing module LM2 name, date, time, and version id.
   (b) Run control verifies that it is proper for LM2 to execute by comparing its name (LM2) as recorded internally within the module with the contents of the run control external file.
   (c) LM2 opens file F1, extracts header information, and compares

Table III—Contents of run control external file after abnormal termination of a job

Module	Execution State	File	Date Last Executed	Time Last Executed	No. of Records
LM1	Executed	F1	8-4-81	13:00	1000
LM2	Not executed	F2	8-3-81	09:30	2500
LM3	Not executed	F3	8-3-81	09:45	5000
LM4	Not executed	F4	8-3-81	09:50	200

Table IV—Association of load modules to files created

Module	LM1	LM2	LM3	LM4
Files Created	F1	F2	F3	F4

it against the audit values in the run control external file recorded when file F1 was created. Let's assume that an earlier version of F1 was mounted and LM2 reads from the file header "8-3-81 9:00." This does not match the time and date logged in the run control external file "8-4-81 13:00." The error condition, file name, and audit values are logged.

(d) Execution of LM2 is terminated by an abend.

4. The operating system JCL processor flushes the job.

This scenario illustrates run control processing under an error condition. If file F1 had not failed its audit, then processing would have continued with modules LM3 and LM4 and actions similar to those recorded for LM2 (without the abend) would have occurred.

Program recovery prior to implementation of the controls typified by this example was often difficult, requiring extensive manual intervention. Run control software has provided dual benefits to the TNDS user community: first, it has reduced the level of direct involvement needed for centralized maintenance personnel to provide operational support; and second, it has made available better diagnostic information for handling problem situations to both the OTC operations personnel and to the central maintenance personnel.

IV. COMMON UPDATE (CU)

4.1 Introduction

Section IV describes the details of the data processing functions carried out by TNDS/EQ components. Because Common Update (CU) provides record base support to a number of the other components, it is described first. It is an example of the integration that has been achieved within the TNDS/EQ and TNDS/TK systems. In particular, CU fully supports the TNDS/EQ systems TDAS and LBS, as well as the Report Distributor facility, and provides partial support of No. 5 Crossbar Central Office Equipment Reports (5XB COER) and ICAN. It also fully supports the TNDS trunking systems TSS and TFS.

Since CU supports both switching equipment and trunking support centers, it has been subdivided into independent subsystems, CU/Equipment and CU/Trunking. This division allows an operating company to have the flexibility of providing multiple installations of TNDS Equipment (TNDS/EQ) systems, including CU/EQ, while maintaining a centralized TNDS Trunking (TNDS/TK) system, including CU/TK. The remainder of this section will focus on CU/EQ, hereafter referred to simply as CU. The article on trunking systems (Ref. 6) discusses CU/TK.

To support the Network Administration personnel who maintain the record base, CU:

1. Validates input record base transactions according to predetermined parameters and ranges, and performs intrarecord and interrecord validations.

2. Outputs activity or addenda reports after each CU cycle to provide updated record base status.

3. Provides messages and reports on transactions found in error.

4. Responds to input requests for full record base listings of a given report type.

CU executes as a series of jobs or runs. It is designed to run on an "as needed" basis, i.e., whenever transactions are created to add, change, or delete information in the record base. In most operating companies, CU executes on a weekly basis with all other dependent TNDS systems running subsequent to its successful completion.

The following four support run groupings are associated with the mainline runs of CU:

1. New TNDS/EQ installation—One-time runs needed during an initial installation of the system at a new site.

2. Release conversion—One-time runs to convert databases resulting from implementation of new TNDS/EQ features in a major new release.

3. Switching machine installations or conversions—Runs to provide CU record base support of changes in the switching or data collection environment. These include automatic database generation and conversions associated with new switching machine generics.

4. Recovery and restart—Runs that allow the database files to be restored to their original content in the event of an abnormal termination.

4.2 System inputs

CU activity is driven by batch input transactions. Each transaction is identified by a unique three-digit number, referred to as the transaction code or type. CU supports 50 different transactions, grouped into series by the TNDS system that they support. For example, transactions in the 700 series (742, 743, 750, 751, $\cdots$) support TDAS, while the 600 series transactions (600, 601, 620, $\cdots$) support LBS. Common information, which is used by multiple systems, is stored in files called tables and their transaction series number is 100 (e.g., 100, 105, 180, 181, $\cdots$).

To accommodate the current CU batch inputs many of the operating companies have developed their own front-end input processors. These systems typically are minicomputers or intelligent terminal configurations that provide screen formats, perform on-line intrarecord validations, and allow geographical dispersion of input sources.

As a future CU replacement, the On-Line Records and Reporting

System (ORRS) discussed in Section 1.4 is being developed. It will have all the benefits of the current front-end systems, as well as immediate interrecord validations, on-line record base inquiry, and on-line report viewing.

4.3 System outputs

4.3.1 General

There are two categories of outputs from CU, record base files and reports. The files support the other TNDS systems, which rely on CU for their record base. The reports support the network administrator's job of maintaining the record base so that it accurately reflects the physical telephone network.

4.3.2 Record base files

User data are entered into CU to establish an up-to-date description of the network. Information relevant to a particular processing system is isolated into one or more record base files. For example, the information that describes the data collection environment needed by the TDAS system is stored in the Traffic Data Administration master file, while the information that describes a switching entity for load balance purposes is stored in the LBS master file.

As updates are introduced to the system in the form of add, change, and delete transactions, the current version of the master file is updated to produce a new version of the file. Generation Data Groups (GDGs) are used to help simplify the updating of these files and the maintenance of adequate backups, as described in Section III.

4.3.3 Reports

CU provides a number of reports to help maintain the record base. The two main categories are error reports and reference listings.

Error reports are issued whenever CU determines that an input transaction is invalid. In addition to these reports of specific errors, other reports provide record base status and error statistics. A number of error conditions do not result in complete rejection of the transaction. Rather, to reduce subsequent data entry, the data are posted in the record base, and an error indicator is turned on. Record base status reports inform the user after each run of all records that remain invalid and require correction. Operating telephone company managers receive error statistics by transaction type and origination, which enables them to analyze error rates and spot training deficiencies. Errors that predominate in a particular geographic region of the company can thereby be recognized and corrected with additional training. Error statistics also enable managers to recognize deficiencies in system documentation.

Reference listings provide a readout of the record base. These reports take the form of addenda and full listings. Depending on the particular report, they can be generated in various formats. For example, the Data Collection Environment report can be sorted and printed three different ways, depending on the needs of the particular user.

As we described in Section IX, CU reports are directed to the report distributor for printing or microfiching and distribution.

4.4 CU processing flow

The processing flow of CU is conceptually quite simple. First, inputs are sorted by transaction type. Invalid transaction numbers are directed to an error file. The sorted transactions are then validated, starting with the 100 series transactions.

Following basic intrarecord validations, the transactions are segregated by series and directed to files for subsequent processing by individual master file update modules.

These modules update the following master files:

- Circuit group master and translation files (CU/TK module)
- Load balance master files
- Traffic unit master file (CU/TK module)
- Traffic data administration and traffic measurement requests master files
- Common access tables.

Errors detected during the update process are directed to error files. Activity and demand requests result in master file records being directed to report files. After all the files have been updated, report and error records are sorted and formatted into standard output reports.

The final step in the CU mainline process is to generate backup tapes for CU table files. These tables are all random access files updated in place in response to their corresponding input transactions. Since destruction or loss of any of these tables would be very detrimental to the entire TNDS downstream operation, the files are automatically copied during the main cycle of the system.

V. TRAFFIC DATA ADMINISTRATION SYSTEM

5.1 Introduction

As shown in Fig. 1, the Traffic Data Administration System (TDAS) is the first of the TNDS/EQ systems to process traffic data. The data are received from upstream collection systems and can be processed immediately, but typically they are processed in batches in weekly cycles. The traffic data shown in Fig. 1 are input to TDAS from a number of different sources. Among the "Other Forms of Data Collection" are vendor-supplied data collection systems and keypunched data collected in a camera/register environment. In response to specific requests for data from the users of various downstream systems, TDAS formats the data as required by those systems. As we described previously, TDAS acts as a centralized warehouse and distribution facility for traffic data.

The primary objectives of TDAS are to:

1. Accept traffic data from any standard data collection source.
2. Edit, validate, and adjust data.
3. Identify the origin of data, i.e., what the measurements represent.
4. Summarize and, if necessary, store the data, to satisfy weekly requests.
5. Transform the data into standard formats that are acceptable to the downstream systems.
6. Provide traffic data reports that primarily satisfy special study requests.

Although Network Administration is the major direct user of TNDS-collected traffic data, there are many other users. Each has different needs and, therefore, may require different subsets of the data in different physical forms and formats. The output reports of TDAS range from machine-readable files for mechanized interfaces to paper and microfiche reports.

Within the total data provisioning process, TDAS is positioned to control the data flow from the collection sources to the end users.

The remainder of this section describes how TDAS functions and its outputs, inputs, and processing flow.

5.2 TDAS outputs

5.2.1 General

TDAS supplies data to a large audience of traffic data users. Some of these are supported by computer systems with which TDAS interfaces via magnetic tape files, while others use reports created directly by TDAS. Thus, the major outputs of the system are interface files and reports.

5.2.2 Interface files

A standard interface file has been defined for each distinct system with which TDAS interfaces. This definition contains file characteristics and record formats. Associated with the record format is a detailed definition of every field within the record and its data attributes. File characteristics include definitions such as file format and block size. As user requests are processed, TDAS directs the appropriate data to the corresponding physical file, which is typically a tape but can be any IBM-compatible secondary storage device. At the completion of a weekly TDAS run, the interface files are transported to their respective data processing sites for further processing by the downstream systems. TDAS has the continuing requirement to ensure that the output formats remain stable as future changes occur in the data collection process. Note that because each class of end user has a distinct and different interest in the data, the formats serving these end users are different. For example, the format of the Trunk Servicing System (TSS) interface file is different from the Load Balance System (LBS) interface file.

5.2.3 Reports

TDAS provides data summary reports for users of traffic data not served by a downstream data analysis system. For these outputs, TDAS can perform some formatting and processing functions on the data to assist the user. These include editing, validating, adjusting for improper sample (or scan) rate, summarizing, and identifying with the appropriate Equipment Measurement Code (EMC) or Trunk Group Serial Number (TGSN) designation.

In addition to reports oriented to end users, TDAS also produces a series of reports to aid in the data provisioning and tracking process. These include:

1. Data edit error reports, tailored to each data collection source,
2. Data log reports, which relate data requests to data availability, and
3. Input data tape processing reports, which account for the many input tapes processed by TDAS in a given week.

All of these reports are written to a report file for input to the Report Distributor (Section IX) for final distribution and printing (or microfiching).

5.3 TDAS inputs

5.3.1 Record base

To achieve its objectives, TDAS relies on CU for maintenance of its record base. The TDAS record base consists of two primary inputs—a definition of the data collection environment, called the Traffic Data Administration (TDA) master file, and a list of all user requests, called the Traffic Measurement Request (TMR) master file. Using these two basic inputs, together with traffic data itself, TDAS can treat the data provisioning job as a basic order inventory problem. Orders, or TMRs, for data are compared with the inventory of traffic data. Based on prescribed rules and the use of the TDA master file for identification purposes, the orders are filled by TDAS.

5.3.2 Traffic data

Concurrent with the evolution of the switching network has been the evolution of data collection hardware. As the means of collecting data have changed, so has the format of the data itself. For example, data that originates from the old camera/register environment, which still serves some electromechanical switches, are in the format of keypunched card input. Another old form of data collection is 1 *ESS* switch data in the form of paper tape. All other modes of data collection are in machine-readable magnetic tape form but in many different formats. *ESS* switch data differ in format from electromechanical data; 4 *ESS* switch data are different from all other *ESS* switch data. Even the same switch can have different data formats based on the software version of the switching machine. One of the important functions performed by TDAS is to mask these format differences and present the data to the end user as though they were collected in a common manner.

Based on its format, the traffic data are read into the appropriate TDAS data entry module for editing, sorting, and converting them into a common format for further processing. The following input data formats are currently defined by TDAS:

1. Keypunch data in support of camera/register data collection.
2. Paper tape data in support of 1 *ESS* switch paper tape.
3. ASCII-encoded data in support of older EADAS machines; the Peripheral Bus Computer (PBC), which collects 4A toll machine data; and various outside vendor collection machines.
4. 4 *ESS* switch formatted data.
5. Binary-coded data in support of current EADAS machines and various outside vendor collection machines.

The last format, the Binary Interface, is a flexible data format intended to standardize the data coming from current stored program control systems. This interface has been documented as a General Trade *Technical Advisory* and has been distributed by AT&T through the United States Independent Telephone Association (USITA). TDAS will process data collected by any data collection system meeting the interface.

5.3.3 System options

To satisfy the variability in data collection environments and other factors that make each operating telephone company unique, TDAS accepts a number of user options, which are stored in its parameter file. One such option is the expiration date offset value for each mode of data collection. It indicates the maximum number of days from time of collection that the system will wait for all data to be received for a given type of Data Collection Unit (DCU). If all data are not received by the calculated expiration date, the TMRs will be processed with whatever data are currently available from that collection machine. This option is particularly important for camera/register collection, since each data item for this type of DCU is received as a separate and independent input. On the other hand, data collected via mechanized methods, such as EADAS, are usually received on a single input tape.

5.4 TDAS processing flow

5.4.1 General

A complete execution of TDAS, called a cycle, is performed on a weekly basis. To simplify its operation and scheduling for the computer operations personnel, TDAS is divided into a number of discrete jobs.

The mainline TDAS programs (that is, those that are run every production cycle) are divided into three functional entities: CU/TDAS File Interface Process (FIPS), Data Entry, and Data Analysis. FIPS extracts the records required from CU and formats them for efficient TDAS processing. It is a multistep run and is executed once per cycle. Data Entry serves as the front-end process of TDAS. There are five Data Entry modules (hereafter called editors), one for each of the data entry formats, as listed in Section 5.3.2. Their frequency of execution depends on the data collection environment and computer scheduling of the individual OTC. Some OTCs run daily editors, while others batch the daily data tapes into a weekly editor run. TDAS allows both or any combination in between. The editors prepare the traffic data for the Data Analysis programs, which are executed once each production cycle. This overall TDAS processing flow is shown in Fig. 4.

The primary responsibility of TDAS is to process traffic data, and that it does—in excess of 50 million data items per week in an average size OTC. It is this magnitude of data that makes TDAS the longest running system in TNDS/EQ. Run time, therefore, has been a continuing area of concern.

One of the techniques that were introduced to improve processing times is data screening. Soon after TDAS was introduced it was recognized that much of the data collected was not needed by any end user. Of the 50 million data items received by TDAS, only about 25 percent are sent downstream—the rest are excess! The varying data requirements of the different end users may result in upstream data collection for almost 24 hours of every day, even though only a small portion is needed for the entire period. Figure 5 shows in a typical example the individual needs of the downstream systems that TDAS serves and the resulting collection of unrequested data. Because of the needs of one downstream system, it becomes necessary to collect the full increment of 1000 data items for 24 hours. The function of data

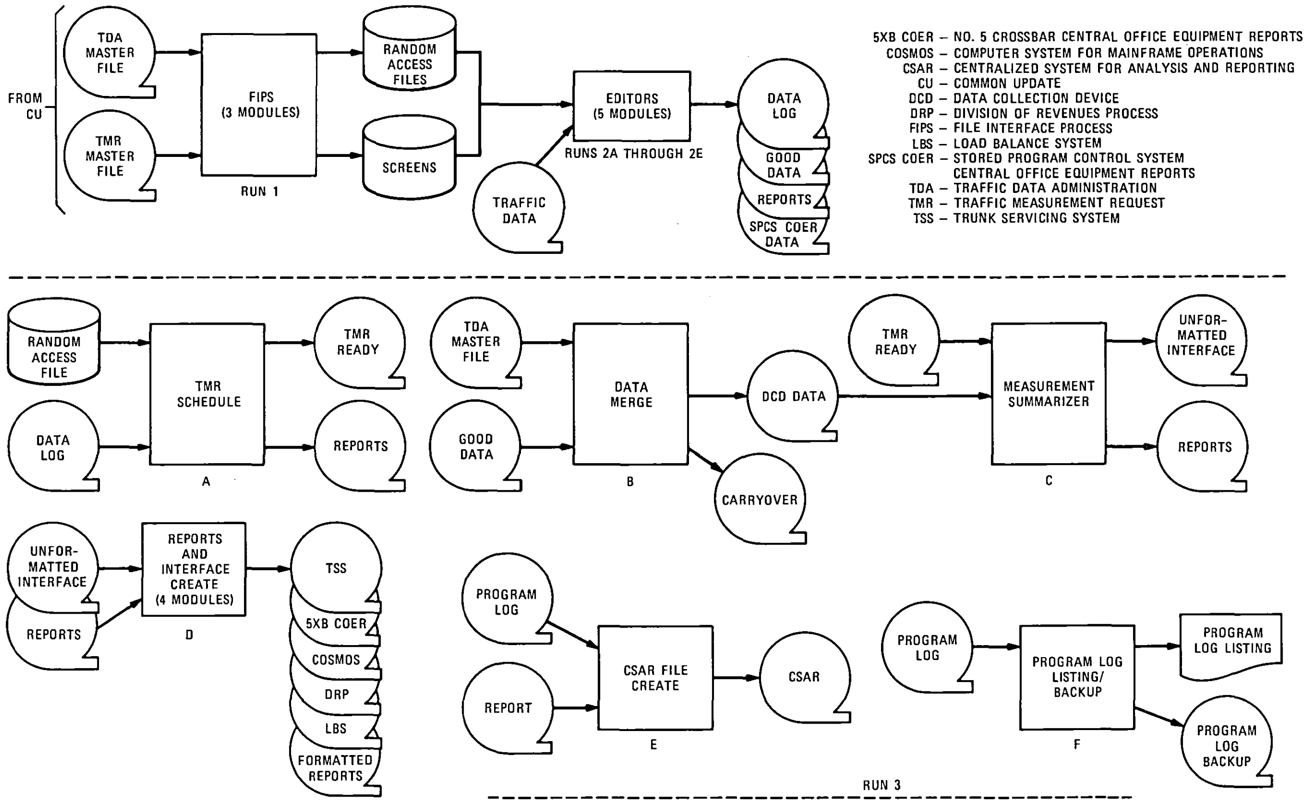

Fig. 4—TDAS processing flow.

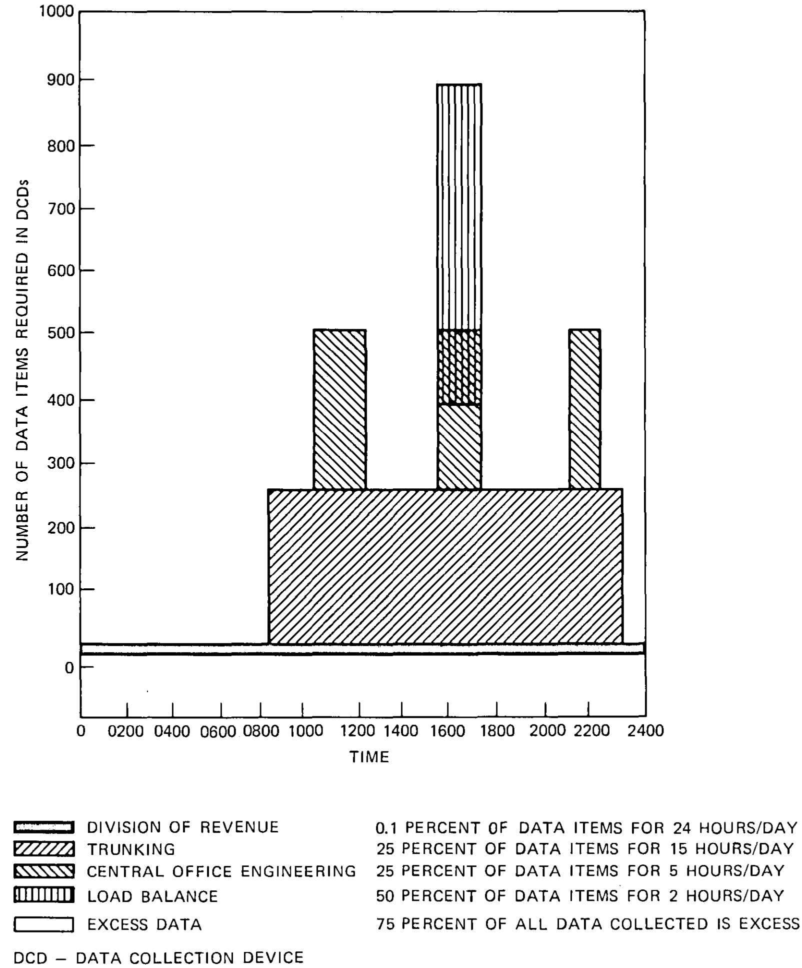

Fig. 5—Example of excess data.

screening, then, is to discard excess data at the earliest possible point in the process (the editors). It should be noted, however, that the function of screening itself incurs an expense, and for that reason it has been implemented only in those editors that process large volumes of data. The camera/register data editor, for example, does not perform automatic screening because much of the screening is done concurrently with the manual keypunching activities.

The remainder of the discussion on TDAS will focus on each of the runs that encompasses the mainline system. A number of additional support runs that perform peripheral functions, like file recovery, will not be discussed here.

5.4.2 File Interface Process

To perform its processing, TDAS requires certain information from the Common Update TDA and TMR master files to be accessed randomly. FIPS, then, is the first run in a TDAS cycle and is responsible for converting the TDAS master files into various random access files for efficient processing. Some files use the Virtual Sequential Access Method (VSAM), where the file is created sequentially based on a defined key and the access method routines retrieve data records randomly using indexes and pointers. Other files use the Basic Direct Access Method (BDAM), where the file is created and randomly accessed using a calculated key.

FIPS creates a two-level screen with which the editors screen the data. The programs first relate TMRs to Data Collection Units (DCUs) on a day and time basis to indicate for a given day and time whether any requests exist for data from the DCU. This first screening level is accomplished by creating a VSAM file keyed by DCU and date. The second level of screening is an interrogation of a BDAM bit map file consisting of bits for each data item, called Data Collection Device (DCD) within the data collection unit. Each bit indicates whether the data item (DCD) has been requested (1) or not requested (0). The VSAM record for a DCU/date contains calculated keys that point to maps of 1000 DCDs in the BDAM file for particular times of the day. Figure 6 shows this screen creation process.

FIPS is run following the weekly CU cycle so that the latest record base changes can be captured for the current TDAS cycle.

5.4.3 Data entry

After FIPS has completed its run successfully, the editors as listed in Section 5.3.2 begin execution. Each of those five editors is run separately. They can be executed in any sequence but not concurrently. An editor is responsible for detecting errors, and screening, sorting, and reformatting the acceptable data into a common output format. The editor also maintains a log of all input data. To account for lost data and the status of individual input data tapes, each editor produces a Data Errors report and a Tape Processing report. In addition, if data are lost because the collection machine goes out of service, the cause of that loss is indicated on the data tape when the collection machine returns to service. This indication is displayed on the Tape Processing report by TDAS and the same information is also forwarded to the Centralized System for Analysis and Reporting (CSAR)[2] for its total system data tracking responsibility.

The editors that process data from Stored Program Controlled Systems also strip off a portion of that data and write it to a file for transmission to the centrally located SPCS COER[1] system. The

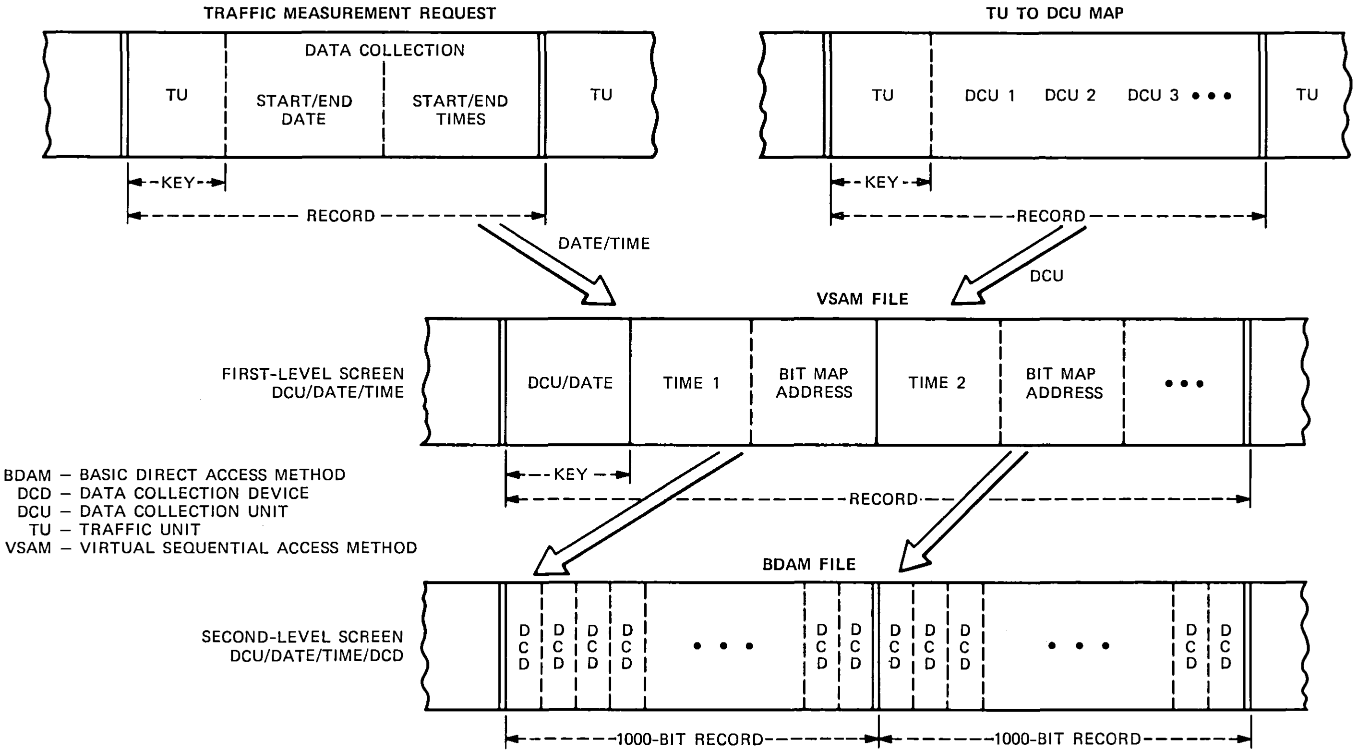

Fig. 6—Data screen creation.

TDAS-to-SPCS COER interface is the same as the collection-machine-to-TDAS interface. As a result data directed to SPCS COER are output by the editors (after being screened). Some OTCs that input daily tapes to the editors transmit the daily outputs to SPCS COER for overnight processing.

A problem inherent in the current collection-machine-to-TDAS interface is the mechanics of tape transport. In a typical OTC, 30 to 50 tapes are transported between these systems weekly. The administration and transportation costs as well as loss of tapes incurred under this arrangement have stimulated investigation into a direct collection-machine-to-TDAS data link. A future release of these systems will include the capability to transmit the data over either a dedicated or dial-up link.

5.4.4 Data analysis

After submitting all the desired data to the TDAS editors, the next and final run is Data Analysis. It is a multistep process that performs most of the work of the system.

The first step, TMR Scheduling, determines which TMRs are ready for processing in the current TDAS cycle. It compares the list of TMRs to the available data, and then based on rules and options governing calculation of expiration dates, releases whatever data are available.

The Data Merge module then merges the data with identification information from the TDA master file. The data that remains (for example, data that was input to TDAS for a partial week and does not fully satisfy the TMRs on file), along with the unexpired TMRs, are held for processing in a subsequent cycle.

With TMRs now ready for processing, the Measurement Summarizer module verifies that the collection equipment was operating correctly at the time the data were gathered. As an example, usage data are expected to be collected at a sampling rate of 36 scans per hour. If the rate falls outside acceptable limits, between 32 and 37 for electromechanical switches and exactly 36 for stored program control switches, it is marked invalid and not passed downstream. Otherwise the data are considered acceptable. If the user requests that the data be adjusted, they are then adjusted to the hourly (36 scan) equivalent. The data are then summarized into intervals as requested on the TMR and written to the appropriate interface and report files. Throughout this process statistics are compiled on invalid data and missing data.

The interface files are then sorted and formatted by the Interface Create modules for delivery to the downstream systems. Next, the Reports Create modules format the various TDAS-generated reports, and the CSAR Create module formats the CSAR interface file. Finally, the run ends with system cleanup functions.

VI. NO. 5 CROSSBAR CENTRAL OFFICE EQUIPMENT REPORTS SYSTEM (5XB COER)

6.1 Introduction

The No. 5XB COER system was one of the earliest TNDS component systems developed. It was designed to assist the network engineers in forecasting future equipment requirements in 5XB central offices. 5XB COER summarizes, over a one-year period, traffic usage data for each individual 5XB equipment component. This interval of time is referred to as the engineering year. Maintenance of this full-year history is one of the distinguishing features of 5XB COER, as compared with other systems, such as EADAS, which provide quick looks at the traffic data over short intervals of time (e.g., weekly).

The summarized data are organized into a report, referred to as the Machine Load and Service Summary (MLSS). This report is the key output of the system. The year's worth of data summarized in this report, combined with past years' data and other forecast information, enable the engineer to determine future equipment requirements.

The following sections will discuss the MLSS reports as they apply to the 5XB engineering function, and will discuss the responsibilities of Network Administration in ensuring the quality of those reports. Finally, a description of the software implementation is also provided.

6.2 MLSS reports in support of engineering

"Load" and "Service" in the title describe the content of the MLSS report: load, referring to the traffic carried by a particular component; and service, the measure of performance associated with that particular load. The engineer determines what equipment will be needed for the probable level of service to meet a projected load.

Data are summarized in High Day (HD), Ten-High-Day (10HD), and Average Busy Season (ABS) results for a Time-Consistent Busy Hour (TCBH). TCBH implies the single hour (e.g., 9 a.m. to 10 a.m.) in which the particular equipment component (e.g., originating registers for *Touch-Tone** telephone) consistently experiences its heaviest traffic load. Within this busy hour, HD, 10HD, and ABS can best be described with the help of Figs. 7 and 8. Assume in these figures that dots represent traffic data values of load on which a component is to be engineered. Figure 7 plots these values for each day (typically only business days) over the 12 months of the engineering year. As expected,

* Trademark of AT&T.

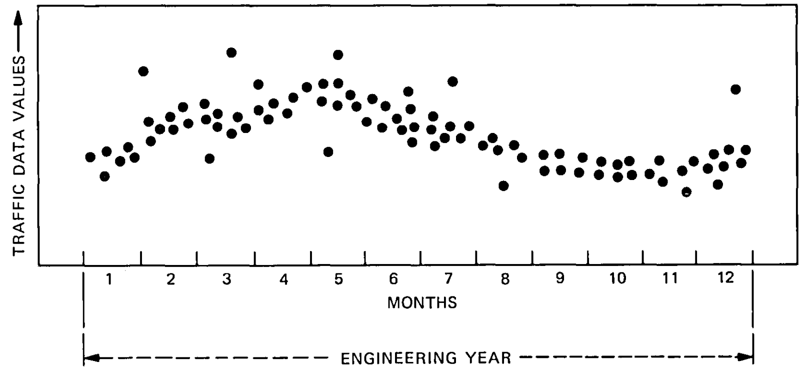

Fig. 7—Typical traffic profile.

some days have higher than normal values, and there are extended periods where the general average is higher than the rest of the year. Figure 8 shows this data summarized in terms of monthly averages for each of the twelve months, and also shows the selection of the ten highest days and three busiest months. The average of the ten highest days forms the 10HD value, and the three busiest months the ABS value.

The common format of the MLSS report used for all 5XB components, shown in Fig. 9, incorporates all the engineering requirements. Header information identifies details such as office, component, and time span. The columns (left to right) list the date, key column value* (used to engineer the component) and support columns (significant calculations such as percent overflow/blocking, holding time, percent occupancy/capacity, along with the actual values used in their calculations). The Ratio to Busy Season (RA/BS) and Fitted Ratio (FIT RAT) enable the engineer to determine if a particular reported high-day key column value is a yearly recurring event. High days typically not expected to recur (e.g., heavy telephone traffic due to a severe storm) should be excluded from the engineering base. RA/BS is the ratio of the key column value to the ABS value. FIT RAT is a statistical quantity relating this key column data point to the total year's values, assuming a gamma distribution. The RA/BS value should be very close in value to the corresponding FIT RAT value if the data point is in fact a legitimate recurring high day. The rows display the highest day (HD), the next nine highest days, the 10HD average, the next 15

* Typically, this quantity is a usage measurement normalized with respect to an office parameter such as main stations (MS) or trunks (TRK), depending on which is applicable.

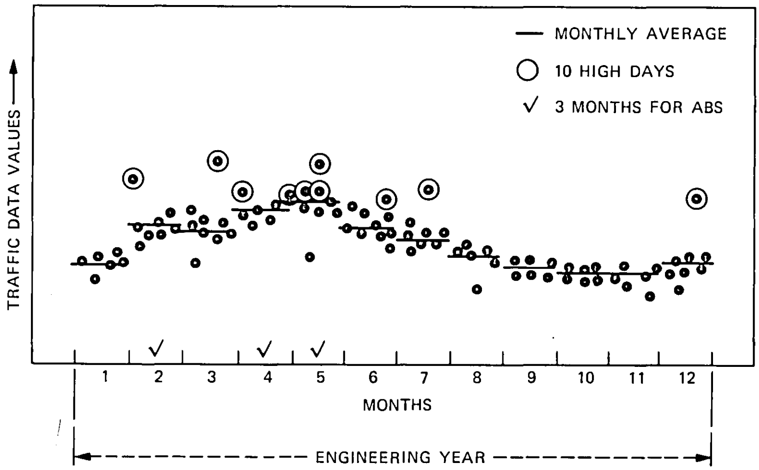

Fig. 8—Summary of traffic data.

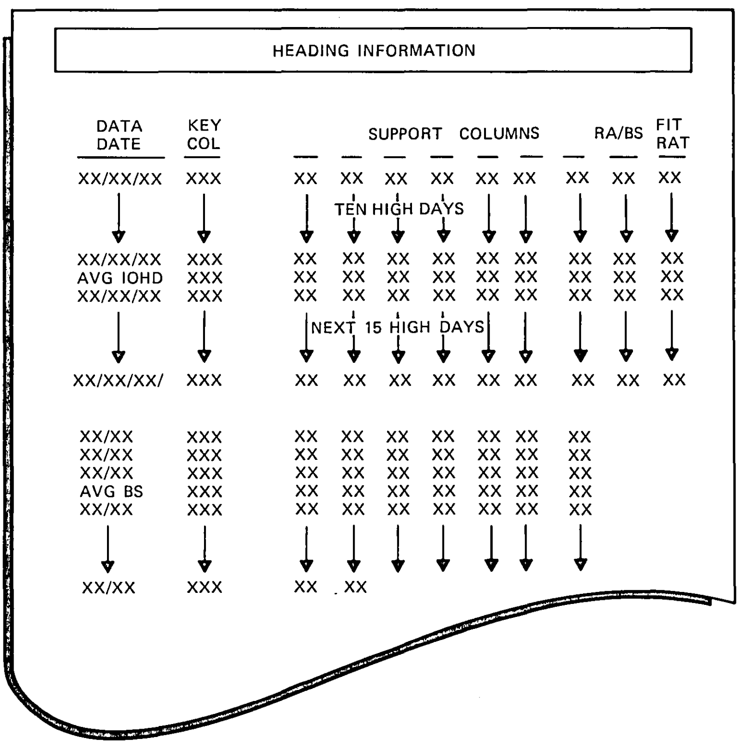

Fig. 9—MLSS report format.

highest days (displayed for their potential use as one or more 10HD replacements), the three highest monthly averages, the ABS calculation, and the monthly averages for the remaining nine months.

As traffic data are received by the system, typically in weekly batches, the MLSS report is automatically developed for each component and hour requested.* Among the components are line link frames, markers, originating registers, outgoing senders, incoming registers, trunks, and transverters.

6.3 Network administrative support

The network administrator is responsible for supplying a complete, accurate, and timely set of MLSS reports to the engineering staff. This job function fits in with the overall responsibility for monitoring the daily performance of the central office to offer the best quality of service possible with the currently installed equipment. To perform this function, equipment utilization information is gathered by coordinating the collection, report generation, and analysis of the traffic data.

To achieve the prime system objective of producing usable MLSS reports, the No. 5XB COER system supplies extensive aids to the network administrator. Among these are:

1. Busy hour determination
2. Data validation
3. Data surveillance
4. Data management.

Each will be described in the following sections.

6.3.1 Busy hour determination

Integrated into the 5XB COER system is the ability to perform a Busy Hour Determination (BHD) study. A BHD study identifies the busiest hour(s) for each engineerable component. Those hours will then be the ones reported in the next year's MLSS. Any particular component may have many busy hours contending for the heaviest load. Busy hour contention normally results from heterogeneous traffic (e.g., business versus residential). The selection of the correct hour(s) is critical because data collection and processing for the coming engineering year will be based on that decision. Selection of the wrong MLSS hour(s) can result in invalid engineering decisions.

A BHD study is conducted at least once a year, usually just prior to entering the busier part of the year. The study period typically spans

* The system will process up to five hours per component. In general, only the busy hour is studied, but one or two side hours may also be studied if uncertainty exists concerning the busy hour.

two weeks. For this study period, up to twenty-four hours per day of data can be collected in half-hour increments. The half-hour totals are summed by TDAS into hourly totals, ending either on the hour or the half hour or both. This results in up to 48 traffic data totals per day for each component measured.

Once all the study data have been collected and processed, a BHD report is produced for each component, as shown in Fig. 10. The key and support columns in this report parallel the MLSS report.

The system-selected busiest hour (based on key column) is identified with a "BH" to the left of the hour. Two additional columns, RA/BH and DAYS = BH/#, are computed to help the user understand and interpret the results, as well as to assess the degree of confidence in selecting the busy hour. RA/BH is the ratio of the individual hour to the selected busy hour. This column reveals hours that are close in load to the selected busy hour. The next column displays the actual number of days in which that hour was the busy hour (DAYS = BH), compared to the number of days that had valid data (#).

6.3.2 Data validation

In order for the summaries in the MLSS reports to help the engineers make equipment forecasts, the basic measurement data must

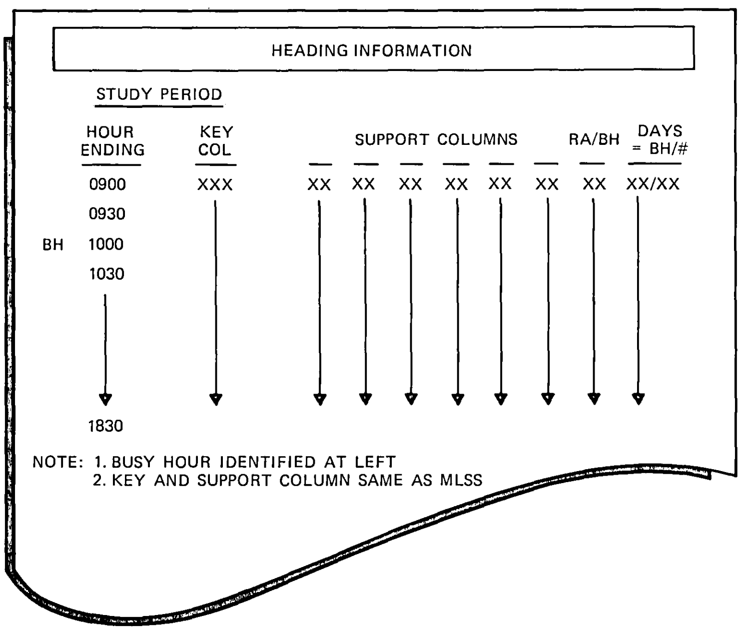

Fig. 10—Busy hour determination report format.

be a true representation of what occurred physically in the central office. An extensive set of data validation tests and associated failure reports are therefore provided.

The tests can be classified into four major categories, and they are usually performed in the following order:

1. Sanity checks
2. Peg count balance tests
3. Holding time tests
4. Historical tests.

The order of performance is based on the relative "strength" of the test. For example, sanity check failures carry a 100-percent confidence factor that the tested data are in fact incorrect. At the other end, failures in the historical tests should be investigated further before the data are declared invalid.

6.3.2.1 Sanity checks. These checks are simple and detect impossible conditions. Two examples of test failure conditions for originating registers (OR) and incoming registers (IR) are:

1. OR total usage > 36 ccs × no. of ORs in group
2. IR maintenance usage > total usage.*

6.3.2.2 Peg count balance test. This type of test defines certain relationships that must exist among several peg count (PC) measurements. For example, the Dial Tone Marker (DTM) peg count (which is scored on each call origination) should be equal to the sum of those available peg counts that indicate call disposition. In this case there are three dispositions:

1. Call successfully seized an originating register, causing the OR peg count to score.

2. Call encountered an "All OR Busy" (AORB) condition, causing the AORB peg count to score.

3. Call was not able to find an idle path through the switch, thereby causing a "Dial Tone Marker Failure to Match" peg count (DTM FM PC) to score.

The form of this test would be:

$$\text{Lower Range} \leq \frac{DTM\ PC}{OR\ PC + AORB + DTM\ FM\ PC} < \text{Upper Range.}$$

Because of the complexity of the 5XB equipment interrelationships, these types of tests often involve a large number of data items. Exact balances generally cannot be achieved because of the lack of complete measurements and the time skew for recording the various peg counts. Therefore, a range is defined around the theoretical value of 1.0.

* Total Usage means traffic plus maintenance usage.

Normally the PC balance test is performed on a component group. However, the precision of the test frequently can be refined by applying diagnostic tests to each individual in a failing group.

6.3.2.3 Holding time tests. Based on the specific equipment design, the average Holding Time (HT) of a component, e.g., Outgoing Senders (OSs), should be within a certain range of theoretical values. The range is usually broad to account for variable central office characteristics. For any given central office, the user can override and redefine the range test to strengthen the validation for that office.

6.3.2.4 Historical reliability test. The Historical Reliability test is used for data that on a statistical basis fall within a range based on past history. An example of this is the office calling load per main station, CCS/MS. The value of CCS/MS will vary according to telephone calling fluctuations but usually within a historically determined range. For these cases, tests must be designed based on past data by using a range tracking procedure.

There are many statistical techniques used for this type of validation. The method implemented by 5XB COER has produced good empirical results, was simple to implement, and required little storage. Three parameters are stored for this test: an estimated value based on past history, a variance, and a residual. The residual helps to adjust the acceptable range as the data value changes. This is particularly important in minimizing false alarms as busy seasons are encountered. The valid range for the next data point is the estimated value plus or minus one standard deviation. Based upon a comparison of the residual to the standard deviation, additional adjustments are applied to the estimated value.

6.3.3 Data surveillance

The ability to detect subtle patterns in equipment performance can mean the difference between the correct diagnosis of an equipment malfunction and the incorrect conclusion that an office is underprovisioned. Real-time surveillance is provided by systems such as EADAS, which produce exception reports based on exceeded thresholds. To perform a comprehensive analysis of total office performance, component information for the entire office must be available for the same hour. 5XB COER produces an optional report package consisting of detailed component data, validation failures, and raw measurement listing for the entire office for one or two individual hours from a week's worth of data. Network administrators select the hour(s) to be reported, and the criteria on which the system will select "the highest hour(s)" from those received to process. One of four criteria may be selected for each hour to be reported: Terminating Peg Count, Originating Peg Count, Originating plus Terminating Peg Count, Originating plus Terminating Usage. The selected reports are produced each processing cycle.

Another data surveillance feature of 5XB COER supports the annual MLSS product. As the system receives and processes the MLSS data, normally weekly, it produces output reports titled Weekly Machine Administrative Reports (WMARs) and Monthly Machine Administrative Reports (MMARs) for each component and hour requested. WMARs are optional reports. The format of these reports parallels the associated MLSS reports, except that it displays daily data for each day of the week or month rather than 10HD and ABS data. This permits a review of the daily results to ensure their appropriateness for later inclusion in the MLSS results.

6.3.4 Data management

To permit user judgment and knowledge to be reflected in the results, several capabilities can alter the contents of the MLSS report. These capabilities are called data management and include the ability to change an input measurement, change a day in the high day list, change a month in the average busy season list, or change a validation failure flag. Any changes made are identified on the MLSS so the engineer receiving the report is aware of the changes and can question their reason. The changes are made by input transactions that direct the system to modify the history for the current engineering year. At the end of the engineering year, the system produces the completed MLSS report, which reflects the entire year's data along with all data management changes. After this report is produced, the system retains the history for up to two additional months to allow for further changes.

6.4 No. 5XB COER software

6.4.1 Overview

No. 5XB COER software consists of two main procedures and four support procedures. The four support procedures perform:

1. File initialization—Establishes the record base for new installations of the system.

2. File conversion—For each new software release that affects the record base, this procedure converts from the old to the new.

3. Parameter file maintenance—The parameter file contains the information that controls the processing order of the system. This procedure can modify that information.

4. Program log list—As described in Section 3.2.4, the program log is a continuous record of system execution activities. This procedure produces a report of that information and is used primarily as a debugging aid.

The two main procedures, Control File Update and Process Traffic Data, will be described in the following sections. An overview of the processing steps and interactions of these two procedures is shown in Fig. 11.

There are two primary files in 5XB COER. The Control File contains all the user-supplied control information and office parameters. It serves a function for 5XB COER that is analogous to that provided by the CU files for TDAS and other systems. The History File contains all the processing results being accumulated over the engineering year for MLSS reporting. The data management transactions are depicted in Fig. 11 from a logical point of view as going directly to the history file; in reality the transactions pass through the Control File.

6.4.2 Control file update

The Control File Update process incorporates numerous input transaction checks to ensure database accuracy. There are two levels of checks: syntactic checks that ensure correct data formats and values, and semantic checks that perform further edits to ensure the records are logically consistent. Errors result in one of two actions: either the transaction is rejected outright, or the affected office is placed in an "error" state. The error state prevents further processing of traffic data for the office until the error is corrected. Transactions may be effective immediately or they may become effective in the future, depending on the date supplied by the user. Since this is a file update process, it is designed to run as often as necessary prior to a single execution of the Process Traffic Data run.

6.4.3 Process traffic data

Process Traffic Data is the main batch procedure of the 5XB COER system, generally executed weekly. The major functions of this process are:

1. Data Edit—Screen out all but the valid data requested for processing
2. Validate—Assess each measurement and assign a degree of confidence tag
3. Calculate—Formulate all required numerical results
4. History
   - (a) Mesh new results with historical base
   - (b) Perform special BHD and data management processing (if applicable)
   - (c) Assemble output results file for reporting
5. Reporting—Format the output results for input to the Report Distributor for printing.

These functions are described in the following paragraphs.

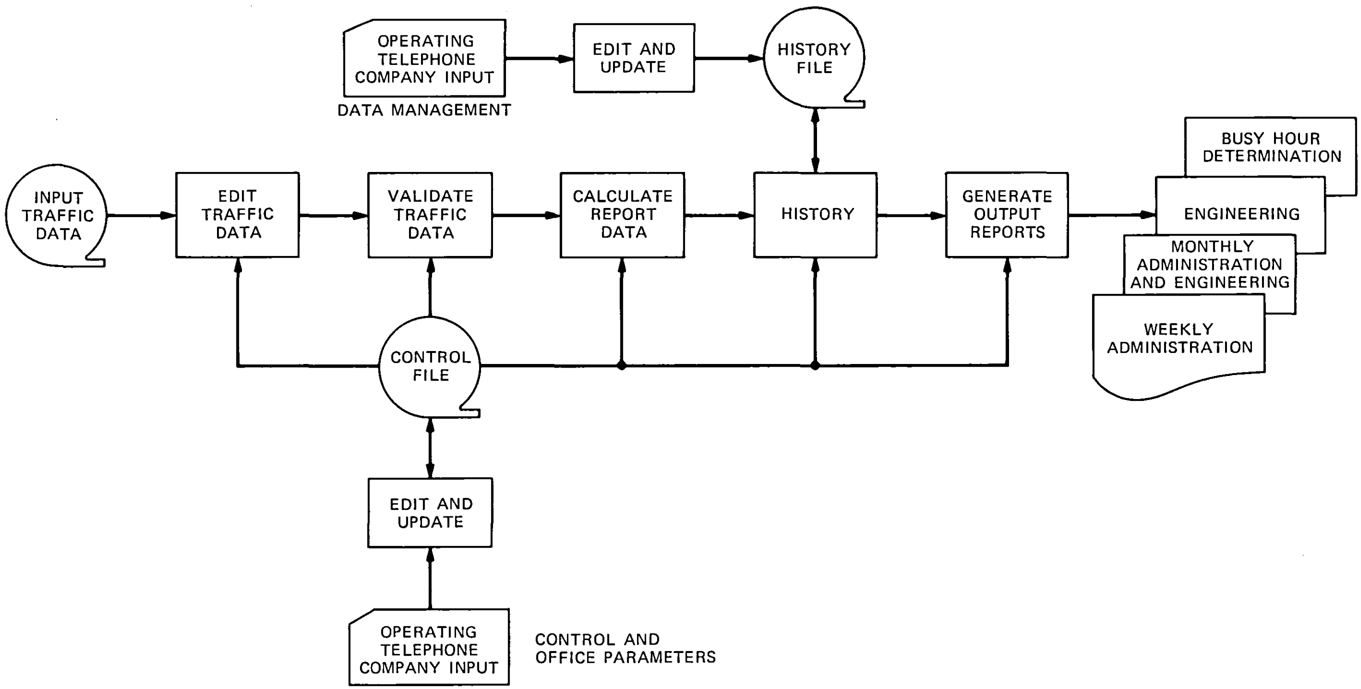

Fig. 11—Processing overview of 5XB COER.

6.4.3.1 Data edit and validate. The Edit function performs syntactical checks on new traffic data received from TDAS, places data in a Save Data File for any office indicated in the control file as being in "error", and combines new traffic data with any saved data for further processing. All input edit errors are detailed in an error report that includes a statistical log reflecting the current year's activity.

The Validation function makes a comprehensive assessment of the quality of the data by employing the extensive set of tests described in Section 6.3.2. All validation test failures are reported, identifying the data as well as test parameters involved.

6.4.3.2 Calculate. The Calculation function is the next major task. The validation tags are propagated into the calculations so that a calculation result receives the lowest tag of any of its input data items.

The calculation module has incorporated a particularly robust design that is described further because of its usefulness in other applications. The design centers on a common computation subroutine driven by tables that define the calculations to be performed. Calling routines are organized to parallel the component report structure, such that for a particular summary there is one calculation routine.

The information in the control tables is stored as a stack in symbolic postfix form. The computation subroutine evaluates the stack in two passes. The first pass checks the stack for syntactic and semantic errors and builds a work table of operand values. The second pass performs the calculation and checks for overflow and transaction errors. The control information is organized in five tables, one of which provides the calculation formulas in postfix form, the other four of which provide operand values.

Seven permissible operators have been incorporated in the calculation table: the five normal arithmetic operations of addition, subtraction, multiplication, division, and exponentiation; and two end of stack symbols, one of which indicates that a percent check should be made to ensure that the resultant value is less than 100.

This design has enabled the calculations to be easily maintained. To correct or modify any calculation, all that is required is a table update. To add a new report requires the addition of a new calling routine in which the most difficult part is the table definition.

Up through the calculation function, all the modules have operated on the current cycle's traffic data. The remaining functions, History and Report, integrate that data with previously processed data and report the results to the user.

6.4.3.3 History. The History Update module performs three major tasks. First, it merges the current traffic data with the stored traffic data in the history file. To do this, the current daily data records are sequentially added to the appropriate section of the file; the high day list is reorganized with respect to the new data with new days inserted and old days deleted; and once sufficient daily records exist, the monthly average is computed (or recomputed if it already exists). Also, if a BHD study is active and new data are received, the merging of the new hourly results by component must be completed. Second, any user data management transactions are applied to the high day and monthly average records. Third data appropriate for the WMAR, MMAR, or MLSS reports are released.

6.4.3.4 Reports. The last function executed in the process traffic data procedure is Report Formation. This module formats and merges all the outputs from the previous modules into a file for input to the Report Distributor System (Section IX). There are currently more than fifty basic report formats and over 200 variations within these formats. A multilevel table-driven design provides a generalized module that is simple to modify, maintain, and test.

VII. INDIVIDUAL CIRCUIT ANALYSIS SYSTEM

7.1 Introduction

Maintaining wiring and assignment records for traffic data collection is a complicated task. It is quite common to sample more than 10,000 points in an electromechanical switching machine. These sampled points are combined to provide accumulated counts in typically a thousand registers or data collection devices (DCDs) used for engineering and administering the office. This accumulation of sampled points into DCDs greatly simplifies the engineering and administration functions by reducing the amount of data to be analyzed. However, the record-keeping and physical wiring associated with this grouping function is a major source of error that can propagate into costly engineering mistakes.

To understand the Individual Circuit Analysis Program (ICAN) and how it assists in this record-keeping process, it is necessary to understand the functions of EADAS/ICUR. This subject is covered in detail in Ref. 4. However, Fig. 12 presents a brief overview of EADAS/ICUR.

With EADAS, traffic data counts occurring in the central office are encoded and sent via a data link to the collection machine. In basic EADAS (i.e., without the ICUR option) the usage data are collected at the output of the Traffic Usage Recorder (TUR), where the measurements have already been grouped by wiring on the cross-connect field of the TUR frame. This grouping function allows the accumulation of related information into a single DCD (e.g., usage data collected from 10 individual trunks into one trunk group). With the ICUR option, which can be applied to electromechanical switching offices, the data are grouped via software in EADAS rather than via wiring

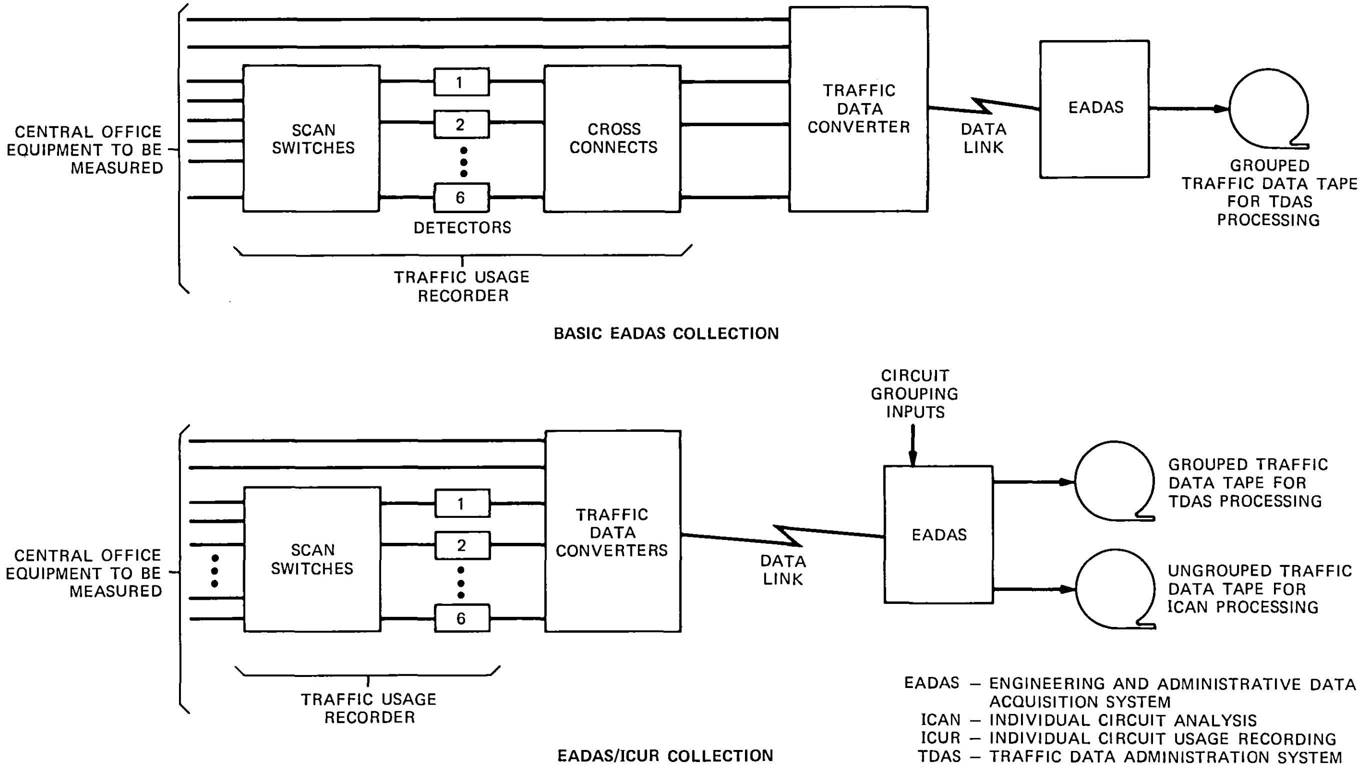

Fig. 12—Overview of EADAS/ICUR.

on the TUR cross-connect field. Software grouping simplifies the TUR administration function by replacing manual TUR cross-connect logs. In addition, the ICUR option allows network administrators to observe the usage being accumulated on each of the individual TUR inputs within a group. These ungrouped data can be analyzed to detect switching equipment problems in the central office, data collection equipment problems, and various database errors.

EADAS/ICUR produces a magnetic tape for ICAN processing. This tape contains the ungrouped usage data and additional information necessary to administer the TURs. The ICAN process is unique among TNDS/EQ components in having access to individual circuit usage information. This makes the system highly valuable in detecting central office operational problems that otherwise could go unnoticed.

This overview of ICUR operation will lend perspective to the description of the inputs, outputs, and processing flow of ICAN that follows.

7.2 ICAN inputs

ICAN's purpose is twofold: to aid in the administration of TUR frames through the use of the EADAS/ICUR circuit grouping map and to help detect faulty central office equipment items. To accomplish these functions, inputs are needed from three sources: EADAS/ICUR, the Common Update record base, and the network administrators.

EADAS/ICUR provides the ungrouped traffic data measurements, the schedules on when these data were collected, a map on how the individual measurements are grouped into DCDs, and the updates to this map since the last ICAN tape was written.

Descriptive information for each of the DCDs is obtained from the Common Update record base. Finally, the network administrators provide control information specifying the reports to be produced and the individual circuit data to be purged from the ICAN history file.

7.3 ICAN outputs

With the basic inputs described above, ICAN can produce a series of output reports to aid network administration. These reports can be produced on either a demand or exception basis. The reports are divided into the categories of record base listing reports, error reports, equipment surveillance reports, and system administration reports.

7.3.1 Record base listing reports

ICAN maintains the record base describing how individual TUR inputs are grouped into DCDs. Various views of this record base can be produced in the form of report listings tailored for a particular network administration task.

TUR administration listings provide a complete layout of the TUR inputs. Each input is identified by physical location of the central office equipment being measured, along with the associated DCD accumulating the information. To aid in the assignment of new TUR measures, a separate report can be produced itemizing those inputs that are unassigned or unequipped.

Other views of this record base list the individual DCDs and specify the associated TUR inputs that are accumulated in each DCD count. This view aids in troubleshooting various validation test failures as well as in making new DCD assignments.

A third view of the record base is organized by the equipment location being measured and identifies the associated TUR input and associated DCD. This view is primarily used in central office maintenance activities for troubleshooting reported data collection problems.

Individual TUR inputs are associated with a primary DCD. However, the scorings can also contribute to a secondary device. These secondary devices, called load grouping codes, typically accumulate total office input equipment utilization broken down by load division (see LBS, Section VIII). This secondary method of accumulation simplifies various office engineering functions by providing a single DCD representing the total usage accumulated over hundreds of other devices. Load grouping code listings show this type of DCD summarization.

In addition to the various TUR and DCD administration listings mentioned above, there are two other ICAN record base listing reports: the TUR schedule report, which identifies the ICUR data collection periods, and the Site Definition listing, which identifies all EADAS/ICUR data collection machines within the company.

7.3.2 Error reports

The ICAN error report series produces various exception reports itemizing discrepancies found in processing inputs and updating the TUR record base maintained in ICAN. The reports specify the various input transaction errors, and the discrepancies found between Common Update, EADAS/ICUR, and ICAN record bases. The record base auditing functions incorporated in ICAN correlate various parts of the TNDS record base and verify that these record bases are kept in synchronization.

7.3.3 Equipment surveillance reports

Of particular interest in the ICAN system are the three central office equipment surveillance reports. Because information is available on the basis of an individual TUR input, additional verification can be performed to assure the central office equipment is operating properly.

*7.3.3.1 **Abnormally Short Holding Time reports.*** The Abnormally Short Holding Time (ASH) report is useful in detecting a "killer trunk." A killer trunk is a circuit that reflects an abnormally short holding time because customers abandon it almost immediately to escape noise, bad transmission, or other problems. Because these trunks are quickly returned to an idle state, the switching equipment is likely to pick them more often than other trunks within a group. One faulty trunk may cause numerous failures in attempts to use the group, even though the group is well below capacity. The ASH tests detect probable killer trunks by a statistical analysis of the individual usage data.

Following is a simplified example demonstrating the working of ASH. Consider the sequence of busy/idle indications for two circuits of the same group connected to a 100-second scan TUR:

Circuit A: 111001100
Circuit B: 101001101,

where: 1 = busy condition when sampled, and 0 = idle condition when sampled. As you can see, both circuits were busy five of the nine times they were scanned by the TUR. However, Circuit A has only three transitions between busy and idle status, while Circuit B has six. If we assume normal call duration of 200 to 300 seconds, it is probable that Circuit A is carrying only two calls of at least a 200-second duration in this time period. It is also probable that Circuit B is carrying at least four distinct calls (busy indication followed by idle indication) of much shorter duration. Circuits displaying these characteristics are likely to be ASH circuits.

As noted above, many busy-to-idle transitions are indicative of a killer trunk. A busy-to-idle transition is considered a positive ASH factor. Similarly, a trunk that is detected as being busy for a number of successive TUR scans is more likely to be operating normally. A trunk appearing busy for two successive TUR scans is therefore considered a negative ASH factor.

To determine which trunks have abnormally short holding times, ICAN first weights these positive and negative ASH factors by the group occupancy rate (total CCS for the DCD divided by the number of TUR scans) and by the individual trunk occupancy for all other trunks. These weighted factors are accumulated until their sum (called the "ASH statistic") exceeds a positive or negative threshold.

If the positive threshold is reached, a message is output on the ASH report indicating that the trunk has an abnormally short holding time.

Thus, ICAN infers through a statistical algorithm that a trunk has an abnormally short holding time. Even though the ASH detection method is statistical, it is designed to be 98 percent reliable in reporting true office problems.

7.3.3.2 Unusual usage report. ICAN analyzes the usage data received and reports on TUR inputs that show usage but are unassigned or unequipped, as well as TUR inputs that are assigned but appear inactive.

Unassigned/unequipped analysis is quite simple. With each transmittal from ICUR, the unassigned/unequipped TUR inputs are analyzed for transitions between busy and idle states. If a transition has occurred, the TUR input is flagged for exception reporting. In the exception report, the amount of circuit occupancy is computed and displayed.

Inactive circuit analysis is also quite simple. If the TUR input shows no transitions between busy and idle, it is flagged for exception reporting. The state of the circuit (100 percent busy or 100 percent idle) along with the date of last activity are displayed on the exception report. Care must be taken to avoid reporting false alarms for those measurements that are typically inactive. The set of measurement types that are typically inactive are automatically identified by the unusual usage software and further analysis is bypassed.

7.3.3.3 Individual circuit usage displays. The individual circuit usage display is a valuable tool available with ICAN. At the user's prerogative, usage on selected circuits can be accumulated and displayed to help analyze problems too subtle to be detected by the ASH algorithms. These displays are also used for tracking specific troubles indicated on central office equipment. For the selected DCDs, the occupancy of each of the individual measurement components that accumulate data for the DCD are displayed in a graphical form. This enables easy comparison of measurement components with similar characteristics. For example, the occupancy over a selected time period of each of the trunks in a given trunk group can be displayed for analysis. In essence, the display report shows a view of the central office equipment operation that is similar to a time-lapse photograph.

7.3.4 System administration reports

The final set of ICAN reports is designed for administration of the ICAN System itself. This set of reports depicts activity that has occurred, and summarizes processing results for the individual ICUR input tapes on an individual DCU basis. The majority of these reports are intended for use by network administration personnel. Additional reports list computer resource usage and record base activity. The latter would be needed in the event of a system recovery.

7.4 Processing flow

7.4.1 General

The ICAN System consists of a single mainline program and several support utilities. This single job approach simplifies OTC scheduling, operation, and control. Based on control inputs, the mainline ICAN process validates inputs, processes ICUR tapes, produces periodic and exception reports, purges unneeded data, and backs up the database.

Because of the large volume of data and the complexity of some of the analysis algorithms, the system has been designed for efficient data processing. A typical company uses ICAN to administer 300 TURs, requiring the weekly processing of over 10 million individual circuit usage measures and over 30 million individual validations.

The flow of mainstream ICAN and the various utilities are described in the following sections.

7.4.2 Mainline process

The functions performed in the mainline ICAN process are shown in Table V. A single execution consists of four phases. The first, the card input processing phase, analyzes and validates the input requests, produces various error reports, and sets up the controls for the additional phases as requested by the input transactions.

The second phase, tape processing, is the most complex of the mainline ICAN and includes:

1. Schedule processing—Analyzes the ICUR tape writing schedules, produces the associated master record base updates, and outputs any appropriate error messages.

2. Card image processing—Analyzes the Circuit Grouping Map (CGM) updates produced since the last ICUR tape was written. The updates are validated and the CGMs are updated. Any discrepancies are noted on the appropriate error reports.

3. Map processing—Compares the updated ICAN CGM with the current ICUR CGM within this function. Again, after analysis, the appropriate error messages are generated and the ICAN CGMs updated to reflect current assignments in EADAS.

4. Circuit data processing—Analyzes individual circuits for ASH and unusual usage, stores data for display reports and long-term analysis, Pending Status Flag (PSF) processes, and deletes old, unneeded data currently resident in the record base.

The reporting phase, three, formulates and produces both requested and exception reports.

In the final phase a backup tape of the record base is made so that restart and recovery steps are simplified if errors are encountered in future runs of mainline ICAN.

Within mainline ICAN processing, a checkpoint feature is imple-

Table V—Mainline ICAN processing phases

	II: Tape Processing					
I: Card Input	Schedule Processing	Card Image Processing	Map Processing	Circuit Data Processing	III: Reporting	IV: Backup
Validate input cards and report discrepancies	1. Compare schedule with existing schedule 2. Report schedules if changed 3. Determine number of tape writes on tape	1. Validate and report discrepancies 2. Update ICAN maps or discrepancies	1. Compare 1-for-1 with existing maps 2. Report discrepancies 3. Change ICAN maps	1. Store data 2. Unusual usage-inactive circuit 3. PSF change 4. Produce ASH report 5. Display report 6. Delete data if requested	Assemble and produce reports	Generate backup tape

mented so that if processing problems are encountered later in the same job, job restart will begin automatically at the point immediately after the last successfully processed ICUR tape.

7.4.3 Support process

To supply needed DCD descriptive information to mainline ICAN, a single support run is periodically executed to extract the DCD information from the Common Update record base, identify discrepancies between the record bases, and produce the necessary record base updates.

7.4.4 Utilities

A number of support utilities are available to be used when needed. These utilities create and modify various portions of the ICAN record base, provide recovery capabilities of the ICAN record base, and serve as a backup in case problems occur in the EADAS/ICUR record base.

VIII. LOAD BALANCE SYSTEM

8.1 Introduction

A principal Bell System goal has always been to offer the best possible service at the lowest possible cost. For the network administrator, this objective translates into making the most efficient use of existing switching facilities while maintaining an acceptable level of service to all customers. To achieve this goal, network administrators must distribute the telephone traffic load evenly across the machine's switching (or load) units. This principle is called load balance. This section will describe a tool available to help achieve that goal—the Load Balance System (LBS). LBS is a integrated component of TNDS/EQ. It relies on CU for maintaining its record base and on TDAS for supplying the weekly traffic data with which it does all its processing. Distribution and printing (or microfiching) of all LBS reports are the responsibility of the Report Distribution Facility (see Section IX).

LBS supports all 1 *ESS* switch, 2 *ESS* switch, No. 1 Crossbar (1XB), 5XB, Crossbar Tandem (XBT), and Step by Step (SXS) machines for which data are collected. Network administration functions are supported by calculating a Load Balance Index (LBI) and providing assignment and balance guides.

8.2 LBS outputs

8.2.1 Index reports

Because of the importance of load balance, an administrative index was devised in the mid-1960s to measure the effectiveness of load balance procedures. LBS mechanizes the implementation of the revised Load Balance Index Plan, as prescribed in AT&T administrative practices. Index reports are produced on a traffic unit (TU) to indicate the level of service and balance achieved by that TU. From these measures indexes are aggregated into the hierarchy of district, division, area, and company scores. On a Service Observing Month (SOM) basis, the company-level report is forwarded to AT&T.

For those offices in which the index is not calculated by LBS, typically smaller, rural SXS machines, the index can be entered into the LBS record base (via CU) for inclusion in higher-level LBI reports.

8.2.2 Data summaries

Data summaries provide detailed information on the data used in calculating the Load Balance Index. These reports require constant attention. Identified on these reports are load units that have rising or falling load trends, unusually high or low loads, or penalty points assigned from the index calculations. These are signs of either current or potential trouble spots in the switching machines and indicate not only where corrective action may be needed but where additional studies may be appropriate. These reports should be used in association with the line assignment aids (see Section 8.2.4) to achieve an increase of traffic in the more lightly loaded load units while diminishing traffic in the more heavily loaded load units. This is typically accomplished through normal disconnect and assignment activity. In drastic cases, reduction of heavy loads is accomplished through specific transfer of customer lines out of heavily loaded units into lightly loaded ones. This procedure to improve line load balance (and thus the Load Balance Index) should result both in increasing the machine's effective capacity and improving service to the customer.

8.2.3 Second Session Studies reports

Many switching entities experience a secondary daily busy hour during some part or all of the year, due to the calling habits of a community of customers with different calling patterns than those responsible for the primary busy hour. While the LBI is calculated using only primary busy hour data, it is also important to provide quality service during the secondary busy hour. To identify secondary busy hour imbalances, LBS produces Second Session Study reports. The same principles apply to use of second session data summaries as to the busy hour data summaries. The reports are similar, with the exception that penalty point and index values are not calculated for second session studies.

8.2.4 Line Assignment/Transfer reports

LBS produces two basic aids for line assignment. Line Assignment

Guides provide lists of load units, with the most lightly loaded ones listed first, thus appearing in order of preference for assigning new lines. Line Equipment Transfer Guides start at the other end of the list; the most heavily loaded units appear first and are thus in order of preference for transfer of lines. The user may request a Line Assignment Guide or Transfer Guide in any length desired. These reports assist the manual process of determining new line assignments and, if necessary, transferring current assignments to other load units. As telephone companies introduce mechanized line assignment processes, such as the Computer System For Main Frame Operations (COSMOS), manual assignments from LBS reports are discontinued. For that reason Line Assignment Guides are produced only on demand on a TU basis.

8.2.5 Nonindexed balance aids

There are some components such as trunks and service circuits that are also sensitive to load balance but are not included in the LBI calculation. To assist in achieving proper balance of those components, LBS produces balance aids that contain corrective load values.

8.3 LBS inputs

LBS requires three categories of input to perform its functions: traffic data, traffic unit characteristics, and a history base. The history base is a file internally maintained by the system itself. The other two inputs are provided by TDAS and CU, which operate with LBS in a tightly integrated fashion.

8.3.1 Traffic unit characteristics

For LBS to judge the degree of balance within a traffic unit and to provide aids in achieving or maintaining proper balance, the system requires a detailed description of certain TU characteristics. This information is provided by the user via CU. Through the input of various transactions, the user builds a Load Balance Master File, which contains a description of all TUs that LBS supports, specifying the traffic measurements to be processed for each TU. Included in such TU descriptive information is common language identification, number of main stations, destination of output reports, assignment and loading division information, theoretical capacities, and average holding time (AHT).

8.3.2 Traffic data

In addition to TU characteristics, LBS requires data that represent the traffic load being offered to the TU. The source of the data is, of course, the switch itself. However, before reaching LBS, the data are first obtained by a data collection system and passed through TDAS. Using the traffic measurement definitions and traffic measurement requests, TDAS validates, adjusts, and summarizes the data before writing them to an interface file for LBS processing. TDAS also applies an Equipment Measurement Code (EMC) label to each data item, permitting LBS calculations to be driven by internally interpreting the EMCs.

8.3.3 History files and system parameters

Since some of the computations involve averaging and trend analysis, it is necessary that history files be maintained for reference and updated on each run of LBS. Included in the history files are the previous week's loads and penalty points, average holding times, percent capacities, and total-to-sample ratios. The latter is a means of determining total usage on a load unit from a sampling of two of its links.

Certain numerical values are required in the LBI algorithms used by LBS. These include default values for average holding time and total-to-sample ratio, Quality Control Limits (QCL), designation of threshold values for "hot-spot" determination, and specifications of LBI correction values. The QCL is a measure of how far the load on a load unit is allowed to deviate, without penalty, from the average load for the loading division.

8.4 LBS processing flow

8.4.1 General

A complete execution of LBS is performed on a weekly basis following the running of TDAS. The system is divided into six discrete runs, one for mainline LBS and the other five serving as support runs. The remainder of this section will address the mainline run of the system, as shown in Fig. 13.

8.4.2 Study formation and validation

The Study Formation and Validation module receives the input traffic measurements from TDAS. It performs validation tests on the traffic data and combines it with office characteristic information obtained from the CU files for further processing. In addition, if the appropriate peg count data are available, average holding time (AHT) and total-to-sample (T/S) ratios are calculated for later use in index and corrective CCS calculations.

Validations are categorized as follows:

1. Study identification—Data are checked to ensure that only valid combinations of studies have been requested.

2. Data volume—The number of hours requested is not always the

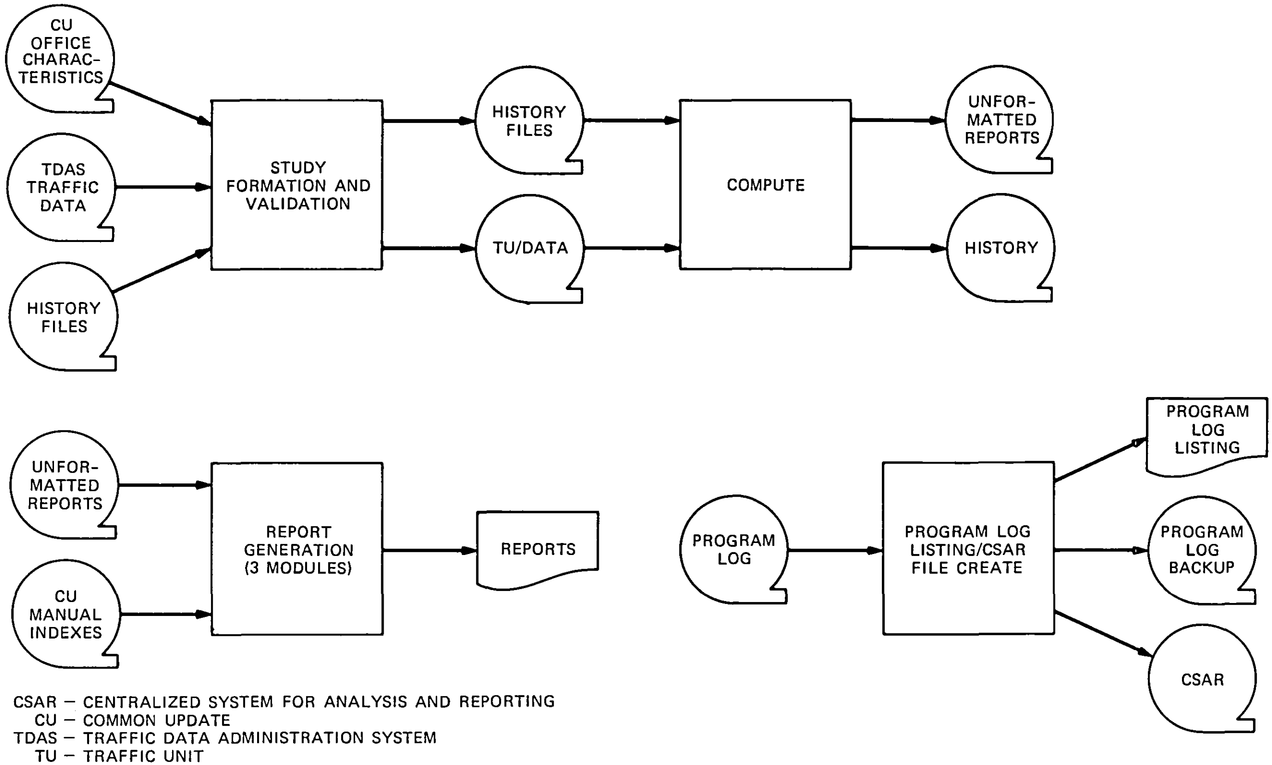

Fig. 13—Load Balance System processing flow.

number received. Acceptable ranges for the different studies have been established to allow for reasonable variation in equipment performance and to retain the statistical reliability of the study. If sufficient data are received, the study is abandoned. In addition, usage data for any hour having either zero or more than 36 CCS are rejected.

3. Data identification—Data are validated to contain acceptable EMC and measurement-type designations.

4. AHT and T/S Ratio—Calculated AHT and T/S ratios are compared to historical values. If they differ by more than a fixed percent, they can be discarded, based on the judgment of network administration personnel.

5. Study criteria—To perform an index study, more than 75 percent of the load units in a TU must have valid data and the loading divisions must be at greater than 30-percent capacity. If both conditions are not satisfied, the study is aborted.

8.4.3 LBS compute

Following the above data validations, the Compute module performs the many calculations necessary to satisfy the study requests. The first function is to react to history change requests to recompute, delete, and reinstate data from history.

Component calculations that contribute to the LBI are performed next. Using AHT and T/S ratio, the quality control limits are determined. Scores are then developed depending upon the indicated deviation of actual usage from the average usage. Load units operating at loads above predetermined thresholds are assigned a "hot-spot" penalty. The index is then calculated using weighted linear regression techniques that utilize up to 12 weeks of past data. Since load imbalances affect service more severely as the total traffic handled by the office increases, the percent load capacity is another factor used in determining the final index.

Finally, corrective CCS values are derived from linear regression estimates of what each load unit's deviation from the average will be at the time of the next study. These values are displayed on the balance guides in the appropriate order to facilitate manual corrections to imbalance.

8.4.4 Reports formations and system cleanup

The remaining modules assemble, sort, and format the various reports produced by the system, and perform cleanup functions.

IX. REPORT DISTRIBUTOR

9.1 Introduction

For the effective utilization of the TNDS downstream systems, it is essential that the reports be distributed in an efficient and timely manner to the users, most of whom are located physically remote from the computation center executing TNDS. The TNDS/EQ components alone can produce several hundred different report formats for distribution to as many as 600 distinct locations within an operating telephone company. To complicate this process further, selected reports are needed on microfiche produced by equipment manufactured by several different vendors, and other reports require multiple copies.

Early in the evolution of TNDS, a Report Distribution and Printing Facility was developed. This capability is utilized today for all the downstream TNDS systems. In addition, it is available as a standalone package and can be used by non-TNDS systems to solve their report distribution needs. The following sections describe the capabilities of the Report Distributor, its inputs, outputs, and processing flow.

9.2 Report Distributor capabilities

The design of the report distribution process is based on two standards implemented throughout the downstream TNDS systems:

1. Every report type is uniquely identified with a five-character report identification specified in the report header.

2. Every report utilizes a four-character responsibility code indicating the specific user who is to receive the report. This responsibility code is also identified in the page header.

With these two basic standards, the report distribution process was designed to produce:

1. Microfiche—For long-term storage, or cost reduction reasons, users are able to select particular reports for output to microfiche. The Report Distributor has the capability of producing outputs that meet the specifications of several microfiche hardware vendors.

2. Multiple copies—Users are able to request multiple copies of selected reports without multiple passes through the distribution process.

3. Proprietary marking—The users can specify that the appropriate proprietary notice be printed on the individual report pages.

4. Monitoring distribution—Although the primary individual is identified with the report responsibility code, alternate distributions can be identified, usually for monitoring purposes.

5. Burst pages—Before a printout can be distributed, it must be broken down by responsibility code. To make this easier and to reduce errors, burst pages are optionally printed (three pages of output with characters printed over the fold) between reports where the responsibility code changes.

6. Remote print—Where the report destinations are physically remote from the computer center running TNDS, time and cost

savings are realized by producing report tapes according to destination and having them printed locally at the remote site. This facility allows selected reports to be routed to numerous locations.

7. Selective retrieval—To recover lost listings, or for historical checking of a report, selective retrieval can be made of an individual report or all reports for a particular responsibility code from a report tape.

8. Print restart—Following a system failure or job cancellation, the report printing program is able to restart at a point that results in minimal duplication of printout. Typically, report print programs are run with a dedicated printer, so that restart points can be communicated to operations through the system console. The frequency of the checkpoints are at the discretion of the OTC. When restarting is needed where multiple print tapes are involved, the restart does not have to spool through complete tapes that have already been printed.

9.3 Report file inputs

The report distribution facility imposes two requirements on every system providing a reporting file for distribution:

1. The report file must be organized in responsibility code/report number order.

2. Key records must be included at specific points within the report file. These records identify the number and responsibility code of the report following the key record. At a minimum, a key record is specified on each change in responsibility code and report number. In the worst case, a key record is provided for every report page. (This is necessary if the report is on microfiche and an index to the fiche is provided on a page basis.)

The report distribution process detects these key records and through the use of Common Update files determines the proper output distribution device.

9.4 Common Update inputs

Inputs through Common Update define the pattern for the distribution process and the types of proprietary messages to appear on reports. Three types of transactions are necessary to define these characteristics:

1. One transaction defines the report output device. These devices are used to break up the distribution process into parts. An individual device may be designated for transmission to a remote computational center where the reports on that device are further subdivided for distribution. A device may be designated for input to a particular type of microfiche equipment, or it may represent the set of reports to be distributed to an individual user.

2. The second type of transaction specifies which report number/responsibility code combinations are to be distributed to each of the numerous output devices.

3. The third transaction specifies by report number the proprietary message to be printed at the bottom of each report page.

The current contents of the various report distribution tables maintained by Common Update can be examined on request.

9.5 Outputs

For problem-tracking considerations, output reports are produced that summarize the information written to each of the output devices during every execution of the report distribution process. These reports identify each responsibility code/report number change, along with any further report content information contained in the key record.

9.6 Report distribution process

The report distribution process is implemented by a combination of three programs (Fig. 14): the merge program, the distributor program, and the printer program.

9.6.1 The merge program

All TNDS report generation programs produce report data sets ordered by destination (responsibility code), similar to the outputs from the distributor and printer programs. When multiple report tapes are to be processed by the distribution facility, a merging of the reports must be made. This is the function of the merge program. It is designed as a separate run because (1) its output may be input to either the distributor or printer program, and (2) to put the merge functions in the distributor program would increase the number of devices required to run it.

9.6.2 The distributor program

The distributor program must be used when there is a need for the generation of remote print tapes and/or microfiche tapes. It uses as input the TNDS keyed report tapes, a key definition file, and the report distribution table. It produces any combination of three outputs: keyed report tapes, nonkeyed report tapes, and microfiche tapes. Keyed report tapes are printed at a local or remote site and can be fed to the printer program directly or further processed at the remote site by another copy of the distributor program before being printed. Nonkeyed report tapes are printed at a local or remote site, and microfiche tapes can be processed by any of the supported microfiche systems.

Report processing by the distributor is key record driven. That is,

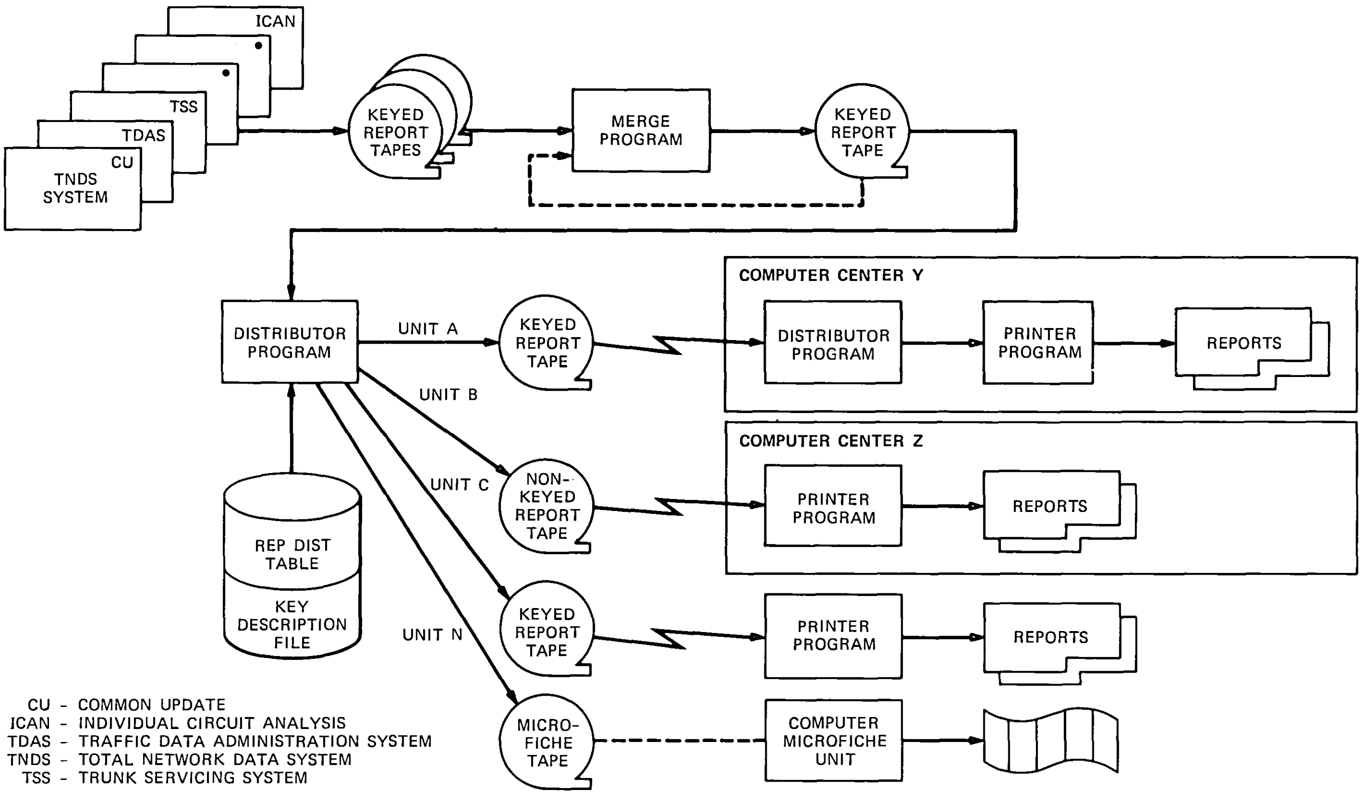

Fig. 14—Report distribution process.

when it encounters a key record on a report tape, it matches the responsibility code and report number against the selection masks in the report distribution table, identifying those report devices that are to receive the report. It then treats all report records up to the next key record as one report, spooling each record to the appropriate data set(s). If no report selection mask matches the report, the report is written to a default report device. When the report is completed, it produces its own reports summarizing the processing that has occurred.

To accommodate changes easily in the meaning of key record fields for reports, allow for addition or deletion of reports by a system, and permit use of the distribution facility by other than TNDS systems, the distributor program has no system-specific internal information. All report-dependent information (i.e., legal report numbers, algorithms for formatting key record data on microfiche, index lines, and distributor reports) is kept external to the program, in a Key Definition file constructed by the using systems.

There are four options supported in the distributor program for microfiche output, which are provided by the device definition record in the report distribution table. These options deal with blank microfiche separators, magnification factors, forms flash, and microfiche sequencing. The blank microfiche option, if used, specifies that a blank microfiche is to be generated when the responsibility code changes. This allows operations to more easily cut the microfiche film for distribution. The magnification factor parameter specifies the magnification to be used in generating the microfiche. The forms flash option enables the user to have a preprinted form superimposed over a microfiche frame. Finally, the microfiche sequencing option determines whether the sequence number as printed on the microfiche is to be reset to one when a new responsibility code is encountered. Each report number or responsibility code change causes a microfiche advance.

9.6.3 The printer program

The printer program can be used both at the OTC central computer center and at a remote print site to produce hard copy listings for distribution. It has options for burst pages and printer restart (options under control of EDP operations) and selective report retrieval and multiple copies (under control of the TNDS EDP coordinator).

X. MANAGEMENT REPORTING SYSTEM (MRS)

10.1 Introduction

In the past, the OTCs originated many requests to develop specialized reports in addition to the current fixed-format reports available through the various TNDS components. The Management Reporting System (MRS) was developed to answer this need. MRS enables both TNDS/EQ and TNDS/TK users to extract information available in various TNDS files and format this information to satisfy their local reporting needs.

Besides customized report generation, MRS offers two other major capabilities for its users. First, MRS can be used to mechanize the production of Common Update input transactions. This capability is especially useful for initializing an office database and for various office conversion activities.

Second, MRS can serve as an interface between software processes. Data may be extracted and made available to OTC-developed systems without the user having to know the actual structure of the TNDS source, nor care about future revisions to this structure.

In its implementation, MRS is an adaptation of the MARK IV™ software package of Informatics, Incorporated, Canoga Park, California. Using this package, Bell Telephone Laboratories personnel provide the necessary file definitions, interfacing software, and associated documentation required to produce reports, transactions, and system interfaces.

10.2 File definitions

More than twenty major TNDS files are defined for MRS use. The file definitions describe the record characteristics and define the structure of each segment within a record. These definitions allow a convenient interface between the TNDS products and the flexible reporting capabilities made available to the user community through MRS. BTL provides the file definitions and updates them with each major release. Because the Mark IV definitions describe the data information items independent of file structure, OTC-designed programs need not be modified as the basic files are restructured with new TNDS releases.

10.3 Mark IV interfacing software

In addition to the basic file definitions and documentation, BTL provides various MARK IV interfacing software. This software ranges from catalogued requests to restructure the TNDS files for MARK IV use to BTL written code that is integrated with MARK IV providing specialized translation operations. Examples of these specialized translation operations include:

1. Conversion of Trunk Group Serial Numbers to the Common Language Circuit Identification (CLCI).
2. Conversion of Location Machine Processing Codes to Common Language Location Identifications (CLLI).
3. Date and time format conversion routines.

10.4 Using the system

To illustrate the use of MRS, consider a request by a TNDS/TK user to identify, list, and count all trunk groups that use a particular group as an alternate leg in completing a call. Using one of the MRS reference manuals, the MARK IV definitions of the required information are identified. In this example seven distinct field definitions are used:

Field definition 1—The Trunk Group Serial Number (TGSN) of the group being investigated.

Field definitions 2 through 4—The TGSN of the first, second, and third alternate legs at the A end of the circuit group.

Field definitions 5 through 7—The TGSN of the first, second, and third alternate legs at the Z end of the circuit group.

A MRS request is now coded on standard MARK IV forms. The forms specify the comparisons to be made on the input files, the sort sequence of the data appearing on the output report, and the format and field titles of the output report.

In the above example a scan is made of the A and Z alternate legs for the particular TGSN in question (DV000189). If DV000189 appears as an alternate leg, the trunk group is listed in TGSN order with all legs identified. Finally, a count of the TGSNs using DV000189 is computed and displayed. Figure 15 shows an example of the resulting output report.

10.5 Processing flow

The generic flowchart for an MRS run is shown in Fig. 16. Standard TNDS files and MARK IV user source are submitted as input to the run. The specified fields are identified and analyzed in Step A, during

03/27/82 GROUPS USING DV000189 AS ALTERNATE LEGS PAGE 1

TGSN	TGSN AL1	TGSN AL2	TGSN ZL1	TGSN ZL2	TGSN ZL3
DV000170	DV000189	DV000173	DV000173	DV000189	
DV000183	DV000196	DV000189	DV000189	DV000196	
DV000187	DV000165	DV000189	DV000189	DV000165	
DV000186	DV000175	DV000189	DV000189	DV000175	
DV000142	DV000168	DV000189	DV000189	DV000168	
DV000143	DV000164	DV000189	DV000189	DV000164	

03/27/82 GROUPS USING DV000189 AS ALTERNATE LEGS PAGE 2

TGSN	TGSN AL1	TGSN AL2	TGSN AL3	TGSN ZL1	TGSN ZL2	TGSN ZL3

GRAND COUNT 6

Fig. 15—Example of a Management Reporting System report.

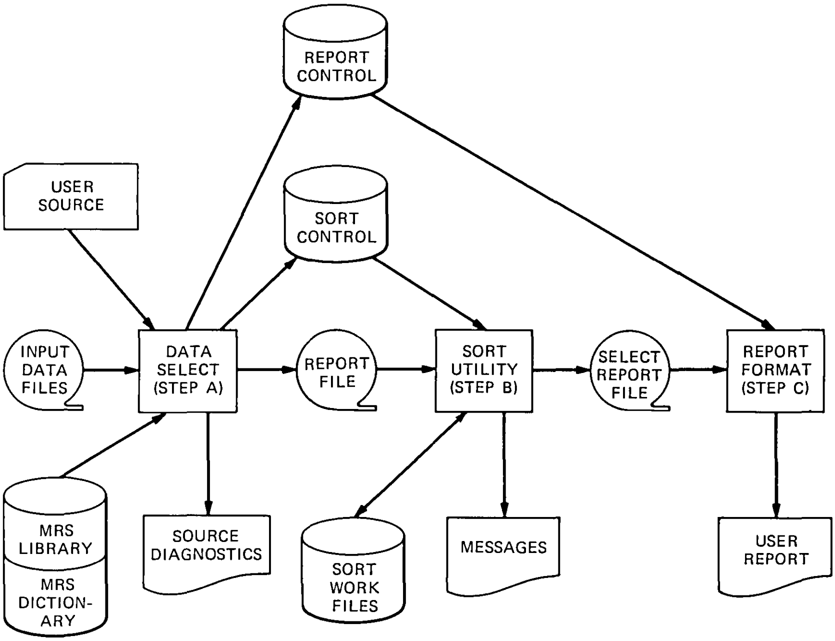

Fig. 16—Generic flowchart for the Management Reporting System.

which the basic report file is produced. Step B is a sort utility that organizes the report file into the sequence requested by the user. Step C takes the organized report file and formats the output report requested by the user. If the objective was to generate a TNDS transaction file or create a file for input to other mechanized systems, the sort output file is used directly, bypassing the third step.

XI. SUMMARY AND FUTURE DIRECTIONS

Components of TNDS/EQ have been in use since the early 1970s. These systems play a key role in the overall administration of network data. The systems are in production use by all Bell System operating telephone companies and Bell Canada. Extensive controls have been implemented to ensure smooth and reliable operation of the data processing center. Through standardization of the data collection environment, interchange of data among companies of common items has become possible. Development and user support activities are carried out by Bell Laboratories and Western Electric central developers, with enhancements to the systems provided by annual releases.

For the future, the major work for TNDS/EQ will be twofold: to maintain support of the continuously evolving central office and trunking network, and to provide more responsive service by using the on-line record base and reporting features of the On-line Records and Reporting System (ORRS). This replacement to Common Update will be developed during the period from 1983 to 1986. The modernization of TNDS/EQ will include establishing telecommunications links between collection machines and TDAS, integrating the No. 5XB COER Control File into the ORRS record base, and assimilating functions currently carried out by CSAR into ORRS. In assessing future telephone company needs to be supported by TNDS/EQ, the key will be responsiveness. The use and expected number of users of traffic data will increase, along with demands for more timely, reliable, and accurate data. The modernization program for TNDS/EQ will be essential for these new business requirements.

REFERENCES

1. R. F. Grantges, V. L. Fahrmann, T. A. Gibson, and L. M. Brown, "Total Network Data System: Central Office Equipment Reports for Stored Program Control Systems," B.S.T.J., this issue.
2. D. R. Anderson and M. J. Evans, "Total Network Data System: Performance Measurement/Trouble Location," B.S.T.J., this issue.
3. D. H. Barnes and J. J. O'Connor, "Total Network Data System: Small Office Network Data System," B.S.T.J., this issue.
4. C. J. Byrne, D. J. Gagne, J. A. Grandle, Jr., and G. H. Wedemeyer, "Total Network Data System: Data Acquisition and Near-Real-Time Surveillance," B.S.T.J., this issue.
5. M. S. Hall, Jr., J. A. Kohut, G. W. Riesz, and J. W. Steifle, "Total Network Data System: System Plan," B.S.T.J., this issue.
6. P. V. Bezdek and J. P. Collins, "Total Network Data System: Trunking Systems," B.S.T.J., this issue.
7. W. S. Hayward and J. Moreland, "Total Network Data System: Theoretical and Engineering Foundations," B.S.T.J., this issue.

AUTHORS

N. Dudley Fulton, B.S.E.E. (Electrical Engineering), 1967, Purdue University; M.S. (Computer Science), 1972, Ohio State University; Bell Laboratories, 1972—. Mr. Fulton has worked on several components of the Total Network Data System family including TDAS, SXS COER, 5XB COER, and TRFS. He is currently Supervisor of the Trunk Forecasting Software Development group. Member, Tau Beta Pi, Eta Kappa Nu, IEEE, and ACM.

James J. Galiardi, B.S.E.E. (Electrical Engineering), 1966, Illinois Institute of Technology; M.S.E.E., 1968, Ohio State University; Bell Laboratories, 1966—. From 1966 to 1970 Mr. Galiardi was located in the Columbus Laboratories and was involved in the initial hardware and software design of the 4A ETS toll switching machine. In 1970 he became involved in mechanizing various switching machine laboratory support functions in minicomputer hardware and software. From 1971 to 1973 his work responsibilities included hardware diagnostic programming for Common Channel Interoffice Signaling. In 1973 he was promoted to a development supervisor position in the Total Network Data System Project and relocated to New Jersey. From 1973 to the present Mr. Galiardi has held various supervisory positions within TNDS, ranging from system planning through the development process to system test. Member, Eta Kappa Nu, Tau Beta Pi.

Edward J. Pasternak, A.B. (magna cum laude, Engineering Sciences), 1957, Harvard University; M.S.E.E., 1962, E.E. (Professional), 1966, Columbia University; Bell Laboratories, 1962—. Mr. Pasternak was initially assigned to an Exploratory Development group which was examining stored program control methods for modernization of SXS. With the inception of the Traffic Service Position System (TSPS) project, he was assigned development responsibilities for the TSPS Audit programs. In 1966 he was promoted to Supervisor of the TSPS Audit and Recent Change Programs Group. Following cutover in 1969 of the first TSPS office in Morristown, New Jersey, Mr. Pasternak was appointed Head of the Stored Program Control (SPC) Development Department with hardware and software design responsibilities for the SPC 1A processor. In 1973 Mr. Pasternak was transferred to the Business Information Systems area of Bell Laboratories as Head of the BISCUS/FACS System Test Department. In 1974 he became Head of the Servicing and Estimating Department, with design responsibilities for several systems that were subsequently incorporated into the Total Network Data System. Mr. Pasternak currently has responsibilities for the TNDS Equipment Systems component of the Total Network Data System. Member, IEEE, ACM.

Stewart A. Schulman, B.S. (Computer Science), 1972, City College of New York; M.S. (Computer Science), 1974, Stevens Institute of Technology; Bell Laboratories, 1972—. Mr. Schulman began his career at Bell Laboratories in the System Test area of the Service and Estimating (S&E) department, which later evolved into the Total Network Data Systems—Equipment (TNDS/EQ) development organization. Following his initial assignment, he was involved in the development of various TNDS/EQ component systems. In 1979, Mr. Schulman was promoted to Supervisor, responsible for ongoing development of the Traffic Data Administration System (TDAS) and the Load Balance System (LBS). His current assignment is development Supervisor for the On-line Records and Reporting System (ORRS), an on-line replacement for an existing TNDS batch system and the vehicle for future TNDS/EQ modernization.

Herman E. Voigt, B.S.E.E., 1962, Polytechnic Institute of Brooklyn; M.S.E.E., 1964, Columbia University; Bell Laboratories, 1962—. Mr. Voigt's initial work was in exploratory development of Common Control for Step-by-Step local switching offices. From 1963 to 1973 he was involved in the development of the first Traffic Service Position System (TSPS). His prime areas of contribution were in the software supporting the operator actions call processing functions, for which he received two patents, systems test and eventually field support of all installations. Since 1973 he has been supervising the development of various operation support systems that provide administrative and engineering equipment utilization reports for local electromechanical and electronic switching offices. Member, Eta Kappa Nu, Tau Beta Pi, IEEE Computer Society.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Trunking Systems

By P. V. BEZDEK* and J. P. COLLINS*

(Manuscript received August 3, 1982)

The Total Network Data System/Trunking (TNDS/TK) systems are those modules of TNDS that support the engineering and administration of the message trunk network. TNDS/TK consists of the Trunk Forecasting System (TFS), the Trunk Servicing System (TSS), and Common Update/Trunking (CU/TK). Using representative trunk group loads and switching system growth data as input, TFS forecasts trunk needs for five future years. TSS fine tunes the first year of the forecast by showing where to rearrange the network to meet current demand. And CU/TK supports TSS and TFS with a record base that contains a description of the network and user-stated parameters. We describe TNDS/TK from the standpoint of its environment, functions, system internals, and future direction. We also present a high-level view of its algorithms. For more detail, the reader is referred to the "Theoretical and Engineering Foundations" article and its references in this issue.

I. THE CIRCUIT ADMINISTRATION CENTER

1.1 General

The Total Network Data System/Trunking (TNDS/TK) systems support the activities of people in the Circuit Administration Centers (CACs). These are work centers in the Bell Operating Companies (BOCs) that are responsible for engineering and administering the

* Bell Laboratories.

message trunk network. Each BOC has at least one CAC; large companies may have several CACs with responsibilities divided geographically.

The CAC functions consist of (1) forecasting the required size and placement of trunk groups over five future years (trunk forecasting), (2) determining how best to modify the existing network to satisfy current demand (trunk servicing), and (3) initiating and monitoring the work orders that result in rerouting traffic from one trunk group onto another (route administration). TNDS/TK directly supports trunk forecasting and trunk servicing. However, it does not directly support route administration, although it does provide a major input to routing by defining, in forecasting, where reroutes are planned.

For a description of trunk servicing and trunk forecasting that is more comprehensive than that given below, the reader is referred to the "Environment and Objectives" article in this issue.

1.2 Trunk servicing functions

By sending Traffic Measurement Requests (TMRs) to TNDS/EQ, the trunk servicer indicates needs for the collection of traffic data on each trunk group throughout the year and frequently during the expected busy season. Traffic measurements that do not satisfy validation tests or represent normal customer behavior must be excluded from use. The servicer must initiate action to correct problems in the data collection process so that subsequent weeks' data can be collected successfully. The accepted measurements are used to estimate each trunk group's offered loads, current service levels, and the number of trunks required to provide objective service.

The servicer issues and tracks Message Trunk Orders (MTOs) to add or disconnect trunks to maintain objective levels of service and utilization. Representative busy-season loads, called "base" loads, must also be selected for each trunk group to support the forecasting process. Annually, Trunk Administration Measurement Plan (TAMP) reports are produced, which describe the adequacy of trunk-group data collection and compare the actual trunk network with a theoretical, low-cost network that produces the desired level of service and utilization.

1.3 Trunk forecasting functions

Trunk forecasting determines the future size and location of trunk groups as a major input to the construction program. Although the generation of the forecast is mechanized by TNDS/TK, the forecaster must manually determine and monitor much of its input. The activities this involves are described below.

The forecaster must define central office growth, typically in terms of main stations and traffic volume per main station. This person must specify the presence or absence of trunk groups, based on switch plans, homing arrangements, switch capacities, costs of facilities, tariff changes, and marketing demand forecast changes. Also based on these sources, the forecaster must indicate where the major shifts in traffic load will occur and select the appropriate load projection and sizing algorithms. The forecaster must perform these functions for scheduled forecasts and, often on short notice, for unscheduled ones in support of new switch plans or other construction program constraints. Finally, the forecaster must make input corrections that will reconcile future forecasts with current load variations that servicing reacts to with unplanned MTOs.

II. A FUNCTIONAL DESCRIPTION OF TNDS/TK

2.1 Overview

Figure 1 shows that TNDS/TK is a batch system made up of three component systems. The Trunk Servicing System (TSS) supports the servicing functions described above. Except for that part of the system related to annual TAMP, TSS is usually run once a week. The Trunk Forecasting System (TFS) performs most of the calculations involved in producing a General Trunk Forecast. It consists of several modules that may operate independently, though all are run between two and four or more times per year. Supporting TSS and TFS is Common Update/Trunking (CU/TK), which maintains a record base that stores a network description and associated parameters.

Separate installations of TNDS/TK are present in each BOC. Depending on the size of the company, an installation may process data for as few as 4000 to more than 100,000 trunk groups.

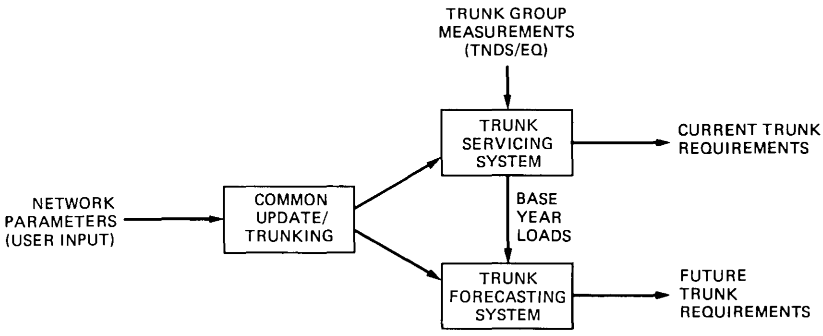

Fig. 1—TNDS/TK components.

2.2 Common Update/Trunking (CU/TK)

2.2.1 General

CU/TK is the interface between the user (the servicer or forecaster) and the rest of the system. Its record base stores a description of the network, engineering and administrative parameter values, some intermediate calculations, revisions to calculations, and some traffic loads. CU/TK provides 27 reports that describe its content. It also validates the input and advises the user of actual and potential errors.

2.2.2 Record base content

CU/TK data fall in two broad categories: those that apply to the company as a whole and those that apply to a portion of it, possibly to a single switch or trunk group. Normally, the first type is created and maintained by a person with company-level responsibility. The second is the obligation of the individual servicers and forecasters.

The company-level information is global in nature and serves several purposes. It may specify company policy or provide a key processing control. A single company input may also obviate numerous identical inputs by individual users. Among other things, company-level inputs define Bell System Common Language name change associations, report formatting and distribution criteria, engineering options, and the forecast years for TFS.

The servicer/forecaster-level information is more detailed and network dependent. *Traffic Unit Records* define the nodes in the network. They give the start and end dates of switching systems or NXX's that terminate trunk groups or for which growth data are to be stored. These records also allow forecasters to override the company-level specification of minimum trunk group size, described later. *Circuit Group Records* define the links in the network. They specify the life spans of trunk groups, load projection and sizing algorithms, trunk servicing criteria and defaults, base year loads, and alternate routes. Figure 2 shows a sample Circuit Group Record (TU500). *Growth Records* contain the main station and traffic forecast data that apply to central offices and are used in projecting trunk group loads. *Traffic Transfer Records* contain a description of the network impacted by area transfers, rehomes and reroutes, and the loads involved. In addition, there are inputs to revise intermediate TFS calculations, request nonautomatic reports to analyze the forecast, and specify that TSS ignore certain weeks of nonrepresentative traffic data.

2.2.3 Input validation

CU/TK verifies that the inputs contain the proper character set (field validation) and that all data on a record are consistent (intrarecord analysis). Consistency *between* records is verified downstream

```
TOTAL NETWORK DATA SYSTEM                                  PROCESS DATE 07 23 81
TRUNKING SYSTEMS                          TNDS             BISP REPORT     TU500
ACCEPTANCE TEST NETWORK                                    RESP CODE        UNF4
SERVICING                              CIRCUIT GROUP
                                          RECORD
                                                              12-31 74      ON
AN000227
            TROY NY 03 MGO   ALBY NY SS CGO   M  PH 55 IE                   I

PROJ CODE      13    PNT CGSN                         ALT RTE AL1    AN000632 100
GRTH CD A      00      PCT 2 WAY                              AL2    AN000178 100
        Z      00    REQD MONTH     01                        AL3
CONV CODE       7    DSTN-CTGY                          COMP TNDM
TRMT CODE            MDLR ENG CD                       A CL FINAL    AN000632
CONF CODE            MIN PRI CCS                      ALT RTE ZL1
PROB BLKG            STIM-FACTOR                              ZL2
NBR SUB GPS          CNST PGM CGY                             ZL3
NBR ACSG EQ                   USE                       COMP TNDM
D/D VAR                 OWSHP A Z                     Z CL FINAL
W-FACTOR             FCST RMKS CD                     T AREA A
ECCS           09    PSEUDO A RTE                     T AREA Z
CCS/TRK                     PCT A                     OTHER CGID    00237
TS CONV CD           PSEUDO Z RTE                     TF-RESP CD    UNF4
TF INFO MSGK1 N    16.6 I    .0 O     .0


SVC OPT 1            DSGN-PK FCTR                     REPORT TIME   E
      3              CKTS   IN SVC  0015              RPTS EXC-IND
PRI-HR         10            PEND                     HIST CGSN
SEC HR         11       TRND DOWN                     SVC RMKS CD
AVG HT        170    TDAS UPDATE                      TS RESP CD    UNF4
TS INFO

* * * * * * * * * * * * * * * * BASE VALUES * * * * * * * * * * * * * * *

BASING PRD      BASE   CNTL   BASE   TOTAL      OFL     PRIMARY       BASE
END DATE        TRKS   BASE   NAME   OFRD CCS   CCS     OFRD CCS   MO  EXT MO

 09 30 80         15    AL2    HBH      267                            3
                               AL1      267                            3
                               AL2      320        8                  11

* * * * * * * * * * * * * * * * * CHANGES * * * * * * * * * * * * * * * *

  S119    ALT RTE AL1     AN121781                  PCT OFL AL1    100
          ALT RTE AL2     AN121484                  PCT OFL AL2    100
          A CL FINAL      AN121781
```

Fig. 2—Circuit Group Record (TU500).

in TSS and TFS. CU/TK rejects inputs that fail its analysis and notifies the user with a Record Update Errors Report (TU001) and several statistical summaries.

2.3 The Trunk Servicing System (TSS)

The following sections outline the major functions of TSS. Representative processes within these functions are described in some detail, illustrating the types of processing that TSS performs.

2.3.1 Cycle setup

Although a TSS cycle can only process traffic measurements taken during a single week, the data received from the Traffic Data Administration System (TDAS) can pertain to several weeks. (The different technologies employed in data collection, ranging from photographic to electronic, produce varied delays before the measurements are made available to TSS.) The TSS cycle-setup function therefore separates the measurements that pertain to a user input "study week" from those that pertain to other weeks. The study-week measurements are passed to the measurement validation function; the others are held for later runs of the system. Measurements that pertain to those weeks before the study week are reported to the users so that back-dated cycles can be run as appropriate.

The CU/TK circuit group records contain (at one time) a time-varying description of the trunk network. The cycle setup function also extracts a static description of the trunk network, appropriate for the study week to be processed.

2.3.2 Measurement validation

The systems that collect traffic measurements and deliver them to TSS do not eliminate machine or human error. TSS must attempt to keep erroneous measurements away from its load-estimation processing. Although certain measurement errors generate absurd data, other errors can create apparently normal data or unlikely but theoretically possible data. The TSS validation function screens out measurements statistically likely to be in error. Rejected and missing measurements are reported to the users, so that problems in the measurement collection systems can be identified and corrected.

A few sample validations are described below. These validation tests are performed for each hour in which the appropriate measurements are available. The measurements discussed here are defined in the "Theoretical and Engineering Foundations" article earlier in this issue.

The peg count (PC) and overflow (OFL) measurements taken at one end of a trunk group are compared. If

$$OFL \geq PC > 0,$$

then both measurements are rejected. It is not valid for OFL to exceed PC. Although they can theoretically be equal and positive, most cases of such measurements are caused by measurement error.

When usage (U) is measured at both ends of a trunk group, the measurements need not agree, because of the sampling technique used. Based on a few assumptions about holding times and random calling patterns, each measurement should be an approximation to the true usage on the group, and hence to the other measurement. Specifically, if

$$(U_1 - U_2)^2 > (200 \text{ seconds})(U_1 + U_2),$$

TSS determines that the two measurements are not sufficiently close for both to be acceptable. It rejects the lesser measurement because typical problems in collecting usage data produce undervalued measurements. If, however, the two measurements are sufficiently close, TSS uses their average as its estimate of true usage.

2.3.3 *Load estimation*

TSS can estimate the average load offered to a trunk group from any of five measurements, or from two particular combinations of these measurements. The specific calculation used for a trunk group depends on trunk group type and the availability of measurements after the validation function. In most cases, several parameters that describe the traffic on a group and the group's performance are calculated, rather than average offered load alone. These parameters are determined from up to five hours' data, representing a fixed clock hour for five days of a business week. Thus, up to 24 sets of estimates, one per measured clock hour, are produced for the business days of the study week.

If, for example, usage, peg count, and overflow measurements are available from the same hour for the appropriate hours, TSS estimates the offered load (a) for each such hour as

$$a = \frac{PC - OFL \times R}{PC - OFL} U,$$

where R is based on trunk group type and accounts for customer retrial of blocked calls. This equation is based on the assumption that calls blocked by the trunk group and not retried would have the same average holding time as completed calls. These estimates of offered load are averaged over five days to produce study week hourly average offered loads. The variance among the daily offered loads in each clock hour is computed and stored, as well as estimates of blocking, average holding time, and peakedness.

If only usage measurements are available for a trunk group because of equipment limitations or lost or rejected data, TSS executes a more time-consuming algorithm that produces fewer, less reliable results. First it averages the usage measurements in each clock hour over the measured days. Then it uses standard numerical-analysis convergence techniques around a program that calculates expected usage, given offered load, to find an offered load that corresponds to the average measured usage. Blocking is computed as a by-product of this process. But average holding time and peakedness cannot be computed in this case. In fact, a user estimate of peakedness (or, if necessary, a TSS default value) is needed as an input to the process.

2.3.4 *Study period formation*

To increase the reliability of its traffic estimates, TSS must produce averaged parameter values that cover study periods of up to four weeks, again for each measured clock hour. The values from the current study week are averaged with values from preceding study weeks, weighted by the number of measured days for each hour in each week. As a result, the measured days are effectively averaged with all equal weights. The function outputs up to 24 sets of averaged estimates, one per measured clock hour.

Many groups are not measured every week. If a group is measured during the study week, TSS normally uses prior weeks' data (at most eight weeks older) to form averages of four measured weeks. Servicers can optionally submit Administrative Period Control inputs to CU/TK, inhibiting TSS from using past data, in cases where the traffic offered to the group is changing significantly (caused, for example, by seasonal variations in demand). If a group is not measured during the given study week, but during one or more of the preceding three weeks, TSS estimates "current" study period loads equal to the previous study period loads, for use in selecting busy hours.

2.3.5 Busy hour selection

From these hourly study-period load estimates, TSS must select a particular hour's load for which to size each trunk group. The network would obviously satisfy service objectives in all hours if each group were sized to satisfy its own service objective in its own busy hour (where "busy" is suitably defined). But because traffic can overflow from one trunk group to another, and unused capacity in one group can relieve a nearby overloaded group, a network engineered in this way would be more costly than necessary. For economic reasons, the hour for which a trunk group should be sized, the administrative hour, depends on loads offered to that group and the surrounding network.

The TSS process for selecting busy hours, Significant Hour Engineering, involves clusters of trunk groups as well as individual groups. A cluster is a final trunk group, together with all the surrounding high-usage groups that (1) share one (fixed) endpoint with the final trunk group, and (2) overflow traffic directly or indirectly to the final trunk group based on the alternate-routing logic of the switch at the shared endpoint.

In cases where all the groups in a nontrivial cluster (i.e., more than one trunk group) have study period load estimates available, TSS computes a cluster load for each hour by summing the carried load on the high-usage groups and the offered load on the final group, in corresponding hours. The hour with the greatest average cluster load per measured trunk is identified as the cluster's busy hour and the "control hour" of the final group. And, the hour with greatest measured blocking on the final group is termed the final group's "service busy hour." These two selected hours may coincide.

Selecting busy hours for high-usage groups is, from a logical view, recursive. Each high-usage group has a set of "significant" related groups defined in the CU/TK circuit group records, consisting of (1) the groups in its alternate route(s) and (2) its cluster final(s) (the plurals here are applicable to two-way high-usage groups only). The set of control hours of the significant groups is computed and then identified as the set of significant hours for the high-usage group itself. Then the significant hour with the greatest load on the high-usage group becomes the group's control hour.

For example, in Fig. 3, there are three clusters: groups AB and AD, BC and BD, and CD. Based on greatest average cluster load per measured trunk, the control hours for groups AB, BC, and CD might be hours 7, 8, and 12, respectively. The significant hours for group BD are now known to be 8 and 12. If the load on BD in hour 8 exceeds that in hour 12, then the control hour for BD is hour 8. The significant hours for group AD are now known to be 7 and 8. If the load on AD in hour eight exceeds that in hour seven, then the control hour for AD is hour eight. An even greater load on AD in hour 9 or 12 would not be considered.

2.3.6 *Trunk group sizing*

TSS computes the number of trunks required to provide objective service on final groups, based on average offered load, peakedness, and day-to-day variation. Required-trunks values are calculated for the traffic in the group's control hour and in the group's service busy hour. TSS chooses the larger required-trunks value as the value for the current study period, and the associated hour is called the group's administrative hour.

High-usage groups are sized for economic, rather than service-based, criteria. Using control-hour traffic data, TSS determines the number of high-usage trunks required to minimize the combined cost of the direct and alternate routes.

The Circuit Group Servicing Record report shown in Fig. 4 is produced after the trunk group sizing process. The bottom part of the report page shows a rough graph of trunks in service and study-period trunks required against time, for the past 65 weeks. The top left gives

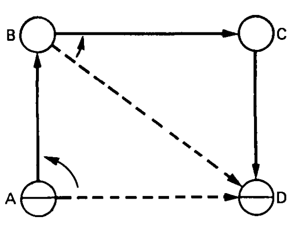

Fig. 3—Sample network for busy hour selection. Solid lines indicate final trunk groups, and dashed lines high-usage trunk groups. Curved arrows show the alternate routing used for calls overflowing a high-usage group.

```
TOTAL NETWORK DATA SYSTEM            TNDS - TRUNK SERVICING SYSTEM                               PROCESS DATE 01 11 81
TRUNKING SYSTEMS                                                                                 BISP REPORT     TB668
ACCEPTANCE TEST NETWORK            NETWORK CIRCUIT GROUP SERVICING RECORD                          RESP CODE     UNF4
SERVICING
CYCLE DATE  12 08 80
                                                                                        CONV TS-  CONF REPORT DESIGN
                                                             PRI SEC  SVC-OPT ECCS  HT  CODE CONV CODE TIME     PF
AN000227    00237     TROYNY03MCO ALBYNYSSCGO M- PH 55 IE     10  11  0 3      09  170  7               EDT    1.00

 LAST STDY-WK            CIRCUITS            SEL     CIRCUIT GROUP ADMINISTRATIVE HOUR          *  SIGNIFICANT HOURS
  OF STDY-PRD  I-SVC REQD FCST TND-DN PND   MEAS  HOUR OFRD-CCS OFL-CCS OFRD-PF VAL-DAYS        *   LEG  HR    CCSM
 ------------  ----- ---- ---- ------ ---   ----  ---- -------- ------- ------- --------        *  ------ --  ---------
 80 OCT 06       15   10   16               UPCO   11    251       0     1.01     18          *    AL1 19 AN000632
    OCT 13       15   10   16               UPCO   11    263       1     1.02     17          *    AL2 11 AN000178
    OCT 20       15   10   16               UPCO   11    260       1     1.02     17          *    AL3
    OCT 27       15   10   16               UPCO   11    252       1     1.06     18          *    ZL1
    NOV 03       15   10   16               UPCO   11    250       1     1.05     17          *    ZL2
    NOV 10       15   10   16               UPCO   11    242       0     1.04     17          *    ZL3
    NOV 17       15   10   16               UPCO   11    246       0     1.04     18          *  A-CLST 19 AN000632
    NOV 24       15   10   16               UPCO   15    239       0     1.00     17          *  Z-CLST
    DEC 01       15   10   16               UPCO   15    245       1     1.10     18          *    CBH 14
    DEC 08       15   11   16               UPCO   11    265       0     1.00     19          *    SBH
 * = STUDY WEEK DELETE APPLIED.----# = CKTS REQD = I-SVC - 1.----ADM PRD CNTL =          ----+ = EXCEEDS SUBGROUP CCS LIMIT.--
-----------------------------------------------------------------------------------------------------------------------

                        79          80
 LAST                       19NOV     28JAN     07APR     16JUN     25AUG     03NOV
 STDY-WK                 15OCT . 24DEC . 03MAR . 12MAY . 21JUL . 29SEP . 08DEC                   REMARKS
 (MONDAY)      CIRCUITS    .    1    .    2    .    3    .    4    .    5    .    6    .
 OF STDY-PRD            12345678901234567890123456789012345678901234567890123456789012345
                     25
                     24
                     23
                     22
                     21
                     20
                     19
                     18
                     17
                     16
                     15 SSSSSSSSSS SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS
                     14
                     13 RRRRRR  RR RRR
                     12       RR      RR                        RR
                     11                 RRRRRRRRRRRR R   R RRRRR  RRRRRRRRRRR          R
                     10                             R RRR R                    RRRRRRRRRR
                      9                                                       R
                      8
                      7
                      6
                      5
  NOTE:               4
  S = CKTS IN SVC     3                                                                          FORECASTING
  R = CKTS RQRD       2                                                                          TROYNY03MCO
  E = CKTS SVC=RQRD   1                                                                          PAGE      2
```

Fig. 4—Circuit Group Servicing Record.

more detail on administrative hour load conditions for the past ten study periods. The top right shows the significant hours identified for the current study period.

2.3.7 *Trunk group banding*

Based on observed blocking on final groups, and comparison of trunks in service with trunks required on high-usage groups, TSS assigns a "service band" to each measured group for each study period. The five bands identify groups that are overprovided (underloaded), possibly overprovided, close to objective, possibly underprovided (overloaded), and underprovided. The thresholds that separate these bands are adjusted for the statistical reliability of the data available for a trunk group. For example, small trunk groups with a few days' data available would be expected to have wide variations in their TSS-calculated load parameters. For such groups, the limits of the "close to objective" band are wider than for other groups.

In each study period groups judged to be overloaded are reported to the servicers. However, groups may appear underloaded simply because they have (currently) spare capacity needed for an upcoming busy season. Therefore, TSS only displays underloaded groups for corrective action if such groups have been underloaded for their entire banded history (up to 65 weeks).

The reports of overloaded or underloaded trunk groups from this function advise the users of significant service or utilization problems. Servicers must then determine what actions are appropriate, if any, or if only a transient condition in loads or measurements has occurred. The Circuit Group Servicing Record, as well as reports of weekly average load data, daily measurements, and measurement validation results, are available to support the servicers in this task.

2.3.8 Base selection

Just as TSS must select an appropriate current load from all the hourly loads in a study period to do trunk group sizing, so must TFS have an appropriate future load from all the hourly loads in a future forecast year to forecast future trunk requirements. The TSS base selection function executes a process similar to busy hour selection, starting with the hourly loads in a user-specified span of study periods. It outputs the most significant trunk group loads and hours in the specified base period, so that TFS can later estimate the loads in corresponding hours of future years.

Servicers can input Extended History Delete transactions to CU/TK to exclude a specified study period's data for a specified trunk group from the base selection function. This feature allows the user to prevent data that represents unusual traffic, or errors in measurement, from affecting the trunk forecast. The user can also replace any program-generated base load with a human-generated load before TFS is run, after reviewing the Base Selection Details Report.

2.3.9 Network optimization

After base selection has determined the significant trunk group loads from a twelve-month historical period, the TSS network optimization function determines a low-cost trunk network that should have been in place to handle these loads. Unlike the trunk group sizing function, which calculates the required size of each group for its actual, measured load, the optimization function estimates the theoretical load on overflow-receiving groups, to adjust for proposed trunk changes in the surrounding network, and sizes these groups accordingly. Optimization can therefore suggest widespread, related changes in trunk-group sizes but it does not suggest adding or removing trunk groups.

2.3.10 Trunk Administrative Measurement Plan (TAMP) reporting

In support of TAMP, the AT&T plan to index the performance of the trunk administration process, TSS produces reports annually from base selection and network optimization outputs. Each final group is assigned to one of the five bands previously described, according to its blocking in its annual busy hour. In addition, every group is assigned to a band, based on a comparison of trunks in service (in the annual busy hour) and optimized required trunks. The distribution of trunk groups among these bands is reported, by administrative responsibility and by trunk group type (end office, operator connecting, tandem, tandem connecting, and auxiliary services). Also, the TAMP reports list and summarize the number of days' data available to compute each group's busy hour load.

2.4 The Trunk Forecasting System (TFS)

The principal output of TFS is the General Trunk Forecast, a document that is a major input to the BOC construction program. A company-created forecast period table controls the span and complexity of the forecast. Through the table, the company defines the five forecast years, the first four of which must be consecutive and are normally chosen to be in the immediate future. For planning purposes, the company may choose a more distant fifth year, perhaps 20 years ahead.

The company also uses the table to define "base year subdivisions" and "forecast periods." These are partitions of the base (current) year and forecast years that tell TFS how much detail to create. Up to 20 forecast periods (normally quarters for each forecast year) and four base year subdivisions may be specified. TFS assumes conditions that exist on the last day of a forecast period exist throughout the period, so the importance of defining more than one forecast period per year is clear. Single annual periods would result in an improperly sized network if major events such as area transfers did not coincide with busy seasons, as frequently occurs. Although a company may opt for numerous forecast periods, it does so at the expense of computer run time which is proportional to the number of periods chosen.

The sections below describe the major functions of TFS, including user interactions, reports, and a high-level view of the calculations. Following these is a description of the sequencing of operations considering both TFS and CU/TK.

2.4.1 Growth factor computation

To forecast traffic load on trunk groups, TFS requires knowledge of the growth in load of the switches that the trunk groups terminate on. For each such switch, the forecaster enters growth data through CU/TK. The data are usually in the form of main stations (MS) and hundred call-seconds per main station (CCS/MS), as of a given date. The data may represent all customers served by the switch, or a subset of the customers according to NXX or class of service, e.g. coin, residence, etc. The data are entered normally by year to cover the base year and all forecast years.

TFS uses these data to compute growth factors in the form of ratios of future load to present load. First, TFS assures that data are available for each base year subdivision and forecast period end date; when the data have not been expressed to coincide precisely with the needed dates, TFS derives what it needs by linearly interpolating from surrounding data. Next, the data are aggregated to the appropriate level for forecasting trunk group loads. For example, if the forecaster input data by NXX but required growth factors for only the total switch (comprising several NXX's), then the NXX data would be summed for the entire switch.

Products of the like-dated individual growth data terms are then formed, e.g., MS and CCS/MS are multiplied to form CCS. Finally, by simple division of a future product by a base year product, growth factors are produced. For each switch, one growth factor is computed for each forecast period relative to each base period (as defined by the base year subdivisions). The forecaster has the option of bypassing this process by manually stating growth factors, if computed ones are inappropriate.

With calculations complete, the system outputs a Growth Factor Report. At this point, errors in the input data are most evident. To avoid the need to rerun the process with new inputs, the system allows the forecaster to substitute new growth factors for the incorrect ones through CU/TK.

2.4.2 Network disassembly

TSS base selection provides TFS with trunk group loads that represent busy hour and busy season conditions during the base year. TFS must convert these loads to a form appropriate for the projection function. For the most part, this amounts to removing overflow traffic from the loads of groups that are alternate routes for other (high-usage) trunk groups. This "disassembly" is necessary, since the overflow is a function of a previous network structure and must not be projected with the (network-independent) first-route traffic.

For each high-usage group, TSS base selection identifies the load offered to the group and its overflow. Using the alternate route information on Circuit Group Records, TFS associates the overflow with the appropriate alternate route groups. (When the overflow is fragmented over multiple alternate routes, TFS allocates a portion of the overflow to each route according to user-specified percentages.) TFS then subtracts the overflow values from the offered loads of the receiving group.

Because of differences in data collection schedules and measurement errors, a group's received overflow may appear to exceed its offered load. In such a case, TFS will assume a 0 value for the first route load. If the forecaster anticipates this and concludes that a 0 value is an understatement, the forecaster may input a "minimum primary CCS" that will override any lower computed value.

2.4.3 *Projection*

With first route loads and growth factors as input, TFS estimates future load. The first step is to compute projection ratios. These are created by substituting the growth factors into a formula that the forecaster has selected for each group. For example, if the AT&T recommended formula $(A + Z)/2$ were selected, the projection ratio would reflect the average growth of the originating end (A-end) and terminating end (Z-end) of the group, as determined by A's and Z's growth factors.

Of the total set of growth factors available for a specific A and Z, the ones used in the formula are those developed for the base year subdivision and forecast period containing the measurement period (busy month) of the base load. So if a trunk group had a base load with a May busy month, and the base year and forecast years were calendar years with a quarterly base and quarterly forecast periods, then the growth factors used would be those that apply to the second quarter of each forecast year relative to the second quarter of the base year.

The multiplication of the base loads by the projection ratios completes projection. In a similar manner, the loads stated on traffic transfer records are projected.

The system allows several options that the forecaster can exercise in advance of a projection run. Instead of relying on growth factors, the forecaster may state manually derived projection ratios or specify an annual percent for compounding. The forecaster may supply a stimulation factor or select a projection formula that accounts for community-of-interest considerations or regional growth. For a specific A or Z, the forecaster may request the use of growth factors that apply to only a part of the switch. Using "pseudo" trunk groups, described later, the forecaster may project separately the individual traffic items on a group. Finally, if all the system-provided methods are inappropriate, the forecaster may state future loads that are externally derived.

2.4.4 *Sizing*

Disassembly, projection, and sizing are run as a single process; the forecaster has the opportunity to review all calculations only after required trunks are computed. The same blocking objectives and economic criteria are used in TFS sizing as in its TSS counterpart. In general, TFS computes one trunks-required value per group per forecast year based on the load in the forecast period that contains the group's busy month. For office(s) affected by major event(s) such as cutovers, however, the forecaster may specify that required trunks be computed as of the date of the event(s). This will result in more than one trunks-required computation per year (except where busy months and event dates fall in the same forecast period).

As a first (logical) step, the process adds together (1) projected first route loads, (2) projected traffic transfer CCS values (+ or −), and (3) user-stated load adjustments (+ or −), if any. Transfer loads are needed since projected trunk group loads alone may not compensate for the effects of reroutes and new or deleted groups that occur after the base year. Next, the process sizes those groups that receive no overflow (only-route and primary high-usage groups). Overflow from the primary high-usage groups is then computed.

The rest of the procedure is the inverse of disassembly. Overflows are added to the offered loads on the alternate routes. When a group receives the overflows from all expected sources, the process computes the peakedness of its load and its required trunks. If the group is high usage, overflow and variance are computed. The previous steps are then repeated. The result is that the network is sized iteratively, bottom-up through that portion of the five-level network hierarchy included in the company's database. The process will add to the computed size of a group a trunk adjustment (+ or −), if any is stated; for high-usage groups this is performed before overflow is computed.

To facilitate more economic purchasing of trunk equipment, TFS will convert the above values to modules of 12 or 24 trunks. The forecaster, however, must request this action on a group-by-group basis. Which of four modular sizing procedures is invoked depends on whether the group is high usage, one way or two way, or uses digital facilities or digital terminal equipment. Modular sizing takes place before computing overflow from high-usage groups.

A major function imbedded in sizing is the determination of where new trunk groups are warranted. Although TFS mechanizes the test for new groups, the forecaster must identify candidates in advance by creating what are called "pseudo" Circuit Group Records in CU/TK, complete with base loads. To test fully all possibilities, the forecaster should input large numbers, perhaps thousands, of pseudo groups. But the practical limitations of deriving base loads and specifying all the necessary parameters reduce the number created to a small subset of those possible.

TFS projects the pseudo trunk group base load as if the pseudo group were an actual group. The sizing process calculates required trunks for it and determines whether a minimum-trunk-group-size threshold is exceeded. If so, the pseudo group is retained as a planned group (one that will appear in some future year of the forecast). If not, the pseudo group's loads are distributed onto existing groups specified by the forecaster. This is necessary since the pseudo group's loads represent actual traffic that is not accounted for elsewhere and would otherwise be lost.

Apart from planning new groups, pseudo groups give the forecaster more flexibility in projecting traffic. The forecaster may remove the load of one or more traffic items from an existing group's base load, define pseudo groups for them, and project them by different methods. If the forecaster indicates that the pseudo groups were created solely for separate projection, sizing will suppress the minimum-size test and add the pseudo groups' projected traffic back to their existing route(s).

The principal user output from sizing is the Preliminary General Trunk Forecast (PGTF). In addition to trunks, it displays several of the major intermediate calculations: base and future offered CCS, projection ratios, and peakedness. For each group, the report indicates, among other things, whether it was a pseudo, is modularly engineered, received overflow, or includes transfer loads. At this point, the forecaster may revise only future offered CCS and required trunks, as none of the other items appear on the final TFS reports. Since the analysis may be complicated, the forecaster may request up to three additional reports. The Forecast Detail Report, Detail Traffic Transfer Summary, and Overflow Summary supplement the PGTF with the information their titles imply.

With revisions complete, TFS produces its final reports. The General Trunk Forecast (GTF) may be issued in several forms, combining the newly computed trunk values with absolute differences or percent differences from previous forecasts. Figure 5 shows a GTF. The Office Trunk Studies display load and trunk values in a way useful for central office equipment engineering, and the Construction Program Worksheets provide a starting point for preparing AT&T construction program reports.

TFS also produces a computer tape of forecast results that is input to another Bell Laboratories product, the Facility and Equipment Planning System (FEPS), and BOC-developed programs. FEPS is a module of the Trunks Integrated Record Keeping System (TIRKS).

2.4.5 Sequencing with CU/TK

Figure 6 shows the relationship between CU/TK and TFS in a complete forecast run. The diagram refers to two TFS functions that have not been described to this point. Growth Factor Analysis (GFA)

```
TOTAL NETWORK DATA SYSTEM                     TNDS - TRUNK FORECASTING                     FORECAST RUN DATE 03 31 81
TRUNKING SYSTEMS                                                                           PROCESS DATE       10 09 81
ACCEPTANCE TEST NETWORK                  ORIGINATING GENERAL TRUNK FORECAST                BISP REPORT           TW530
SERVICING                                                                                  RESP CODE              EF30

GLLD PA GN CGO   1ES   CLASS 5

                                    ORIG PROB RQD BASE                  REQUIRED TRUNKS               RMKS SERIAL
TO                 TRUNK TYPE       UNIT ECCS MO  TRKS  81/82  82/83  83/84  84/85    85/86 CODE NUMBER  OTHER CGID
....           ................... .... .... .. ....  .....  .....  .....  .....    ..... .... ....... ...........
PHLA PA SL 42T M- IB 5- DD            11   02                          98    101      102  DC AA074262 00200 GLLD
           42T M- DF 54 PA           010   02     7      7      7      11     11       11  DC AA084304 00200 GLLD
           42T M- DF 54 TO CR        010   03            2      2       2      2        2  DC AA092097 00200   GLLD
           42T M- DF 54 SP NC2     W010H   01                          19     19       19  DC AA092380 015C1 1B006D
           42T M- DF 54 SP CN3     W010L   01                          60     61       61  DC AA092382 20701 1B006D
PHLA PA TH CGO M- PH 55 IE            07   02    11     12     13                          XW AA055919 00200 00410
           CGO M- PH 55 IE            12   02                          21     21       21  XW AA055919 00200 00410
PHLA PA WV CGO M- PH 55 IE            07   02     6                                        XW AA055920 00200 00410
           CGO M- PH 55 IE            14   02            9     10      18     19       19  XW AA055920 00200 00410
RDPK PA RP DSO M- DF 55 IE         W010L   11                                          59  XW AA094563 00200
           52C M- PH 55 IE            05   02    29     24     26                          XW AA055921 00200 B61M
           52C M- DF 55 IE         W010L   02                          58     59     59 X XW AA055921 00200 B61M
           52F M- PH 55 IE            09   02                           8      9      9 X XW AA055922 00200 B164M
SPFD PA SF DSO M- DF 55 IE         W010L   09                                          87     AA094565 00200    SES
           MGO M- PH 55 IE            08   02    28     30     32                          06 AA055923 00200 B220M
           MGO M- DF 55 IE         W010L   02                          88     88     88 X 06 AA055923 00200 B220M
           MG1 M- PH 55 IE            08   02    17                                        03 AA055924 00200 B60M
           MG1 M- DF 55 IE            08   02           13     15                          03 AA055924 00200 B60M
           MG1 M- DF 55 IE         W010L   02                          41     42     42 X 03 AA055924 00200 B60M
WAYN PA LA 42T M- AF 51 DD         W010H   02    77     61     67     145    149      149  DC AA062071 00200 GLLD
           42T M- IF 51 DD            25   03    27     63    131     138    140      140  DC AA067372 00200 GLLD
           42T M- DF 51 DD CR        010   03     2      2      3       4      4        4  DC AA073149 00200 GLLD
           42T M- DF 51 SP CN3       010   01                           2      2        2  DC AA085604 20701 1B006D
           42T M- DF 51 SP NC2       010   01                           3      3        3  DC AA085805 015C1 1B006D
WAYN PA WY CGO M- PH 55 IE            14   02                           4      4        4  XC AA082120 00200  00406
WCHS PA WC CGO M- PH 51 TE            14   02                          11     11        9  XC AA086392 00200  00406
WLMG DE WL 41T M- DF 53 SP CN3     W010L   07    26     22     22      22 X                XC AA055886 20701 D779SP
           41T M- PH 53 DD            11   02     6      7      8       8 X                XC AA055929 00200 A75640
           41T M- DF 53 SP NC2     W010H   07    10     11     11      11 X                XC AA057168 015C1 D779SP

                       TOTAL INTRA BLDG TRKS    192    181    191
                       TOTAL INTER BLDG TRKS    611    648    756    1566   1569     1513
                         * = MANUALLY REVISED                                                       GLLD PA GN CGO
```

Fig. 5—General Trunk Forecast.

and Circuit Group Analysis (CGA) extend the validations of CU/TK to include interrecord error checks. The philosophy behind this sequence, referring to the numbered steps in the diagram, is as follows:

1. Computer runs of CU/TK and GFA are needed to establish or correct the database for this forecast view. The forecaster receives error reports and, with this step and the ones below, is allowed time to submit corrections.
2. CU/TK is run to accept the corrections.
3. CU/TK is run to allow corrections to erroneous inputs in Step 2, GFA is run to freeze the database, and growth factors are computed.
4. CU/TK and CGA are run to accept growth factor revisions and identify any remaining interrecord errors.
5. Similar to Step 2.
6. Similar to Step 3, concluding with all remaining calculations.
7. CU/TK accepts requests for additional analysis reports, issued by TFS.
8. CU/TK accepts the forecast revisions. These are merged with the preliminary results and final outputs are produced.

Rarely, however, does a company execute the full sequence as described. To do so, allowing for forecaster interactions, would require two or three months. Typically, the company will omit Steps 2 and 5 and will merge GFA and CGA into a single step. A company may also

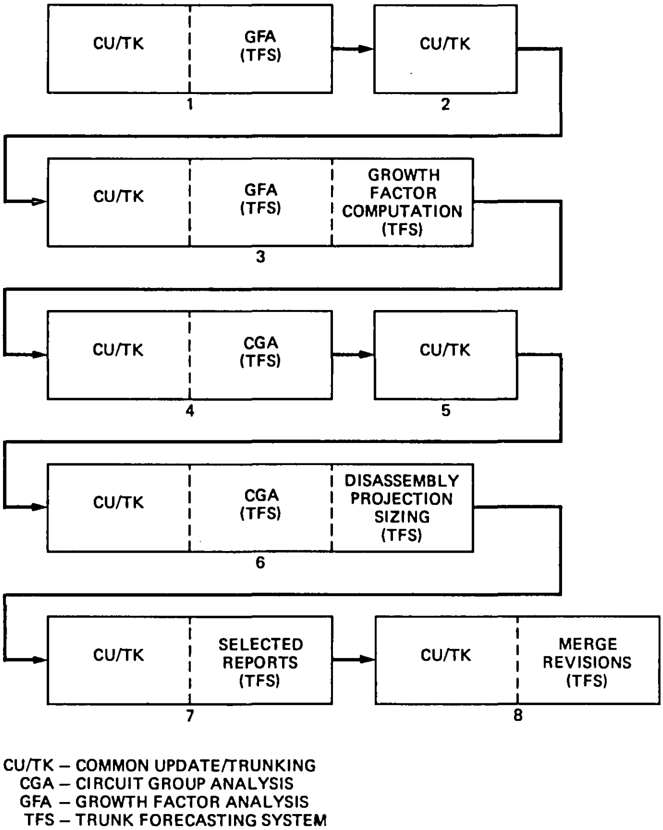

Fig. 6—Sequencing TFS with CU/TK. Arrows show the sequence of operations, not flow of data, and indicate pauses in computer run. During pauses, the forecaster reviews the intermediate output, prepares collections to the output and database, and inputs collections to CU/TK.

overlap the early steps of one forecast run with the concluding steps of the previous run. In some cases, too, only the forecast that pertains to a few switches is of interest. The company may then run TFS in the "Area Forecasting" mode, against only a selected subset of the database, and speed up the process accordingly.

III. SYSTEM INTERNALS

The TNDS trunking systems are designed to run in batch mode on an IBM 370-series computer system or equivalent, supported by the Bell System Standard Operating Environment software. The systems are coded primarily in COBOL, with some programs in FORTRAN or PL/I. Table I describes the size of these systems in several dimensions.

The source code and development documentation for these systems

Table I—Approximate system size

	Lines of Source Code	Compiled Source Parts	Delivered Programs	Data Files
CU/TK	34,000	25	14	15
TSS	77,000	100	58	170
TFS	85,000	110	68	135

are created and maintained using the *UNIX**/Programmer's Workbench (PWB) operating system and text-processing facilities on a minicomputer system. Multiple releases of a system, in simultaneous use at different installations, are maintained conveniently using the Source Code Control System. The Change Management Tracking System monitors the status of Modification Requests through various phases of investigation and resolution.

The programs are converted to executable format, tested, and transmitted to the users' computation centers by use of a telecommunications software system for computer-to-computer data exchange with operating telephone companies (TTRAN).

IV. LOOKING TO THE FUTURE

4.1 With TSS

With a planned trunk demand servicing policy feature, TSS will more effectively separate service problems that require corrective action from transient conditions associated with normal fluctuations in load. This feature will also select the best groups for corrective action within overloaded clusters, subject to user-input facility constraints. A second feature, providing a short-term forecast of loads, will then allow TSS to recommend cost-effective solutions to service problems before they occur.

An extensive revision of the TSS load estimation and study period formation functions will provide more accurate load data and trunk requirements through (1) calculation of loads from individual rather than averaged measurements, (2) inclusion of more days' data in weekend study-period averages, and (3) greater use of calculated parameter values (such as holding time) in preference to user-estimated or theoretical values.

TSS reports will be restructured so that a lesser amount of data is output automatically to the servicers. Servicers will be able to receive more detailed supporting data on request. Currently, TSS is designed to produce all of its reports automatically, thereby providing servicers with an excess of data from which they must extract the information of interest.

* Trademark of Bell Laboratories.

4.2 *With TFS*

Over the next few years, significant attention will be devoted to improving the user interface for forecasting. There are major implications in this for both TFS and CU/TK. One thrust will be to replace a large portion of the batch operation with on-line capability. Candidates for on-line use include database update, validation and inquiry, report retrieval, revisions to calculations, and the calculations themselves.

Another major enhancement that would be particularly useful in an on-line environment is a "what if" capability. Currently, TFS produces a single forecast with a database that gives a single consistent (though evolving) picture of the network at any given time. A "what if" capability would allow the user to produce several forecasts, each for a different switch and homing configuration. The best forecast, by criteria to be defined, would result in a database update with the preferred configuration.

The companies will also be able to exchange network information in a more mechanized way than is now possible. The exchange will include CU/TK records, intermediate TFS calculations, and forecasted trunks. To do all three will require the companies to coordinate their forecast schedules. The design of this feature will be robust enough to apply to the current network partition or to other partitions required by Dynamic Non-Hierarchical Routing (DNHR) or legislation.

Other proposed enhancements include the Sequential Projection Algorithm (SPA) and a Trunk Implementation Plan (TIP). SPA will replace the current projection method with one based on Kalman filter prediction theory. TIP is a method to convert the demand forecast from TFS into an administered one, considering forecast uncertainty, expense and capital costs, trunks in service, and facility availability. DNHR is also under study as it applies to TFS.

AUTHORS

Paul V. Bezdek, B.S. (Mathematics), B.A. (Physics), 1969, Wayne State University; M.S. (Computer Science), 1976, Steven Institute of Technology; Bell Laboratories, 1972—. Mr. Bezdek is currently Supervisor of the COEES/DOPS Systems Engineering Group. Previously he was responsible for developing functional requirements for the Total Network Data System/Trunking Systems. Member, Phi Beta Kappa.

Julius P. Collins, B.S., M.S. (Mathematics), 1974, Polytechnic Institute of New York; Bell Laboratories, 1974—. Mr. Collins has worked on design, implementation, and testing of several operations support systems for engineering and administration of the public telephone network. He is currently working on an engineering system for corporate telecommunications networks. Member, Pi Mu Epsilon.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Central Office Equipment Reports for Stored Program Control Systems

By R. F. GRANTGES,* V. L. FAHRMANN,* T. A. GIBSON,* and L. M. BROWN*

(Manuscript received February 16, 1983)

Stored Program Control System Central Office Equipment Reports (SPCS COER) is one of the earliest of the family of Operations Systems. Deployed on a centralized time-shared system, it produces engineering and administrative reports for Stored Program Central Office Equipment. This paper describes the user needs that motivated the design of SPCS COER. It traces the development of the system from a joint experiment of Bell Laboratories and the New York Telephone Company's Manhattan Engineering Department to the present system serving over 2800 electronic central offices from three centralized Amdahl V6 computers. We explain the functional system design from the user's point of view and outline the structure of the software, emphasizing the "tools" approach that is central to the development. The project management featured a research and development style, which we call the "whole-job" approach and discuss in some detail. Finally, we explore possible future directions for the project.

I. INTRODUCTION AND HISTORICAL PERSPECTIVE

When the first 1 *ESS*† electronic central office cut over into service on May 30, 1965, at Succasunna, New Jersey, the practice of traffic

* Bell Laboratories.
† Trademark of Western Electric.

data collection took a large step forward. Peg count, overflow, and usage measurements were made by the stored program control of the switching machine and kept in memory until their scheduled printout on a remotely located traffic teletypewriter.[1] No longer was there any need for the old electromechanical traffic registers, for clerks to read them, or for cameras to photograph them. Since the traffic teletypewriter could easily perforate a machine-readable paper tape with the measurement data, the tapes could be collected and read into a computer whenever convenient. It seemed clear to AT&T and Bell Laboratories planners that telephone company written computer programs would soon be processing the measurements into statistical information that would be unprecedented in accuracy, reliability, and timeliness.

1.1 Traffic data needs

Central office traffic measurements tell a Bell Operating Company (BOC) what is going on in the switch and the network. The company needs to know about call load, customer service, and equipment utilization. Information needs vary from short to long term. Network administration people are responsible for ensuring that valid, complete, and timely data are properly collected and reported for administration, business, engineering, maintenance, and other purposes. If something is wrong with the data, or the switching office, the administrators need to know quickly so they can fix it quickly, thus minimizing the effect of the problem on data and service quality.

Network administrators use the data to trend and balance loads and determine exhaust dates for central office equipment components (i.e., dates when the expected call loads will exceed the traffic capacity of the installed equipment). Network design personnel have more long-term needs, using the data to determine equipment capacities and the size and timing of future office additions. Marketing people help determine how well customers' needs are met by existing equipment arrangements.

Because telephone use varies widely (even wildly) with time, most data needs are best served by statistical information (e.g., averages) rather than by individual measurement readings. This means arithmetic operations (data processing) must be done on the measurements collected. It also means that—unlike, for example, the banking industry—a few missing measurements may not be serious. Some data may need to be excluded from the averages to keep the results representative. A good example of this was the day it snowed in Miami in 1977. Or the day a telephone building caught fire in New York City. Or the time when an equipment trouble caused all calls to certain areas to be set up so that no one could talk. Thus, data cannot be accepted uncritically but must pass the test, "Is this measurement both accurate and representative?"

To meet these needs in the years before electronic switching had been difficult and labor intensive. Data systems were semiautomatic at best. For example, measurement registers (mechanical counters that resemble automobile odometers) could be grouped together in large arrays and photographed periodically by automatic cameras. But, after development, the film had to be manually read and keypunched before it could be used by data processing machines. It was particularly difficult to obtain timely information in quantity. Thus, the electronic switching systems' promise of an economical, completely mechanized collection and data processing capability was truly an important and exciting step forward. Increased mechanization was going to permit improved information quality and reduced processing time and save a lot of manual effort as well.

The expected increase in mechanization was not intended to permit force reduction but to prevent a large force increase. Electronic switching systems were going to produce a lot more traffic measurements than previous electromechanical systems. This was because of the increased complexity of the switch (requiring new measurements), the ease and low cost of making the measurements, and the plans of the Bell System to offer new and special services using the capabilities of the stored program control. Indeed, within 10 years a typical 1 *ESS* central office was estimated to be producing perhaps 18,000 register readings daily or some 4.5 million per year. As the Queen said in Lewis Carroll's *Through the Looking Glass,* it was going to take "... all the running you can do, to keep in the same place."

1.2 An experiment begins

By 1969, little of the expected telephone company development had taken place. A few companies were running programs on time-shared computers that rearranged the traffic measurements into neatly labeled columns. Many companies were using transparent overlays to locate a few key measurements in the teletype printout, and were then recording and processing those manually. Some 50 *ESS* central offices were in service and new offices were being installed at an increasing rate that would soon reach one a week.

In mid-1969, work to design an experimental traffic data processing system for 1 *ESS* offices was begun as a joint project of Bell Laboratories Traffic Studies Center and the Manhattan Engineering Department of the New York Telephone Company under the auspices of the Traffic Division of AT&T. The system was experimental because, from the outset, the intention was to use technology radically new for the time in an attempt to build a practical and economical data system uniquely responsive to the needs of the telephone company users. Those needs, of course, were incompletely understood at the time, a further reason for considering the system experimental.

After two years of initial development, overlapped with about a year of trial use in several companies, AT&T in October 1971 announced a new time-shared computer program to process and manage traffic measurement data in 1 *ESS* offices. The new system, "PATROL" (Program for Administrative Traffic Reports On Line), provided an uncommon combination of features of significant help to the network administrator and network design engineer in the management of data. The announced features included the following:

1. On-line access to large quantities of traffic data.
2. Daily, weekly, monthly, high day, and busy season reports on demand in nearly real time.*
3. Automatic selection of high-day data.
4. Built-in data validation and exception flagging capabilities.
5. Single end-user control of traffic data entry, data retention, report generation, and data management in general.
6. A forgiving, interactive user dialogue that makes it easy to use with minimum computer knowledge and little or no formal training.

(Later experience proved two more features important:)

7. Reliable operation with rapid, low-cost system maintenance and feature development.
8. No hardware investment required of the user community for the experimental system that used only widely available teletypewriters with paper tape readers.

PATROL was initially implemented in a combination of FORTRAN and EXEC language on an IBM 360/67 computer operating under a modified Cambridge monitor. The vendor, Computer Software Systems, Inc. (CSSI) of Stamford, Connecticut, was chosen for the unique capability of its service at the time. There were no plans to make PATROL portable.

1.3 The first system

After initializing the system with a detailed description of a particular switching system, a network administrator would dial up the CSSI machine over the regular telephone network and transmit hourly ("Block H") traffic register counts from the paper tape punched by the traffic teletypewriter. PATROL would immediately make data validation tests, and items exceeding predetermined tolerance limits would be called (flagged) to the administrator's attention, along with

* In this context, "near-real time" refers to reports obtainable with little delay (perhaps minutes) after the end of the period being measured.

a brief summary of office performance data. The measurements would be accumulated in history files of data on an hourly basis. As many hours of daily data as desired could be processed, with costs increasing in proportion. Because of cost, most offices processed only one or two hours of data a day.

The history files were arranged for the generation of reports on a daily, weekly, and monthly basis, or for any desired span of dates at user request for each individual switching office. The reports were usually printed off-line and mailed to the administrator, but when time value justified higher cost, the reports could be printed on-line at the administrator's (or network design engineer's) terminal.

Probably little in PATROL was startling or new from a computer science point of view (a term just then beginning to come into vogue). What was new was the application of the technology to assure maximum capability and ease of use for end users with little training. A single time-shared machine serving thousands of offices spread over a continent was a concept still considered a bit adventurous (perhaps even foolhardy) in the early 1970s, although there were several precedents in the 1960s.[2,3] An on-line, interactive development facility was still a wonderful novelty for program designers and, of course, proved efficient and cost effective.

One design approach that was not recognized as new at the time was to provide each switching office with a database distinct from every other office and to conduct almost all processing as though only the one office and the one administrator existed. The (perhaps concurrent) processing of other offices for other administrators was generally hidden in the design, it being left to the operating system to sort out. This single-office approach greatly simplified development and maintenance by reducing the complexity of the software compared to other Total Network Data Systems (TNDSs) [e.g., Traffic Data Administration System (TDAS) and No. 5 Crossbar Central Office Equipment Reports (5XB COER) described elsewhere in this issue] whose designers had to plan for the sequential processing of perhaps hundreds of different offices in the same batch run.

1.4 Growth and evolution

In the years that followed the initial introduction of PATROL, most of the original features were preserved, while many features and systems for other types of electronic switching systems were added. At the end of 1971, 27 offices in 9 companies were using PATROL.

During 1972, the participation of the New York Telephone Company was gradually phased out and Bell Laboratories continued the experiment, working closely with AT&T and representatives of the BOC network administrators.* Many report types were added, an on-line documentation feature, Lessons, was produced, and processing costs were reduced. By March 1973, there were more than 165 offices on the system.[4] In subsequent years, versions of PATROL were developed for 2 *ESS*, 3 *ESS*, Remote Switching Systems (RSS), Voice Storage System (VSS), and 5 *ESS* machines.[5]

When an in-house (AT&T) computer service offering nearly the same software features [Virtual Machine/Conversational Monitor System (VM/CMS)] as our original vendor became available in 1974, PATROL was the first large project to move to that facility. It has remained there since.

In 1975, when the Engineering and Administration Data Acquisition System/Traffic Data Administration System (EADAS/TDAS)† data collection portion of TNDS was available in most BOCs, PATROL developed a major new feature called Batch Data Entry (BDE). Data could be collected by EADAS, sent by magnetic tape to TDAS, and relayed from TDAS to the PATROL computer over the high-speed corporate data network (see Fig. 1). Compared to paper tape data entry, BDE significantly reduced computer expense at the sacrifice of two initial PATROL features. A single end user could no longer control the whole system alone, as coordination with EADAS, TDAS, and the corporate data network was required. And, because of the delays inherent in the new batch process, results from PATROL could no longer be obtained in "near-real time." However, because of the data processing and reporting capability of EADAS, the near-real-time feature of PATROL was no longer needed where BDE could be used.

In 1977, a major rewrite of PATROL changed its basic orientation from processing clock hours for all office components to processing selected busy and side hours for each office component. This component busy hour feature was followed in 1979 by a Busy Hour Determination System. By the end of 1978, there were more than 1600 offices using PATROL.

1.5 The experiment ends

During its early years PATROL was enthusiastically received by network administrators and network design engineers. However, upper BOC management, often advised by computer experts that time shar-

---

* As has often been our experience, the collaboration with a BOC, here New York Telephone, had been richly rewarding. The "real world" expertise and "get it done now" attitude of the New York people really helped move the project along. By the end of 1972, however, the on-line system was providing feedback from the whole country.

† EADAS is the subject of the article "Data Acquisition and Near-Real-Time Surveillance," this issue.

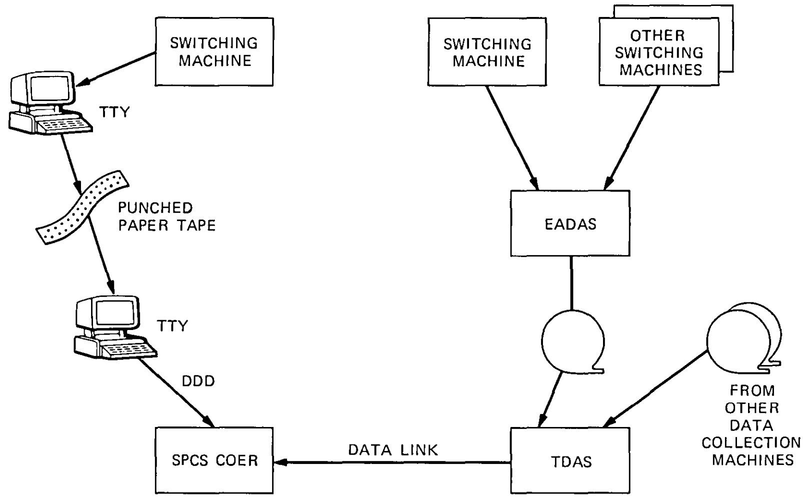

DDD – DIRECT DISTANCE DIALING
EADAS – ENGINEERING AND ADMINISTRATIVE DATA ACQUISITION SYSTEM
SPCS COER – STORED PROGRAM CONTROL SYSTEM CENTRAL OFFICE EQUIPMENT REPORTS
TDAS – TRAFFIC DATA ADMINISTRATION SYSTEM
TTY – TELETYPEWRITER

Fig. 1—On-line and batch entry of data in SPCS COER.

ing was always more expensive than batch, feared that perhaps the system was too expensive. In 1976, these fears were laid to rest by two events. A major survey of computing costs showed that batch operating costs were much higher than was widely believed—sometimes 10 times higher. Finally, an extensive multilevel, multidepartment survey of every BOC, including three that at that time were building or operating their own traffic data systems, was conducted. It found that the PATROL system, of all other candidates, had the best match of features to the future data needs of the BOCs. With these and other studies concluded, it was finally clear that the PATROL experiment could be considered at an end. Continuing standard development under the Business Information Systems (BIS) funding agreement was begun in January 1978.

1.6 Continuing development

The BIS Advisory Committee asked that the name of the system be changed from PATROL to *ESS* COER to better agree with other TNDS component system names. This was later amended to SPCS COER when we realized that the system would be used for Stored Program Control System Reports not involving switching. However, because the term PATROL had become extensively embedded in Bell

System literature and had been used by the BOCs for many years, and because two syllables are much easier to say than six,* the term PATROL is still frequently used by the field forces.

The change from experimental to standard status made it possible to emphasize system design considerations in SPCS COER over simple feature availability because long-term planning and investment of resources were now warranted. The original version of PATROL, with substantial amounts of interpretive code, had been replaced by a much more efficient FORTRAN version by the end of 1972. As an experimental system, the software emphasis was on achieving features quickly, while steering a prudent course between run-time inefficiency and high performance. With the prospect of an assured future, the development effort began to emphasize new approaches to the design in 1977. A modular, generic-tools approach would achieve both high performance and reduce development and maintenance effort in the long term. This new approach is described in Section III.

In 1979, a new experimental feature was designed for AT&T. The Data Profiles feature consolidates selected information from all the individual 1/1A *ESS* office databases for the whole Bell System. Elaborate retrieval, computation, and reporting facilities are provided in an English-like language for the use of personnel engaged in studies of various types.

During the 1980–81 busy season (January to March), SPCS COER began to strain the resources of its host (Amdahl V6) computer, with the traffic data load from nearly 2400 *ESS* machines. Therefore, in September 1981, after about a decade of continuous operation on a single computer, the user community was divided up among three computers (at the same location).

In December 1982 a new feature was introduced that allows printing of reports in the user's local company data center. In addition to existing options that allow reports to be either printed on-line or off-line and mailed, the user may now have the reports sent to the nearest BOC data center over the corporate data network or via magnetic tape. This new option was provided to allow users willing to incur additional delay in the distribution of reports the opportunity to reduce centrally billed expense for SPCS COER by as much as one third.

During its first decade of operation, SPCS COER compiled an excellent record of availability. While few detailed records were kept, the longest recorded nonscheduled total system outage (no access to reports) was about six hours. Some individual users or BOCs experienced longer outages while software errors were being repaired.

At the end of 1982, there were more than 2800 offices using SPCS

* The preferred pronunciation of SPCS COER is S-P-C-S KO-AIR.

COER. The program had grown to 225,000 lines of FORTRAN code with an additional 105,000 lines of comment. There were about 1000 pages of on-line user lessons for all the component systems. Users were logging on to the system at a rate of about 25,000 per month and issuing about 1,000,000 commands a year. About 3,000,000 hours of data were being handled annually and 1.2 million pages of reports produced per month. One third of the report output is produced on fiche, the rest on paper.

II. FUNCTIONAL OVERVIEW

This section presents a functional overview of SPCS COER from the user's point of view.

SPCS COER is really a family of on-line data management systems that process traffic data from the class of switching systems called sorted program control systems. This includes 1/1A *ESS,* 2/2B *ESS,* 3 *ESS*, 5 *ESS* and RSS switching entities.

Each switching entity, or office, served by SPCS COER has its own distinct, independently managed database. This implies separate data collection, reporting, and data management for individual offices. This single-office view is incorporated in the software design and the user interface. It reflects the usual approach taken by network administrators and network designers in examining central office traffic data. A consequence of this design is that most of the SPCS COER systems cannot summarize data across multiple entities. There is, however, one SPCS COER system, described in Section 2.6, that provides data summaries across offices.

2.1 User interface and documentation

A paramount goal in designing the SPCS COER systems is that they be easy to use by network administration clerks as well as managers. An interactive English-like dialogue, on-line documentation, and a generally forgiving design help support this goal. Since SPCS COER is an on-line system, it gives the user direct control of the system functions. When something unexpected happens, for example, an invalid request is submitted, the user is notified and given an opportunity to take immediate corrective action.

The user interface takes the form of a question and answer type of dialogue. The user types a command in response to a REQUEST = prompter. Any parameters required before the command can be executed are then prompted for by the system. The user need type only a carriage return to ask for help in responding to a particular prompter. The system will then print a list of valid responses or response formats. This list is also printed if the user types an invalid reply. The user can type CLEAR in response to any prompter to abort the command and return to REQUEST =.

A complete set of on-line user documentation is available with each SPCS COER system. The user may access the documentation on request and have one or more sections of it typed at the terminal or printed off-line and sent by mail. This allows changes in the documentation to be made available to the users quickly and easily by the developers. It also assures the user of an easily accessible source of the latest information. The combination of readily changeable on-line documentation with a hot line open 24 hours a day (by answering machine at night) has proven effective in keeping the documentation error free and the users informed, confident, and satisfied. One finding of many studies is that software system problems are often due to the unavailability of up-to-date documentation at the user sites. SPCS COER totally avoids this problem by placing the most current information in the hands of the user whenever the user asks.

News messages may be printed on request as well. News items are put on the system by the developers to announce new features, revised documentation, or the existence of (or, preferably, the solutions to) newly discovered common problems. Users are alerted to news items by one-line flashes whenever they log on the system. Two levels of flashes and news items are used depending on the intended audience for the item—all users or only users of a particular component system.

Since costs of any operations support system need to be monitored, SPCS COER provides estimated expenditures to the user for all on- and off-line requests. This feedback has helped users manage their use of the system effectively.

An important part of the SPCS COER user interface is a network of "SPCS COER coordinators," one or more in each BOC. These coordinators directly assist the BOC users in operating and managing the SPCS COER system. Usually it is the company coordinator who contacts the SPCS COER staff over the hot line about problems. The coordinators help the users identify and rectify problems. Over the years, the coordinators have asked for and been given the tools to identify and correct system problems that initially only the PATROL central staff could fix.* These tools have been grouped into a separate interactive system, modeled on SPCS COER, called COMP for Coordinators' Management Package. COMP even has its own set of on-line lessons.

2.2 Office description file management

Before any SPCS COER system can process traffic data for an office, an Office Description File (ODF) must be established for that

* That is, when the generally forgiving design had stopped being forgiving to some user.

office. The ODF contains information describing the office and the traffic measurements that are to be processed. This information includes:

1. The size of the office
2. The equipment or service components of the office
3. Characteristics of each component
4. Measurements to be processed and components to which they correspond
5. Hours of traffic data to be processed for each component.

An SPCS COER user establishes, or activates, an ODF through an on-line request. The user must supply to COER all the information needed for the ODF. This information is used by the system to interpret, validate, and make calculations on the traffic input data and to store and retrieve processed data. Therefore, it is essential that the ODF be current and correct.

A validate function ensures that the ODF information is internally consistent and that the user-specified values fall within appropriate bounds (e.g., the information represents a valid *ESS* office configuration). Validation is automatically done when an ODF is first activated and whenever the information in it is updated. If errors are found in the ODF, the user is alerted and COER will not process any new traffic data until the errors are corrected.

The on-line update command allows the user to examine the ODF and change it. All changes are validated, and any problems immediately reported to the user on-line. The user can then fix any errors by issuing update commands. Commands are also available to print the contents of an ODF on-line or off-line, and to deactivate an ODF (remove the ODF and its associated database from the system).

2.3 Data entry

To enter data into the database for an office, an SPCS COER system must first read the raw data values put out by the SPCS on what is known as a traffic schedule. A traffic schedule contains measurement data for a particular office, the collection date, collection interval, and end-of-interval time. The system uses ODF information to associate measurement values with their corresponding components. Format checks are made on the values, and traffic statistics are calculated for each component based on the raw data values and ODF information. These processed component values are then stored in the database for the office.

As part of the data entry process, SPCS COER makes various data reliability tests. These tests are intended to pinpoint instances of traffic measurements that are invalid, for whatever reason. Invalid measurements are caused by equipment or software malfunctions in the switching machine or data collection system, highly unusual and unrepresentative traffic conditions, or (most of the time) errors in the ODF information. If data for a component fail one or more reliability tests, they are flagged, or marked as questionable, in the database. Flagged data are processed specially during data retrieval. These data are marked on all reports on which they appear and are excluded from report averages. The user has a recourse if data are inappropriately flagged, as described in Section 2.5. The results of the data reliability tests made during data entry are saved and made available to the user as exception reports.

Traffic data may be entered into SPCS COER in one of two modes, on-line or batch, as shown in Fig. 1. On-line data entry, the original PATROL data entry mode, involves using punched paper tape generated at the traffic teletypewriter connected to the switch. These tapes, containing the required traffic measurements, are transmitted to SPCS COER on-line over the switched telephone network via a terminal with a paper tape reading facility. The user must monitor the process while the tapes are being read. For large volumes of data, this is a tedious and time-consuming process that also involves substantial computer connect time charges.

The newer batch data entry facility reduces the cost and clerical effort associated with on-line data entry. It is now used for most of the offices served by SPCS COER. In designing BDE, the idea was that a user could set things up once, and then data would flow from the switch into SPCS COER almost unattended.

Batch data entry may be used by offices supported by EADAS and TDAS,[6] or their equivalents. Under this method of data entry, data for many offices and for several days are transmitted from a telephone company TDAS computer site to the SPCS COER computer site over the corporate data network. The SPCS COER system automatically places this data into a common disk storage area as pending files. An individual pending file is created for each block of data that corresponds to a particular office, date, and collection time. These pending files are organized on a per-user basis.

At night, an automatically scheduled job is run for each user who has pending files. This job processes the data from the pending files into the databases for the user's individual offices. It also generates messages for the user that indicate the status of the batch data entry job. These messages may be printed on request when the user next logs on to the system. In this way, a user knows exactly what data were entered into the database, or why any data were not entered.

Pending files that are not successfully processed are retained on the common disk area for two weeks. They are also available for 45 days from a backup tape. The user can obtain a listing of all pending files for a given office. Users are provided with on-line commands for managing these pending files. Pending files can be erased from the common area or restored from backup tape. Certain errors in the header information for pending files or in the traffic data values can be corrected interactively to enable the system to process those data.

2.4 Reports

To retrieve data from an office database, a user requests traffic reports from the system on-line. The user has the option of having them run on-line and typed directly at the terminal, or run in a batch mode, printed off-line, and mailed to the user.

There are two basic types of traffic reports produced by the SPCS COER component systems: daily reports, which produce individual daily data and data averaged over a given span of days; and Machine Load and Service Summary (MLSS) reports, which make available year-to-date summaries of monthly averages and the highest traffic days.

Daily reports may be requested in weekly, monthly, or intermediate form. Weekly and monthly reports provide data for all central office equipment components and all hours. Intermediate reports may cover any span of days and any selected components and hours. Each section of a daily report is devoted to one component in the office. It shows the daily statistics as well as the averages of the values over the report period. Flagged data are excluded from the averages. Monthly reports are useful to the network administrator for analyzing the overall service in the office, for trending office loads, and for data validation (as described in the next section).

The MLSS reports are used by the network designer in capacity determination and in planning for future central office equipment needs. An MLSS report contains, for each component and hour requested, data for the 15 high load days of the year to date, together with the average of the 10 highest unflagged days. Also included are the monthly averages for the year to date, together with the average of the three highest months. At the end of the data collection year, this last average represents the Average Busy Season (ABS) value for the component.

The formats of the traffic reports produced by the SPCS COER component systems follow those produced by the 5XB COER system wherever possible. The 5XB COER reports are covered in some detail in the article "Equipment Systems," in this issue.

The Busy Hour Determination (BHD) systems (1 *ESS* BHD and 2 *ESS* BHD) of SPCS COER produce a different set of reports from the ones just described. The BHD systems help the network administrator determine, for each component in an office, which hour has, on average, the highest traffic load. This is called the time-consistent busy hour for the component. A BHD study for an office is usually done once or twice a year, for two to four weeks. During that time, data are collected and processed by the BHD systems for 12 to 24 hours per day. The output of this study helps the network administrator determine which hours of data should be processed for each office component in 1 *ESS* COER or 2 *ESS* COER. There are five sets of reports produced by the BHD systems. They range from detailed hourly reports to a high-level summary report listing the suggested COER collection hours. Because RSS and 5 *ESS* switches are Extreme Value Engineered (EVE),[7] instead of time-consistent busy hour (or average busy season busy hour) engineered, they do not require busy hour determination studies.

2.5 Database management

SPCS COER makes available to the user several data management capabilities. There are on-line requests that tell the user what data are available in the database and allow the user to flag and unflag data, remove unwanted data, and initiate (or terminate) a data collection year.

The on-line flag and unflag requests allow the user to override the data reliability flagging decisions that were automatically made by the system during data entry. A user can selectively flag or unflag data for any component, hour, and date. This is usually done before requesting a monthly report.

After a monthly report is run, the monthly averages are stored with the MLSS data, and users have received their paper or microfiche reports, there is usually no need for SPCS COER to keep the detailed daily data for that month. The SPCS COER has a remove command that lets the user remove daily data for any span of days from the database. Data can also be removed selectively from the MLSS high day lists. A cycle function lets the user remove all MLSS data from the database at the end of the year.

The database management features, along with ODF validation, are key parts of the success SPCS COER has achieved in producing high-quality results. However, the degree to which user actions are required to manage the database might seem inefficient. Why not further mechanize these data management processes to require less user intervention? Let us look at why this user control is so important.

The user, usually a network administration clerk, is responsible for producing accurate and timely information. But it appears to be a fact of life (based on some 20 years of data processing experience) that bad data often does not become apparent until after it is used to produce reports. And every new type of report produced seems to uncover new instances of bad, or at least suspicious, data.

When this happens in SPCS COER, the user, after suitable investigation to learn the facts, simply flags or unflags data as appropriate and reruns the report. Only after the reports have been thoroughly examined is it generally safe to remove the data from the system. Data processing systems without these features often leave their users with no usable results and/or the task of reprocessing the data by hand.

It is also a fact of life that users make mistakes. While this may be regrettable, it is no excuse for producing poor results. Thus, almost every action that a user can take in data management can be reversed if the action proves to have been a mistake.

User control of the data management has a cost to the user in data processing and storage charges, of course, But if the user feels that exercising the recovery features is not worth the cost, then the feature is not used and the cost avoided.

There is an overhead cost in software for having some of these options available, even if they are not exercised by individual users. However, the system design has assumed that these are far offset by several types of savings that their presence makes possible. First, manual data processing by clerks is eliminated for this purpose. Second, results are sometimes achievable that could not possibly be obtained in other designs (for example, when a serious error is discovered some months after the data has been discarded and a large monetary decision hangs in the balance). Third, formal training expense is avoided by virtue of the new system design (e.g., users can learn by doing with minimal risk of data loss through user error).

Since data retention affects storage requirements (and costs), the system monitors storage usage and alerts the user when the amount of storage consumed exceeds a predetermined threshold. The system also provides the user with a storage command that reports the current storage consumption.

2.6 Data profiles

The functions and capabilities described so far in this section apply to most of the family of SPCS COER systems. However, one system, called Data Profiles, serves a special function. Data Profiles, unlike the other systems, is not restricted to a single-office view of data. It provides a facility for summarizing traffic data across many offices (and also across more than one service year).

This system collects preselected MLSS data from each SPCS COER office database and stores this information in a single hierarchical database. In a hierarchical database, data are divided, subdivided, and sub-subdivided in a treelike fashion as many times as needed to permit convenient access to its parts. Data Profiles divides and subdivides the MLSS data by BOC, type of *ESS* switch, particular office, and collection period (current year, latest complete year, and archived older years).

The data for an office are collected by Data Profiles when the SPCS COER user requests a monthly report for that office. The database currently contains data for all 1/1A *ESS* offices. Data exist for each year, starting with 1978. Users can access any or all these data on an on-line demand basis.

Data Profiles uses a generalized database system/tool called the Off the Shelf System (OTSS),[8] which provides data management routines and a query language. This query language allows Data Profiles users to retrieve any items or combination of items stored in the database. The query language statements are simple. They consist of a few verb and modifier expressions that specify the part of the database to be traversed and what action should be taken.

By typing a single OTSS sentence, a user can ask, for example, "What was the average busy hour CCS per main station for *Touch-Tone** dialing digit receivers in New York Telephone offices during 1980?" Users can set screens on data so that the system will retrieve only data that satisfies certain conditions. For example, including the phrase, "when average busy season percent capacity is greater than 90," in the above query would cause the system to include in the average only data for receiver groups that were operating at more than 90 percent of capacity. The query language has commands that distribute, plot, print, alter, and make statistical computations on the data. Because queries may become complex and may be useful to many users on many occasions, a query library is maintained by the system.

Data Profiles provides a powerful tool for studies since it contains data for an entire area and so can be used by telephone engineers for many different planning purposes.

III. SOFTWARE STRUCTURE

The first SPCS COER system was hard coded with data structures and algorithms specifically designed for the 1 *ESS* central office application. The resulting system, although remarkably popular, was also inflexible and hard to enhance. What might seem a small conceptual change (the component busy hour feature) required a major rewrite of the system. Even the seemingly trivial change to compress the batch reports seven spaces to allow for hole punching required an embarrassing amount of reprogramming. In this section, we describe a software design method that led to SPCS COER modules that are easier to develop and maintain.

---

* Registered service mark of AT&T.

3.1 The view from the back room, or a hacker's view of history

One motivation for this design philosophy was the rapid proliferation of SPCS COER subsystems (serving different switching systems) during the mid-1970s. The 2 *ESS* COER was the first offspring of the initial 1 *ESS* system. This system was produced by the venerable red pencil technique. Engineers poured over the ponderous listings of 1 *ESS* COER (60K lines) deleting and changing sections of code that were not quite right for the new application. The resulting system worked but became difficult to enhance. Then 3 *ESS* COER was produced from 2 *ESS* COER in much the same way. The prospect of making parallel enhancements in these systems was awesome in much the same way that the Okefenokee Swamp is awesome. The laws of genetics suggested that this inbreeding would have to stop or the species would become extinct.

The crux came in the summer of 1977 when the SPCS COER team was faced with the task of developing four new subsystems (Busy Hour Determination, Voice Storage System, Remote Switching System, and Advanced Communication System). It was clear that the red pencil technique was no longer tenable. It was not so clear how these ambitious systems could be produced without a large development effort and substantial duplication of code. Fortunately, this crisis coincided with the emergence of the software-tools approach to system development. This method advocates developing larger systems out of a kit of small general-purpose tools rather than building monolithic special-purpose systems. This philosophy was an enormous help in managing the large development task. Rather than rushing off in four separate development efforts, time was taken to define a set of tools that could be useful in all four systems. The result was a marked reduction in development time and lines of code. In addition, the resulting systems are simpler and easier to maintain.

3.2 Tools plus

As valuable as the tools approach was, it was not the whole answer to the problem. Each SPCS COER system contains a fair amount of domain-specific knowledge both about traffic engineering and about a particular switching system. Without further philosophical underpinnings, it would be difficult to capture this knowledge in simple, generically useful tools. What was needed was some way to factor this specialist's knowledge out of the tools. The goal was to make the tools simple enough to be flexible and yet sophisticated enough to be knit into a useful system. The solution was to capture the knowledge of traffic engineering in a data model that can be used by the tools and to capture the switch specific knowledge in a set of tables that drive the tools. As a result, the systems developer's job changed. Previously, it was to write or change large amounts of code to meet the requirements of the system. With the new method, it is to first select the proper tools for the system, then define the tables specifying the new system to the tools, and finally to write the controlling code which welds the individual tools into an effective system. The advantage is that the developer spends less time in low-level development and is free to devote more time to the design of an integrated system meeting the customer's need. In effect, with intelligent tools we have raised the developer's perspective from that of programmer to that of system designer.

The heart of the new approach is the traffic data model. The model consists of both a view of the traffic data and the routines that manipulate the data in accordance with this view. The traffic data that SPCS COER deals with come in two varieties—daily and MLSS.

3.3 Daily data manager

The SPCS COER Daily Data Manager (DDM) captures the notion of daily data in an abstract model. To DDM, an individual traffic measurement is represented by a tuple of five elements (day, hour, equipment tag, measurement tag, value). The quadruple (day, hour, equipment tag, measurement tag) is an index to the data values. Any measurement in an office's database can be accessed by specifying the identifying index.

The DDM provides standard database operations. The user of DDM can create and destroy databases, and add, delete, retrieve, and update items in an existing database. But DDM is a powerful tool because it understands the particular needs of an SPCS COER system. For one thing, it understands traffic data. It understands that the data are collected as a block of measurements for all the equipment in the office for a given interval. The data structures and algorithms of DDM are carefully tailored to efficiently organize and store data collected in this way. The DDM also understands that these measurements are subject to error and provides routines to flag questionable data items.

The DDM is particularly useful to SPCS COER systems because it understands how the systems are likely to use the data it manages. Report generation provides a convenient example. An SPCS COER report is a formattted output of the traffic statistics for a specific piece of equipment at a particular hour over a span of days. The DDM provides two special interface routines to aid in report generation. The first takes an hour, an equipment tag, and a range of days as an argument to prepare DDM to retrieve data for the report page in a systematic way. Each time the second routine is called, it returns a new line of data for the report which includes the day on which the data were collected.

3.4 MLSS data manager

The MLSS database contains the same measurements as the daily database, but instead of keeping a day-by-day account, the MLSS database contains rank-ordered lists of the busiest days and the monthly averages on the equipment in the office. The MLSS Data Manager (MLSM) is an SPCS COER tool for managing MLSS databases. It views an MLSS data item as a five element tuple (hour, rank, equipment tag, measurement tag, value). The MLSM provides SPCS COER systems with the same sorts of services for managing MLSS databases as DDM provides for daily databases. In particular, MLSM maintains the rank-ordered lists.

DDM and MLSM free the individual system developer from worrying about the complexity of data management. They take care of storing and retrieving data on disk and worry about such issues as concurrency control and recovery. But more importantly, they enforce a standard data model across all the SPCS COER subsystems. This has two important ramifications. First, since all the traffic data are stored in a common manner, they can be analyzed by common software. Second, since the data are stored by a common data manager, higher-level tools can be built that use DDM and MLSM to fetch and store the data.

3.5 RSL

The SPCS COER Report Specification Language (RSL) is a good example of a high-level tool that uses DDM and MLSM. It also illustrates how a system designer can customize a generic tool to fit specific subsystem needs. The RSL is a simple but powerful language which allows the user to specify the format of a wide class of SPCS COER reports. Figure 2 shows a sample RSL program and Fig. 3 shows a report that was generated from this specification.

A brief study of Fig. 2 illustrates how RSL works. The first section of the report specification is the heading. This is the information to be printed at the top of the report. Information in quotes is printed here as it is written on the report. The key words offnam, datnam, schnam, and stunam specify translations to be made when a specific report is printed. The key word offnam will print the name of the office for which the report is being generated. Key word schnam will print the hour for which the data were taken, stunam will print the name of the equipment, and datnam will print the Gregorian representation of the first or last day in the report.

The supertitles section presents titles that span more than one

```
daily report 2 for (s52, s53, s54, s55, s56);
        flag line on position 9;
        heading
            offnam;
            '';
            'NO. 1 ESS SERVICE CIRCUITS AND SPECIAL TRUNKS';
            'DETAILED DATA FROM', datnam (fstday), 'TO', datnam (lstday);
            'ALL COMPONENT HOURS';
            '';
            '';
            'ITEM NO.', %d stutag, ':', stunam, '          ', schnam;
            '';
        end;
        supertitles
              'PER STATION' 14-24;
              'USAGE' 29-33;
        end;
        columns
            datnam (curdat) 1-8%s'' no average;
            i1  12-18%6.4f* 'CCS' exclude on '*%&';
            i2  20-25%5.3f* 'CALLS';
            i3  26-31%5d* 'TRFFC';
            i4  32-35%3d* 'MB';
            i5  36-41%5d* 'PEG COUNT';
            i6  43-48%5.2f* 'HT SEC.';
            i7  50-54%4.2f* 'OVFL %';
            i8  55-58%3d* 'CAP %';
            i9  59-63%4d* 'NCI';
            i10  64-69%5d* 'CAP CCS';
            100* (i3+i4) / 36*i9  70-72%3d  'OCC %';
        end;
end report;
```

Fig. 2—Sample RSL program.

0000013H MLDNMAELCGO SP CTX-8 3.5 COMP. ID: 001
NO. 1 ESS SERVICE CIRCUITS AND SPECIAL TRUNKS
DETAILED DATA FROM 07/01/79 TO 07/09/79
ALL COMPONENT HOURS

ITEM NO. 54: T-T MEMRY P 10:00

	PER STATION CCS	CALLS	USAGE TRFFC	MB	PEG COUNT	HT SEC.	OVFL %	CAP %	NCI	CAP CCS	OCC %
07/01/79	0.0029	0.058	40		792	5.04	0.00	3	55	1256	2
07/02/79	0.0041	0.047	55		634	8.72	0.00	4	55	1256	3
07/03/79	0.0493+	0.047	66		643	10.26	0.00	5	55	1256	3
07/04/79	0.0077	0.052	105		705	14.89	0.00	8	55	1256	5
07/05/79	0.3333*	0.039	62		531	11.73	0.00	5	21*	1256	3
07/06/79	0.0070	0.066	2*		897	10.64	0.00	8	55	1256	5
07/07/79	0.0050	0.001*	68		547	12.36	0.00	5	55	1256	3
07/08/79	0.0054	0.051	73		693	10.59	0.00	6	55	1256	4
07/09/79	0.0060	0.043	81		581	13.96	0.00	6	55	1256	4
AVERAGE	0.0109	0.050	69		669	10.91	0.00	6	55	1256	4

Fig. 3—Sample report generated from an RSL program.

column in the report. The quoted string gives the title, and the numbers give the character positions over which the title should be centered.

The columns section specifies the main body of the report; each line in this section specifies a column in the report. The first entry is the name of the data item to be printed in this column and the numbers show the character position boundaries for the column. Information after the percent sign tells how the value should be printed (s = string, d = integer, f = floating point, n.n = how many total digits in the number and how many digits after the decimal point. Compare with the C language function, printf). The asterisk shows that the data item can be flagged, and the string in quotes is the column heading. The key words no average indicate that there should be no average taken over this column. The key words exclude on followed by a list of flags indicate the data items flagged with any of the listed flags should be excluded from the column averages.

The system developer writes the report specification when designing the system. The RSL compiler converts the report specification into intermediate code to be interpreted at report generation time—that is, when the SPCS COER user requests a report of this type. At report generation time, the SPCS COER subsystem calls the RSL run-time interpreter with a short description of the desired report. This description includes the report type, the office, the equipment tag, the hour, and the range of days. The RSL interpreter then prints the report according to the specification and data. First, RSL prints the heading, supertitles, and column headings in an attractive fashion. It then calls DDM or MLSM to get the data for each line of the report, which it prints according to the column specifications. Finally, it prints the appropriate averages for each column.

RSL takes the drudgery out of report generation. Instead of writing code to fetch, format, and print reports, the system designer writes a few simple specifications that describe the required reports. RSL is able to do the rest because it understands what is generic in SPCS COER reports.

This section has described only a few of the tools used in the construction of new SPCS COER subsystems. The present arsenal contains many other tools both of the type providing some generic service, such as DDM, and of the table-driven type, such as RSL. Moreover, each new component SPCS COER system that is built suggests new tools so that the arsenal continues to grow. The net result is that the system designer can concentrate more and more on the problems of each system, with less and less bogging down in the details of coding.

The savings can be considerable. Studies on the RSS COER system indicated that tools have resulted in a 65-percent reduction in the code required. Tools have also made SPCS COER an easier project to manage. Since a large portion of the code is generic, it is well established and tested. This substantially reduces the maintenance burden.

In addition, better service is given to system users since new features and component systems can be developed more quickly.

IV. THE MANAGEMENT APPROACH TO SPCS COER SPECIFICATION AND DEVELOPMENT

Because of its origin as an experimental system, the SPCS COER project was managed with a research and development style,[9] instead of the classical software design style. Although many software projects seem to use this approach, it does not appear to be discussed anywhere in the software management literature. Weinberger hints at it in various places in his excellent book.[10] On the other hand, there is vigorous literature on research and development management methods.[9,11-14] We find this literature applicable to software projects and have coined the term "whole-job software design method" to describe it.

4.1 Whole-job team

The "whole-job" software design method puts all responsibility for a project, including requirements and development, on a small design team. The principal advantage of this approach is that the developers themselves have a close relation with the clients. This results in other derivative advantages: high motivation, flexibility, and low overhead. The team is responsible for:

1. Studying and understanding previous work.
2. Meeting with clients and understanding their views.
3. Proposing a solution. This includes selling the proposal to the BOC's, to AT&T, and to Bell Laboratories affected organizations and management.
4. Designing the system.
5. Implementing and testing the system.
6. Writing user documentation.
7. Conducting a field trial.
8. Managing the general release.
9. Following up on user problems.

Steps 1 to 3 are systems engineering functions, and 4 to 9 are development. In the whole-job approach, they are all done by the same team.

4.2 Structure and responsibility in the whole-job team

A team consists of one to a few members of technical staff. A team member may belong to more than one team. Typically, we have used people with Master's degrees in computer science. The teams are flexible; they reform from time to time as people come and go, but maintaining continuity of people assigned to a given project is a primary goal of whole-job management. This has not been difficult to achieve.

For multiperson teams, the supervisor sometimes appoints a "team leader." Other times, team members are free to organize themselves as they see fit. They devise and tell the supervisor their plans and a method for reporting progress. The supervisor must choose between imposed and free organization. This choice is based on a variety of factors: the needs and expertise of the team members being the principal ones. In any case, the team plans its own work.

The team does not make commitments on its own to AT&T or BOC clients; management does that. With experience, this hard line can be relaxed somewhat. Team members learn management's views and make minor commitments. This smooths client relations because clients get quick answers.

The whole-job team tells management what it believes it can deliver and when. Management may adjust these estimates a little, sometimes a lot. But even if the dates are adjusted, the team members are usually confident they own them. This has led to much hard work and overtime, not because it was demanded by management, but because the team members assumed a responsibility to deliver to *their* clients what *they* had promised.

4.3 High team motivation

The whole-job approach places virtually all project responsibility on the shoulders of a few people. It is not divided among several organizations (e.g., a requirements group, development group, test group, maintenance group, etc.). The team members can immediately learn about and satisfy customer needs. This is a motivator.

High motivation arises from all the advantages of this approach: low overhead, closeness to clients, career flexibility, and broadening of team members. Small teams are easier to motivate than large teams. Self-organization is a motivator, a required motivator among professionals.

Team members develop a sense of pride in their project that we call "parental motivation." This carries into parts of the project that are less enjoyable, for example, documentation and maintenance. In SPCS COER, there have been few complaints about documentation and maintenance; they just get done. User documentation is delivered to end users at least a month before general availability of the software.

4.4 Low overhead

Low overhead arises for two reasons:

1. There are no "handoff ceremonies" from requirements to development, and from there to other follow-on organizations.

2. There is only one chain of command within Bell Laboratories that is responsible for the technical project management.

The whole-job approach allows the team to begin some design work even before requirements work is completed. There is a tight feedback loop from design to requirements because few people are involved. When a plan is generally accepted by clients, and details of formats, formulas, thresholds, etc., are all that remain, the planner's mind can gradually shift to design. Coding of stable or generic parts may even begin. Thus, a smooth and natural transition occurs.

4.5 Better understanding of the client's problem by developers

Requirements writers are usually close to the client. As they work their part of the problem, they can understand the client's point of view. They can see possibilities for trade-offs that will not compromise real objectives.

If the requirements people are the developers, then it follows that the developers understand the client's problems, have empathy, etc.

If there is a formalization of requirements and a handoff to developers, then the developers are once removed from the client. Many developers minimize this distance by attending some requirements meetings and getting to know the customers. Nonetheless, there is a distance between developers and clients when developers are once removed from the clients by project structure.

In general, the whole-job approach avoids the problem of distance between developers and clients. There are only two groups of people who must thrash things out, see each other's points of view, and agree: the whole job team and the client. One is trying to provide something that the other needs. Relations are simple. Add a third organization, requirements, and you must choose between the two communications networks shown in Fig. 4. The left plan has distance between client and developer, and slow feedback loops; is inflexible; and has difficulty in exploiting opportunities. The right has overhead, tri-team committee meetings, and is a little more flexible and capable of exploiting opportunities. The whole-job diagram is shown in Fig. 5. It has direct feedback and is maximally suited to exploiting opportunities.

Some developers see distance between development and client as helpful. It gives the developers more freedom to pursue design without reference to the customers' opinions. It protects developers from the well-known human trait of clients asking for more and more features on the original development schedule. Experience on the SPCS COER project indicates that neither of these conditions is a problem if the whole-job team works closely enough with the client.

The whole-job approach does have at least two problems. Close ties to the client tend to emphasize short-term over long-term goals and

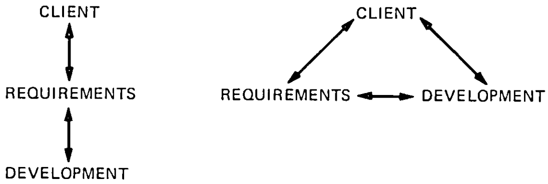

Fig. 4—Communications in a multiorganization software project.

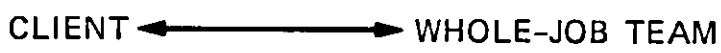

Fig. 5—Communications in a whole-job software project.

activities. Overcoming this calls for very close ties and a willingness to invest time in explanations and discussions. When a client is obdurate, perhaps it really is better to emphasize the short-term goal. A second potential problem is that projects staffed by one or two people would seem extremely vulnerable to sickness and accidents. SPCS COER experience over 10 years shows that this is not a serious problem. Because of the multiply overlapping team design and because people develop friendships with their peers and like to talk, several members of other teams always seemed to know what a missing team member had been doing so that gaps could be quickly closed.

4.6 *Career flexibility*

System developers often view systems engineers as plodding writers of dull requirements. Systems engineers in turn frequently perceive development folk as narrow crank turners, plagued by self-imposed inflexible deadlines. While there may be something to be said for both points of view, both are distorted views and do not encourage communication and personal growth.

The whole-job approach ensures that individuals are exposed to the realities of both systems engineering and development—both the good and bad of each. This is a broadening experience. Entering the project with a development orientation, as most fresh college graduates do, many whole-job people develop a strong desire to move into full-time systems engineering. In our experience, this rarely happens in conventionally structured developments. Graduates of the whole-job experience have gone on to do well in both systems engineering and development areas. We view this as being greatly to the benefit of both the individual and the organization.

4.7 Extending the scope of whole-job possibilities

Whole-job projects are better for the customer than large multiorganization projects. Customers' needs are met more quickly and at lower cost. Long planning/development cycles are not needed. Customers' needs often change rapidly, and the whole-job approach is well suited to a rejuggling of priorities.

This conclusion seems evident. It is common knowledge that small, highly motivated teams can hit quicker and harder and more accurately than large organizations. In the military, commandos are considered elite. But commandos cannot win wars alone. And probably the whole-job approach cannot work on all software projects.

For small projects of two or three persons, most would agree it can be used. But size alone is not an adequate criterion. The method has worked on one medium-sized (10 to 20 person) project, SPCS COER, but may not work on others of this size. To find a criterion for where the whole-job approach can be applied, let us examine the structure of the project.

SPCS COER is organized horizontally instead of vertically, as shown in Fig. 6. SPCS COER is designed as a set of largely independent facilities. Of nine current and one future facility, their interdependencies graph is shown in Fig. 7. The asterisks represent customers. Most of the facilities depend only on BDE, which is a source of data from the world outside of SPCS COER. The 1 *ESS* subsystem supplies data to the world outside. Remove these two interfaces, and the graph reduces to Fig. 8.

What is remarkable here is the strong parallel between the customer's view of the project and its division into loosely coupled subsystems. Let's examine what is meant by the customer's view. Consider one of these subsystems, VSS COER. One person developed this subsystem. He consulted with the AT&T Network Design and Network Admin-

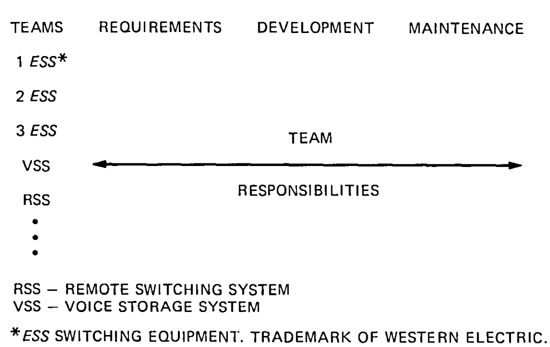

Fig. 6—SPCS COER team organization.

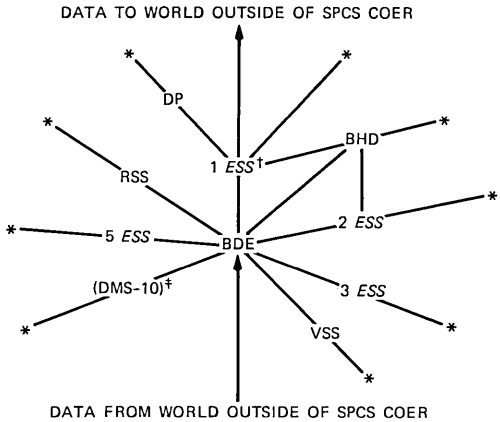

* – CUSTOMERS OF SPCS COER
BDE – BATCH DATA ENTRY
BHD – BUSY HOUR DETERMINATION
DMS – DIGITAL MESSAGE SWITCH
DP – DATA PROFILES
RSS – REMOTE SWITCHING SYSTEM
VSS – VOICE STORAGE SYSTEM

† *ESS* SWITCHING EQUIPMENT. TRADEMARK OF WESTERN ELECTRIC.
‡ TRADEMARK OF NORTHERN TELECOM.

Fig. 7—SPCS COER subsystem interdependencies.

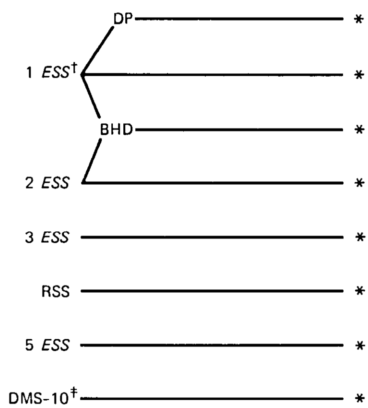

* – CUSTOMERS OF SPCS COER
BHD – BUSY HOUR DETERMINATION
DMS – DIGITAL MESSAGE SWITCH
DP – DATA PROFILES
RSS – REMOTE SWITCHING SYSTEM

† *ESS* SWITCHING EQUIPMENT. TRADEMARK OF WESTERN ELECTRIC.
‡ TRADEMARK OF NORTHERN TELECOM.

Fig. 8—Simplified SPCS COER subsystem interdependencies.

istration departments in a small meeting. There was no representation from 1 *ESS*, 2 *ESS*, or any other parts of the COER project. The independence of the VSS module from other SPCS COER modules was clear from the customer's point of view. And this is reflected in the structure of this meeting. There are no customer requirements that the VSS database or processes interact closely with any other SPCS COER subsystem.

Consider 1 *ESS*, 2 *ESS*, and 3 *ESS* central offices, and the role of a real network administrator in a Bell Operating Company. One administrator may be responsible for all three types of offices. But this person usually works on offices one at a time. A network administrator may use 1 *ESS* COER for an hour or so examining a problem in one particular office, then switch to 2 *ESS* COER to examine the validity of some new data, and request a report. Although one person may use several modules of SPCS COER, the independence among these modules is still there from the customer's point of view. The independencies derive from the method of use.

So from both requirements and usage points of view, the customer's view of the project is a set of almost independent facilities. And each of these facilities is small enough to be planned and developed by a small team.

Underlying all SPCS COER, however, are common system design decisions: the on-line nature of the system, method of data entry, dialogue style, style of reports, etc. These can be viewed as system design decisions that affect all subsystems; but they were all made some years ago and evolve at a slow pace. It is these underlying design criteria that give SPCS COER a cohesiveness from the customer's point of view, i.e., make it one system.

So our hypothesis is this: It is not the size of a project, per se, that makes the whole-job approach applicable or not. The whole-job approach is applicable whenever:

1. A system can be divided into parts that are independent, from a customer's viewpoint.
2. The resulting parts are loosely coupled.
3. Each part can be planned and developed by a small, whole-job team.

Extending the whole-job concept to large projects requires that they be structured according to these criteria. Large projects that cannot be structured this way can probably only be done by the multiorganization approach.

4.8 Whole-job approach summarized

The whole-job approach to software specification and development is a project method. It is described in the research and development literature, but apparently not in the software literature. It is a proven technique and has been in use on software projects for at least 10 years. Its central advantage is that developers and customers build a close relation. This leads to realistic design by the designers, and high customer acceptance of the software products. Structure of a project, not size, determines whether this method can be used.

V. FUTURE DIRECTIONS

SPCS COER, well into its second decade of service, is still growing vigorously in features, component systems, and entities served in response to the expressed needs of its users. But this is only to be expected of a successful product serving the needs of a growing market. We close this paper briefly citing four new and challenging areas toward which SPCS COER thinking is being directed. It is not possible to tell now which of these areas, if any, will bear fruit.

As mentioned earlier, an SPCS COER system was developed for VSS. While VSS COER is not now in operation for reasons having nothing to do with the present paper, this development was an interesting first application of SPCS COER technology outside the realm of switching systems. There appear to be many other potential applications beyond switching that require the capabilities of SPCS COER. Work is presently proceeding to expand our applications in this arena.

The reports produced by SPCS COER are laboriously worked out with user representatives to meet the reporting needs of the BOCs. But surveys have shown that fixed-format reports seldom meet all the needs of individual users. This leads to wasteful manual posting from different mechanized reports to the desired report format. Planning is well advanced, under the title of "Flexible Reporting," for new features that will allow users to alter the format and content of existing reports and create new reports from scratch.

While the development of generic tools, described in Section III, has greatly simplified and speeded the development of new SPCS COER systems, developing a system for each new application is still a significant and time-consuming effort for highly skilled and experienced program designers. Under the name of "Generic COER," also referred to as "User Programmability," planning is under way to allow users to (in effect) create their own component systems for new switching applications. Unlike flexible reporting, the user in "User Programmability" need not be the ordinary SPCS COER user. Most hands-on users of SPCS COER are clerks with little formal training. The user in "User Programmability" could well be a highly trained computer program designer employed by an individual BOC as long as such a resource were readily available to the operating users in the BOC.

Finally, in the 12 years SPCS COER has been operational, there has been a continuing advance in the computing facilities available to the BOCs in their data centers and elsewhere. It now seems possible and economically attractive to move the SPCS COER operation from its single central location for all BOCs into one or more installations in each BOC. Planning work is in progress on this future direction.

REFERENCES

1. H. J. Dougherty, H. Ragg, P. G. Ridinger, and A. A. Stockert, "No. 1 ESS Master Control Center," B.S.T.J., *43*, No. 5, Part 2 (September 1964), p. 2286.
2. E. P. Gould and J. W. Mosior, "TELPORT - Time Shared Information Systems," Bell Lab. Rec., *46*, No. 6 (June 1968), pp. 197–202.
3. N. R. Sinowitz, "DATAPLUS - A Computer Language for English Speaking People," Bell Lab. Rec., *46*, No. 11 (December 1968), pp. 362–9.
4. M. S. Hall and G. P. Patti, "PATROL Surveys Traffic for No. 1 ESS Offices," Bell Lab. Rec., *51*, No. 3 (March, 1973), p. 72.
5. V. L. Fahrmann, "Reporting Central Office Traffic Information," Bell Lab. Rec., *57*, No. 4 (April 1979), pp. 99–104.
6. M. S. Hall, J. A. Kohut, G. W. Riesz, and J. W. Steifle, "System Plan," B.S.T.J., this issue.
7. K. A. Friedman, "Precutover Extreme Value Engineering of a Local Digital Switch," Proc. of the Tenth Int. Teletraffic Cong., Montreal, June 8–15, 1983.
8. L. E. Heindel and J. T. Roberto, "The Off-The-Shelf System—A Packaged Information Management System," B.S.T.J., *52*, No. 10 (December 1973), pp. 1743–63.
9. E. B. Roberts, "Stimulating Technological Innovation—Organizational Approaches," Res. Manage., *12*, No. 6 (November 1979), pp. 26–30.
10. G. M. Weinberger, *The Psychology of Computer Programming*, New York: Van Nostrand Reinhold, 1971.
11. T. J. Allen, D. M. S. Lee, and M. L. Tushman, "R&D Performance as a Function of Internal Communication, Project Management and the Nature of the Work," IEEE Trans. Eng. Manage., *27*, No. 1 (1980), pp. 2–11.
12. J. D. Goldhar, L. K. Bragaw, and J. J. Schwartz, "Information flows, management styles and technological innovation," IEEE Trans. Eng. Manage., *23*, No. 1 (1976), pp. 51–62.
13. T. J. Allen and S. I. Cohen, "Information flow in research and development laboratories," Admin. Sci. Quart., *14* (1969), pp. 12–19.
14. T. J. Allen and A. R. Fusfeld, "Design for Communication in the Research and Development Lab," Technol. Rev. *78*, No. 6 (May 1976), pp. 65–71.

AUTHORS

Richard F. Grantges, B.S., B.E.E., 1953, University of Minnesota; Bell Laboratories, 1953—. Mr. Grantges was first assigned to long range engineering studies of transmission systems where, in 1954, he demonstrated the economic feasibility of Pulse Code Modulation for exchange carrier. After military service he was transferred to *ESS®* Systems Engineering in 1958 and promoted to Supervisor in 1961, responsible for switching network engineering and the mechanization of central office engineering practices. He was promoted to Head in 1968, responsible for appraisal studies in the Traffic Studies Center. In 1969, he went to AT&T as Traffic Operations Manager—Electronic Switching, responsible for the traffic engineering and administration of local and toll electronic switching machines. Returning to BTL in 1971 as Head of the Operator Services and Traffic Engineering Department, he was active in planning the mechanization of operator services and traffic data collection and processing, particularly the SPCS COER, Small Office Network Data System and EADAS systems. In June 1981, he became responsible for the systems engineering of the Total Network Data System (TNDS) except for Network Management and Trunking. Member, Tau Beta Pi, Eta Kappa Nu, IEEE.

Virginia L. Fahrmann, B.A. (Mathematics), 1963, St. Mary-of-the-Woods College; M.S. (Mathematics), 1965, Purdue University; Bell Laboratories, 1965—. Ms. Fahrmann initially worked on traffic load simulations of 1 *ESS* and No. 1 Traffic Service Position System (TSPS). In 1971 she joined the AT&T Switching Engineering Department, with responsibility for providing technical support in long-range planning studies for local wire centers to Operating Companies. She returned to Bell Laboratories in 1973, and was involved in systems engineering and development for SPCS COER. In 1980 she became Supervisor of one of the two SPCS COER groups. Currently, Ms. Fahrmann supervises a group that works on requirements and development for the Plug-in Inventory Control System/Detailed Continuing Property Record (PICS/DCPR) system, which provides inventory control and maintains investment records for central office equipment.

Thomas A. Gibson, B.S.E.E., 1960, Kansas State University; Bell Laboratories, 1960–. At Bell Laboratories, Mr. Gibson did systems engineering studies on Electronic Switching Systems and developed a variety of software systems including the Master Links database management system and SPCS COER. He is currently Head of the Plug-in Inventory Control System/Detailed Continuing Property Record Development and Support Operations Department. Mr. Gibson has chaired a subcommittee of the American National Standards Institute. Member, ACM, Sigma Xi, Eta Kappa Nu.

Laurence M. Brown, B.A. (Classics), B.S. (Fundamental Sciences-Computer Science), 1975, Lehigh University; B.A. (Literae Humaniores), 1977, Balliol College, Oxford University, Oxford, U.K.; M.A. (Electrical Engineering), 1978, Stanford University; Bell Laboratories, 1977—. Mr. Brown first worked on the software tools of ESS COER. He is presently Supervisor of the Software Systems Design and Support Group.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Small Office Network Data System

By D. H. BARNES* and J. J. O'CONNOR*

(Manuscript received December 3, 1982)

The Small Office Network Data System (SONDS) is a part of the Total Network Data System (TNDS) designed specifically to meet the network data requirements for the small step-by-step offices of the Bell System. Its ability to accurately relate to the load and service characteristics seen in these small, mostly rural offices has ensured its success. It maintains low cost by measuring and recording a single observation per day for each traffic measurement. This is the daily peak hour load per group. A moving average of the mean and variance of peak hour measurements is interpreted by using a modified-Gumbel extreme value distribution function. This has been demonstrated to fit well within the actual distributions of daily peak hour loads. It is used to construct traffic capacity tables that better reflect customers service experience in these offices than do existing methods. It uses field-tested parameter assumptions of the underlying distribution from which peak values are sampled to estimate equivalent time-consistent busy hour loads relative to the observed peak hour loads. This allows service to be compared among SONDS and non-SONDS step-by-step offices. SONDS uses on-site data collection units, a single central computer for all Bell System offices, and provides local printing of output reports. Data are collected each night and daily, weekly, and monthly output reports are distributed over the switched network.

* Bell Laboratories.

I. INTRODUCTION

The Small Office Network Data System (SONDS) is a miniature Total Network Data System (TNDS) designed to meet the unique and special needs of the Bell System's smaller, rural-type offices. In respect to SONDS a small office is one serving fewer than 5000 main stations (MSs). These comprise almost 40 percent of the Bell System's total switching entities but serve less than 8 percent of its customers. The offices that are candidates for SONDS are predominantly of the noncommon control, step-by-step type. They are generally located in remote rural locations and are unattended.

The amount of traffic and the service vary among these offices and among the serving groups within each office for several reasons. First, as random traffic theory predicts, there is a greater proportional variation in traffic with small versus large groupings of traffic sources. Observations have shown that the actual variations of source traffic are, in fact, much greater than assumed by the random theory underlying the probability tables normally used to size central offices. Second, the mix of customer types varies widely among these offices and this further affects the within-a-day and day-to-day traffic distributions in often unpredictable ways. In addition, as a result of the relatively small originating line and terminating connector groups of the small step-by-step offices, load balancing problems are accentuated. These offices are unattended, which means less frequent surveillance of central office equipment problems such as busy switches or trunks. On the average, there will be more such maintenance problems and they will be longer in duration than in larger, attended offices.

The small offices require special attention if both service and cost objectives are to be met. They are, however, too small and remote to justify support by conventional, continuous measurement systems. Typically, a few days of traffic data per year are available in these offices. Present standard dimensioning techniques and design service criteria are thus inadequate in the presence of minimal data and large traffic variation. We have demonstrated with SONDS how large office standards can be applied to these small offices at an affordable price.

II. DIMENSIONING OF TRAFFIC SYSTEMS

Design service criteria for standard Bell System engineering are based on the concept of a Time-Consistent Busy Hour (TCBH). It is by definition that clock hour which has been observed by study to contribute the greatest amount of total traffic over some defined period of time. The sizing of a group of traffic servers is based on meeting service criteria for the highest controlling TCBH loads for the average busy season, average ten high day, or highest day.

The traffic theory used to calculate engineering capacity tables assumes the behavior of a constant system of chance causes ("statistical equilibrium"). To account for different levels of departure from truly random behavior, traffic capacity tables are adjusted to allow for options of low, medium, or high levels of day-to-day variation.

Time-consistent busy-hour practice is a compromise for data collection and processing costs. Simplifying assumptions are used, however, which lead, particularly in small offices, to overprovisioning to protect for variations not reflected in an average statistic. Due to lack of data, maintenance problems, and forecasting errors, underprovisioning can also occur.

TCBH practices assume that provision of objective service in the average busy hour will provide adequate service in all hours. Average service is not, however, the service seen by individual customers nor by groups of customers in offices with traffic imbalance. While TCBH service criteria have been shown by experience to be a reasonable approach to provide an adequate overall service control for large groups of customers, they are less valid for small groups because of the greater deviation from truly random behavior of such groups. For example, SONDS office studies reveal that the TCBH is the actual busiest hour of the day on only about 20 percent of the days.

III. THE SONDS DEVELOPMENT

SONDS is the outcome of a series of Bell Laboratories' traffic studies and computer experiments that were initiated originally to explore the power of computers to completely control data collection from a central location. Included were studies of traffic characteristics and the evaluation of design service criteria that might be used to better control service levels and at the same time provide a more efficient allocation of resources. In recognition of the largely unsolved small office data and design problem, we initiated an experimental network data system designed to meet that need. These relatively simple offices, largely unmeasured, were judged to be a good environment to explore a completely mechanized system that would use time-shared computers to interface with smart microprocessor-controlled central office terminals. We were looking for a system free from the control of people other than the direct user. These offices were not directly coupled to TNDS and thus gave us a simpler environment in which to study innovative traffic systems.

The SONDS development, the formulation of analysis requirements, and the specifications of output reports were designed to incorporate:

- Low cost
- Service measurement
- A compact database

- A nearly paperless system
- Switched-network data communication
- An interactive computer interface
- Automatic report generation and distribution
- Complete access by the authorized user for:
  —Nonscheduled data or special study requests
  —Database control
  —Trouble reports and data for analysis
  —Remote testing of terminal equipment.

SONDS has met these desired objectives. Its success is due to a combination of the following features:

1. An Extreme Value Engineering (EVE) method has been formulated and adapted to the studied characteristics of the candidate offices. EVE uses order statistics of an extreme value distribution formulated from daily peak hour data observations. The need for simplifying assumptions using other traffic dimensioning models is largely reduced. Through use of the order statistics the computer is able to detect questionable data and to issue timely exception reports, including backup data, which are useful for analyzing and correcting trouble conditions.

2. New EVE design service criteria are based on a broad study of office characteristics and are selected to match the objectives of current dimensioning methods. On the average, with the same or less switching equipment and trunks, overall service will improve. With full-time data surveillance available for the first time and with a good system of trouble detection, there will be less need to play safe by overprovisioning.

3. A compact database is realized because only a single hourly value per day, the daily peak hour value, is recorded in the database for each traffic measurement. The extreme value distribution parameters are estimated from a moving average of the mean and variance of this measurement.

4. Automatic computer polling of office terminals during early morning hours is an efficient serial use of computer ports and the switched network.

5. A largely paperless system has been realized. The standard reports required in TNDS are automatically distributed immediately following the end of each reporting interval. Database updates and requests for special data and reports are made via an interactive interface with the computer and require no other (human) interactions.

6. No special monitoring or scaling is required to compensate for holiday or "odd-ball" traffic conditions in the office engineering because automatic outlier tests reject such data.

7. Dial tone delay service measurements and their reporting as a service index are contained in the system. Separate and costly dial tone test call facilities are not needed. The result is a more accurate measurement and includes a measure of group-to-group service differences not available otherwise.

IV. FORMULATION OF THE EVE METHOD FOR SONDS OFFICES

The following is a short history of work that helped generate an interest in the EVE approach used in SONDS.[1-4]

4.1 Early work

Karlsson[1] claimed that, particularly for international routes, the (time-consistent) busy hour is both unstable in amplitude and unfixed in time. He proposed the dimensioning of traffic routes from integrated information of all traffic peaks taken over a sufficiently long time interval T, typically six hours a day for five days, to give good statistical stability. This was an early attempt to recognize the importance of differences in the variances of traffic groupings with equal loads.

Gershwin, Laue, and Wolman[2,3] solved a problem related to the load administration of Subscriber Line Concentrators (SLC) using extreme value statistics. The problem was to administer line assignment far enough ahead so that concentrated lines could see nearly zero blocking in this additional switching stage. Extreme value methods were used to keep peak hour blocking less than 0.005 except for a few times per year.[4]

Unlike Karlsson, Gershwin et al. retained hourly measurements but dropped the concept of a time-consistent busy hour. They recorded only a weekly peak load. The mean and variance of these weekly peak loads were found to be good estimates of the fitting parameters of the Gumbel extreme value distribution function.[5] The plan was simple and inexpensive and served well to direct line assignments and to provide capacity estimates for planning purposes.

4.2 Daily peak values

The initial work that led to extreme value engineering as used by SONDS is published in Refs. 6, 7, and 8. In SONDS it is important that traffic reports and service measurements be integrated into and be consistent with the rest of TNDS. It was thus important, in contrast to weekly values used for the SLC, to follow daily busy hour and busy season concepts and to furnish monthly traffic reports consistent with other TNDS reports. The use of daily peak values for design and administration of small offices was of particular interest because we had experimented with techniques to measure, collect, and summarize such data at a very low cost.

Day-to-day variation of TCBH originating traffic loads as observed in step-by-step offices is illustrated in Fig. 1. The coefficient of variation in percent is shown as a function of offered load in erlangs. Curve A is an approximate fit to the TCBH loads observed in many step-by-step offices. The data showed considerable scatter, indicating that day-to-day variation is more than just a function of the offered traffic. Curve C is a theoretical curve based on the assumption of constant system of chance causes ("statistical equilibrium") which is basic to the probability theory used to construct traffic capacity tables. Curve B depicts a possible location for daily peak-hour data. Its location between the theoretical curve C and the TCBH curve A will be a function of the number of hours in the day that have a load high enough to possibly be the peak load for the day. For calculation purposes we will develop an equivalent number of equally loaded hours with which to describe the underlying distribution from which peak hours come. We designate these as "Candidate Busy Hours".

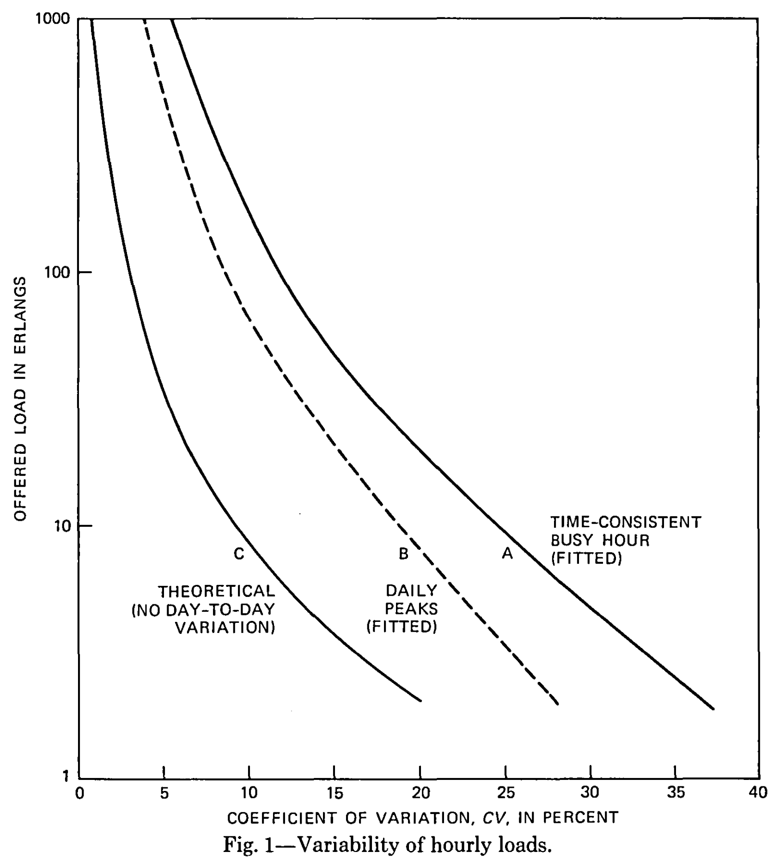

Fig. 1—Variability of hourly loads.

Figure 2 illustrates a comparison between TCBH and peak-hour concepts using actual data. These data were taken from a typical line finder group in a SONDS office, which had 13 line finders serving the group. The data are plotted on probability paper constructed to match the normal-to-the-sixth extreme value probability function derived in Section 4.4.

Curve A is a theoretical curve based on purely random, no day-to-day variation; it is the underlying distribution of the TCBH capacity tables. We could not plot this curve exactly as a straight line on this graph. To simplify the illustration we have approximated it as two line segments plotted between the 0.05 and the 0.95 probabilities points and the average value. This average value was scaled to be equal to the 242 hundred call seconds (CCS) measured average of the TCBH data. This gives the 0.95 point, equivalent to a 20-day return period load, a load of 286 CCS.

Curve B is a plot of the 20 TCBH loads observed in this office. It is seen that these data points significantly deviate from the theoretical curve A, and this is due in part to actual day-to-day variations. Note that the observed 20-day return period load, the 0.95 point, is 326 CCS during the time-consistent busy hour.

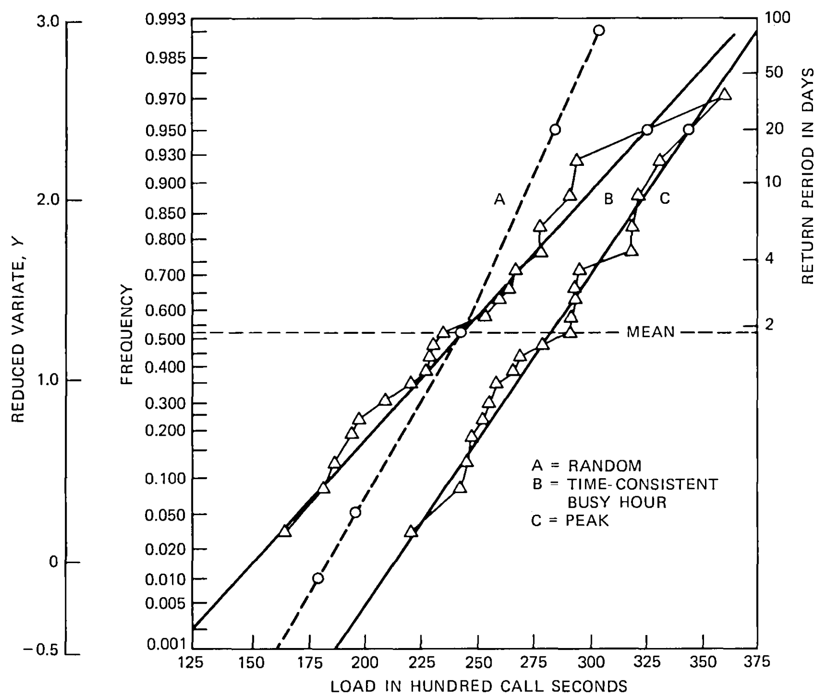

Fig. 2—Extreme value probability plot for originating loads.

Curve C is a plot of the 20 daily peaks taken for the same days as the TCBH data. The straight-line character of the data on this graph is typical of the agreement between daily peak data and the normal-to-the-sixth distribution function. Note that the mean value is 283 CCS and the 0.95 value is 346 CCS. While the mean and 0.95 values of the peak curve are larger than those of the TCBH curve, the larger slope of the peak curve indicates less variance, that is, daily peak values are less volatile than TCBH values!

These several load values obtained from Fig. 2 can be used to compare the dial tone delay service as given by current engineering capacity tables with the range of expected values that a customer will experience. This is given in Table I, which shows the service values for both mean and 20-day return period loads. Since the 20-day return period load X_{20} is also interpreted as the load to be exceeded once a month, on the average, we can examine the service implications on this basis. Table I shows that a line finder group properly designed to give 1.5-percent dial tone delay during the time-consistent busy hour would be expected to exceed 4.3 percent once a month, even with purely random traffic. The day-to-day variation observed in Fig. 2 would increase this to 8.7 percent, and the use of the daily peak-hourly loads of Fig. 2 would predict a service level of 12.5 percent to be exceeded once a month.

Extreme value statistical methods can be applied to load balance. Loads are balanced through line assignment, which takes into account both additions and disconnects. Table II shows an office TCBH and peak load averages and their coefficients of variation, CV, using a month of data. The last column, X_{20}, is the 20-day return period load calculated from the mean and standard deviation of the daily peak values. As expected, we can see that the CV of the daily peak loads (12 percent) is less than that of the TCBH loads (22 percent). Also note that the CV of the peak loads (6 percent) across the loading division is less than that of the TCBH loads (7.3 percent). Even the X_{20}, once-a-month load, which includes the additional parameter of variance, has only a CV of 7.9 percent.

Table III shows how the corrective action needed to balance load against load main station movement depends on the type of data used

Table I—Load service characteristics

	Load in CCS		Expected Dial Tone Delay (13 Line Finders)	
Distribution	Mean	X_{20}	Mean	X_{20}
A–Random	242	293	1.5	4.3
B–TCBH	242	326	2.1	8.7
C–Peak BH	283	346	4.2	12.5

Table II—Originating load distributions

Line Finder Group No.	TCBH CCS	TCBH CV (Percent)	Daily Peaks CCS	Daily Peaks CV (Percent)	X_{20} CCS
1	159	21	185	11	220
2	155	19	199	8	226
3	139	17	185	8	211
4	179	17	214	9	249
5	169	25	220	14	273
6	147	21	200	13	244
7	162	22	196	12	238
8	164	26	193	18	252
9	162	29	207	13	253
Avg	160	22	200	12	241
Total percent of *CV*		7.3		6.0	7.9

Table III—Originating main station balance

Line Finder Group No.	TCBH	Daily Peaks	X_{20}
1	−1	−15	−18
2	−6	−1	−13
3	−27	−15	−25
4	+24	+14	+7
5	+12	+20	+27
6	−17	0	+3
7	+3	−4	−3
8	+5	−7	+9
9	+3	+7	+10
Avg.	11	9	12

from Table II. The best service balance will be accomplished using X_{20} data. Note the very large differences by individual groups between the movement indicated by TCBH and that indicated by X_{20} data.

4.3 Gumbel distribution

Preliminary work on SONDS, following the direction of the previous work, used the Gumbel[5] extreme value distribution, a double exponential distribution, shown below.

$$G(x) = \exp\{-\exp[-\alpha(x - \mu)]\}. \tag{1}$$

$G(x)$ is the probability that the largest observation in a large set of observations is less than the value x. In this expression μ is a location parameter (the mode), and α is an inverse measure of dispersion. The mean of the distribution is:

$$E(x) = \mu + \gamma/\alpha, \tag{2}$$

where $\gamma = 0.57721$ (Euler's Constant) and the variance is:

$$V(x) = \frac{\Pi^2}{6\alpha^2}. \tag{3}$$

The Moments Method may be used to estimate μ and α by replacing $E(x)$ and $V(x)$ with their sample values from a sample of N hourly peaks $x_1, \cdots, x_N$. The sample mean is:

$$\bar{x} = \frac{1}{N} \sum_{i=1}^{N} x_i \tag{4}$$

and the sample variance is:

$$\bar{s}^2 = \frac{1}{N-1} \sum_{i=1}^{N} (x_i - \bar{x})^2. \tag{5}$$

Substituting eqs. (4) and (5) into eqs. (2) and (3), we have:

$$\hat{\mu} = \bar{x} - \gamma/\hat{\alpha} \tag{6}$$

and

$$\hat{\alpha} = \frac{\Pi}{\bar{s}\sqrt{6}}. \tag{7}$$

Data points fitting a Gumbel distribution function should plot as a straight line on Gumbel extreme value probability paper. Figure 3 shows an attempt to do this with hourly peak data obtained in a Private Branch Exchange (PBX) office serving a single business customer.[6] Notice that the data points do not fall within the control limits. A number of the low peaks were identified as holidays and days preceding holidays. By applying rejection test statistics to one low point at a time and recalculating the order statistics until the lowest point fell within control limits, it was possible to plot the data as shown in Fig. 4, after rejecting only nine of the original 80 data points. This method appeared to properly represent the high data points of most interest and automatically reject data of little interest, such as holidays.

If we reject the lowest few points to force a straight line fit of data to the double exponential probability plot, which is useful for eliminating holiday traffic, we also might eliminate some low data points necessary to identify problems in SONDS. This forced rejection of the lower data points makes the estimate of the mean of the daily peaks biased on the high side and the estimate of the variance biased on the low side, making the effect on the resulting estimated 20-day return period load indeterminate.

We realized that the problem arises because the daily peak loads do not satisfy the large sample assumption that leads to the double

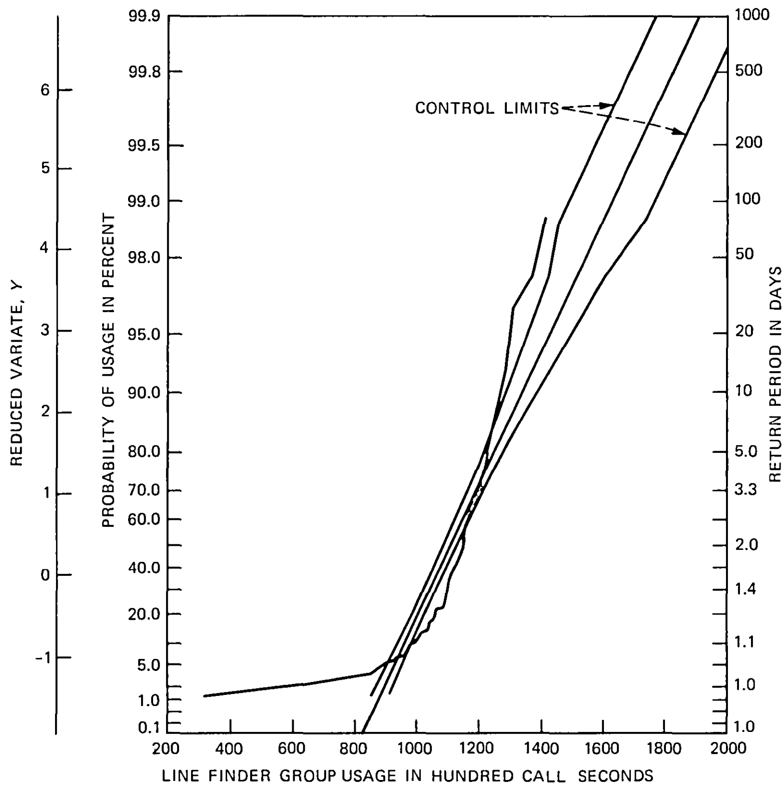

Fig. 3—Extreme value probability plot for 80 peak hours of data.

exponential distribution function of eq. (1). Instead of selecting a peak value from among a large number of candidate hours, we find that we are typically selecting daily peaks from a few to perhaps ten candidate busy hours.

4.4 Normal-to-the-sixth distribution

An improved extreme value distribution function model was needed to fit the small office traffic data. A set of statistical test procedures was needed in the initial start-up testing of a SONDS office (see Section 9.1) to resolve installation problems and other central office problems not previously detected for lack of data.

The first step was to test the distribution of equipment loads across days for the busier hours of each of many offices. The findings were that the distributions satisfactorily tested for normality. The second step was to test whether round-the-clock hourly data can be modeled adequately by h equally loaded candidate busy hours even though the actual clock hours do not have equal loads. These tests led to the

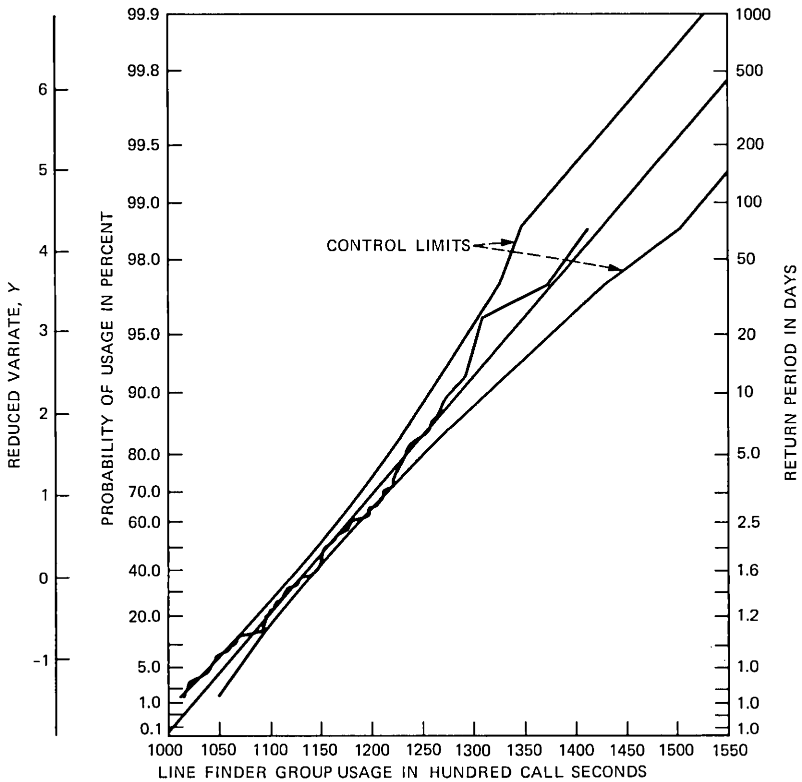

Fig. 4—Extreme value probability plot for 71 peak hours of data.

conclusion that a modified extreme value distribution function suitable for SONDS data is the normal probability distribution function raised to an optimum power of h. Table IV gives the results of three tests that resolved the optimum value of h as 6.

Therefore, the modified extreme value distribution function for describing daily peak-hour usage data, x, is the normal distribution raised to the sixth power, $F(x)$:

$$F(x) = \left[\frac{1}{\sqrt{2\pi}\sigma}\int_{-\infty}^{x} \exp[-(t-\mu)^2/2\sigma^2]dt\right]^6, \tag{8}$$

where μ and σ are the mean and standard deviation of the underlying equal "candidate busy hours" load distributions.

Tippett[9] numerically evaluated the probability of finding the largest value of h samples drawn from a normal distribution. In particular for $h = 6$, in terms of the normalized variate, the mean is:

$$E(x) = \mu + 1.267\sigma \tag{9}$$

Table IV—Selection of number of candidate busy hours, h, in the (NORMAL)^h model

Tests	Ranges of h Indicated by Tests
Distribution of daily peak loads among hours during the day	6–10
Validation and monitoring of lower daily peak loads	4–6
Estimating highest of 20 daily peaks	6–9
Collectively (all of the above)	4–10
Value selected	6

and the variance is:

$$\mathrm{Var}(x) = 0.416\sigma^2. \tag{10}$$

In SONDS using the moments method, μ and σ are estimated by matching the first two moments of the distribution of eq. (8) to their unbiased sample estimates. Replacing $E(x)$ and $\mathrm{Var}(x)$ by the unbiased sample estimates of the mean, $\bar{x}$, and variance, $\bar{s}^2$, the moments method equations are obtained:

$$\bar{x} = \hat{\mu} + 1.267\hat{\sigma} \tag{11}$$

and

$$\bar{s}^2 = 0.416\hat{\sigma}^2. \tag{12}$$

Manipulating eqs. (11) and (12), we have the parameters of the normal-to-the-sixth-distribution in terms of the underlying daily peak parameters as

$$\hat{\mu} = \bar{x} - 1.96\bar{s} \tag{13}$$

and

$$\hat{\sigma} = 1.55\bar{s}, \tag{14}$$

where $\bar{x}$ and $\bar{s}$ are the unbiased sample estimates, calculated directly from the daily hourly peaks x_i, as shown in eqs. (4) and (5).

The 95th percentile of $F(x)$ can be calculated directly from eq. (8) by setting $F(x) = 0.95$ to give

$$x_{20} = \hat{\mu} + 2.39\hat{\sigma}, \tag{15}$$

which by eqs. (13) and (14) can be stated in terms of the daily peak values as

$$x_{20} = \bar{x} + 1.74\bar{s}. \tag{16}$$

This x_{20} is the load that is exceeded once, on the average, every 20 business days. It is also called the once-a-month load.

V. EVE SERVICE CRITERIA

The use of daily peak-hour measurements to estimate a once-a-month load reflects to a greater degree the service viewpoint of customers than do present average time-consistent busy hour practices. The estimated once-a-month load is based on continuous monitoring of all hourly data, and as interpreted in SONDS by the $F(x)$ model of eq. (8), takes into account the variability of loads.

In order to use a once-a-month load for engineering and administration, new peak service criteria and corresponding load objectives must be determined. For small step-by-step offices, peak service criteria have been selected for line finders, connectors, selectors, and only-route trunk groups.

Peak service criteria are determined by calculating the expected blocking or delay that would occur in offices if the once-a-month loads were offered to the amount of equipment required by time-consistent practices. Continuous collection of hourly traffic loads in the study phase of SONDS provided the data needed for this comparison. The following are the objectives for selecting peak service criteria:

1. Provide overall service that is at least as good as that resulting from use of the current time-consistent criteria
2. Provide better service during peak load periods than results from using the current time-consistent criteria
3. Provide service to customers that is independent of the size of the serving office or equipment group
4. Require no more total equipment in the Bell System than is in current use.

Having resolved by field study the EVE blocking objectives for the various components relative to their present TCBH blocking objectives, it is a simple matter to extend the existing capacity tables to the new blocking values and calculate equivalent tables to replace existing TCBH tables. The existence of these extreme value capacity tables is necessary to apply EVE techniques to central office engineering, as discussed in the next section.

VI. TRAFFIC ENGINEERING USING EVE

In the traffic engineering of step-by-step offices, it is necessary to distribute the forecasted originating, local, and incoming traffic loads and to provide equipment quantities to meet objective service criteria.

Customers are associated with a group of line finders that return dial tone and connect the customer to a first selector, which will respond to the first dial pulses. A selector steps vertically in response to each digit (0–9) dialed. It then searches horizontally up to ten terminals for an idle path to a following stage of selectors or for an outgoing trunk. Successive selector stages respond to successive dialed digits to direct traffic toward outgoing trunks or intraoffice destinations. In the last switching stage a connector completes the call to the customer line corresponding to the telephone number. A connector responds vertically and horizontally in response to the last two numbers dialed. Incoming traffic from other offices is served by incoming selectors associated with incoming trunks. The number of selector switching stages and how the traffic mixes prior to switching to the customer's connector group depend on the numbering plan and size of the office.

To provide measurements for each part of the traffic flow would be too expensive. Using time-consistent practices it has been customary to estimate unmeasured traffic by adding and subtracting traffic measured in other parts of the network. With EVE this is not feasible because the peaks in the various parts of the traffic flow may not occur at the same time.

There are two relatively simple answers to estimating equivalent time-consistent traffic loads, which then may be apportioned as is done in existing step-by-step practices. Both depend to a degree on empirical relationships observed in small offices and relate to the candidate busy hour concept.

The first approach is based on analysis of field data, as was done to determine $h = 6$ of Table IV. In this case, the best fit between the daily peak-hourly values and the TCBH values is for $h = 2.5$. The mean and the standard deviation of the normalized variate for $h = 2.5$ are 0.725 and 0.790, respectively.[7] Using these values in eqs. (11) and (12) we calculate the TCBH equivalent load as:

$$\mu_{\text{TCBH}} = \bar{x} - 0.918\bar{s}, \tag{17}$$

where $\bar{x}$ and $\bar{s}$ are defined in eqs. 4 and 5.

A second method—more appropriate, we think, for estimating the TCBH mean value—is the use of the SONDS peak traffic capacity tables. These tables were obtained by deriving the equivalence over many offices between time-consistent busy hour service criteria and peak service criteria. Traffic capacity tables are available in the SONDS computer for both the peak loads and equivalent time-consistent loads. Using the load values in these tables, we can convert an observed peak load to a time-consistent load.

Table V shows an example comparing the estimates of time-consistent loads derived from once-a-month values with actual measured time-consistent loads for a typical mix of traffic. Both the normal-to-the-2.5 and the capacity table methods give satisfactory results. The table method is preferred, however, because it will come closer to correcting the unmeasured traffic to SONDS objective service levels.

Table V—Intraoffice* traffic estimates (in CCS)

	Once-a-Month Actual	Time-Consistent Estimates		TCBH Actual
		h=2.5	Table	
Terminating	2789	2503	2418	2512
Incoming	1961	1737	1670	1757
Local		766	748	755

* Intraoffice = Terminating-Incoming

VII. SONDS DIAL TONE DELAY ANALYSIS

For SONDS offices, a dial tone delay service index is required for peak measurements, and it should be able to be compared directly with non-SONDS offices using TCBH measurements. The following paragraphs describe the procedures used to make this conversion.

In current practice the service index relates monthly average service results to objective capacity levels. SONDS introduces the concepts of the monthly Average Measured Peak Service (AMPS)[8] as the average of the peak values of the Dial Tone Delay (DTD) results for all the line finder groups in a loading division. An AMPS objective is chosen for the peak service criterion, thereby including the measured volatility of loads within the service objective.

The coefficient of variation, CV, of the current once-a-month load is found from eq. 16 to be:

$$CV = \bar{s}/\bar{x} = (x_{20} - \bar{x})/1.74\bar{x}. \tag{18}$$

An empirical relation between CV and load has been determined.[8] The CV of the capacity load is estimated from that of the measured load by:

$$CV_{\mathrm{CAP}} = CV\left(\frac{x_{20}}{x_{\mathrm{CAP}}}\right)^{0.432}. \tag{19}$$

The value x_{CAP} is the x_{20} value at capacity. Given the CV_{CAP} and the number of line finders per group, the AMPS objective can be selected from Table VI.

These peak values of AMPS and AMPS objectives can be converted[8] to DTD TCBH equivalents from the following empirical relationship derived from data studies.

$$\frac{TCBH\ DTD}{TCBH\ DTD_{\mathrm{OBJ}}} = 0.85\left(\frac{AMPS}{AMPS_{\mathrm{OBJ}}}\right) + 0.17\left(\frac{AMPS}{AMPS_{\mathrm{OBJ}}}\right)^2. \tag{20}$$

SONDS uses this equation to convert the measured AMPS values to equivalent TCBH DTD values. These values of TCBH DTD can be entered directly into current, accepted index tables.

Table VI—AMPS objective (percent) coefficient of variation of daily peak loads at capacity

		0	0.05	0.10	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5
Number of line finders per group	1	9.0	8.3	7.6	7.2	6.7	6.3	5.9	5.6	5.3	5.1	4.8
	2	9.2	7.9	6.9	6.1	5.5	5.0	4.6	4.2	3.9	3.7	3.5
	3	9.3	7.8	6.6	5.6	4.9	4.4	3.9	3.6	3.3	3.1	2.9
	4	9.4	7.4	6.0	5.0	4.3	3.8	3.4	3.1	2.8	2.7	2.5
	5	9.6	7.4	5.8	4.8	4.0	3.5	3.1	2.9	2.7	2.5	2.4
	6	9.7	7.1	5.4	4.4	3.7	3.2	2.8	2.6	2.4	2.3	2.2
	7	9.7	7.0	5.3	4.2	3.5	3.0	2.7	2.5	2.3	2.2	2.1
	8	9.9	6.9	5.1	4.0	3.3	2.9	2.6	2.4	2.2	2.2	2.1
	9	9.9	6.8	4.9	3.8	3.2	2.8	2.5	2.3	2.2	2.1	2.1
	10	10.0	6.6	4.8	3.7	3.0	2.6	2.4	2.3	2.2	2.1	2.0
	11	10.1	6.6	4.7	3.6	3.0	2.6	2.4	2.2	2.1	2.1	2.1
	12	10.2	6.5	4.5	3.5	2.9	2.5	2.3	2.2	2.1	2.1	2.1
	13	10.3	6.4	4.4	3.3	2.8	2.5	2.3	2.2	2.1	2.1	2.1
	14	10.3	6.4	4.3	3.3	2.7	2.4	2.3	2.2	2.1	2.1	2.1
	15	10.4	6.3	4.3	3.2	2.7	2.4	2.3	2.2	2.2	2.1	2.1
	16	10.5	6.2	4.2	3.2	2.7	2.4	2.3	2.2	2.2	2.2	2.2
	17	10.5	6.1	4.1	3.1	2.6	2.4	2.3	2.2	2.2	2.2	2.2
	18	10.6	6.0	4.0	3.0	2.6	2.3	2.2	2.2	2.2	2.2	2.2
	19	10.7	6.0	3.9	3.0	2.6	2.4	2.3	2.2	2.2	2.2	2.2
	20	10.7	6.0	3.9	3.0	2.5	2.4	2.3	2.2	2.2	2.3	2.2
	21	10.8	5.9	3.8	2.9	2.5	2.3	2.3	2.2	2.2	2.3	2.3
	22	10.9	5.8	3.8	2.9	2.5	2.3	2.3	2.3	2.3	2.3	2.3
	23	10.9	5.8	3.7	2.8	2.5	2.3	2.3	2.3	2.3	2.3	2.4
	24	11.0	5.8	3.7	2.8	2.5	2.4	2.3	2.3	2.3	2.4	2.4
	25	11.1	5.7	3.6	2.8	2.5	2.4	2.3	2.3	2.4	2.4	2.4
	26	11.1	5.7	3.6	2.8	2.5	2.4	2.3	2.3	2.4	2.4	2.4
	27	11.2	5.7	3.6	2.8	2.5	2.4	2.4	2.3	2.4	2.4	2.5
	28	11.2	5.6	3.5	2.7	2.5	2.4	2.4	2.3	2.4	2.5	2.5
	29	11.3	5.6	3.5	2.7	2.5	2.4	2.4	2.4	2.5	2.5	2.5
	30	11.3	5.6	3.5	2.7	2.5	2.4	2.4	2.4	2.5	2.5	2.6

VIII. SONDS LINE ASSIGNMENT AND LOAD BALANCE

Line assignment is an important task of the network administrator and a load balance index is part of TNDS. The procedures for both of these are well established but are based on TCBH techniques. Farel[10] derived new calculations to use the once-a-month loads obtained by SONDS. Consistent with the previous techniques, the SONDS line assignment recommendations are smoothed over time so that stable line assignment procedures are obtained. Also the load balance index reflects the level of load carried by the office; in particular, it recognizes that bad load balance does not produce bad service in a lightly loaded office.

Farel's line assignment algorithm[10] is based on the load deviation of each line finder group from the average load of the loading division, as follows:

$$\bar{u}_i(t) = w_i(t)F \tag{21}$$

and

$$w_i(t) = a[\bar{u}_i(t-1) + \lambda(k_i - \bar{k})] + (1-a)y_i(t), \tag{22}$$

where:

$\bar{u}_i(t)$ is the smoothed once-a-month load deviation for the ith line finder group for month t,

$\bar{u}_i(t-1)$ is the same quantity for the previous month,

λ is the CCS/MS derived for the current month by dividing the sum of the once-a-month loads in hundreds of call seconds (CCS) for the loading division by the sum of the main stations (MS) in the loading division,

k_i is the change in the number of main stations from last month for the ith line finder group,

$\bar{k}$ is the average of k_i,

$y_i(t)$ is the deviation of the once-a-month load for the ith line finder group from the average of the loading division,

a is a constant that weights the previous month's data to those of the current month,

and

$$F = 1 - (N-3)(.034\bar{z})^2/\sum w_i^2, \tag{23}$$

where:

N is the number of line finder groups in the loading division,

$\bar{z}$ is the loading division average of the line finder group once-a-month loads,

w_i is an intermediate variable, given by eq. (22).

The line assignment recommendation for the ith line finder is:

$$q_i(t) = -\bar{\mu}_i(t)/\lambda \tag{24}$$

rounded to the nearest integer. A positive value of $q_i(t)$ indicates that lines should be added to line finder group i.

Examination of eq. (22) shows that the inclusion of the k term accounts for the line assignment activity (connects and disconnects) of the previous month.

The value of a has been set at 0.32 from data studies. This means that 68 percent of the weight of the current line assignment recommendation is based on the current month's data; 32 percent of the weight is based on history. This leads to consistent month-by-month line assignment recommendations.

The quantity F of eq. (23) is the James-Stein factor and is a measure of the confidence to be placed on the calculations based on statistical volatility of the data. It is always less than unity and typically runs about 0.9.

The same quantities used to generate the line assignment recommendations are also used in the load balance index calculation.[11] First the nonstatistical deviation of the current month's loads $D(t)$ is calculated as:

$$D(t) = \sum w_i^2/N - (.034\bar{z})^2(N-1)/N. \tag{25}$$

This current month's value is used with that of the previous month $D(t-1)$ to derive a smoothed value L by a constant b as follows:

$$L = bD(t-1) + (1-b)D(t). \tag{26}$$

The value of b has been set at 0.7 from data studies. This means that 30 percent of the weight of the smoothed deviation is based on the current month's data; 70 percent is based on history.

Finally, an imbalance factor is derived as:

$$(V/M)_{\mathrm{OAM}} = L/\bar{z}. \tag{27}$$

While this factor is a direct measure of the load imbalance, the calculation of the index has to account for the level of load being carried by the office, called percent capacity C. It is calculated as:

$$C = \bar{z}/z_{\mathrm{CAP}}. \tag{28}$$

The value of z_{CAP} is obtained from the SONDS capacity table for the number of line finders in the line finder group.

The imbalance factor and the percent capacity are used to derive the load balance index, as shown in Table VII. This table shows that a given imbalance factor results in a lower index as the carried load approaches capacity. Therefore, the index is a measure of service as seen by the subscribers, while the imbalance factor is a message to the network administrator to bring the office into balance as the load builds up.

The imbalance factor is calculated for each loading division because the lines are assigned by loading division. The office load balance index is derived by combining the imbalance factors of all the loading

Table VII—SONDS load balance index

Imbalance Factor V/M_{OAM}	Percent Capacity							
	≥96	95–86	85–76	75–66	65–56	55–46	45–36	≤35
0–0.8	100	100	100	100	100	100	100	100
0.8–1.3	99	99	100	100	100	100	100	100
1.3–1.7	98	99	99	100	100	100	100	100
1.7–2.1	97	98	99	99	100	100	100	100
2.1–2.6	96	97	98	99	100	100	100	100
2.6–3.3	94	96	97	98	99	100	100	100
3.3–4	92	95	96	98	99	99	100	100
4–5	90	93	95	97	98	99	100	100
5–6	85	92	94	96	98	99	99	100
6–7	75	90	93	95	97	98	99	99
7–8	63	87	91	94	96	98	99	99
8–10	45	82	90	93	96	97	98	99
10–13	18	68	87	92	95	97	98	99
13–16	11	57	84	91	94	96	97	98
16–25	1	20	64	87	92	95	97	98
25–33	0	11	48	83	91	94	96	97
33–50	0	6	28	74	89	93	95	97
≥50	0	1	14	61	86	91	94	96

divisions according to the relative number of main stations within each loading division.

IX. ANALYSIS OF EVE DATA

One of the advantages of EVE techniques is that one can determine if a set of data points follows an extreme value distribution by establishing a set of quality control limits. It is seen that the data points of Fig. 3 violate their quality control limits. When nine low-side outliers were removed, the remaining data points, shown in Fig. 4, fall within their (updated) quality control limits.

This feature is used in SONDS in a powerful data validation test to determine if a new data point is an acceptable member of an established distribution. But first one has to establish the reference distribution for each usage measurement.

9.1 Start-up tests

The reference distribution is established in SONDS by accepting 20 days of peak data with no screening. At the end of this time, called the start-up period, three low-side, followed by two high-side, outlier tests are applied to the data. These are called censored outlier tests because the contribution of the data point being tested is removed from the estimates of the mean and the variance. This is done so that a contaminated data point will not be accepted because of its influence on the test conditions.

The following equations are given by Friedman[8] for estimating the mean and the variance of the equal "candidate busy hour" load distribution from the censored sample set of peak data points via the moments method. The censored mean is

$$\hat{\mu} = \bar{x}_{n-1} - \frac{\hat{\sigma}}{n-1}\{1.267n - E[y_{(n)}]\} \tag{29}$$

and the censored variance is

$$\hat{\sigma}^2 = \bar{s}^2_{n-1}\,(n-2)/(0.416(n-1) - \frac{n}{n-1}\{E[y^2_{(n)}] - 2.534E[y_{(n)}] + 1.605\}) \tag{30}$$

where

$\bar{x}_{n-1}$ is the sample mean and $\bar{s}^2_{n-1}$ is the sample variance of the remaining $n-1$ data points as each highest or lowest data point is removed

and

the reduced variate $y = (x - \mu)/\sigma$.

Table VIII—Low-side and high-side censored means and variances

n	Low-Side		High-Side	
	$E[y_1]$	$E[y_1^2]$	$E[y_n]$	$E[y_n^2]$
17	—	—	2.515	6.50
18	0.182	0.114	2.535	6.60
19	0.168	0.108	2.554	6.70
20	0.156	0.102	2.572	6.79

$E[y_{(n)}]$ is the expected value of the largest of the n data points, each of which has been selected as the largest of six samples drawn from a normal distribution with zero mean and unity variance. $E[y_{(n)}^2]$ is the expected value of the largest square of n such data points. $E[y_{(1)}]$ and $E[y_{(1)}^2]$ are the equivalent values for the lowest of n data points. Values for both sets of parameters as used in SONDS are listed in Table VIII as given by Friedman.[8]

Using the censored mean and variance of eqs. (29) and (30) with $E[y_{(1)}]$ and $E[y_{(1)}^2]$, we compute $F(x_L)$ from eq. (8). We then compute the low threshold L as:

$$L = 1 - [1 - F(x_L)]^n \leq 0.06. \tag{31}$$

If the value of L is less than 0.06, the data point is rejected and the next lowest data point is tested. This value of the threshold will reject six percent of valid lowest-of-twenty data points.

Using the censored mean and variance of eqs. (29) and (30) with $E[y_{(n)}]$ and $E[y_{(n)}^2]$, we compute $F(x_H)$ from eq. (8). We then compute the high threshold H as:

$$H = 1 - [F(x_H)]^n \leq 0.01. \tag{32}$$

If the value of H is less than 0.01, the data point is rejected and the next highest one is tested. This value of the threshold will reject one percent of valid highest-of-twenty data points.

If either three low-side data points or two high-side data points fail the threshold tests, the complete set of start-up data is rejected and a new start-up interval is automatically initiated. Restarts almost completely disappeared after the normal-to-the-sixth replaced the Gumbel as the assumed distribution.

If both thresholds are passed, the particular measurement is considered operational and a different set of data validation tests are applied, as described in the next section.

9.2 Operational tests

When the start-up data set for a particular measurement is accepted, its mean and variance are used for the initial operational data validation tests. The values of the mean and the variance used for subsequent operational tests are continually updated as new data are accepted. The failures of certain of these operational tests will cause an exception report to be issued. Other test failures will increment a counter, which may or may not result in an exception report. If a given day's data contain too many exception conditions on individual measurements, the whole set of daily data is rejected and an exception report is issued.

The following are the operational tests applied to EVE data.

9.2.1 *Upper physical bounds test*

The received data should not be greater than 36 CCS times the number of components. (This test is also applied to non-EVE usage measurements.)

9.2.2 *Lower physical bounds test*

The received value must be greater than zero. A zero value for the daily peak usage measurement of a group of components to be subjected to EVE analysis would indicate a trouble condition or a mistake in measurement assignment. (Zero is allowed for both non-EVE usage measurements and peg count measurements.)

9.2.3 *Once-a-month load exceedance test*

This test compares each new data point with the once-a-month load computed from eq. (16) with the current values of mean and variance. A data point higher than that value will increment an up-down counter; a data point lower than that value will decrement the counter (but not below zero). An exception report is issued at the count of three. In this way traffic trends or bad data can be detected.

9.2.4 *Coefficient of variation test*

If a register or piece of equipment is inoperative, the coefficient of variation of the moving average of the mean and the variance falls below the threshold value of 0.025, and an exception report is issued.

9.2.5 *High and low outlier tests*

Before a new EVE data point is accepted, it must pass both a low and high outlier test. The current values of the mean and the variance are used to compute $F(x_i)$ from eq. (8).

The low threshold L is computed as:

$$L = 1 - [1 - F(x_i)]^{20} \leq 0.06. \tag{33}$$

If the value of L is less than 0.06, the new data point is rejected and it is not used to update the mean or the variance. This threshold value will reject only about one in 300 valid daily peak values.

The high threshold H is computed as:

$$H = 1 - [F(x_i)]^{20} \leq 0.01. \tag{34}$$

If the value of H is less than 0.01, the new data point is rejected and not used for updating. This threshold value will reject only about one in 2000 valid daily peak values.

By solving eqs. (33) and (34) numerically and using eqs. (13) and (14), the acceptable range for a new data point in terms of the current values of the mean and standard deviation can be found as:

$$\bar{x} - 2.442\bar{s} \leq x_i \leq \bar{x} + 3.889\bar{s}. \tag{35}$$

X. SYSTEM CONFIGURATION

A block diagram of SONDS is shown in Fig. 5. The most prominent feature is the single, central computer for the Bell System. The source of the traffic data is a data collection unit in each step-by-step office. Output reports are delivered to the user by shared access to a local reports terminal. The user controls the system by means of the user interface. All of these elements are connected together by the switched network.[12]

One of the advantages of the system is its use of the low-cost switched network instead of dedicated data links for the data collection and report distribution. The use of the switched network and computer resources is optimized by middle-of-the-night polling, followed by early morning reports transmission. The scheduled reports are available to the user at the start of business on the due date, and often contain results as recent as the previous day.

Consolidating the software in a single computer led to reduced development costs and timely revisions to correct bugs or to incorporate enhancements. The time-shared nature of the computer imple-

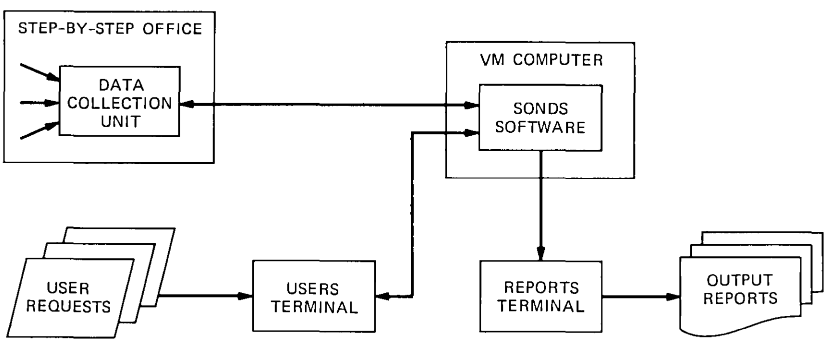

SONDS – SMALL OFFICE NETWORK DATA SYSTEM
VM – VIRTUAL MACHINE

Fig. 5—Block diagram of SONDS.

mentation made it natural to put the user manual on-line. This not only reduced the cost of documentation but made it more timely, user-oriented, and easily accessible. This type of documentation has reduced costs further by precluding the need for any formal training of the (many and casual) users.

Individual sections of the user manual can be obtained on-line. Also, the entire manual or any of its sections can be printed off-line and mailed to the user as the result of an on-line request. The date of the user manual is the date of its most recent change and the user is notified of this each time he or she logs onto the computer. The table of contents contains the date of the most recent change of each section. Each page contains the date of the most recent change incorporated thereon and the change itself is noted by marks in the margin.

To enhance the fully mechanized and less paper aspects of the system, the computer contains the capacity tables for the components to be engineered by use of the system's data. This allows the system to produce fully analyzed results, requiring no further manual manipulations by the user.

The user interface is the key to making the system easy to operate and completely under user control. The most important function of this interface is entering the office information that SONDS needs to produce the output reports. This information includes the component names that SONDS presents in the output reports. It also includes the number of the components that SONDS needs to access the capacity tables for those components. Further, SONDS must know about the units of measurement to convert the received data to the proper units: CCS for usage, units for peg counts, and percent for dial tone delays.

The interface contains a series of prompts that help the user enter the information. Out-of-bound values (probably typographical errors) are detected immediately and the user can correct them at this time or add them later if the proper value is unkown. Whenever the user does not know what to do next, typing a question mark will produce a list of acceptable inputs.

Another application of the user interface is to request previously generated reports that were either misplaced or not received because of a reports terminal problem. Typically, the most recent three or four scheduled reports are kept in this on-line archive; the (daily) exception reports are kept for two weeks. Also, the user can demand a current version of any of the scheduled reports. He will receive a partial version of the next scheduled report. This request, however, will not interfere with the regular generation of the scheduled report.

Perhaps the most powerful use of the interface is to command the system to go into hourly data collection. Without interfering with the regular nightly polling and report generation, the computer will initiate a data poll each hour between 6 a.m. and midnight, or as specified by the user. Thus, time-consistent hourly data can be generated. This feature is needed for limited use to verify data for a newly equipped office, to troubleshoot a data problem, or to conduct special studies.

10.1 Data collection unit

To obtain the desired data from the step-by-step offices, SONDS has specified a set of features, some of which are optional, for a Data Collection Unit (DCU). The three vendors and the four models of DCUs used on SONDS are Western Electric (PDT2A) (13), Tele-Sciences (TE450), and Alston (Model 565 and 566). Each microcomputer-controlled data terminal is connected to a data set that has an automatic answer capability. This is used for dial-up access by the SONDS computer to obtain the measurement data. This dial-up access can also be used by maintenance people for remotely running diagnostics, as specified by the vendor.

The DCU collects both peg count and usage data, the latter for either one-second or ten-second scan rates.

The output register set for each measurement contains three registers. One is the active register, which accumulates the counts during the measurement hour. At the end of each hour the active registers are reset automatically and the contents are made available to the previous-hour and long-term registers.

For each measurement, the new hourly value from the active register overwrites the value contained in the previous-hour register where it is retained for the next hour. The new hourly value also goes to the long-term register, which has been specified to be either a peak or an accumulative register. If the long-term register is peak, the new hourly value is compared to that stored in the peak long-term register and the higher of the two is retained in the peak long-term register. If the long-term register is accumulative, the new hourly value is added to its contents.

The long-term registers are reset after reading by the SONDS computer during the nightly poll. The values obtained from the long-term registers are either the highest hourly value experienced since the last resetting or the accumulations of the hourly values since the last resetting.

The DCU reads out the register contents on command from the SONDS computer. During the nightly poll, the computer requests the contents of both the previous-hour and long-term registers and then the computer resets the long-term registers. For the collection of hourly data during the daytime, the computer only requests the contents of the previous-hour registers and disconnects without any resetting. The SONDS computer will accept up to 200 register readings in response to a given request.

Consideration of mid-train engineering and the possible use of SONDS data for the pre-installation engineering of an electronic replacement office has led to the definition of two optional features. One is called the peak of the sum. What is desired here is the (24-hour) peak value of a measurement that itself is the sum of other measurements. Since the individual measurements might peak at different hours, the sum of the (24-hour) peaks will, in general, not be equal to the (24-hour) peak of the sum of the hourly values. Therefore, the hourly summing and (24-hour) peaking must be done in the DCU. The output is available in one of the regular peak long-term registers.

The other optional feature is called the peak of the difference. The desired result is the peak value of a measurement that itself is the difference between two sets of other measurements. Again, the peak of the difference is, in general, not equal to the difference of the peaks. The output is available in one of the regular peak long-term registers.

Besides these measurement features, the DCU satisifes several other SONDS requirements. Each response has a specific format that the SONDS computer expects. It includes the DCU time of day, which is resettable by the SONDS computer, and a status word to indicate several conditions within the terminal that might make the data suspect and that would aid in the trouble location of a DCU problem. The response also contains a checksum value of the response itself. The SONDS computer calculates the checksum of the received data and compares it to the received checksum. In this way the SONDS computer can be assured that it has an accurate copy of the register readings as they exist in the DCU.

10.2 Data calculations

To generate the several types of output reports, SONDS has four different types of data analysis. In entering the office information, the user selects one type of analysis for each measurement.

10.2.1 EVE analysis

EVE analysis is used for the usage measurements in the main part of the SONDS monthly report and for selected measurements in the miscellaneous section of that report.

As seen in eq. (16), the once-a-month load is calculated from the mean and the standard deviation (the square root of the variance) of the daily peaks. Therefore, it is not necessary to retain all the received values; only an estimate of the mean and the variance is required. This is accomplished in SONDS as a continual updating of these values, once they are derived at the end of the start-up period.

The equations for this updating are as follows:

$$\bar{x}_i = \mathrm{p}x_i + (1 - \mathrm{p})\bar{x}_{i-1} \tag{36}$$

$$v_i = p(x_i - \bar{x}_i)^2 + (1 - p)v_{i-1}, \tag{37}$$

where

x_i is received value for the ith day,

$\bar{x}_i$ is the mean at the ith day,

$\bar{x}_{i-1}$ is the mean from the previous day,

v_i is the variance at the ith day,

v_{i-1} is the variance from the previous day,

and

p is a constant.

This updating method reduces the size of the database. Only two values need to be stored for each measurement; no raw data are retained. It gives an updated value for the mean and the variance, which is the exponential weighting of all the previously received data. The value of p was selected as 0.095 to make the average age of the weighted data the same as if the latest month of data itself had been used. One effect of this calculation is that 87 percent of the weight of the current estimate is derived from data received during the most recent month; 13 percent of the weight is due to data received in prior months. This approach, besides minimizing data storage requirements, is also compatible with current TCBH practice of reporting monthly results.

Another advantage of this updating technique described in eqs. (36) and (37) is that it makes the calculation robust for (occasional) losses of data. For any day that data are not obtained, the estimate is not updated, that is, the estimate is stale by that amount. The exponential weighting of previous data reduces this staleness, and the effect of a lost data point is soon eliminated.

10.2.2 PA analysis

Periodic Average (PA) analysis is used for dial tone delays and for those usage measurements in the Miscellaneous section of the SONDS monthly report that do not lend themselves to the EVE analysis.

The equation to update the estimate of the mean is:

$$\bar{x}_j = x_j/j + \bar{x}_{j-1}(j - 1)/j, \tag{38}$$

where

j is the jth day during the current period.

This analysis type is called Periodic Average because it calculates the true, straight-line average during the period of the monthly report; the value is reset when the monthly report is generated. PA also retains the single largest value received during the calculation period.

10.2.3 SMA analysis

The third type of analysis is called the Service Month-to-date Accumulation (SMA). This analysis type is used for the Automatic Number Identification (ANI) measurements needed for the Network Switching Performance Measurement Plan (NSPMP). It basically gives the month-to-date summation of the daily values of the (24-hour) accumulative long-term registers. Similar to the PA analysis, the SMA analysis also produces the periodic average and the single highest received value.

10.2.4 TOPC analysis

The fourth type of analysis is called Total Office Originating Peg Count (TOPC). It is specified during the entry of the office information to be applied to originating PC measurements. These measurements are stored for the generation of the (calendar) month TOPC report in which certain weekday and weekend ratios and averages are calculated. Also certain adjustments are made for the loss of individual days of data.

10.3 Output reports

As part of its goal to be a fully mechanized system, SONDS automatically calculates, generates, and transmits several output reports. The monthly report, the (monthly) Total Office Originating Peg Count Report and the weekly report are regularly scheduled. Exception reports are issued daily as needed. Only the hourly data report depends on user initiation (to request hourly polling). It is generated and transmitted daily to give the previous day's polling results.

10.3.1 Monthly report

The heading of the monthly report gives the office name and the date of the report, which is typically the 22nd of the month. The report heading also gives the date that the office information file was last updated; this is a message to the user because his/her monthly updates of office component counts can affect the numbers presented on the report.

Another part of the header information is the number of days of valid EVE data. Although SONDS collects data seven days a week, EVE data are only used for five days. Typically, these are the weekdays, but the user can include Saturday, Sunday, or both in the five days if weekends are high traffic days in a particular office. There are only 20 to 23 EVE days in a month and the number of valid days indicates the weight of the EVE results due to data collected during the current month. Fifty percent of the weight occurs in the current month if seven days of valid EVE data are obtained. Data for measurements with fewer than seven days in the current month are shown on the monthly report, but they are flagged with a ?. Incidentally, if a measurement is in start-up, it is flagged with an asterisk.

The first part of the Originating Results section of the monthly report contains information on the Line Finder Groups (LFGs). This includes the items discussed in prior sections, such as the current load as percent capacity, main station assignment guide, and maximum value of Dial Tone Delay (DTD), all on a per-LFG basis. Also there are total office results in terms of the Dial Line Index (DLI) and Load Balance Index (LBI).

The Originating Results section of the monthly report also contains an outgoing trunk group section. Certain columns containing information not previously discussed were specifically designed for the trunking people. One such column gives the number of circuits required to carry the measured load at the once-a-month blocking objective; this quantity is equivalent to one of the main outputs of the Trunk Servicing System (TSS). The second column can be used as the TSS (TCBH) study period load; it is derived from the proper TCBH trunk capacity table, depending on the trunk accessing arrangement, for the number-of-trunks required specified in the previous column. SONDS retains these study period values for the most recent 12 months and each month prints the highest value in the third column. This is the base period load, which can be the input to the Trunk Forecasting System (TFS). The fourth column prints the month when this highest load occurred.

The Terminating Results section of the monthly report contains similar information for selectors, connectors, and incoming trunks.

The Miscellaneous Results section contains any other measurements the user has decided to add to the system. The user can select peg count or usage measurements, peak or (24-hour) accumulating registers, and has a choice of EVE, PA, or SMA analysis types. One of the main uses of this section has been to obtain the ANI counts for NSPMP reporting. Another use has been to obtain the results of the summing and differencing registers to obtain data for pre-installation engineering of an electronic replacement office. Of course, the user could also choose to obtain Last Trunk Usage (LTU) or Subscriber Line Usage (SLU) measurements.

10.3.2 Weekly report

The Weekly report is scheduled for Monday mornings; it gives daily results for seven days up to the previous Friday. The main part of the report gives the "always busy" data obtained by the SONDS computer during the nightly polling by reading the previous-hour registers of the data collection unit. Integer multiples of 36 CCS in these data are indicative of busied-out or stuck components. The report also contains the daily peak values of dial tone delays so that the user can determine if high values of dial-tone delays were caused by maintenance conditions indicated by the always-busy data. The report contains the daily values of total office originating peg counts to make these data available to the Division of Revenues people before the end of the (calendar) month. The report contains the month-to-date results of the DTD (TCBH) and NSPMP measurements so the user can monitor the reportable service results during the month.

10.3.3 Exception reports

Exception Reports (ER) are issued daily as needed when abnormal conditions are experienced during the night's poll. These conditions might be a failure to connect (after four attempts), error bits set in the status word, and failed data validation tests. These ERs are short messages giving the ER code, the descriptive name, and the received data, if any. There are some 52 different types of ERs and each has a one- to three-page description in the SONDS User Manual to help the user locate the problem. Also, the user can access the SONDS computer to obtain a list of the ERs for the past two weeks.

10.3.4 Total Office Originating Peg Count report

The Total Office Originating Peg Count report has a calendar-like format showing the daily values with weekday values separated from weekend values. Weekday-to-weekend ratios are calculated as well as values of average business days and calendar days. The computer contains algorithms to adjust for isolated days of missing data and it computes adjusted averages.

10.3.5 Hourly data report

Hourly data reports are generated only when hourly polling has been requested by a user. The reports are received in the morning and contain up to 16 hours of polling results of the prior day. This produces time-consistent hourly data, including dial tone delays, to track the flow of traffic into, out of, and across the office. For each hour, the report also contains a traffic balance calculation to determine if permanent signals are inflating the measured value of line finder usage.

10.3.6 Retention periods

As mentioned previously, SONDS retains on-line for user access two weeks of exception reports and the last three or four reports of other types. The user can also receive on-line a demand version of any of the reports, which can supply data as recent as that of the previous day.

XI. COMPUTER AND SOFTWARE DESCRIPTION

The SONDS software consists of approximately 260 programs residing on an AT&T Information Systems computer located at the Corporate Computer Center in Piscataway, N.J. SONDS shares with many other corporate users an Amdahl 470/V8 processor under the Virtual Machine/Conversational Monitor System (VM/CMS) operating system.

11.1 VM/CMS computer facility

VM/CMS manages the resources of a real-time computing system in such a way as to make them available to many users at the same time. Through VM/CMS, each SONDS user has a virtual machine composed of a virtual system console (a terminal), virtual CPU, virtual storage (disk), and virtual channels and Input/Output (I/O) devices.

The VM/CMS Auto-dial Facility allows application programs to initiate a direct distance dialing connection to a remote terminal properly equipped with an auto-answer feature. SONDS uses the Auto-dial Facility in two ways: It extracts data at a 300-baud rate from the data collection unit in each SONDS office by nightly polling, and it prints output reports on users' printing terminals at the rate of either 300 or 1200 baud. This report distribution process is faster and less expensive than mailing reports directly and can make use of terminal equipment outside normal working hours.

11.2 SONDS software

Each of the major modules of the SONDS software is described in the following sections.

11.2.1 User interface

The user enters and maintains office information on the user interface. It can also be used to change the polling mode of the office: start (nightly) polling, stop polling, and schedule hourly polling.

The user accesses the Virtual Machine (VM) system over the switched network via any ASCII-compatible terminal at either 300 or 1200 baud. After entering a user identification (userid) and a password, the user is connected to the SONDS user interface. The system gives any short news items that the user should know, and the date of the last update to the User Manual. Whenever the user doesn't know what to do next in the interactive dialogue with the system, he or she can enter a question mark. The system will respond with a list of appropriate entries.

11.2.2 Users' database

The office information file entered by the user via the user interface is stored in the users' database. Also the traffic information obtained from the data collection unit is stored here. All the data are stored on an office-by-office basis. Some administrative reports, however, do combine information on how SONDS is working for all the offices on a userid and all the offices for a company.

11.2.3 System files

The system files primarily contain office status information to determine which activities should be started and when it is appropriate to start them. They are updated as the scheduled activities are performed. They thereby always contain the current information and can be queried by the SONDS Support Team to provide a current report of system status.

11.2.4 Master scheduler

The master scheduler controls four main activities: nightly polling, hourly polling, report generation, and report transmission.

Nightly polling occurs seven days a week, including holidays, during low traffic periods, usually between midnight and 6:00 a.m. local office time. The polling load is distributed over as many hours as necessary, taking advantage of the different time zones across the country.

Hourly polling may be requested by the user through the user interface for as many as six days of polling. The user specifies the earliest hour and the latest hour for polling, between 6 a.m. and midnight. Hourly polling can be started on the next clock hour. Hourly polling increases SONDS costs and charges, which in turn have prevented abuse of this feature.

The master scheduler activates the report generator and report transmitter processors according to the schedule for each of the five output report types, described in Section 10.3. The master scheduler also establishes an efficient order of nightly polling, hourly polling, report generation, and report transmission by considering the number of Automatic Calling Units (ACUs), the number of offices per time zone, estimated time to perform each function, etc.

The master scheduler runs each hour, usually starting at one minute past the hour and running for 15 seconds to 2 minutes. In addition, one of the major objectives of the master scheduler is to invoke "batch jobs" as needed for the poller, the report generator, and the report transmitter. Once initiated, these batch jobs will run on their own without master scheduler intervention.

11.2.5 Poller

The principal poller is activated during the night to place a call over the switched network to the data collection unit using the ACUs. The data as well as the status information of the DCU are sent to the users' database and the system files are updated. The action of the hourly poller is similar except that it is activated during daytime hours.

11.2.6 Report generator and transmitter

When the report generator is activated by the master scheduler, it obtains the traffic data from the users' database, applies the necessary algorithms to calculate the numbers that appear on the report, and stores the report data. The report transmitter then places a call over the switched network to the user-specified reports terminal via the ACUs, retrieves and formats the report data, and transmits the report to the users' terminals over the switched network.

11.2.7 Map Management System

The PDT2A data collection unit contains a software map to cross connect the input leads to the output registers.[13] This software map is created by the Interactive Map Assembly Program (IMAP), which exists on a time-shared DEC 10 computer at a Western Electric facility. Using the VM batch facility, the Map Management System will accept a user-generated request to place a call over the switched network via the ACUs to access IMAP and obtain the most recent edition of the software map, which is then stored in the users' database for down-loading to the PDT2A via the poller.

11.2.8 Administrative System

The Administrative System is a group of programs that monitor the activities of all the other modules, recording their successes and failures. The system generates several types of detailed logs and summary reports, which are available to the SONDS Support Team to determine the health of the SONDS and VM systems and to identify and resolve problems in a timely manner. The user has access to a (pertinent) subset of these administrative reports via the user interface.

11.2.9 Billing system

Billing for VM/CMS is done once a month and a separate statement is produced for each userid. Each statement includes a detailed listing of the charges for each off-line activity the userid had during the month, and a breakdown of the usage on the userid by accounting information, which identifies the initiator of the on-line activity and the type of off-line activity. To increase the user's cost consciousness, the user interface prints an approximate charge, upon logoff, for that user session.

XII. ACKNOWLEDGMENTS

The authors are indebted to K. A. Friedman, who was the main analyst throughout this project. He proposed and analyzed the normal-to-the-sixth distribution, the outlier detection tests, and the dial tone delay conversions. He collected and analyzed the vast quantities of field data needed to derive extreme value service criteria. He derived the new engineering tables used in SONDS for all the major components of the step-by-step office. R. A. Farel's work on the use of extreme values for line assignment and load balance, the results of which are published for the first time herein, was a significant contribution towards making SONDS a member of the TNDS family. We are particularly appreciative of the contributions, suggestions, and critiques of J. G. Kappel throughout the project. The experimental version of SONDS, which was used to develop many of the concepts reported herein, was possible due to the talent of the software team led by G. B. Foley and S. Chen and to the dedicated field support of R. R. Plum. The one thousand SONDS offices now operating on the VM computer would not have been possible without the dedicated effort of J. Schwartz and his software development team at AT&T Information Systems. Their generous review of the pertinent parts of this paper is gratefully acknowledged; the significant contribution of K. L. Dziedzic is specifically acknowledged.

REFERENCES

1. S. A. Karlsson, "The Dimensioning of Telephone Traffic Routes for Measured Integrated Peak Traffic," Fifth Int. Teletraffic Cong., New York, June 14–20, 1967, pp. 7–17.
2. S. B. Gershwin, R. V. Laue, and E. Wolman, "Peak-Load Traffic Administration of a Rural Line Concentrator," Seventh Int. Teletraffic Cong., Stockholm, Sweden, June 13–20, 1973, pp. 146/1–7.
3. S. B. Gershwin, R. V. Laue, and E. Wolman, "Peak-Load Traffic Administration of a Rural Multiplexer With Concentration," B.S.T.J., *53,* No. 2 (February 1974), pp. 261–81.
4. R. V. Laue, "Extreme-Value Engineering Designs for Peak Traffic," Bell Lab. Rec., *56,* No. 2 (February 1978), pp. 38–42.
5. E. J. Gumbel, *Statistics of Extremes,* New York: Columbia University Press, 1958.
6. D. H. Barnes, "Extreme Value Engineering of Small Switching Offices," Eighth Int. Teletraffic Congr., Melbourne, Australia, November 10–17, 1976, pp. 242/1–7.
7. D. H. Barnes, "Observations of Extreme Value Statistics in Small Switching Offices," Ninth Int. Teletraffic Cong., Torremolinos, Spain, October 16–24, 1979, 9 pages.
8. K. A. Friedman, "Extreme Value Analysis Techniques," Ninth Int. Teletraffic Congr., Torremolinos, Spain, October 16–24, 1979, 7 pages.
9. L .H. C. Tippett, "On the Extreme Individuals and the Range of Samples Taken from a Normal Population," Biometrika, *17* (1925), pp. 364–87.
10. R. A. Farel, unpublished report.
11. R. A. Farel, unpublished report.

12. G. B. Foley, J. J. O'Connor, and J. Schwartz, "SONDS Says 'Sound Off' to Small Offices," Bell Lab. Rec., *58,* No. 10 (November 1980), p. 333–8.
13. J. Arthur Grandle, Jr., and Richard R. Plum, "A Versatile Data-Gathering Tool," Bell Lab. Rec., *57,* No. 8 (September 1979), pp. 227–31.

AUTHORS

Douglas H. Barnes, B.S. (Industrial Engineering), 1938, Lehigh University. Mr. Barnes joined Ohio Bell Telephone Company in 1941 as an Assistant District Traffic Superintendent, where he worked on a variety of traffic assignments including operator force management, traffic staff, service observing, local and toll office traffic engineering, trunk engineering, and division of revenue studies. He came to Bell Laboratories in 1951 to work on requirements for traffic measurement needs and systems. He was appointed Head of the Traffic Systems Engineering and Analysis Department in 1961. He was responsible for systems engineering of all traffic data collection and processing systems. Some of these are: Traffic Usage Recorders, Traffic Data Recording System, Engineering and Administration Data System, Mechanized Service Observing, Force Administration Data System, Subscriber Line Usage System, and Small Office Network Data System. He was also responsible for various improvements in traffic methods, such as the adaptation of Extreme Value Engineering. He retired from Bell Laboratories in 1979 and has since been a Teletraffic Consultant.

John J. O'Connor, B.S.E.E., 1951, M.S.E.E., 1952, Columbia University; Ford Instrument Co., 1952–1954; Electronics Research Laboratories, C.U., 1954–1958; Hathaway Instrument Co., 1959–1961; Martin Co., 1961–1964; Bellcomm, 1965–1972; Bell Laboratories, 1972—. At Ford Instrument Co. Mr. O'Connor worked on terrain synthesis for an automatic radar map-matching guidance system. At Electronics Research Laboratories he studied and worked on simulating low-frequency (0 to 0.1 Hz) noise of radar glint in the Nike missile system. He worked on the man-rating of the Titan III launch vehicle at the Martin Co., and on the spacecraft attitude control propellent budget and structural oscillations (POGO) of the Apollo-Saturn launch vehicle for Bellcomm. In 1972 Mr. O'Connor joined Bell Laboratories, where until 1973 he worked on the maintenance aspects of electromechanical switching machines. From 1973 to 1976 he was involved with various traffic-related problems of toll switching machines. From 1976 until 1982 he worked on SONDS, and from 1982 until the present has been studying the project management aspects of the Dynamic Nonhierarchical Routing network.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Performance Measurement/Trouble Location

By D. R. ANDERSON* and M. J. EVANS*

(Manuscript received November 24, 1982)

The Total Network Data System (TNDS) includes a subsystem that measures TNDS performance and assists network administrators and those responsible for network traffic data collection in finding TNDS troubles. This interactive subsystem provides the telephone company managers with a versatile tool for a performance analysis of the many systems and organizations associated with TNDS. The troubleshooting features usually provide sufficient detail from which to specify corrective action. The performance information can be aggregated across all organizational levels from the single traffic unit to the entire Bell System.

I. INTRODUCTION AND SUMMARY

Good service to the customer has always been a hallmark of the Bell System. This service is ensured by a number of measurement plans used by Bell System managers to monitor the quality of service to the customer. Operations such as those associated with the Total Network Data System (TNDS) do not all have an immediate effect on the customer, but if neglected for a long period of time, could result in costly engineering mistakes. Thus, a TNDS performance measurment plan is needed so that the efficiency and integrity of TNDS operations are maintained at a level sufficient to support continued good service

* Bell Laboratories.

to the customer at reasonable cost. A TNDS Performance Measurement Plan (TPMP) was introduced in 1980. This measurement plan is highly automated, using a computerized support system called the Centralized System for Analysis and Reporting (CSAR).

In addition to highlighting performance, measurement plans often provide supplementary data useful to those responsible for assuring good performance. This function is especially important in TNDS operations because of the wide diversity in systems, organizations, geography, and timing associated with TNDS. Thus, assisting troubleshooting is an important feature of CSAR. It often allows the analyzer to identify the specific source of a data problem while still at the terminal.

II. TNDS PERFORMANCE MEASUREMENT PLAN (TPMP)

2.1 Need and objectives

In 1976, as the TNDS systems were being widely deployed in the Operating Telephone Companies (OTCs), TNDS project managers at AT&T and systems engineers at Bell Labs undertook a field study of TNDS Operations in South Central Bell. This company offered an ideal environment for observing TNDS operations. Its applications included rural areas and large metropolitan areas across five states.

With the help of South Central employees, personnel at Bell Laboratories and AT&T were able to obtain valuable insight into TNDS performance and typical operations. Over several months the AT&T/Bell Labs team observed just how TNDS traffic data reports were obtained, and how accompanying TNDS administrative information was analyzed and used in a variety of operating company offices. It was observed that managers often had difficulty obtaining an overview of performance from the administrative reports generated by the various systems of the TNDS. Also, the geographical diversity of these systems, the time interval over which traffic data progress through the TNDS, and the volume of the reports from the systems all made it difficult and time-consuming to detect and pinpoint data problems. These difficulties were the result of the increasing size and complexity of TNDS and its wide deployment.

This study clearly pointed out a need for a comprehensive, automated performance measurement plan that could help to localize traffic data problems. The reports associated with the measurement plan should provide a concise management overview with enough detail about the type and location of problems to allow managers to apply resources to bring about solutions. The system generating the measurements should also make available to the TNDS administrators selected, detailed information about data as the data pass through the various systems in the TNDS. The measurement system should allow enough investigation, in increasing detail, to identify the location and time of problems. Such investigation should also identify enough about their cause to determine either the corrective measures or the testing and analysis that is required. The traffic data problems should be identified as quickly as possible so that critical traffic data studies are not jeopardized.

All of these requirements are best met by a system that obtains performance data directly from the other TNDS systems and provides a variety of reports on an on-line, interactive basis.

In considering measurement plans for TNDS, we noted a significant trend in other Bell System measurement plans. A measurement task force has recommended that measurement plans should avoid averaging results since they are aggregated together for several measured entities. This is important so that poor performance in one or two entities is not masked by good performance in a number of other entities. This principle is extremely valid in TNDS operations measurements where consistently poor performance in obtaining traffic data for an office could deny engineers the data needed to plan for its growth.

The method being adopted in Bell System measurement plans to preserve isolated indications of poor performance is to associate the results obtained for each measured entity into one of four bands: H, O, L, or U. The bands are designated as Higher than objective, Objective, Lower than objective, and Unsatisfactory. Then, aggregated results give the number of entity measurements in each of the bands. Thus, a poor performer is never lost in the crowd. In TPMP, there is usually no extra expenditure for flawless performance so there is no need for the Higher than objective category. Thus, TPMP only includes bands O, L, and U.

The measures important in monitoring TNDS performance are related to data completeness, data validity, and the accuracy of the record bases that associate the data with the telecommunications equipment. The TNDS Central Office Equipment Reports (COER) systems,[1,2] Trunk Servicing System (TSS),[3] and Load Balance System (LBS)[1] all present traffic data to end users (engineers or administrators) and, thus, can supply comprehensive information about TNDS performance. They all provide performance data to CSAR. The Trunk Forecasting System (TFS),[3] also an end-user system, does not provide data to CSAR at this time, but its inputs are directly dependent on the (CSAR-measured) TSS. The Traffic Data Administration System (TDAS),[1] which is an intermediate system in TNDS data flow, also makes available performance data to CSAR. This is done because TDAS is the final TNDS system for certain special study information and equipment types, and because analysis of TDAS results can allow early detection and isolation of data-gathering problems.

The above systems generate performance data pertinent to data completeness, data validity, and record base accuracy. In addition, the Individual Circuit Analysis System (ICAN)[1] was designed primarily to help administer the record base for usage data acquired on individual circuits through the Individual Circuit Usage Recording (ICUR) version of the Engineering and Administrative Data Acquisition System (EADAS).[4] As such, ICAN is an important source of CSAR record base performance data for those switching systems served by EADAS/ICUR.

2.2 Performance measures

Each system that furnishes CSAR with performance data has its own validity checks, record bases, and schedules. Each has its own way of measuring data completeness, data accuracy, and record base accuracy. Hence, the first level of the hierarchy for CSAR reporting is by the TNDS. Table I shows the TNDS performance measures generated by CSAR and gives a brief explanation for each. The results are reported as "problems," i.e., as percentage of missing data, validity failures, or record base errors. Two thresholds are assigned to each measure to distinguish between the three bands, O, L, and U. The threshold settings are determined differently for each measure because each is associated with a different part of the TNDS process, different operations, or different equipment. There is a guiding rationale, however, applied to all. The Objective band starts at zero and includes performance levels where minor flaws occur but are being attended to. The first threshold marks the transition to Band L, where the equipment errors and record base errors become significant and systematic. The upper threshold, which marks the beginning of band U, is set to indicate equipment problems that persist longer than is normally necessary for their remedy or to indicate the accumulation of record base errors.

2.3 Reporting

TPMP is designed to assist operating company TNDS managers at all organizational levels. Their access to reports is through the on-line, interactive computer system, CSAR. As mentioned earlier, the first level in the CSAR reporting hierarchy is the TNDS, from which CSAR has obtained the performance data. The second level is the organization for which the manager is responsible. For a given request, results from all reporting units, switching offices, trunk services, etc., in that organization are summarized or delineated according to the request. The dialogue with the computer enables the manager to get a

Table I—TPMP/CSAR performance measures

Measures	Performance Data Source	Description
DCU—Intervals missing	TDAS	The half-hour DCU* intervals that are required to satisfy all of the requests for obtaining data through TDAS, but where data are not recorded.
DCU—Intervals with any verification-device failures	TDAS	The half-hour DCU intervals where data are obtained but where one, or more, of the devices used to verify data accuracy have an imperfect score.
Detector test failures (only for TUs obtaining usage data through a TUR-traffic usage recorder)	TDAS	The number of discrepancies encountered when comparing a CU record base containing the number of circuits in each measured group with the usage inputs during a test scan when all circuits are simulated as busy.
Invalid detector test scans	TDAS	Discrepancies detected by the verification devices when the detector test is run.
ICAN/CU count mismatch	ICAN	The number of discrepancies encountered when ICAN compares the number of individual circuits in each circuit group in its record base with that in the CU record base.
ICAN/CU other error	ICAN	Any other discrepancies found when performing the record base comparison between ICAN and CU (e.g., circuit group type).
Suspect SCHV assignment	ICAN	The number of cases when ICAN reports an individual circuit assignment to a circuit group to be "suspicious" (these should be resolved by data administrators).
UA/UE SCHV usage	ICAN	The number of individual circuits that have measured usage but are designated in the EADAS/ICUR record base as being unassigned (UA) or unequipped (UE).
CGMT with insufficient data	TSS	The number of circuit group measurement (CGMT) types for which there was insufficient valid traffic data to use as a basis for trunk servicing or forecasting.
CGMT other invalid	TSS	Trunk group validity failures in CGMTs with sufficient valid data.
CGMT unassociated	TSS	The number of CGMTs for which traffic data were received but cannot be associated with a group in the record base.
LU invalid for LBI	LBS	The number of load units [(LUs) concentrators to which circuits are assigned] for which no valid data were provided for the load balance index (LBI) calculation.
LU measurements invalid	LBS	The number of load units where some data were obtained in the loading division but for which the LU had insufficient valid data (excludes cases where entire loading division missed data).
5XB COER control file with error	5XB COER	Indicates that data are not processed by 5XB COER because 5XB COER has detected an unresolved error in its record base.
5XB COER end hour missing	5XB COER	The number of hourly intervals for which 5XB COER expected data but did not receive it.
5XB COER validation failures	5XB COER	The number of failures detected when 5XB COER checked data validity.
SPCS COER missing data	SPCS COER	The number of hourly intervals for which 1 SPCS COER expected data but insufficient data were received.
SPCS COER sanity failures	SPCS COER	The number of data sanity check failures (gross data inconsistencies—tight thresholds apply).
SPCS COER cross-measure type test failures	SPCS COER	The number of cross-measure type tests that have failed (inconsistencies between data items—looser thresholds apply because peculiar switching environments may result in false alarms).

* The physical embodiment of DCUs (data collection units) differ from one type of switching machine to another. However, a DCU may generally be thought of as an aggregate of traffic data items that are scheduled and collected as a set.

breakdown by the next lower organizational level. More will be said later about these reports and the manager's use of them. In addition, CSAR furnishes detailed information for trouble analysis.

An overall summary of results for each TNDS operating company is made available to AT&T every month. The percentage of reporting-unit measures (see Table I) in the bands L and U are included in a Bell System measurement summary called "Network Results," which is published monthly for Bell System use.

III. THE CENTRALIZED SYSTEM FOR ANALYSIS AND REPORTING

CSAR is an on-line, interactive computer system that provides dial-up access six days a week to centralized databases holding over 250M bytes of information. The system combines the cost effectiveness of batch processing large volumes of data with the speed and convenience of on-line database access. CSAR can be accessed using almost any 110-, 300-, or 1200-baud asynchronous terminal. The daily operation of CSAR combines the resources of OTC computer facilities in 24 states, a Bell System data transmission network, a centralized computer system at AT&T, and the Direct Distance Dialing (DDD) network to support the TPMP and generate extensive information to assist in the operation and administration of TNDS. The reporting capabilities offer timely assessments of the performance of the traffic data collection and processing tasks, as seen by each system in the TNDS. Currently, eighteen OTCs use CSAR. Each TNDS installation within these companies becomes a remote source of TNDS performance data analyzed by CSAR and maintained in the central databases. Authorized users in each of these companies have dial-up access to the interactive component of CSAR, allowing flexible retrieval of information pertinent to their own company and certain Bell System results.

3.1 System configuration

The major component in the CSAR system configuration resides at the centralized computer site and receives the majority of its data from the distributed OTC TNDS processing sites via the Bell System telecommunications software system for computer-to-computer data exchange with OTCs (T-TRAN) data transmission network. The CSAR users gain access to the central computer system from existing remote terminals using the standard DDD telephone network. Thus, the system includes five major components that effectively combine existing Bell System facilities:

1. Distributed OTC computer sites
2. A T-TRAN data network
3. A Central AT&T computer site

4. A DDD telephone network
5. Existing remote terminals.

Figure 1 shows the CSAR computer system configuration. The CSAR implementation utilizes and coordinates a number of different hardware and operating system environments. Brief descriptions of the first three system components listed above illustrate the overall CSAR operating environment.

3.1.1 OTC computer facilities

The CSAR measurement process begins at the OTC's computer facility when the batch TNDS subsystem modules are run. It uses the Standard Operating Environment of IBM 370 series mainframe or a compatible computer running the Multiple Virtual Storage (MVS) operating systems. The performance data are gathered by software embedded in the individual TNDS subsystem modules. Thus, the individual TNDS subsystems accumulate the majority of the raw data required for CSAR performance measurement and analysis. A stand-alone CSAR module executes under this environment at the remote computer facilities to merge the data files from the separate TNDS components and prepare a single standard interface for data transmission.

3.1.2 T-TRAN data network

Each OTC uses the T-TRAN network as a data link for the transmission of its TNDS performance data to the central computer site at an AT&T Corporate Computer Center in New Jersey. Weekly CSAR transmissions are an integral part of the TNDS operations in the OTCs. Backup data are retained in each company on disk or tape.

The receipt of a transmission at the AT&T T-TRAN site (an IBM 370 system running MVS with the time-sharing option) automatically invokes a CSAR program that initiates the CSAR batch processing sequence. This first step identifies the incoming CSAR data files and sends all necessary information to a second computer system that is the actual host for the primary CSAR batch and on-line interactive processes. The program under the Multiple Virtual Storage/Time-Sharing Option (MVS/TSO) communicates to the CSAR host machine via a local Inter-Processor Communication (IPC) network.

3.1.3 Central host computer

The heart of the CSAR system also resides on one of the many computers that comprise the AT&T Corporate Computer Center. The CSAR software is developed and centrally executed on an Amdahl 470 series mainframe running under the Virtual Machine (VM) operating system using the Conversational Monitor System (CMS). This VM/

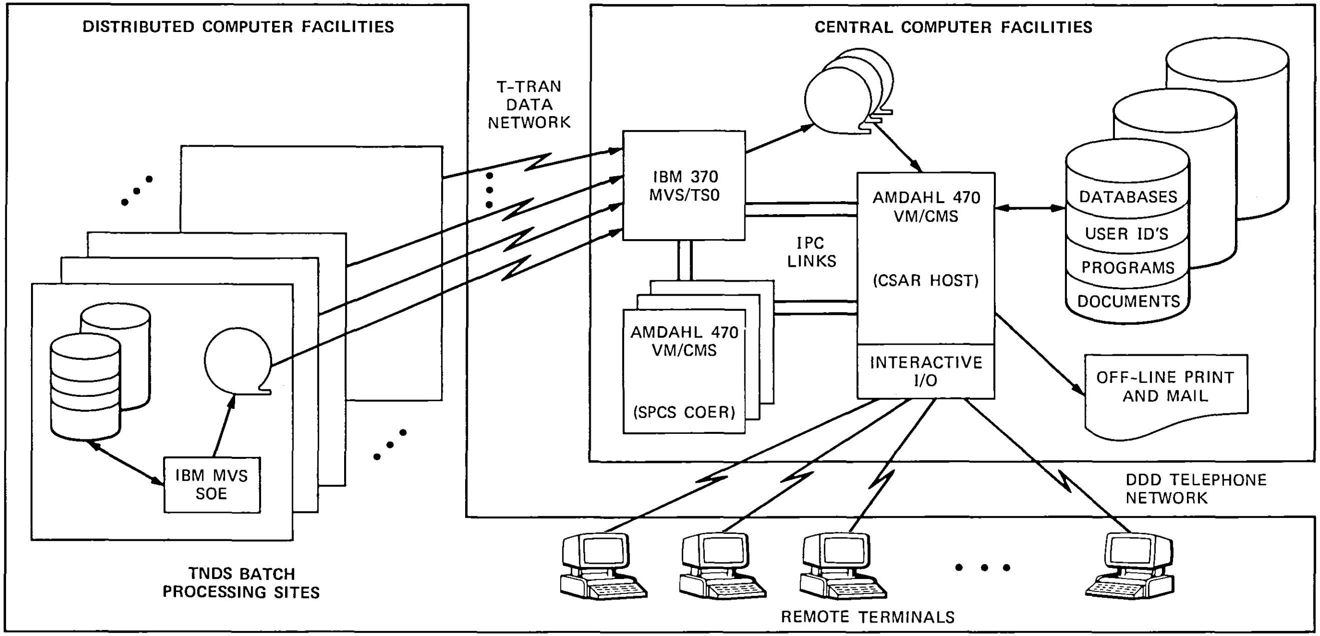

CSAR – CENTRALIZED SYSTEM FOR ANALYSIS AND REPORTING
ID – IDENTIFICATION
IPC – INTERPROCESSOR COMMUNICATION
MVS – MULTIPLE VIRTUAL STORAGE
SOE – STANDARD OPERATING ENVIRONMENT
MVS/TSO – MULTIPLE VIRTUAL STORAGE/TIME-SHARING OPTION
SPCS COER – STORED PROGRAM CONTROL SYSTEMS CENTRAL OFFICE EQUIPMENT REPORTS
TNDS – TOTAL NETWORK DATA SYSTEM
VM/CMS – VIRTUAL MACHINE/CONVERSATIONAL MONITOR SYSTEM

Fig. 1—CSAR computer system configuration.

CMS computer facility is the home of two other TNDS centralized operation support systems: the Stored Program Control System Central Office Equipment Reports system (SPCS COER)[2] and the Small Office Network Data System (SONDS).[5] Performance data are sent directly from SPCS COER to CSAR through a VM/CMS remote spooling interface.

This central host computer system offers the facilities and the processing capacity required by CSAR functions. The large volumes of data received from the 26 remote TNDS processing sites currently sending data to CSAR are scheduled for overnight processing using the VM/CMS batch facility. The analysis, correlation, and database loading occur during off-business hours at a reduced cost and without interfering with the typical interactive user.

IV. CSAR ON-LINE FEATURES

The primary purpose of CSAR is to provide on-line interactive access to central databases that contain extensive TNDS performance information. The on-line CSAR features can be logically grouped into four general categories:

1. On-line user documentation
2. TNDS performance reporting
3. Database and processing controls
4. Administrative information reporting.

The many features work together, thereby enabling each OTC to tailor the database capacity and reporting strategy to its own data requirements, organizational structure, and management needs.

The following paragraphs present an overview of the CSAR features categorized above. The description will also highlight the on-line reporting techniques and selected features that distinguish CSAR as an innovative approach to special-purpose on-line database access and reporting.

4.1 Interactive user dialogue

The CSAR interactive dialogue begins with the highest-level prompt: REQUEST =, which allows the user to select a system feature. The dialogue proceeds in a question and answer fashion until the user enters a specific request or inquiry.

The CSAR interactive user dialogue is designed to offer a user-friendly interface. The dialogue design gives the inexperienced user specific direction and guidance when questions and problems arise. For the experienced user, an abbreviated dialogue sequence can be strung together on a single line to save typing and interaction time. The user dialogue allows flexible inquiry and reporting. For frequent requests of tailored reports, the dialogue sequence can be defined once, saved, and activated when the need arises.

4.2 On-line documentation

A fundamental goal of CSAR is to be a totally on-line interactive tool that includes all necessary user documentation. Within CSAR, user documentation is available on-line in the form of 40 lessons that cover all aspects of CSAR use and administration. A particular lesson can be retrieved on-line at the user's terminal or printed off-line and mailed to specified address. This form of documentation is timely and easy to maintain, and eliminates the many problems associated with distribution lists and central reproduction.

4.3 Performance reporting

CSAR performance reports present data in predefined formats. Report content, however, is dynamically retrieved and prepared to satisfy specific user needs. The flexibility of the CSAR reporting mechanism lies in four levels of data selection available to the user. The first level, Report Type, identifies a particular enumeration or summarization of the performance data and implies an associated physical format. The second level, System, selects one, or possibly all, systems in the TNDS. The third and most versatile level of selection is Organization. The user can select one or more reportable entities, traffic units, by naming the organization to which they belong. The company organization structure (known as the organizational map) is defined to the system by the CSAR company coordinator or administrative user as part of start-up procedures and can be dynamically modified whenever necessary. The CSAR map is more fully described later in this section. The fourth level, Time, specifies the calendar dates for which results are desired. The performance report request, and the underlying database retrieval, are defined by the composite of these selection criteria: Report Type, System, Organization, and Time. Any additional specifications are supported by a list of Options that are different for each report type.

There are four general types of performance reports:

1. Performance Indicator reports
2. Performance Summary reports
3. Performance Monitor reports
4. Unsatisfactory Results Display reports.

The following paragraphs describe the purpose and basic characteristics of these four report types. Specific examples of reports, their use by the OTC personnel, and the benefits to overall TNDS operations are discussed.

4.3.1 Performance indicators

The Performance Indicator Reports (PIRs) are intended for weekly troubleshooting of TNDS operational problems and data abnormalities. Those personnel most immediately responsible for data administration for the various TNDS systems access CSAR on a regular basis to detect these problems and isolate their causes. A PIR exists for each system in the TNDS scored under the TPMP. A PIR request retrieves weekly traffic unit data pertinent to one system in TNDS for any specified company organization. The performance measurements that appear on these weekly reports include those listed in Table I. The measurements typically quantify indications of missing data, validation failures, and record base inconsistency or errors. The report content is formatted as one line per traffic unit so the data for many traffic units can be scanned quickly. All performance measurement values that fall into the Unsatisfactory band are highlighted on the report.

Options are available with the PIR that make it even more powerful as a troubleshooting tool. These include "history" and "exception" options, which can be included in the initial PIR request, and the "detail" option, which can be selected after the PIR has been printed and perused. The history option allows inclusion of up to 15 previous weeks of information on the PIR. The exception option causes CSAR to report information on only those units for which one or more of the week's performance measures exceed the band U threshold. The history and exception options may both be exercised on the same PIR request to highlight offices consistently performing unsatisfactorily.

The detail option is somewhat open-ended. After the PIR is printed for a given week, organization, and TNDS system, the user can request details for any Traffic Unit (TU) in the report. After the detailed report for that unit is produced for the given week and TNDS system, the user is asked if details are desired for the same unit and week for another system in the TNDS. Thus, where applicable, the user can look upstream in the TNDS to see where problems first appeared in the TNDS process or downstream to see their ultimate effect. This may be very helpful in tracking down subtle problems.

There are too many combinations of PIRs and options to describe here, but a typical example of analyzing results and problems using a few reports should be indicative of the assistance provided by CSAR. Figure 2 shows a PIR for the No. 5 Crossbar Central Office Equipment Reports System (5XB COER)[1] for one district. An administrator responsible for monitoring 5XB COER results can obtain this report through the dialogue shown on the top two lines of the figure. In the dialogue, the questions from CSAR are in the large type to the left of "=:" followed by the user inputs in the smaller type. No options were selected in this example.

The report header identifies the report type, organization, and study date. As we can see in Fig. 2, this is the 5XB COER Performance

```
REQUEST = :repo indi 5xb nod district 111581
OPTIONS("D"BLSPC,"E"XCPTN,"H"ISTORY,"P"AGING)=:


        SCB                          CSAR 5XBCOER                DATE 12/17/81
MOD           DISTRICT        PERFORMANCE INDICATOR REPORT        RPT TC034
                                     STUDY DATE 11/15/81


  TRAFFIC      -----INPUT DATA ENDHOURS-----  --VALIDATIONS--  CF
   UNIT                    DEF  DEF  CM  NOT                  ERR
              EXP  MISSING TMR  SCH  OUT EXP  PERFORM  FAILED

  NOD1        MANAGER
NWORLAMUMGO   80    0-  0.0%   0    0   X   0   1590    75- 4.7%   N
  NOD2        MANAGER
FKLNLAMAMGO   25    0-  0.0%   0    0   X   0   1634    83- 5.1%   N
MRCYLAINMGO   30    5- 16.7%*  5    0   X   0   2570    88- 3.4%   N
  NOD4        MANAGER
BURSLAMAMGO   10    0-  0.0%   0    0   X   0    401    29- 7.2%   N
NWORLAARMGO  115    0-  0.0%   0    0   X   0   2902   126- 4.3%   N

* =UNSATISFACTORY PERFORMANCE
X =COLLECTION MACHINE OUTAGE NOT YET CORRELATED


DETAILS? ENTER TUNAME OR NO:mrcylainmgO
```

Fig. 2—Performance indicator report for the 5XB COER for one district.

Indicator Report for the New Orleans District of South Central Bell for the traffic study week of 11/15/81. A data header (feature ① of the figure) identifies the performance measurement data below. Under TRAFFIC UNIT, feature ②, five 11-character TU Common Language codes are shown. These are grouped together by the various organizations directly subordinate to the district being reported. Three subdistricts (NOD1, NOD2 and NOD4) are shown. Thus, any performance pattern associated with the subordinate organizational level would be obvious to the district level.

The next two columns, feature ③, pertain to the 5XB COER performance measure, MISSING DATA. The first of these columns shows the number of one-hour intervals expected for each TU by the 5XB COER system. These expectations are a part of the 5XB COER record base. The next column indicates missing traffic data, both the expected percentage and the number of one-hour intervals. The asterisk for MRCYLAINMGO indicates the percent of missing data for this week exceeds the band U threshold. Quick response to correct the deficiency will prevent it from persisting long enough to be band U for the official reporting period of one month.

The next four columns, feature ④, provide supplementary information resulting from CSAR automatically attempting to indicate the source of any missing data. In this case, the five one-hour missing intervals for MRCYLAINMGO are accounted for in the first column by a deficient Traffic Measurement Request (TMR). (TMRs are used to request that intervals be passed from TDAS to the downstream systems such as 5XB COER, if perhaps an interval was missing from the TMR.) The next column of this feature contains supplementary information indicating that there were no deficiencies in the scheduling of this data from the collection machine to TDAS. The "X" in the following column, Collection Machine Outage (CM OUT), indicates that there were disturbances in the collection phase of TNDS sometime during the intervals when data were gathered for 5XB COER. But since all data have been accounted for, these must have been minor. The last of these four columns indicates that no unexpected data (NOT EXP) was received for any TU. That is, TDAS did not forward to 5XB COER any data that were not expected (excess TMR). Thus, all 5XB COER data missing for this district are accounted for by the DEF TMR, which is the difference between what COER was told to expect and what TDAS was told to forward by their respective record bases.

The next two columns of the PIR, feature ⑤, give the number of validation tests performed by COER and the number and percent that failed. Validation failures are performance measurements, but the absence of asterisks indicates that none of the TU's are in band U for that measure. If the administrator, however, sees the failures exceeding reasonable rates for any TU, detailed information can be obtained as described later.

The last column of the PIR indicates where the COER internal record base validation tests have detected a record base inconsistency (control file error). None has been detected in this district.

Thus, perusal of this week's PIR indicates that the only severe problem is a data loss for one TU caused by a scheduling mismatch between the TDAS and 5XB COER record bases.

After the PIR and its accompanying footnotes are printed, CSAR asks the user to enter a TU name if further details are desired. In this typical example, the last line of Fig. 2 shows that the user responded with MRCYLAINMGO, the TU having the missing data. A 5XB COER Traffic Unit Details Report, Fig. 3, is generated from that request. The top section of this report, feature ①, is a repeat of the PIR format for the one TU. This redundant information is included in the details report so that the report is complete for that TU without requiring reference to the PIR. A few lines below, feature ② shows the status of the data input to 5XB COER. The interpretation of this information, guided by the appropriate on-line lesson, is that six one-hour intervals (ending at 10:00, 10:30, 11:00, 11:30, 16:00, and 16:30) are expected by COER on all of the weekdays. However, the data expected for the one hour ending at 11:30 (between 11 and 12 on the

SCB — CSAR 5XBCOER — DATE 12/17/81
MRCYLAINMGO TU — TRAFFIC UNIT DETAILS REPORT — RPT TC036
STUDY DATE 11/15/81

TRAFFIC UNIT	INPUT DATA ENDHOURS EXP	MISSING	DEF TMR	DEF SCH	CM OUT	NOT EXP	VALIDATIONS PERFORM	FAILED	CF ERR
MRCYLAINMGO	30	5- 16.7%*	5	0	X	0	2570	88- 3.4%	N

①

INPUT DATA SUMMARY:

	# RECORDS	EARLIEST – LATEST
NEW	7475	11/16/81 – 11/20/81
SAVE	0	
PURGED	0	

INPUT DATA STATUS CODES ON AN ENDHOUR BASIS:

```
                  1 1 1 1 1 1 1 1 1 1 2 2 2 2 2
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4
11/15/81 SUN
         MON        ...D     ..
         TUE        ...D     ..
         WED        ...D     ..
         THU        ...D     ..
         FRI        ...D     ..
         SAT
```

KEY TO THE INPUT DATA STATUS CODES USED ABOVE:
. =AVAILABLE, >90% DATA RECVD — M =MISSING, <=90% DATA RECVD
D =DEFICIENT TMR — I =INVALID IN TDAS
E =EXCESS TMR — U =UNAVAILABLE TO TDAS
(^) BELOW ENDHR WITH CM OUTAGE — BLANK =DATA NOT REQUIRED

②

DATA VALIDATION TEST RESULTS ON AN EQUIPMENT COMPONENT BASIS:

COMPONENT	# EXPECTED	# PERFORMED	# FAILED	% FAILED
AM	100	75	0	0.0%
CM	250	248	1	0.4%
CS	0	0	0	0.0%
CX	0	0	0	0.0%
DM	231	214	11	5.1%
DT	75	75	0	0.0%
ET	0	0	0	0.0%
IR	525	437	19	4.3%
LL	438	383	4	1.0%
MC	75	75	0	0.0%
OR	177	165	0	0.0%
OS	541	428	53	12.4%
TK	500	470	0	0.0%

③

* =UNSATISFACTORY PERFORMANCE
X =COLLECTION MACHINE OUTAGE NOT YET CORRELATED

ENTER OTHER SYSTEM OR NO:tdas
ENTER DETAIL TYPE(S) <DCU,SCH,TMR,VD,ALL,NONE>: :tmr

Fig. 3—A 5XB COER traffic unit details report.

figure) were not provided to 5XB COER because they were not scheduled through TDAS on a TMR.

Thus, the 5XB COER details report can serve as a comprehensive trouble indication to organization(s) responsible for scheduling 5XB traffic data for this TU. In this case, the analysis and trouble referral would be complete after a brief on-line computer session.

The bottom of Fig. 3, feature ③, gives details on the data validation tests. When the TU validation failures are higher than an experienced administrator thinks is normal, they could use these details to determine which component tests have abnormal failure rates.

The last two lines of Fig. 3 show the dialogue for optionally obtaining access to further information from other systems in the TNDS for this same TU. In this example the user continues the analysis by requesting complete TDAS details on all TMRs for this TU week.

4.3.2 Performance summaries

The Performance Summary Reports (PSRs), derived from weekly results in the database, evaluate performance in the TNDS operation for management at all levels of the telephone company organization. The summaries are available for one or all systems in the TNDS. The summarization technique determines individual traffic unit results for each performance measurement and places each traffic unit measurement into the appropriate band: Objective, Lower than Objective, or Unsatisfactory. As we discussed earlier, this banding technique avoids averaging the results across entities and masking specific cases of poor performance. As shown in Fig. 4, the summarized results reported are the number of traffic units and the percentage of traffic units in each performance band for each measurement. Totals are included on a system basis in the TNDS, along with grand totals. The on-line summarization is computed for the time period specified by the user (one or more study weeks) and for a requested organization (one or more traffic units).

As part of the regular CSAR processing, an official performance summary is run for monthly results of the total company. These results for the Service Observing Month (SOM) are then retained as part of the CSAR database. The grand total results for each company constitute the official input to TPMP.

As an aid to isolating and correcting problems CSAR also generates an optional list of all traffic units that appear in the unsatisfactory category (Fig. 5). This list can serve as a trouble list, and appropriate remarks regarding disposition and/or remedy can be made in the comments column.

4.3.3 Performance monitors

In addition to traffic unit banding and summary aggregation, CSAR also monitors other aspects of TNDS operations that are not within the scope of TPMP. Performance Monitor reports serve this important purpose and include the following areas of interest:

1. Common Update transaction statistics
2. Collection machine downtime
3. TNDS processing timeliness
4. TNDS software abnormal terminations and run times

REQUEST = :repo offsum all miss area 1081
OPTIONS ("B"RIEF)=:b

SCB — CSAR ALL — DATE 12/17/81
MISS — AREA — OFFICIAL PERFORMANCE SUMMARY REPORT — RPT TC900
STUDY PERIOD 09/27/81-10/24/81

MEASUREMENT CATEGORY	BAND "L" RANGE	TU'S TEST	NO. OF OFC-WK	NUMBER AND % TESTED IN BAND "O"		"L"		"U"	
DCU-INT MISSING	7%-15%	95	291	86	91%	7	7%	2	2%
DCU-INT ANY VER FAIL	10%-25%	95	291	85	89%	8	8%	2	2%
DET TEST FAILURES	3%- 5%	0	0	0	0%	0	0%	0	0%
INVAL DET TEST SCAN	20%-30%	24	54	24	100%	0	0%	0	0%
TOTALS-TDAS		214		195	91%	15	7%	4	2%
ICAN/CU COUNT MSMTCH	3%- 5%	25	96	25	100%	0	0%	0	0%
ICAN/CU OTHER ERROR	.5%- 1%	25	96	25	100%	0	0%	0	0%
SUSPECT SCHV ASSIGN	1%- 2%	25	96	23	92%	1	4%	1	4%
UA/UE SCHV USAGE	3%-- 5%	25	96	24	96%	1	4%	0	0%
TOTALS-ICAN		100		97	97%	2	2%	1	1%
CGMT INSUFF DATA	5%-10%	93	146	85	91%	5	5%	3	3%
CGMT OTHER INVALID	15%-25%	93	146	83	89%	3	3%	7	8%
CGMT UNASSOCIATED	5%-10%	93	146	90	97%	1	1%	2	2%
TOTALS-TSS		279		258	92%	9	3%	12	4%
LU INVALID FOR LBI	10%-25%	69	143	66	96%	3	4%	0	0%
LU MEAS INVALID	5%-15%	69	143	66	96%	3	4%	0	0%
TOTALS-LBS		138		132	96%	6	4%	0	0%
5XCOER CNTRL FILE ER	25%-75%	35	138	34	97%	1	3%	0	0%
5XCOER ENDHR MISSING	10%-15%	35	138	33	94%	0	0%	2	6%
5XCOER VAL FAIL/PERF	5%-10%	35	138	35	100%	0	0%	0	0%
TOTALS-5XBC		105		102	97%	1	1%	2	2%
TOTALS		836		784	94%	33	4%	19	2%

Fig. 4—Performance summary report.

5. TSS data by servicer responsibility.

These reports present an additional view of the effectiveness of data collection equipment and various company organizations.

4.3.4 Unsatisfactory results display

The management users of CSAR often prefer results that more clearly reflect general performance levels, trends over time, and direct comparison between organizations or between the component systems of the TNDS. The Unsatisfactory Results Display reports plot the poor performance (band U) in a series of horizontal bar graphs. The graphs are simple in appearance and offer a wide variety of options that reveal performance trends at a glance. The results can be plotted with subtotals by any system in the TNDS, or by the organizations that are directly subordinate to the one requested. For example, as we

```
REQUEST = :repo summ tss ma district 102581-111581
OPTIONS ("B"RIEF OR "L"LIST)=:1

      SCB                              CSAR TSS                               DATE 12/17/81
MA              DISTRICT        PERFORMANCE SUMMARY REPORT                      RPT TC601
                                 STUDY PERIOD 10/25/81-11/21/81

-------------------------------------------------------------------------------------------
:     MEASUREMENT        BAND "L"   TU'S    NO. OF        NUMBER AND % TESTED IN BAND      :
:       CATEGORY          RANGE     TEST    OFC-WK          "O"            "L"         "U" :
:==========================================================================================:
:CGMT INSUFF DATA         5%-10%     18       67       17     94%     0     0%     1     6%:
:CGMT OTHER INVALID      15%-25%     18       67       18    100%     0     0%     0     0%:
:CGMT UNASSOCIATED        5%-10%     18       67       17     94%     0     0%     1     6%:
:                                   ----           ---------------------------------------:
:  TOTALS                            54                52     96%     0     0%     2     4%:
============================================================================================

      SCB                              CSAR TSS                               DATE 12/17/81
MA              DISTRICT               BAND U LIST
                                 STUDY PERIOD 10/25/81-11/21/82

-------------------------------------------------------------------------------------------
:  TRAFFIC                                                                                 :
:   UNITS              CATEGORIES                VALUE               COMMENTS              :
:------------------------------------------------------------------------------------------:
:BLZNMSMA247       CGMT UNASSOCIATED           ( 22.6% )                                   :
:------------------------------------------------------------------------------------------:
:CLEVMSMA84A       CGMT INSUFF DATA            ( 33.3% )                                   :
:------------------------------------------------------------------------------------------:
:                                   2 BAND U TRAFFIC UNITS OF   18                         :
:------------------------------------------------------------------------------------------:
```

Fig. 5—Performance summary report showing unsatisfactory traffic units.

see in Fig. 6, an area manager can request a five-week performance graph containing an area total and subtotals for each of the divisions in that area. The official summary results are also available as graphs.

4.3.5 Bell System summaries

Analysis of TNDS performance can extend to the level of the entire Bell System when the performance results for all OTCs are aggregated together. Here, there is an opportunity to assess the need to modify TNDS, systems or operating methods in the CSAR itself. There are four types of reports available at this level. One gives Bell System total results in the form of the "Brief" version of the Performance Summary report (no band U list). The "Company" option, when exercised, produces a report of the number of performance measurements falling into each of the bands O, L, and U for each company. The "graph" option generates a graph of the band L and band U results by company. The remaining option is especially valuable for TNDS and CSAR project managers at AT&T and Bell Laboratories. This "Measure" option generates a list of each company's percentages and the Bell System totals in bands O, L, and U for each of the performance measures. Thus, problems common to a large cross sec-

```
REQUEST = :repo offdisp all miss area 1081(5
OPTION("SU"BTEND OR "SY"STEM)=:su

      SCB                             CSAR ALL                          DATE 12/17/81
MISS        AREA      OFFICIAL UNSATISFACTORY RESULTS DISPLAY             RPT TC905

TU-MEASURES  INTVLS
   TESTED                                 % IN BAND-U
                    0    5    10   15   20   25   30   35   40   >40%
                  -+----+----+----+----+----+----+----+----+----+----+-
                   .
                   .
AREA    :MISS      .
  876  05/24-06/27 .=====*
  865  06/28-07/27 .===*
  910  07/26-08/22 .=====+*
  874  08/23-09/26 .====*
  836  09/27-10/24 .==*
                   .
                   .
DISTRICT:MA        .
  317  05/24-06/27 .====*
  306  06/28-07/25 .=*
  341  07/26-08/22 .=====*
  291  08/23-09/26 .==*
  292  09/27-10/24 .==*
                   .
                   .
DISTRICT:MB        .
  257  05/24-06/27 .=====+*
  265  06/28-07/25 .====*
  269  07/26-08/22 .=====*
  259  08/23-09/26 .=====*
  244  09/27-10/24 .==*
                   .
                   .
DISTRICT:MC        .
  300  05/24-06/27 .====*
  294  06/28-07/25 .=====*
  300  07/26-08/22 .=====+===*
  324  08/23-09/26 .=====+*
  300  09/27-10/24 .==*
                   .
                   .
                   .
                  -+----+----+----+----+----+----+----+----+----+----+-
                    0    5    10   15   20   25   30   35   40   >40%
```

Fig. 6—Unsatisfactory results display.

tion of users can be examined to determine what system design or training efforts might be appropriate.

4.4 Database controls

CSAR creates and maintains separate central databases for each OTC using the system. On-line database controls permit the CSAR coordinator or Headquarters staff personnel to decide on the size of the database, the retention periods of three classes of performance information, the timing of certain automatic CSAR processing, and the organizational structure to be used for reporting purposes. Each of these controls may be exercised at any time to adapt to changes in data needs, company reorganization, storage costs, and other considerations.

4.4.1 Data retention

While the current week or current Service Observing Month (SOM) performance results are of primary interest, past results can also be examined for evidence of trends or persistence. To provide this historical perspective, CSAR databases contain fundamental performance data organized on a TNDS subsystem and traffic unit basis for distinct study weeks. In addition, summarized performance results are stored for each SOM. Because different classes of performance data serve distinct purposes and make significantly different demands on the physical database storage capacity, three data retention parameters are available to the user:

1. Short-term Retention (STR) (from 6 to 16 weeks)
2. Long-term retention [from (STR+5) to 52 weeks]
3. Official results retention (from 12 to 30 months).

The most voluminous data, deleted after the short-term retention period, are the supporting details from which the weekly results are derived. These details reflect TNDS processing at the half-hour, or hourly, level, and include measurements in terms of TDAS Data Collection Units (DCUs), switching office equipment components, Trunk Servicer responsibility codes, etc. This fine level of detail helps the CSAR user to track and isolate specific problems. The data analysis algorithms also require these data to correlate cross-system effects, and accurately calculate the weekly traffic unit measures.

Weekly performance results are kept for a period of several months (long-term retention), and SOM results are retained for more than a year (official results retention) so that performance can be compared under similar seasonal conditions.

During the CSAR Batch processing the data retention restrictions are applied, all outdated database information is automatically deleted, and the storage space is freed for future use.

4.4.2 Organizational map

The user-defined CSAR organization map provides a direct and dynamic control over database retrievals. The CSAR user is presented with a hierarchical logical view of the performance database for on-line retrieval purposes. Unlike most hierarchical database management systems, the hierarchy is not a static structure defined at database generation time; instead CSAR provides a flexible structure that can be modified easily on-line without database reorganization.

The basic reporting unit for most CSAR information is the traffic unit. Data tracking and problem diagnosis for the TNDS is most successfully accomplished by analyzing the scheduling, collection, processing, and validation activities at a fundamental traffic unit level. Data reported at this low level reveal problems that can be isolated to specific physical devices, user transactions, record base items, or operational procedures. For performance measurement of complete systems, functions, or company organizations, and for more effective management use of the information, aggregate views of traffic unit data are more appropriate. Responding to these needs, CSAR provides flexible multi-level views that permit data access at any organizational level, from the elementary traffic unit up to the entire company.

The user defines a hierarchical structure (up to nine levels) that establishes a company organization desired for CSAR reporting purposes. The first or top level is reserved for the company name, and the lowest level is designated as the traffic unit level. The remaining seven hierarchical levels can be named to represent the company structure and levels of management responsibility (i.e., district level, division level, area level, etc.). Not all seven intermediate levels have to be defined. At each of the levels two through eight, one or more entities can be named and placed subordinate to an entity named in the level above it. Figure 7 shows an organizational map with a total of five levels.

The CSAR map logically associates each of the many company traffic units to the desired organizations identified in this hierarchical structure. The physical database storage of traffic unit performance data is completely independent of this logical hierarchy. The CSAR dialogue then allows groups of traffic units to be collectively requested by any designated organization at any level.

The map's structure and the association of traffic units to higher-level organizations is defined using special data-entry features of the CSAR dialogue. A complete set of update commands allow the user to add, delete, change, rename, and move individual traffic units, complete organizations, or levels. The changes are performed on-line with the resulting organization in effect immediately for all subsequent report requests. CSAR features also include the ability to list the map in a variety of ways to determine or verify the current company organizational structure.

The direct control over a retrieval hierarchy and the ability to modify it easily make the CSAR organizational map a most versatile and innovative system feature.

4.5 Processing controls

CSAR on-line features can override certain automatic batch processing functions. These controls are enacted through the normal CSAR dialogue, but only by the privileged headquarter's users. The overrides exist to meet two types of needs: (1) a management demand for an early assessment of official monthly results; and (2) an administrative need to respond quickly and conveniently to batch processing problems.

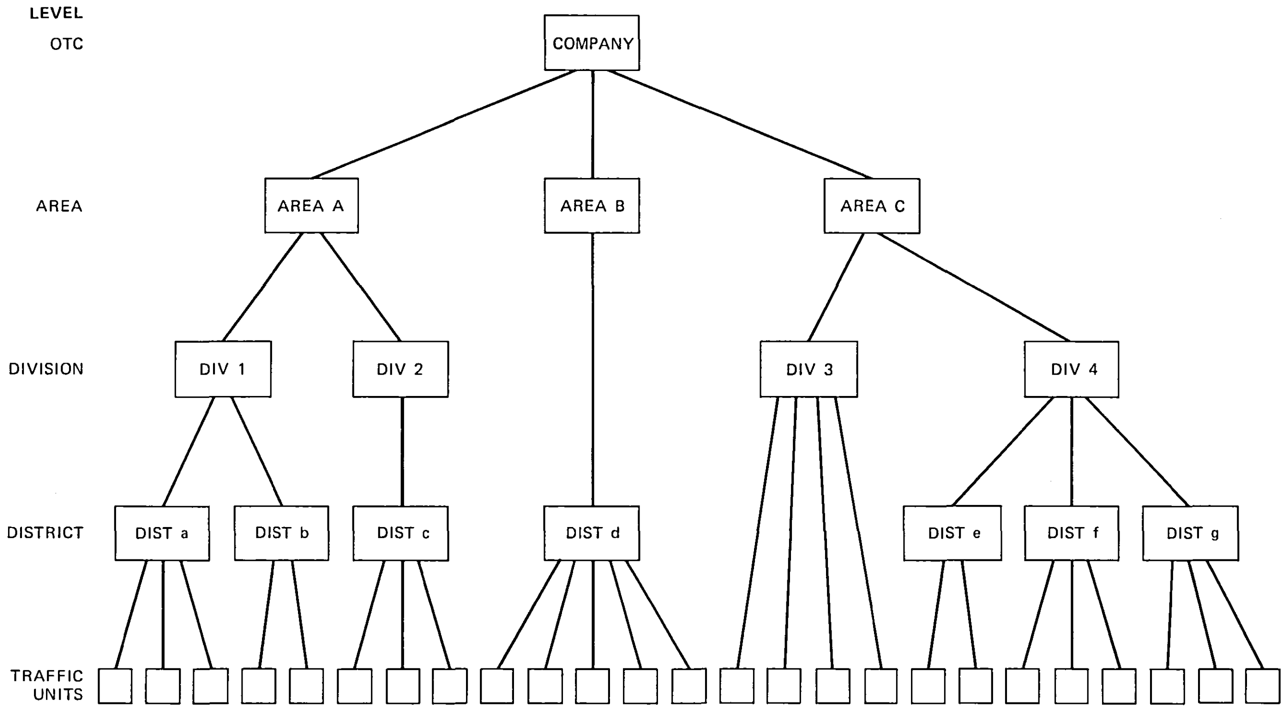

Fig. 7—Organizational map structure.

4.5.1 Generating official results

A process delay parameter, entered as a company option, defines the longest expected delay between the collection of traffic data and the delivery of the TNDS performance data to CSAR. The database loading process closely monitors actual data delay, particularly to determine late data impact on official TPMP monthly results. The official performance summary results are generated automatically during the CSAR batch process after the appropriate period of time has elapsed. In cases where management personnel require official results prior to the expiration of the process delay period, CSAR permits on-line interactive generation of the official performance summary. This on-line feature enables the administrator to satisfy immediate needs or to correct past results by regenerating summaries in the event of unexpected, very late data.

4.5.2 Batch job control

Process control features also exist to simplify overseeing the data transmission and CSAR batch database loading process, described in Sections 3.1.2 and 3.1.3. Each OTC's administrator responds to abnormal completion of CSAR batch jobs. Interactive features include Rerun and Remove commands allowing direct control over the disposition of individual data transmissions that were received but not successfully processed. These controls are exercised after the administrator determines the nature of the problem using diagnostic and error recovery information supplied by the system.

4.6 Administrative reporting

The telephone company administration of the CSAR software system requires activities such as: initial implementation, organizational map definitions and maintenance, selection of company options, coordination of data transmissions, and monitoring of the ongoing data processing. Each OTC has a designated CSAR administrator. Some of CSAR's features simplify the administrator's job. The data transmission and database entry operation have been automated to the fullest extent possible. The CSAR database provides on-line interactive access to status information and various reports that assist in system administration. The administrator relies on several sources of information to monitor the overall operation of CSAR:

1. Local and global login messages
2. News items
3. Tape processing status
4. Batch data processing summary
5. Map Data Discrepancy report

6. Interactive feature usage statistics
7. Operational cost information.

When an administrative user logs on, the Messages, News Items, and Tape Processing Status display important happenings without the need for database query. Based on any abnormal terminations indicated on the tape status, the corresponding batch job is investigated by accessing the Batch Data Processing Summary. The summary serves as a log of all processing activities and contains any error diagnostics identified by the CSAR software. Many conditions can be effectively identified and corrected by the administrator. Other problems may require investigation by the central development organization, and the error message would instruct the user to initiate the trouble referral.

The entry of new data into the CSAR database may involve traffic units that had not been anticipated, or for some reason not properly defined in the company organization map. The performance data are retained for these traffic units and are entered into the database. In addition, each previously unknown traffic unit is temporarily added to the map directly subordinate to the company level, where it remains until the administrator uses the CSAR dialogue to move the traffic unit to the appropriate position in the organization hierarchy. CSAR alerts the users of any such unexpected or unusual conditions via a Map Data Discrepancy report, and by Input Data Anomaly sections of the Batch Data Processing Summary. The software system design incorporates similar error detection and supporting diagnostic information throughout to simplify system administration and maintenance.

V. CSAR SOFTWARE DESIGN

In the design of a software system several conflicting factors are addressed and balanced to provide an efficient, flexible system that is responsive to the user needs. Three aspects of CSAR software design deserve specific mention and are covered in the following paragraphs.

5.1 On-line response time

A primary concern in the design of CSAR was flexibility in selecting the set of traffic units, TNDS systems, and time periods over which performance can be summarized. Depending on the particular request, tens of thousands of data items would have to be summarized on line and reported to the user. To accomplish this summarization with good on-line response time required special attention in the CSAR design. The physical database organization chosen reflects the high-level summarization options of "system" and "time period." Unique files exist for each system/week that contain the key performance information for all reporting traffic units. CMS files are organized and accessed with a physical block size of 800 bytes. The key performance indicators for many traffic units are stored together in logical records whose lengths are an integral multiple of 800 (one TU requires much less than 800 bytes) to make efficient use of the operating system's file structure. File pointer tables exist to identify the CMS record (block) and location within the block for each active TU. The internal traffic unit identification scheme utilizes a hashing function together with the file pointer table to efficiently access the data for a particular TU within a file. The logical structure of the map efficiently converts the OTC organizational level to a set of reportable TU hash indices. Finally, the TU access order is optimized based upon the file pointers to minimize the amount of I/O necessary to read the raw performance data. These techniques used together result in worst-case summarization response times typically less than six seconds in an average size OTC.

5.2 Flexibility in reportable measures

The measurement plan consists of 19 individual performance indicators. These basic indicators have been refined several times since the measurement plan was first introduced to provide equitable reporting across companies and to encourage continued performance improvement.

The design of the software to support these changing requirements had to be robust and easy to modify. The identification of distinct measures, the start and stop effective dates, the banding thresholds, and other key information are all specified in a single file external to the software programs. During system execution this file is accessed to load tables that control the summarization process. This table-driven summarization approach is also beneficial in that the effects of proposed changes to the mesurement plan can be observed by editing a development copy of this file and using this modified version when accessing the OTC data. Many such modifications can be studied and later implemented without changing or recompiling any software.

5.3 Overlay structure

User chargeback generated by using this AT&T VM/CMS facility is highly sensitive to the core size of the virtual machine required for the application program. To minimize this aspect of chargeback, and to improve efficiency by reducing virtual storage paging, the system was designed with an overlay structure consisting of over 120 separate modules. The design of these modules reduced the total amount of code generated by combining approximately 560 functional routines as building blocks.

The overlay structure maintains common routines and information required throughout the on-line session in a root area that is never overlayed. Specific software required to respond to a particular user request is overlayed in multiple levels below the root. In the most complex request, eleven individual overlays are used, reducing the amount of incore storage required from 500 to 130 kilobytes.

VI. CONCLUSION

The goal of bringing performance measurements and operations analysis information to the TNDS has been achieved by implementing the TPMP through CSAR. TNDS managers and administrators now receive reports that are comprehensive, easy to obtain and use and on time. CSAR can be readily adapted and documented as TNDS changes. In fact, method changes are currently under way to improve the effectiveness of CSAR in isolating TNDS data collection and provisioning problems. Recent changes and future plans, not reflected in this article, include the revision or elimination of certain performance measurements. These changes will shift the emphasis from strictly end-system analysis to measurements that detect and quantify problems of data availability earlier in the flow of data through TNDS. As data responsibilities have become better focused organizationally, TPMP has been reexamined to ensure the plan meets the changing needs of OTC managers. This article reflects the CSAR software system and TPMP methods as of the middle of 1982.

CSAR became available to the OTCs in the last half of 1979. No official TPMP reporting was required by AT&T, however, until July of 1980. During the introductory interval, the OTCs began to use CSAR reports as effective tools for traffic data administration and management. At the start of the official reporting, the Bell System average for band U measurements was about 10 percent. (This is thought to be about one-half of what it was at the beginning of the introductory period.) The current Bell System average of band U measurements is about 5 percent. Thus, CSAR is providing effective in improving the delivery of valid traffic data so that the Bell System network can respond efficiently to customer needs.

REFERENCES

1. N. D. Fulton, J. J. Galiardi, E. J. Pasternak, S. A. Schulman, and H. E. Voigt, "Total Network Data System: Equipment Systems," B.S.T.J., this issue.
2. R. F. Grantges, V. L. Fahrmann, T. A. Gibson, and L. M. Brown, "Total Network Data System: Central Office Equipment Reports for Stored Program Control Systems," B.S.T.J., this issue.

3. P. V. Bezdek and J. P. Collins, "Total Network Data System: Trunking Systems," B.S.T.J., this issue.
4. C. J. Byrne, D. J. Gagne, J. A. Grandle, Jr., and G. H. Wedemeyer, "Total Network Data System: Data Acquisition and Near-Real-Time Surveillance," B.S.T.J., this issue.
5. D. H. Barnes and J. J. O'Connor, "Total Network Data System: Small Office Network Data System," B.S.T.J., this issue.

AUTHORS

Dennis R. Anderson, B.S.C.S., 1976, Michigan State University; M.S.C.S., 1978, Purdue University; Bell Laboratories, 1977—. Mr. Anderson initially worked on software development for the No. 5 Crossbar Central Office Equipment Reports system. He later worked on performance measurement of the traffic data collection and processing functions performed by the Total Network Data System. He defined new measurement methods and designed interactive software for efficient data analysis, database access, and performance summarization. In addition to development responsibilities, he wrote feature requirements and participated in project planning. He is currently working on the Digital Ordering and Planning System for No. 5 *ESS*® switching equipment. Member, ACM, Tau Beta Pi.

Melvin J. Evans, B.S.E.E., 1957, University of Utah; M.E.E., 1959, New York University; Bell Laboratories, 1957—. After joining Bell Laboratories, Mr. Evans worked on guidance equations used by the Bell Laboratories Command Guidance System for rocket launch vehicles. He has supervised guidance equation work, lunar mission studies, and military system effectiveness studies. He has also supervised work on operations support systems that assist Bell System network and data administrators. Starting in 1976, his group took part in a study of Total Network Data System performance, and then developed the Centralized System for Analysis and Reporting (CSAR). Currently, he supervises the Planning Mechanization group, which provides computerized support for operations planning.

THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Operating Company Perspective

By J. PFEIFFER, Jr.*

(Manuscript received June 21, 1982)

Implementation and management of the many individual systems comprising the Total Network Data System (TNDS) is a monumental undertaking, with far-reaching impact on a telephone company's organization and methods. This paper describes how one such company—Southern Bell—planned for and converted to TNDS, extending over a number of years. It also discusses the current operations and benefits of TNDS.

I. PRE-TNDS IN SOUTHERN BELL

Before examining TNDS as it exists now within an operating telephone company, it should be helpful to review data-collection procedures before the introduction of TNDS.

1.1 Manual switch counts

Starting with the introduction of dial Step-by-Step (SXS) switching machines in the 1920s, traffic studies were made by manually counting the number of individual switches that were in use, over fixed intervals of time. A large group of clerks would descend upon an SXS central office to make these studies, typically once or twice a year. Counts were made on a sample basis, generally every three minutes, of all

* Southern Bell.

busy switches in specific groups, such as trunking groups or a switching stage. A supervisor would control the timing of each count with a whistle, and record the number of busy switches for subsequent summarization and calculation. (See Fig. 1.)

The deficiencies inherent in these studies were severe and obvious. They lacked accuracy; there was significant data "skew" due to inherent difficulties in starting and stopping these counts with precision; they were labor intensive; and since they could only be made infrequently, there was a high probability of missing the busy hour or busy season for a given office. Although these traffic studies were the only basis for engineering and administering the telephone network, there was a general lack of confidence in their accuracy. A judgment factor was frequently added to the study results in order to protect service. For this reason, a great many offices tended to be over-engineered in at least some of their switching stages. No one really knows how much capital was wasted in providing unnecessary equipment, but it had to be considerable.

1.2 Introduction of the traffic usage recorder

A major improvement was made in the mid-1950s with the introduction of the Traffic Usage Recorder (TUR). The TUR permitted individual switches to be wired into a "grouped" register, each register

Fig. 1—A manual count being recorded of number of switches engaged with calls. The count was made to determine usage in SXS dial offices.

representing a specific trunking group or switching stage. Also, the TUR made it possible to study traffic in crossbar offices, which had not been practical with the manual count method. (See Fig. 2.)

The TUR had one distinct advantage. It needed to be read only hourly, rather than every three minutes. This sharply reduced the clerical load, and it also improved the data-skew problem. It then became practical to schedule traffic studies more frequently, and confidence in data quality increased. As office sizes grew, however, the traffic-register readings and data-summarization chores grew as well.

1.2.1 Camera added to TUR

A further improvement came in the late 1950s. A camera and flash unit were positioned over groups of up to 160 traffic registers. Under the control of a pre-set TUR timer, the camera took pictures of these banks of registers for as many hourly periods as study needs dictated. The advantage gained was to virtually eliminate data skew, and further reduce the clerical effort in gathering the data and manually summarizing it. Computer programs were created in Southern Bell to process this raw data and produce a family of user-oriented reports downstream.

However, other problems were created. The network administrator who was responsible for doing the traffic study now had to get involved in the manual camera and film-processing procedures. A substantial investment in personnel and equipment became necessary in order to view and keypunch data from the film strips. A key verify procedure was also introduced that virtually eliminated keypunch errors, but at the cost of doubling the key-entry costs. Finally, the work load on the network-administration staff was highly variable since traffic studies were not evenly distributed over time. Since studies are scheduled to capture peak loads, they tend to be concentrated in seasonal and monthly peaks. As a result of the uneven loads, some delays occurred in the manual processing activities.

1.3 OCR experiments

Southern Bell, and one or two other Bell System operating companies, experimented with methods to mechanize this key-entry function during the early 1970s with Optical Character Recognition (OCR) techniques. While our trial was a technical success, it did not prove to be economical and was abandoned after two years.

1.4 Pre-TNDS problem areas

Despite the progress made over the years, our data-collection situation just prior to the introduction of TNDS had four serious problem areas.

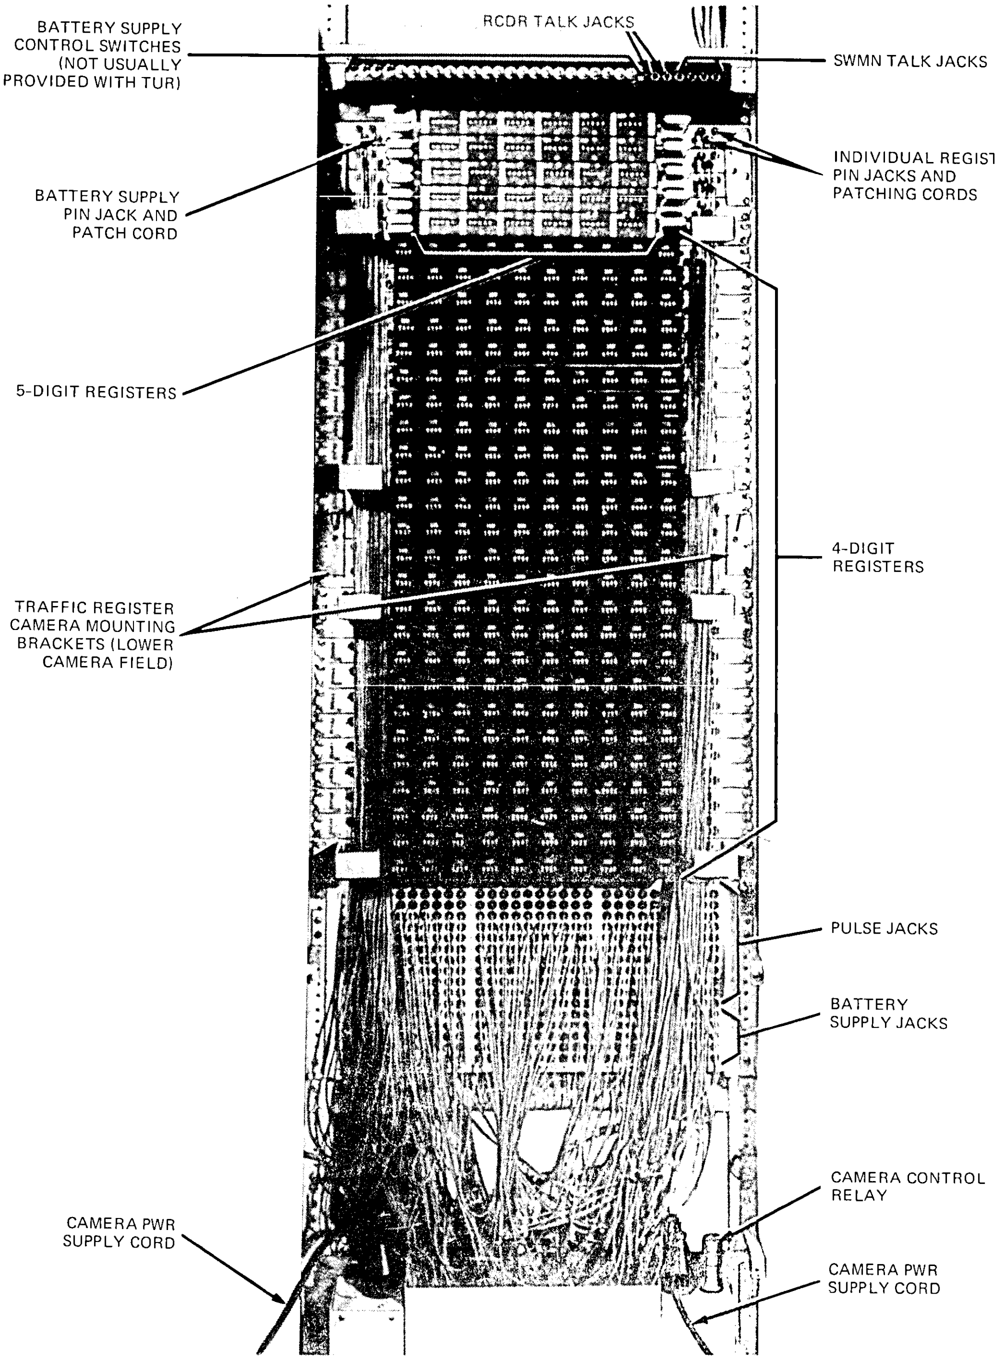

Fig. 2—TUR traffic register cabinet. Camera and hoods are not shown.

First, it was difficult to administer because it required extensive and timely coordination between several departments, each having different priorities.

Second, the data quality was lower than desired. This was due to inherent problems with the cameras, lighting, and film processing.

Third, the completed studies took too long to reach the ultimate user. This was due to the complex sequence of steps necessary, as well as the uneven study schedules. Frequently, the studies had lost much of their value by the time they were received.

Finally, each operating area within Southern Bell had evolved its own set of procedures for data collection. There was a lack of central direction and standardization. The stage was set for TNDS in Southern Bell.

II. TNDS IMPLEMENTATION STRATEGY

2.1 Criteria for TNDS

TNDS was approved for implementation in Southern Bell as a corporate-wide project in November 1973. The following criteria were basic to the implementation plan:

1. To realize the maximum benefits, all TNDS modules would be implemented (both available and planned) unless there were compelling reasons for doing otherwise.
2. The "downstream" data processing modules would be implemented on a centralized basis.
3. The "upstream" data-collection (EADAS) modules would be implemented on a decentralized basis. There would be one Engineering and Administrative Data Acquisition System (EADAS) installation in each administrative area, and administrative boundaries would not be crossed.
4. There would be a small, strong centralized staff to provide support to the areas for planning, implementation, ongoing technical assistance, and methods.
5. There would be no downstream processing of key-entry, filmed (non-EADAS) data in Southern Bell, even though TNDS could accommodate it. This would avoid slowing the downstream weekly-coordinated process to the slower pace of key-entry data.
6. The administrative areas would be cut over to TNDS on a phased basis, with each area's schedule determined by need and capital availability.
7. All central offices would interface with TNDS unless they were below a designated minimum size or were to be replaced within about three years.
8. Commercial courier service would be used to transmit EADAS

data-collection tapes to the single centralized processing location, and output reports back to the areas. Figure 3 illustrates the deployment.

2.2 Deployment schedule

The "downstream" software of TNDS was installed shortly before the first EADAS installation to allow for a proper system test. The "upstream" EADAS systems, and their associated administrative staff, were installed according to the schedule shown in Table I.

This sequence was dictated principally by the rapid growth in Southern Florida, which emphasized the need for valid and prompt traffic data. The last three areas were more rural in nature, with slower growth and therefore less critical data needs, at least initially.

Depending upon the number of eligible central offices, one to two years were required to fully implement TNDS within each Southern Bell area. As was expected, the first few offices in each area required a greater amount of time to cut over, until the learning curve was mastered. Each area was given strong support from the centralized TNDS staff, particularly in the planning and early implementation stages.

TNDS was considered fully implemented in Southern Bell by 1980, with the addition of only growth central offices thereafter.

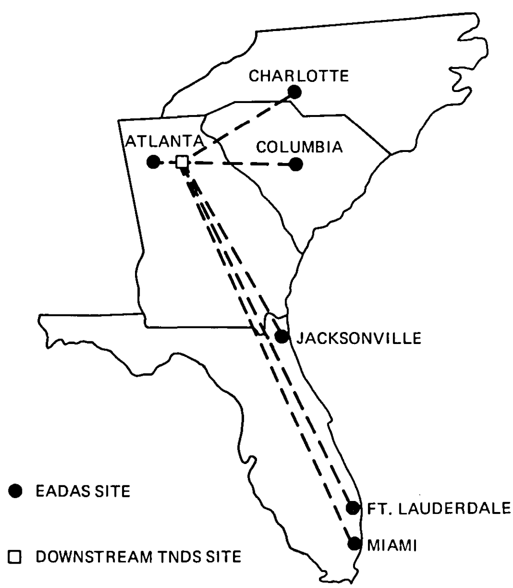

Fig. 3—Routes by which EADAS data-collection tapes and outputs are transmitted between centralized processing location and administrative areas.

Table I—EADAS installation schedule

Area	City	Date
South Florida	Miami	October 1974
Southeast Florida	Ft. Lauderdale	October 1974
Georgia	Atlanta	June 1975
North Carolina	Charlotte	July 1976
North Florida	Jacksonville	June 1977
South Carolina	Columbia	June 1977

III. ORGANIZATION FOR TNDS

3.1 Work centers involved

Southern Bell follows the guidelines of the Total Network Operations Plan (TNOP) quite closely. TNOP organizes all work functions within the Network Segment into clearly defined centers. The centers involved in our TNDS data-collection and report-distribution process are the:

1. Circuit Administration Center (CAC)
2. Network Administration Center (NAC)
3. Network Data Collection Center (NDCC)
4. Minicomputer Maintenance and Operations Center (MMOC)
5. TNDS Coordination Center (TCC)
6. Corporate Data Center (CDC).

3.2 TNDS data flow

The technical description of the data flowing through TNDS is described in detail in other articles of this issue. Figure 4 illustrates how the data flow is managed by these centers in Southern Bell. It should be recognized, however, that there are significant variations in this process within the Bell Operating Companies (BOCs) due to differences in geography, demography, and corporate structure.

1. There are twenty-eight Network Administration Centers (NACs) in Southern Bell, ranging from three to seven per area. Each NAC serves a geographic subset of an area. All requests for TNDS-equipment studies originate in, or are transmitted through, the appropriate NAC to the Network Data Collection Center (NDCC).

2. There are six Circuit Administration Centers (CAC) in Southern Bell, one per area. All requests for TNDS-Trunking studies are originated by the appropriate CAC, and are transmitted directly to the NDCC serving that area. The databases for trunking studies (TSS and TFS) are maintained by the CAC.

3. There are also six Network Data Collection Centers (NDCCs) in Southern Bell, one per area. The NDCCs are responsible for monitor-

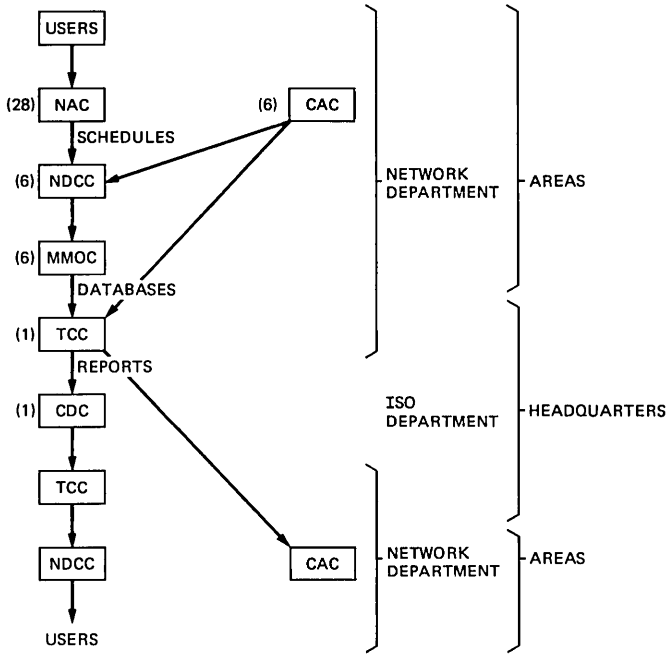

NOTE: NUMBERS IN PARENTHESES INDICATE THE NUMBER OF CENTERS OF EACH TYPE IN SOUTHERN BELL.

CAC – CIRCUIT ADMINISTRATION CENTER
CDC – CORPORATE DATA CENTER
MMOC – MINICOMPUTER MAINTENANCE AND OPERATIONS CENTER
NAC – NETWORK ADMINISTRATION CENTER
NDCC – NETWORK DATA COLLECTION CENTER
TCC – TNDS COORDINATION CENTER

Fig. 4—TNDS data flow among administrative centers in Southern Bell.

ing the operation and maintenance of the EADAS Central Control Units (CCUs) in their area, monitoring the operation of the central office data-collection apparatus and associated data links, maintaining accurate record bases, and coordinating as required with the interfacing work centers to resolve problems. In Southern Bell, the NDCC is the principal focal point for data-collection activities in an area, and has no responsibilities other than successful data collection for TNDS.

4. The Minicomputer Maintenance and Operations Centers (MMOCs) are responsible for the operation and maintenance of the clustered minicomputers in their respective areas. EADAS is one of these clustered minicomputers. The MMOC performs preventive maintenance for EADAS on a schedule mutually agreed upon with the NDCC. It also is responsible for transmitting the magnetic tapes generated by EADAS to the TNDS Coordination Center (TCC) on a timely basis, via commercial courier service.

5. The TCC is a single, company-level work center. It serves as an interface between the various network users and the Corporate Data Center (CDC) of the Information Services Organization (ISO), which operates our downstream programs as part of a utility service for all corporate data-processing requirements. The TCC receives magnetic tapes with study data from all area MMOCs, plus all off-line database updates originated by both the NDCCs and CACs. These are formed into data-processing jobs and submitted to ISO by the TCC. The TCC is responsible for tracking the successful completion of the TNDS job through the CDC. A measurement plan has been instituted to determine the timeliness of CDC production for TNDS output.

6. The CDC is responsible for producing the downstream TNDS reports for Southern Bell. Three types of output media are used, depending upon the wishes of the individual users. These are impact print, laser print, and microfiche. Since TNDS processing is organized into stages, the output is spread over the week to level the load on all centers and resources. After each processing stage, the TNDS output is delivered to the TCC promptly.

7. The TCC then reorganizes the TNDS reports into individual user "batches" destined for each area, and transmits the reports back to the six NDCCs and six CACs via commercial courier service. The EADAS tapes are also returned to the MMOCs for reuse on a monthly cycle.

8. The CACs are end users, and upon receipt of their output reports, can immediately put them to use. The NDCCs are the end users for administrative reports that relate to the accuracy and health of the data-collection process; these are used directly by the NDCC. The other user reports are transmitted to the ultimate users in that area by the NDCC.

IV. MANAGEMENT OF TNDS

4.1 Centers involved in TNDS management

Although each center involved in the TNDS data-collection and distribution process is essential for overall success, the "up-front" centers—NAC, NDCC, and CAC—have an especially important role. These three centers are assigned primary responsibilities in the data-flow process through TNDS, including record-base, scheduling, surveillance, and distribution functions. The NAC is assigned the basic central office assignment functions for equipment measurement needs; the CAC maintains trunking databases; and the NDCC provides a centralized center with the technical expertise to interface directly with the TNDS subsystems. The NDCC performs the surveillance of the EADAS CCUs and is an area distribution point for TNDS output.

4.2 Use of TPMP and CSAR

The TNDS Performance Measurement Plan (TPMP) was introduced in Southern Bell in 1979 to improve the management of TNDS. This plan, along with the Centralized System for Analysis and Reporting (CSAR), has added two greatly needed features. As an index plan, TPMP has increased emphasis on the performance and the resolution of problems in the area of traffic-data collection. Additionally, CSAR, as a reporting tool, gives management at all levels specific indications of the success or failure of the TNDS process.

The pre-TPMP environment often found TNDS problems competing for attention with other indexed items. The TNDS problems often were given much lower priorities, and lengthy time intervals could occur before problems would be resolved. Since TNDS results are now being reported along with other indices in Area, Southern Bell, and AT&T results books, TNDS problems are given a higher priority. The weekly reporting and monitoring capability of CSAR also maintains emphasis on any problems until they are resolved.

4.2.1 CSAR reports

The CSAR program itself has proven a valuable reporting tool for TNDS management. By providing reports tailored for any organization or level from an individual traffic unit up to a total company report, CSAR provides management with increased insight into the problems that may need attention. The level of detail available in CSAR makes it useful for first-level managers in the NDCC, NAC, or CAC to monitor and resolve weekly problems in TNDS processing for their specific area of responsibility. It also provides district-level summaries for local management and provides area and company summaries for reporting results. All of these functions are accomplished without user effort other than requesting reports via a dial-up terminal. All CSAR processes are fully mechanized.

4.2.2 Center responsibility

Based on the functional responsibility of the three centers, the nineteen TPMP measurement categories have each been assigned to a specific center for initiating corrective action. Six are the responsibility of the NAC manager, three of the CAC manager, and the remainder belong to the NDCC. Although coordination among these and other centers is frequently required, the specific centers assigned each category are responsible for initiating the necessary action to correct any problems reported by CSAR. In addition, Southern Bell has established an objective level of performance for TPMP results, which is that 90 percent of all traffic units perform in the objective range.

Steady progress has been made in Southern Bell toward this objective since TPMP became an official Bell System measurement plan in July 1980. Both goals of TNDS performance measurement and performance improvement have been achieved.

V. BENEFITS OF TNDS

There are a variety of benefits of TNDS from the operating telephone company perspective, particularly when compared to the pre-TNDS environment.

5.1 Timeliness

TNDS produces reports to the user far more rapidly than the old camera/register arrangements. Surveillance reports are generated by EADAS/NORGEN in near-real time. Further, they are provided directly to the appropriate Network Administration Center, or Circuit Administration Center, only for those offices that these centers manage. In this way, the Company personnel can detect and resolve problems in switching or trunking very early—often before customers are aware of any troubles.

The bulk, downstream reports of TNDS (COER,* load balance, trunk servicing) are also made available in a relatively short time. In general, these are delivered to the user a week after the end of any study period. Prior to TNDS, it typically required four to eight weeks for any meaningful network data to be made available to the user.

5.2 Quality

In the pre-TNDS environment, errors were commonplace. These were caused by keypunching mistakes, camera problems, and database synchronization problems. The quality of TNDS data is far superior, owing to mechanization of database validation procedures and the availability of a standard measurements plant. The shorter turnaround time of TNDS also facilitates prompt detection and resolution of database or hardware problems. Those data errors that do occur are not institutionalized over long periods.

5.3 Quantity

The pre-TNDS environment was labor intensive. This essentially limited the number, scheduling, and quality of studies. The quantity of studies produced by TNDS today would have been impossible just a decade ago. TNDS has also given us studies that are far more sophisticated and detailed than had been possible previously.

---

* Central Office Equipment Reports.

5.4 Network investment

The family of TNDS reports permits us to engineer and provision the network more efficiently. This improves the existing network utilization and helps us to better forecast growth. Maintenance of the network is also enhanced, allowing further improvement in use of existing facilities. TNDS, therefore, allows fuller use of our present investment and better planning for new capital expenditures.

5.5 Credibility

TNDS lends a high degree of credibility to network data. The improved timeliness, quality, and quantity of data have eliminated the uncertainties of yesterday that resulted in overprovision of facilities in an attempt to protect service.

VI. FUTURE TNDS DEVELOPMENT NEEDS

A formal enhancement procedure is available to the operating telephone companies by which future needs can be submitted, evaluated on a system basis, and, when approved, designed and developed. Enhancements are grouped into annual releases that are accompanied by training and documentation updates.

TNDS is a relatively mature system. Most enhancements now consist of refinements in report content, or in data-processing capabilities.

Southern Bell is a large, geographically dispersed company, and has needs that differ somewhat from those of smaller, more concentrated companies. Our TNDS is logistically dispersed. We see a real need for a system-supported data-transmission capability that would be quick, accurate, yet relatively economical.

There are two "loops" to this need. The first would be a capability of transmitting magnetic-tape data from the deployed EADAS sites to the central CDC for processing. This would avoid the expense of shipping these tapes by courier, would reduce the turnaround time by at least one full day, and would eliminate occasional tape problems caused by the uncertain environment of courier service.

The second "loop" would transmit the TNDS output to each NDCC (or CAC), or alternatively to the individual user. This would also eliminate substantial handling and shipping costs and reduce the turnaround time by at least another full day.

TNDS, although relatively mature, is dynamic. New data needs and improved technology will guarantee that we will always have opportunities for enhancement as long as there is a need for TNDS.

VII. ACKNOWLEDGMENT

The author wishes to acknowledge the contributions of J. S. Rainwater, Staff Manager, Southern Bell. Mr. Rainwater provided invaluable assistance by reviewing this article and making valuable comments.

AUTHOR

John Pfeiffer, Jr., B.A., 1951, Wesleyan University; M.B.A., 1968, Georgia State University; Southern Bell, 1954—. Mr. Pfeiffer was initially assigned to the Commercial Department, involved in market forecasting and business office operations. Subsequently, he transferred to the Traffic Department and was responsible for Traffic Planning in the Georgia Area. He is currently responsible for Total Network Data System (TNDS) support as well as producing the Total Network Operations Plan (TNOP) and Six-Year Plan in Southern Bell.

ACRONYMS AND ABBREVIATIONS

5XB COER	No. 5 Crossbar central office equipment reports
10HD	ten high day
ABS	average busy season
ACP	action point
ACU	automatic calling unit
AER	apparatus exception reports
AHT	average holding time
AMPS	average measured peak service
ANI	automatic number identification
AOOC	automatic out-of-chain
AORB	all originating registers busy
ASH	abnormally short holding time
ATB	all trunks busy
ATEMIS	Alston Traffic, Engineering and Management Information System
AUTRAX	Automatic Traffic Recording and Analysis Complex
BIS	Business Information Systems
BDAM	basic direct access method
BDE	batch data entry
BHD	busy hour determination
BOC	Bell Operating Company
bpi	bits per inch
C	capacitor
CAC	circuit administration center
CCIS	common channel interoffice signaling
CCS	hundred call seconds
CCSA	common control switching arrangement
CCU	central control unit
CDC	Corporate Data Center
CDO	community dial office
CDR	customer digit receiver
CGA	circuit group analysis
CGM	circuit grouping map
CGMT	circuit group measurement
CLCI	common language circuit identification
CLLI	common language location identifier
CM OUT	collection machine outage
CMP	completions
CO	central office
COER	central office equipment reports
COSMOS	computer system for mainframe operations

CPU	central processing unit
CRT	cathode ray tube
CSAR	Centralized Systems for Analysis and Reporting
CU	central unit
CU/EQ	common update/equipment
CU/TK	common update/trunking
CV	coefficient of variance
DAS	Data Access System
DCD	data collection device
DCU	data collection unit
DDD	direct distance dialing
DDM	daily data manager
DGU	detector group usage
DIS	Data Inquiry System
DIXC	data interchange
DLI	dial line index
DMA	direct memory access
DNHR	dynamic nonhierachical routing
DOC	dynamic overload control
DOC-N	district operations center – network
DP	dial pulse
DRE	directional trunk reservation
DRP	division of revenues process
DTD	dial tone delay
DTM	dial tone marker
DTP	data transfer point
E2AP	E2A process
EADAS	Engineering and Administrative Data Acquisition System
EADAS/NM	Engineering and Administrative Data Acquisition System/Network Management
ECCS	economic hundred call seconds
EDP	electronic data processing
EM	electromechanical
EMC	equipment measurement code
EMS	electromechanical system
EQ	equipment
ER	exception reports
ETDC	electronic traffic data collection
EVE	extreme value engineering
FEPS	Facility and Equipment Planning System
FIPS	file interface process
FIT/RAT	fitted ratio
FM	failure to match

GDG	generation data group
GFA	growth factor analysis
GP	graphics process
GTF	general trunk forecast
HD	high day
HT	holding time
HTR	hard-to-reach
ICAN	individual circuit analysis
ICUR	individual circuit usage recording
ID	identification
IMA	ineffective machine attempts
IMAP	Interactive Map Assembly Program
INA	ineffective network attempts
INU	incoming usage
INWATS	Inward Wide-Area Telephone Service
I/O	input/output
IPC	incoming peg count
IPC	interprocessor communication
IR	incoming register
ISO	Information Services Organization
KDP	keyboard, display, printer
LBI	load balance index
LBS	Load Balance System
LEX	lexical
LFG	line finder group
LFN	logical file number
LFS	Logical File System
LL	Long Lines
LTB	last trunks busy
LTU	last trunk usage
LU	load unit
MBC	maintenance busy count
MLSM	MLSS data manager
MLSS	Machine Load Service Summary
MMAR	monthly machine administrative report
MMOC	Minicomputer Maintenance and Operations Center
MON	Monitor System
MPRT	monitor print process
MPS	Message Processing System
MRCS	Modification Request Control System
MRS	Management Reporting System
MS	main station
MTO	message trunk order
MTS	message telecommunications service

MTU	maintenance usage
MUX	multiplexor
MVS	multiple virtual storage
NAC	network administration center
NCP	network control point
NDA	NORGEN data analyzer
NDCC	network data collection center
NM	network management
NMC	network management center
NOC	network operations center
NOCS	Network Operations Center System
NORGEN	network operation report generation
NPA	numbering plan area
NSEC	network switching engineering center
NSPMP	Network Switching Performance Measurement Plan
O	overflow
OCR	optical character recognition
ODF	office description file
OR	originating register
ORRS	On-line Records and Reporting System
OS	Operations System
OS	outgoing sender
OTC	operating telephone company
OTSS	Off the Shelf System
OVF	overflow
PA	periodic average
PATROL	Program for Administrative Traffic Reports On Line
PBC	peripheral bus computer
PBS	Playback System
PBX	private branch exchange
PC	peg count
PECC	product engineering control center
PGTF	preliminary general trunk forecast
PIR	performance indicator report
PSF	pending status flag
PSR	performance summary report
PWB	Programmer's Workbench
QCL	quality control limits
RA/BS	ratio to busy season
RCMAC	recent change memory administration center
RDUS	Reference Data Uptake System
ROC	regional operations center
ROP	receive-only printer
RSL	report specification language

RSS	Remote Switching Systems
RTS	reroute to seizure
SCC	switching control center
SCCS	Source Code Control System
SDOC	selective dynamic overload control
SLC	subscriber line concentrators
SLU	subscriber line usage
SMA	service month-to-date accumulation
SOE	standard operating environment
SOM	service observing month
SONDS	Small Office Network Data System
SPA	sequential projection algorithm
SPC	stored program control
SPCS	Stored Program Control System
SPCS COER	Stored Program Control System Central Office Equipment Reports
STP	signal transfer point
STR	short-term retention
STR	selective trunk reservation
SXS	step by step
TAMP	Trunk Administration Measurement Plan
TCBH	time-consistent busy hour
TCC	TNDS coordination center
TCT	telemetry computer translator
TDA	traffic data administration
TDAS	Traffic Data Administration System
TDF	tape data formatter
TDRS	Traffic Data Recording System
TFS	Trunk Forecasting System
TGSN	trunk group serial number
TIP	Trunk Implementation Plan
TIRKS	Trunks Integrated Record Keeping System
TMR	traffic measurement request
TNDS	Total Network Data System
TNDS/EQ	Total Network Data System/Equipment
TNDS/TK	Total Network Data System/Trunking
TNOP	Total Network Operations Plan
TNPC	traffic network planning center
TOPC	total office originating peg count
TPMP	TNDS Performance Measurement Plan
TRK	trunk
T/S	total to sample
TSO	time-sharing option
TSS	Trunk Servicing System

TTL	transistor-transistor logic
TTY	teletypewriter
TU	traffic unit
TUR	traffic usage recorder
U	usage
UA	unassigned
UE	unequipped
UPCO	usage, peg count, and overflow
USG	usage
USITA	United States Independent Telephone Association
VM/CMS	Virtual Machine/Conversational Monitor System
VSAM	virtual sequential access method
VSS	Voice Storage System
WBP	wallboard process
WCPC	wire center planning center
WMAR	weekly machine administrative report
XPRT	exception print process
YACC	yet another compiler compiler

THE BELL SYSTEM TECHNICAL JOURNAL is abstracted or indexed by *Abstract Journal in Earthquake Engineering, Applied Mechanics Review, Applied Science & Technology Index, Chemical Abstracts, Computer Abstracts, Current Contents/Engineering, Technology & Applied Sciences, Current Index to Statistics, Current Papers in Electrical & Electronic Engineering, Current Papers on Computers & Control, Electronics & Communications Abstracts Journal, The Engineering Index, International Aerospace Abstracts, Journal of Current Laser Abstracts, Language and Language Behavior Abstracts, Mathematical Reviews, Science Abstracts (Series A, Physics Abstracts; Series B, Electrical and Electronic Abstracts; and Series C, Computer & Control Abstracts), Science Citation Index, Sociological Abstracts, Social Welfare, Social Planning and Social Development,* and *Solid State Abstracts Journal.* Reproductions of the Journal by years are available in microform from University Microfilms, 300 N. Zeeb Road, Ann Arbor, Michigan 48106.

Bell System