THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 56, No. 7, September 1977
Printed in U.S.A.

No. 4 ESS:

Network Management and Traffic Administration

By T. V. GREENE, D. G. HAENSCHKE, B. H. HORNBACH, and C. E. JOHNSON

(Manuscript received July 30, 1976)

The No. 4 Electronic Switching System has been planned to provide operational ease in the efficient management of traffic-sensitive equipment for the purpose of maintaining a high quality of service. The network management function has the goal of optimizing the completion of calls during periods of traffic stress. In No. 4 ESS, innovative real-time control and surveillance features are provided to meet this goal. Traffic administration involves the activities of personnel in managing the traffic-sensitive equipment in an efficient, timely, and economical fashion. These activities depend upon the collection of data which reflects the operating characteristics of the No. 4 ESS, and the reporting of information in a form tailored to the user functions.

I. INTRODUCTION

High-quality service requires the proper planning for the quantities of central office and trunking equipment and the effective utilization of this equipment in the operational environment. The network management and traffic administration functions in the No. 4 ESS are planned with these objectives in mind. The network management function is directed at maintaining a high level of service during unusual traffic situations. The traffic administration function has two aspects: network engineering and machine administration. Basically, network engineering encompasses the planning function associated with the dimensioning of equipment, and machine administration consists of the ongoing activities aimed at meeting service objectives. A fundamental distinction between these activities is their differing time frames:

(*i*) The network management activities deal in real time with the unexpected or unusual situation.

(*ii*) Machine administration activities deal with maintaining service objectives on a daily and monthly basis.

(*iii*) Network engineering activities are concerned with future planning for customer service.

1.1 Network management

It is a property of modern telephone networks with common control switching and alternate routing arrangements that they are highly efficient and economical under engineered load conditions but deteriorate under overloads. Once this decline sets in, it is difficult to recover even if the load level is reduced to normal.[1]

The reasons for this decline in efficiency are excessive alternate routing[2] and regenerative switching (queuing) delays. The latter has the more direct impact on performance and its prevention is a prime objective of the No. 4 ESS network management system. Regenerative switching delays, if left uncontrolled, can quickly spread throughout the network, causing the type of decline in carried load shown in Fig. 1. This figure is from Burke who first analyzed this phenomenon in detail.[3] The mechanism at work here is waste usage of common control equipment which, for instance, can be caused by transmitters in one office waiting for receivers in another. A similar throughput decline occurs also within a single heavily overloaded No. 4 ESS because of loss in internal efficiency due to increasing real-time overhead.

The network management system must provide the real-time control and surveillance capabilities that are needed to realize the economical advantages of alternate routing and modern switching systems, without compromising service during stress. In the mid-1960s the dynamic overload control system[3–5] and manual trunk group controls were introduced, which have proven to be reasonably effective during peak-day overloads such as Christmas and Mother's Day. However, these controls are not code selective and, therefore, not very effective during focused overloads which are characterized by a surge of traffic from all parts of the network to a small number of offices or destination codes. In No. 4 ESS, a major innovation in surveillance capabilities is the automatic determination of destination codes which have a low probability of completion and are said to be "hard to reach." This information reveals equipment failure and traffic congestion at points far removed from the No. 4 ESS or its trunking field.

The integration of hard-to-reach code data with automatic controls accomplishes the objective of making this system responsive to a wide range of trunk facility and machine overloads. This enhanced responsiveness of the automatic controls should minimize the need for manual control interactions. Human reaction times are often too slow to prevent

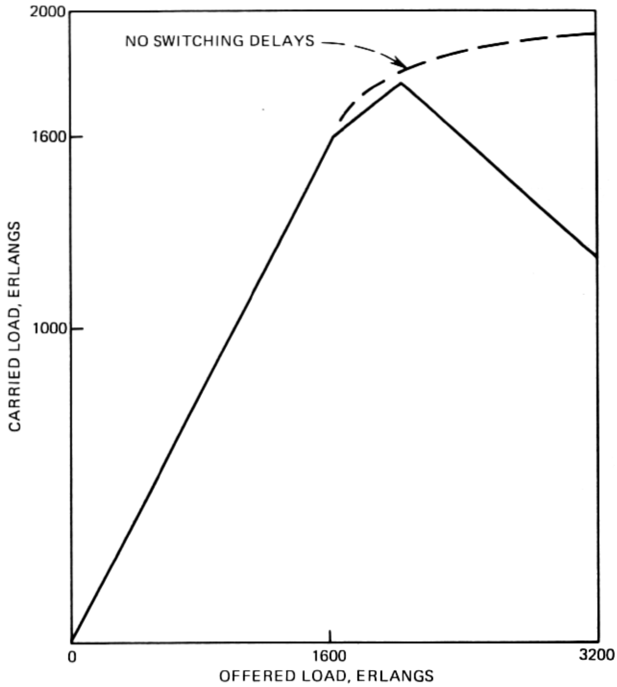

Fig. 1—Network performance under overload.

the earlier-mentioned decline in network efficiency during unexpected overloads.[6] The system has been planned so that network managers can concentrate their efforts on the general supervision of real-time network performance, analysis of network weak spots before they become service-affecting, preplanning of control strategies, fine tuning of the automatic system, and intervening with manual control actions only in problems requiring human judgment. For the surveillance of real-time network performance, a display system is provided which furnishes the network manager with preprocessed data. In the past, network management problems had to be determined from traffic registers, trunk-group busy lamps, and a few machine-status lamps. In No. 4 ESS, potential problems are automatically reported and detailed problem investigation is accomplished using CRT pages programmed for efficient data analysis.

1.2 Traffic administration

The No. 4 ESS, serving as a switching node in the telecommunication network, has to have its central office equipment and trunking to other offices properly engineered and administered to meet its service objectives.

The quantity of provided equipment is determined by the size and character of anticipated traffic loads. In planning the growth of the telephone network, the network engineer must determine how much equipment will be needed so that it can be ordered and installed in time to provide an objective level of service throughout the engineering periods ranging from 6 months to 2 years in length.

In the No. 4 ESS, as in other systems, the network engineer is responsible for both central office and trunk equipment. Central-office engineering involves the establishment of configurations for new offices and office additions, and the specification of quantities of switching subsystems to be installed. Trunk engineering involves the determination of which new trunk subgroups* will be established, the number of trunks each should contain and the changes to be made in the size of existing trunk subgroups.

Once the equipment is installed and working, the machine administrator is concerned with potential and real service-affecting problems occurring within shorter intervals of time than those of interest to the network engineer. However, the machine administrator is not expected to react as quickly as the network manager to service problems, since the network manager is primarily responsible for maximizing the flow of traffic through the network during unexpected events of a transient nature (such as earthquakes and storms).

In preventing and solving service-affecting problems, the machine administrator is responsible for the processing, collection, analysis, and distribution of traffic data. The performance and loading of central-office equipment and trunk equipment should be periodically analyzed and trended to assure sufficient capacity for handling existing and future traffic loads. The machine administrator is also responsible for maintaining the objective service level during office transitions when equipment is repaired, added, or replaced.

The basis for satisfying the information needs of these engineering and administrative functions is a comprehensive set of measurements. The measurement data which is maintained for the previous hour reflects both service as seen by the customer as well as performance of the switching equipment. Subsets of this data are extracted and maintained for extended intervals to tailor several information bases in a cost-effective fashion.

Through sorting and processing the collected measurement data, the information needed by network engineers and machine administrators

* The term "trunk subgroup" is No. 4 ESS nomenclature for the set of trunks in a given route with certain common features such as directionality, pulsing type, and transmission delay characteristics, i.e., satellite versus terrestrial; each trunk subgroup is an entity in itself for traffic measurements and network management control.

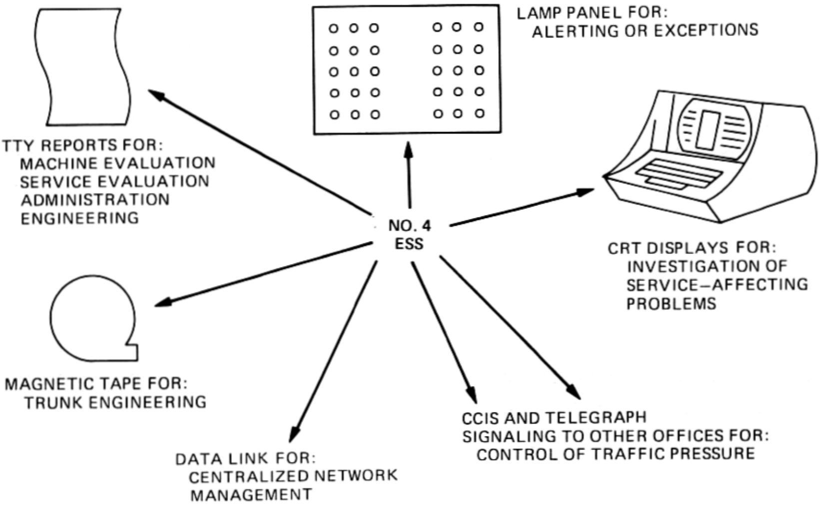

Fig. 2—System interface.

is provided in the form of hard-copy reports which can be flexibly scheduled and directed to work centers.

The type of the report can vary from an overview report obtained by extensive processing of data to a detailed printout of the individual measurements. This structuring of reports, from overview to detail, is intended to minimize the volume of paper generated. In general, overview reports should satisfy the information needs of office personnel and will be supplemented with detailed reporting only when exceptions are indicated by the overview report.

Along with hard-copy reports, a magnetic tape can be scheduled for downstream processing. This tape contains the statistics for the trunk subgroups in the No. 4 ESS which, along with similar information from other toll offices, are analyzed by a downstream process to produce forecasts of message-trunk requirements.

II. SYSTEM INTERFACES

The network management and traffic administration features provide interfaces with office personnel and other machines (Fig. 2).

The main personnel-machine interface is realized with Model 40 teletypewriters which are serviced through the I/O channels of the 1A Processor. The reports for general administration and engineering the office can be scheduled to be output on any of the standard I/O channels. Printed on fanfold paper, each page of report is contained on a separate sheet for assembly into a report folder.

The real-time surveillance system for network management provides for both local and remote network management centers. The Model 40 teletypewriter terminals for providing CRT displays can be remoted through a standard arrangement using 202 data sets. The lamp-driving circuitry for the network-management display panel can interface with an E2A telemetry system for the remoting of the lamp display. The panel consists of Nixies to provide numeric information and lamps for indicating YES/NO type of exceptions which are color coded to denote problem severity. Since the panel is an exception-reporting medium, it has been designed to have a darkened appearance if the system is operating normally.

Besides direct communication with office personnel, three machine-to-machine interfaces are provided. For network management purposes, communication interfaces are provided to pass control information between switching offices and to provide data to centralized network-management systems. For traffic administration, a magnetic tape containing trunk utilization statistics is generated for processing by a downstream system.

The network management control plan calls for the transmittal of machine congestion signals between interfacing offices. There are three levels of dynamic overload control, or DOC, signals that can be transmitted to control the traffic on a trunk subgroup between the offices. To control multifrequency and dial-pulse traffic, signals are transmitted over telegraph channels or data channels between the offices. For integrity purposes, the office receiving a DOC signal will send an acknowledgment signal to the sending office. In this way, impairments to the signaling channels can be detected either through failure to receive an acknowledgment or the receipt of a false acknowledgment. For trunk subgroups that are served by Common Channel Interoffice Signaling[7] (CCIS), the DOC information is passed between offices using the CCIS channel. DOC acknowledgment signals are not used in the CCIS case since the CCIS signaling channels have built-in self-checking features.

A data-link interface is provided for centralized surveillance and control. In particular, standard arrangements will be available for transmitting network management data to a remote centralized network management system called EADAS/NM (Engineering and Administrative Data Acquisition System/Network Management). EADAS/NM is presently being introduced into the Bell System to provide network management surveillance and manual control of an entire geographical area, such as a state, not just a single office. The No. 4 ESS will supply EADAS/NM with a subset of its measurement and control status data for processing and display through EADAS/NM.

The final machine-machine interface provides trunk engineering data to a downstream processing system. The No. 4 ESS generates a standard label tape which contains hourly statistics for all the trunk subgroups in the office. An identification map for the trunk subgroups in the office is provided daily on this tape to furnish a means for checking the consistency of the trunk assignment data contained in the data bases of the No. 4 ESS and the downstream processing system.

III. NETWORK MANAGEMENT CONTROL

The goal governing the application of network management controls is to give the best possible service during overloads by utilizing all available facilities as efficiently as possible. These control actions, whether manually or automatically instituted, modify the routing of traffic. Such actions, however, are not a substitute for proper engineering and equipment provisions.

3.1 Hard-to-reach code analysis

A Hard-To-Reach (HTR) code is a 3- or 6-digit destination code to which successfully outpulsed calls have a very small chance of completing. If the probability of completing through the distant network is very low and the outgoing trunk groups or connected switching offices are congested, the waste usage of these overloaded network resources for traffic to HTR points should be prevented.

To this end, the No. 4 ESS automatically monitors the completion rate on all calls to each numbering plan area (NPA) code. It also monitors call completion to each central office (NXX) code in its home NPA and in as many as six other NPAs which can be selected by the network manager in real time. Every 5 minutes, the No. 4 ESS counts the number of "Network Attempts" (NA) that are successfully forwarded to the next office, and of those, the number of "Ineffective Network Attempts" (INA) which abandon without receiving answer supervision. The INA count represents the number of calls that fail in the distant network for any reason, including line busy and line doesn't answer conditions. The ratio of abandonments, INA, to network attempts, NA, is called the INA rate. Normally, the INA rate will run about 20 to 40 percent depending on the incidence of line busy and doesn't answer conditions. In focused overloads, the INA rate may be in excess of 90 percent.

The No. 4 ESS places an NPA or NXX code on the HTR list if for a statistically sufficiently large NA:

$$\frac{\text{INA}}{\text{NA}} > T_1$$

where the on-list threshold level, T_1, is typically 70 percent. Once a code is declared hard to reach, it will be automatically restored to normal if:

$$\frac{\text{INA}}{\text{NA}} < T_2 < T_1,$$

where the off-list threshold level, T_2, is typically 60 percent. The threshold value, T_2, is chosen to be less than T_1 to avoid oscillation when the control of HTR codes results in a reduction in the INA rate but the basic hard-to-reach cause remains in the network.

By definition, HTR codes are those codes with a very low probability of completing after they have been forwarded to the next office. They do not reflect completion problems the No. 4ESS may encounter prior to outpulsing, such as no-circuit conditions and time-outs. These failures prior to outpulsing are called "Ineffective Machine Attempts" (IMAs). The No. 4 ESS gathers IMA data for the same codes for which it gathers INA data. The IMA counts are available to the display system and provide the network manager with valuable per-code switching machine performance data. However, a No. 4 ESS will not use the IMA counts for entering codes on its own HTR list. This is done because a high IMA count for a given code is not an indication that calls to this code will have a small chance of completing once they succeed in seizing outgoing trunks.

The No. 4 ESS's ability to determine HTR codes anywhere in the network can benefit not only the No. 4 ESS but also other switching offices that are able to administer an HTR code list for use with controls but which are unable to make an HTR code determination on their own, for example, No. 4A toll crossbar offices arranged for CCIS.[7] For these offices, the No. 4 ESS can serve as a network information center by informing them of HTR code problems.

3.2 Automatic controls

3.2.1 Protective automatic controls

Protective network management controls are employed during overloads to prevent the decline in carried load due to switching delays shown in Fig. 1, and to achieve best possible use of available trunk facilities. Two new types of protective controls are employed automatically by the No. 4 ESS: selective dynamic overload control (SDOC), which responds to machine congestion, and selective trunk reservation, which responds to trunk congestion.

The stimulus to SDOC comes from a connected switching office that is sensing machine congestion. Two levels of machine congestion, a "low" and a "high" level, can be sensed with present nonselective DOC as well as by the No. 4 ESS overload-control program. As shown in Fig. 3, the No. 4 ESS will typically respond to the "low" congestion signal by restricting HTR traffic that is headed for the congested machine. Only when a "high" congestion signal is received will other traffic be restricted, typically all alternate-routed calls. The No. 4 ESS's ability to respond to "low" congestion signals with control only if an HTR problem exists,

SELECTIVE DYNAMIC OVERLOAD CONTROL		SELECTIVE TRUNK RESERVATION
CONGESTION LEVEL OF REQUESTING MACHINE	TRAFFIC CONTROLLED	TRUNKS REMAING IDLE IN GROUP
NONE	NONE	$n > n_1$
LOW	HTR ONLY	$n_1 \geqslant n > n_2$
HIGH	HTR + OTHER ALTERNATE ROUTE	$n_2 \geqslant n$

Fig. 3—Protective automatic controls of the No. 4 ESS.

coupled with its ability to transmit and receive such signals via CCIS without the extra cost of separate signaling facilities, allows a much wider deployment of SDOC than was possible with nonselective DOC.

The No. 4 ESS overload-control program continuously looks for shortages in real time and common equipment, and will initiate the "low" and "high" congestion signals when shortages are detected. It also governs the allocation of its own resources during overloads to prevent loss in internal efficiency. For instance, it will postpone nonessential tasks and, as a last resort, reduce the accepted traffic load if a sufficient load reduction does not occur through the use of DOC controls. The No. 4 ESS is also arranged to transmit a unique signal when it is no longer able to process new calls because of a major software or hardware failure resulting in a system outage. This signal is broadcast to connected offices to protect them from wasting resources while attempting to forward calls to the failed No. 4 ESS.

The other type of protective automatic control, called selective trunk reservation (STR), responds to trunk congestion as measured within the No. 4 ESS by the number of trunks remaining idle in the associated trunk subgroup. The response to trunk congestion, also shown in Fig. 3, is quite similar to that with SDOC. When the number of idle trunks, n, is less than n_1, access will be denied to the few remaining idle trunks only for HTR traffic that is headed for the congested trunk subgroup. If this HTR definition does not apply, no controls are taken. When trunk congestion builds up and less than n_2 idle trunks remain, $n_2 < n_1$, trunk access can also be denied to other alternate-routed traffic. However, this second control step is usually applied only to the last-choice trunk subgroups for call routing, called finals, with the objective of protecting service of direct-routed (nonalternate-routed) traffic during heavy alternate routing. The No. 4 ESS will be able to meet this objective without the trunk penalty associated with the present practice of splitting the final trunk group into two groups, one dedicated to direct-routed traffic, the other to alternate-routed traffic.

Calls denied access to an outgoing trunk subgroup by SDOC or STR can be canceled and routed to an announcement, or can be skipped over that trunk subgroup to an alternate route with idle capacity. SDOC and STR are more powerful than nonselective DOC because they are able to respond automatically to a much wider range of overload problems. They can be applied without the hierarchical restrictions inherent in the older controls. They are compatible with existing controls and need not be deployed extensively to be of benefit.

These claims of increased effectiveness of SDOC and STR have been verified by computer simulations under conditions of peak-day and focused overloads. One set of results, derived from a simulated 24-switching-office network[4] is shown in Fig. 4. The results apply to a focused overload during which the load offered to a given office increased 8-fold compared with normal levels. This type of overload is equivalent to that experienced by offices during the 1971 earthquake in southern California. The figure shows the number of successful messages in progress versus time, measured from the beginning of the overload. Without controls, we see the transient decline in carried load, which is similar to the decline shown in Fig. 1 for the static case. Present automatic controls improve call completion but are unable to maximize network utilization. STR, without SDOC, is able to keep network completions high for about an hour. Eventually, however, the buildup of ineffective short-holding-time attempts reduces trunk subgroup occupancies below the point where STR will trigger. Nevertheless, STR alone can keep the network operating efficiently for a sufficiently long time to allow the network manager to intervene with additional manual control. This is particularly significant since STR, unlike SDOC, is a completely autonomous protective control which requires no signaling channels from distant offices.

STR combined with SDOC provides close to optimum call-carrying capacity[8] and furnishes a marked improvement over present nonselective automatic controls. The actual improvement of SDOC and STR over present controls amounted to over 60 percent in the scenario presented in Fig. 4. While these improvements are greatest for focused overloads, simulation results indicate that similar benefits are obtained for peak-day overloads, such as those encountered on Christmas Day.

Further analysis of the simulation results indicate that selective blocking of HTR traffic with SDOC and STR increases network efficiency sufficiently to also benefit HTR traffic and, therefore, improve service into overloaded offices. An even greater service improvement results for customers calling out of the overloaded offices. This means that calls

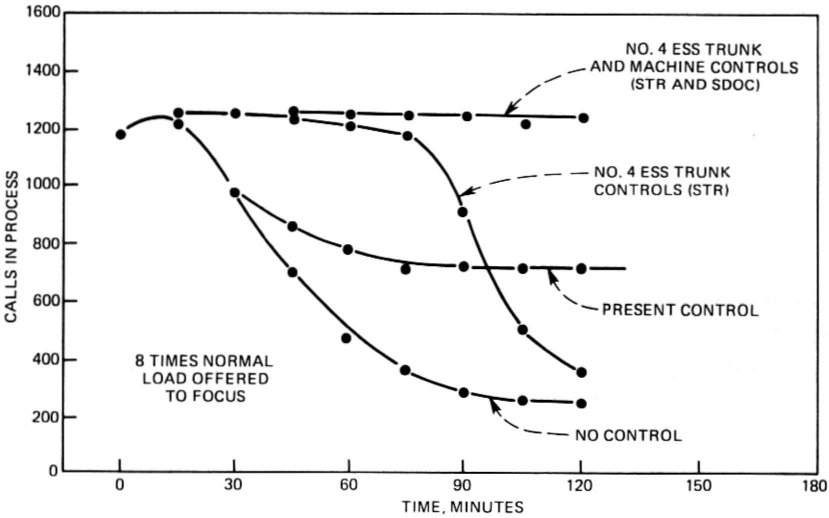

Fig. 4—Control of focused overload.

from the overloaded offices are automatically given the desired preference over incoming calls. Simulation results also show that worthwhile improvements can be attributed to SDOC and STR during their initial deployment when only a small portion of the network consists of No. 4 ESSs.

3.2.2 Expansive Automatic Controls

In contrast to the protective controls which restrict access to overloaded facilities, expansive controls take advantage of idle capacity on "out-of-chain" routes, i.e., on routes which are not within the design of the hierarchical routing structure. Whereas expansive controls have been exclusively manual up to now, the No. 4 ESS permits out-of-chain routing on an automatic basis.

The new feature which accomplishes this is called Automatic Out-Of-Chain (AOOC) routing. It permits calls which overflow the final in-chain trunk subgroup to be offered to out-of-chain trunk subgroups. Out-of-chain trunk subgroups are preassigned according to destination codes so that a call overflowing the final in-chain trunk subgroup will be offered to one of up to seven out-of-chain trunk subgroups that have good connectivity toward the code's destination. When more than one out-of-chain trunk subgroup is provided, traffic will be spread evenly over all trunk subgroups.

AOOC routing is allowed only if there is ample idle capacity in the out-of-chain trunk subgroup, in the "via" office at the far-end of that trunk subgroup, and in the outgoing trunking field from the via office towards the call's destination. This controlled use of out-of-chain routes is accomplished by turning off the availability of an out-of-chain trunk subgroup for a timed interval when any of the following conditions occur: less than a predetermined number of trunks idle towards the via office; the receipt of a "low" or "high" machine congestion signal from the via office; or the receipt of a CCIS trunk congestion message indicating that the last call encountered an all-trunks-busy (no-circuit) condition at the via office or a subsequent office.

The full protection and discipline inherent in AOOC routing will be obtained if the out-of-chain trunk subgroup to the via office is equipped for CCIS. With CCIS, each call forwarded to the via office will have a special traveling classmark attached. This classmark instructs the via office to restrict outgoing trunk selection to the most direct route toward the call's destination. If no idle circuits are found in a direct route, the via office will return the above-mentioned CCIS trunk-congestion message, even when alternate routes are idle. This forces the temporary turn-off of out-of-chain routing for traffic that the via office cannot complete on a direct route. The dynamics inherent in AOOC routing promises to be the key to a more effective utilization of idle capacity on out-of-chain routes than has ever been possible with manual methods.

3.3 Manual controls

Even though the emphasis is on improved automatic controls, manual controls are needed for solving problems which are not recognized by the automatic system because they require human judgment and interpretation of environmental data. In addition, the network manager must be allowed to override the automatic system and fine-tune its response.

The extensive set of manual controls provided with the No. 4 ESS is quite similar to that available in other major toll switchers. It includes the ability to block calls to any 3-, 6-, 7-, and 10-digit destination code specified by the network manager and to route blocked calls to a specially worded announcement. It also includes protective and expansive controls on a per trunk subgroup basis.

Traffic offered to a trunk subgroup can either be canceled and routed to announcement, or forced to skip the trunk subgroup and advance to an alternate routing choice. Traffic overflowing a trunk subgroup can be canceled or rerouted to an out-of-chain route. The application of manual trunk subgroup controls can be limited to HTR codes which is a new capability and provides a high degree of selectivity. Other options provided are the ability to specify control percentages, ranging from 25 percent to 100 percent, and a choice between direct and alternate-routed traffic for trunk subgroup skip and cancel controls. Manual override capability is provided for all automatic controls as well as for the HTR code detection systems. A code can manually be declared as hard to reach—or excluded from HTR treatment—for any outgoing trunk subgroup selected by the network manager.

IV. NETWORK MANAGEMENT SURVEILLANCE

The No. 4 ESS network management display system provides the network manager with the data needed for the surveillance and manual control of the network. The system has access to a large No. 4 ESS-maintained data base consisting of both traffic and plant measurements as well as measurements specially collected for network management. The latter includes 5-minute per code completion data, 5-minute trunk subgroup performance measurements, and measurements of the number of calls affected by each automatic and manual control action taken by the No. 4 ESS. The display system also has access to control and equipment status data which are updated on a per-event basis, as well as to information that shows the outgoing trunk subgroup choices for each destination code.

Since the data base is large, the surveillance function is done in two steps. The first step involves the automatic detection of potential network management problems and the alerting of the network manager to these problems through visual and audible indicators. The second step consists of problem investigation through data analysis in which the network manager employs the capabilities of an interactive CRT display system.

4.1 Surveillance arrangements

Network management problems are diagnosed by the No. 4 ESS through "exception calculations" and recognition of critical events. Exception calculations consist of a set of calculations that the system performs typically every 5 or 15 minutes on a large number of machine and trunk subgroup performance measurements, which are compared against preset thresholds. This machine processing of large volumes of data permits all potential problems (called "exceptions") to be compressed and displayed on about 160 ON-OFF indicators and about 20 numerics which make up the exception panel. The onset of an exception indicator shows that an associated exception calculation exceeded its threshold or that a critical event has taken place.

The onset of an indication on the exception panel will direct the network manager to one of ten CRT pages which provide an overview of the problem. These overview pages are supported by about 50 detailed pages which are organized so that the network manager searches in a pyramidical fashion from a gross indication of a network problem to a detailed description of the problem.

The CRT system is supported by software programs which allow data to be retrieved, processed, and formatted for display. No longer will the network manager have to search through raw data to investigate a problem. The method of presenting data and the page layout has been carefully optimized from a user-oriented point of view. All CRT pages are interactive and consist of a fixed background and data windows, as described in Section 4.3.

Manual controls are activated and deactivated through a standard set of ten interactive CRT displays called "CRT control pages." Control pages offer greater flexibility over hardware control consoles and provide the opportunity for an immediate feedback of data pertinent to the control action.

4.2 Exception reporting

As shown in Fig. 5, the exception panel contains functional groupings of lamps. The exceptions represented on the panel are indicating either the occurrence of events or the analysis of measurement data that reflect key aspects of system performance. For statistical reliability, classes of performance data are analyzed at different rates, namely 30 seconds, 5 minutes, and 15 minutes.

Normally, the only lighted indications are the numerics in the upper left of the panel which are indicating the volume, measured in thousands, of incoming and outgoing traffic in the past 15 minutes. As exceptions occur, green/red/amber/white lamps are illuminated. Green lamps indicate that manual controls or manually initiated overrides are in effect. These lamps serve not only to remind the network manager who invoked the action but also to inform the other center in a local/remote arrangement that manual control actions are taking place. Exceptions related to system performance are indicated by red/amber/white lamps. The color of the lamp is keyed to the significance of the exception. To draw attention to an important change on the exception panel, a spurt audible alarm accompanies the lighting of a lamp.

The majority of exceptions are determined by checking various aspects of traffic-handling performance against a set of acceptable criteria or thresholds. Since exceptions are meant to alert the network manager to potential problems which may warrant manual intervention, the selection of thresholds is an ongoing process. For instance, network performance differs for an average business day and a peak day, such as Christmas. Consequently, CRT pages were designed to allow threshold values to be easily changed.

An intrinsic part of the exception-reporting system is an exception printer. This printer provides a chronological record of the key machine status exceptions, network performance exceptions, and manual and automatic control actions. These printouts provide details on the ex-

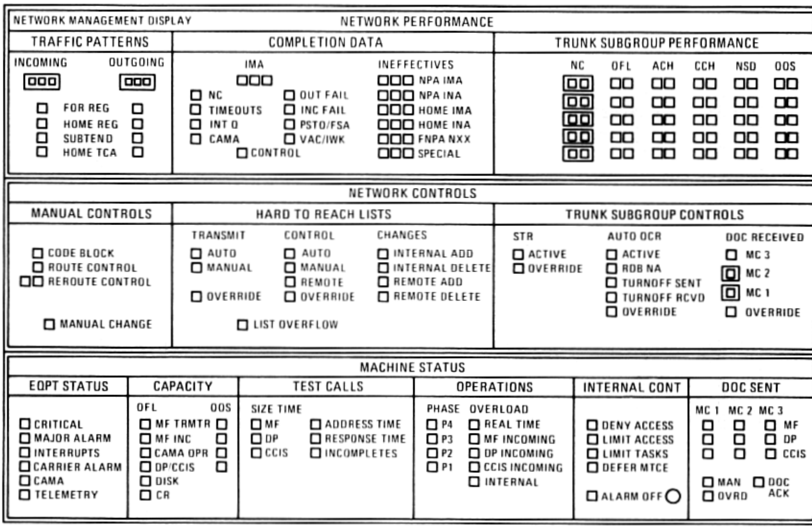

Fig. 5—Network management exception display panel.

ception analysis which resulted in a lamp indication. For instance, the lighting of an OFL lamp in the TRUNK SUBGROUP PERFORMANCE section of the panel would be accompanied by a printout identifying the trunk subgroups which exceeded the "percent overflow" threshold. In a similar manner, other printouts detail the exceptions for the network manager. Besides supplying information for the real-time analysis of problems by network managers, the exception printouts provide a historical record of exception conditions and control responses for post-problem critiques.

4.3. CRT display system

The CRT display system provides a flexible capability for investigating problems that are brought to the network manager's attention by the exception panel. Since effective problem solving is based, to a large extent, on empirical knowledge, the design is aimed at providing a flexible set of general-purpose capabilities for the analysis, presentation, and interactive querying of system data rather than a set of generically imbedded investigative sequences. The specified displays of information which are presented on a CRT screen are determined by specifications which are stored in the No. 4 ESS as sets of data that can be interpreted by the generic software system. These page specifications can be readily changed to accommodate modifications to problem-solving approaches based upon experience or upon changing characteristics of the network.

The displays presented on the CRT screen can be categorized into three

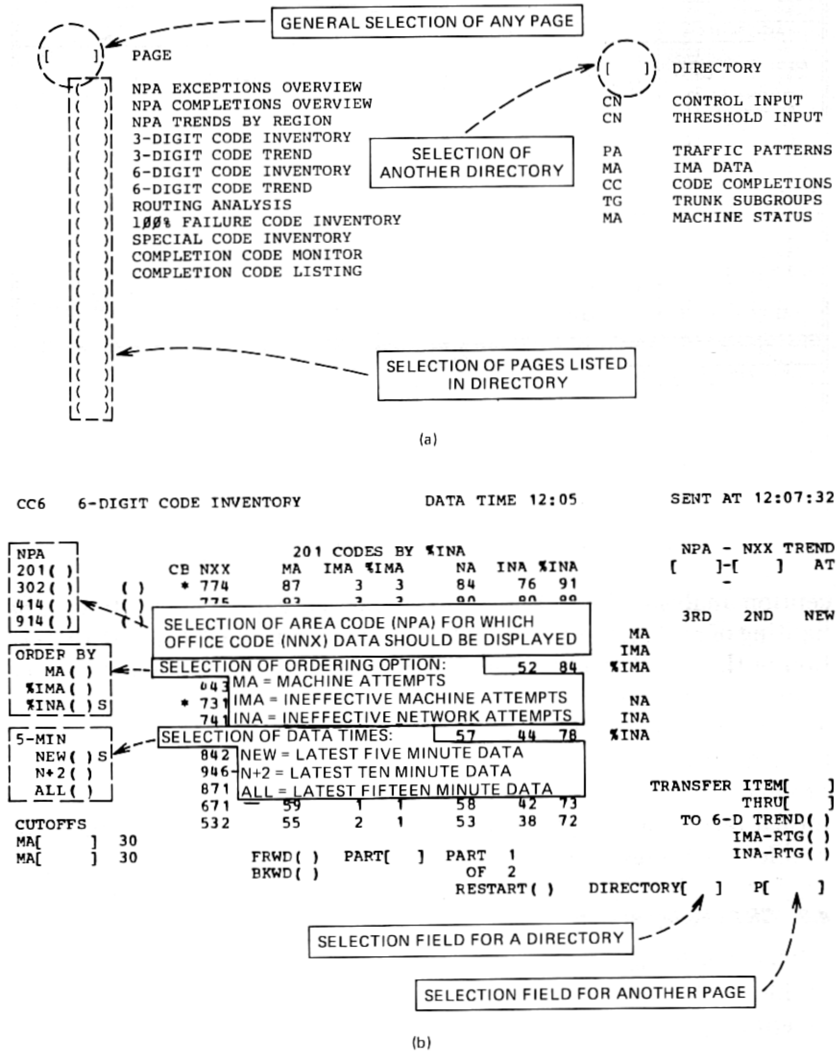

Fig. 6—(*a*) Example of a directory page. (*b*) Example of a display page.

general types: directory pages, display pages, and control pages. The directory pages provide a listing of the control and display pages that are contained within the system. The display pages provide processed system information on the CRT screen while the control pages allow the selection of the controls discussed in Sections 3.2 and 3.3.

CRT terminals are dedicated to this display function and the user operates in a "closed" system in the sense that the only valid input requests are either for another page or an interaction on the present page. The user makes selections of displays using the mobility capabilities

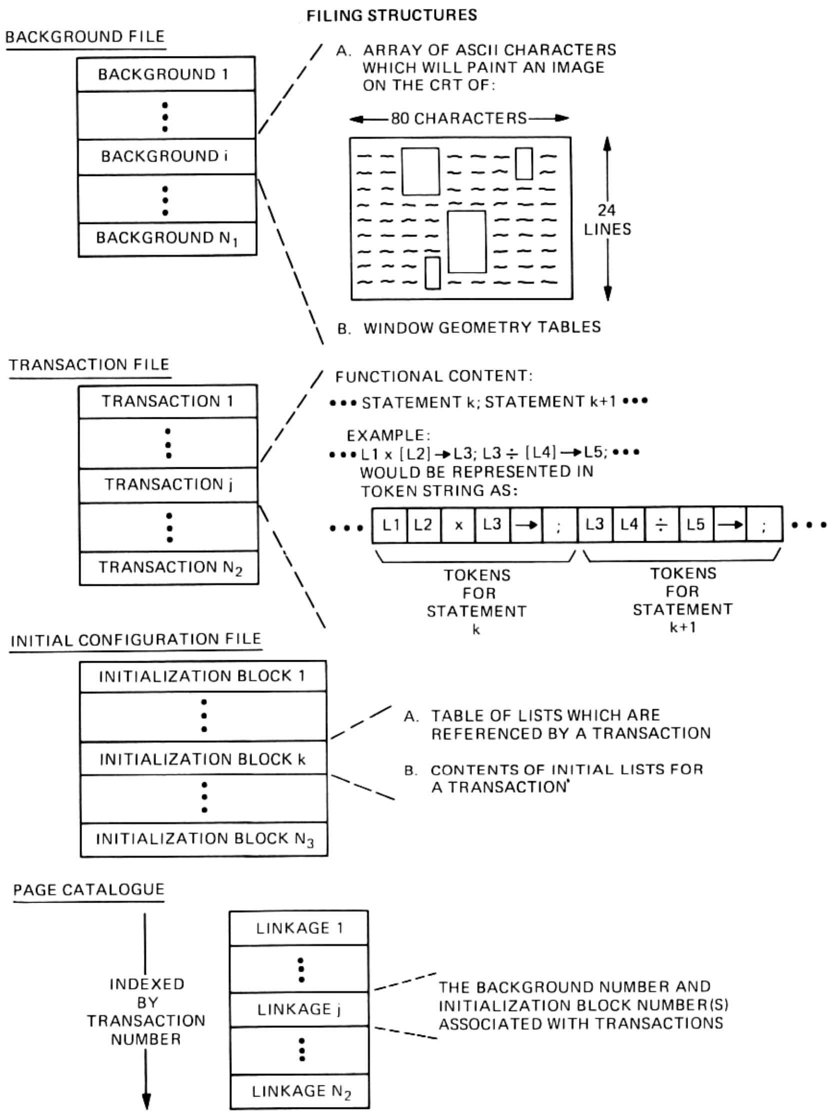

Fig. 7—Page data.

shown in Fig. 6. On a directory page, the user can select a page listed in the directory by placing a check mark next to the desired page and transmitting the request into the system or by specifying the page directly, e.g., CC6. On any nondirectory page the user can make a direct selection of another display/control page or a directory page. Whenever the system cannot interpret a request that has been transmitted from the CRT, the first directory page is displayed.

Display and control pages are generally designed to allow user interactions in order to realize a simpler and more orderly investigation of system data. As shown in Fig. 6, the user can request the same functional type of system data to be analyzed in several ways.

4.3.1 Page data

The information describing a page is contained in a disk filing structure as shown in Fig. 7. The specification for a page has three basic constituents: background information, a set of processing instructions called a transaction, and initialization information. In addition, the filing structure for pages contains a catalog which provides the basic linkage information to locate the page data for a user request.

The background data contain a description of the display as it is to appear on the face of the CRT. The ASCII image of the display background is stored along with window geometry tables. This background information will be displayed on the CRT screen as protected characters while the window areas will be in the unprotected mode. (Protected characters on the CRT screen cannot be overwritten by keyboard operations at the terminal). Two window geometry tables are provided. The first table provides the data for positioning the cursor of the Model 40 teletypewriter at window positions when displaying system data in window areas. The second table describes the sequence of window elements which will be received on a transmission from the Model 40 teletypewriter.

A transaction describes the processing actions associated with a CRT display. The description for the actions is contained in encoded statements of list-processing instructions which will be interpreted by the display system software for execution. The encoded form of a statement has each element represented by a byte called a token. The set of instructions in a transaction compose a program for retrieving system data, processing this data, and controlling the sequencing of actions to provide interactive responses. The instruction set available for encoding into the transaction includes arithmetic, list manipulation, conditional branching, and data retrieval operators. For example, the processing statement, L1 X[L2] $\rightarrow$ L3, will cause an element-by-element multiplication of the contents of List 1 (L1) and List 2 (L2) and the results stored List 3 (L3).

At the start of execution for each transaction, the dynamic storage area for list information is initialized.

4.3.2 Display processing

In order to provide an understanding of the processing associated with a CRT display, a typical sequence for display processing is presented.

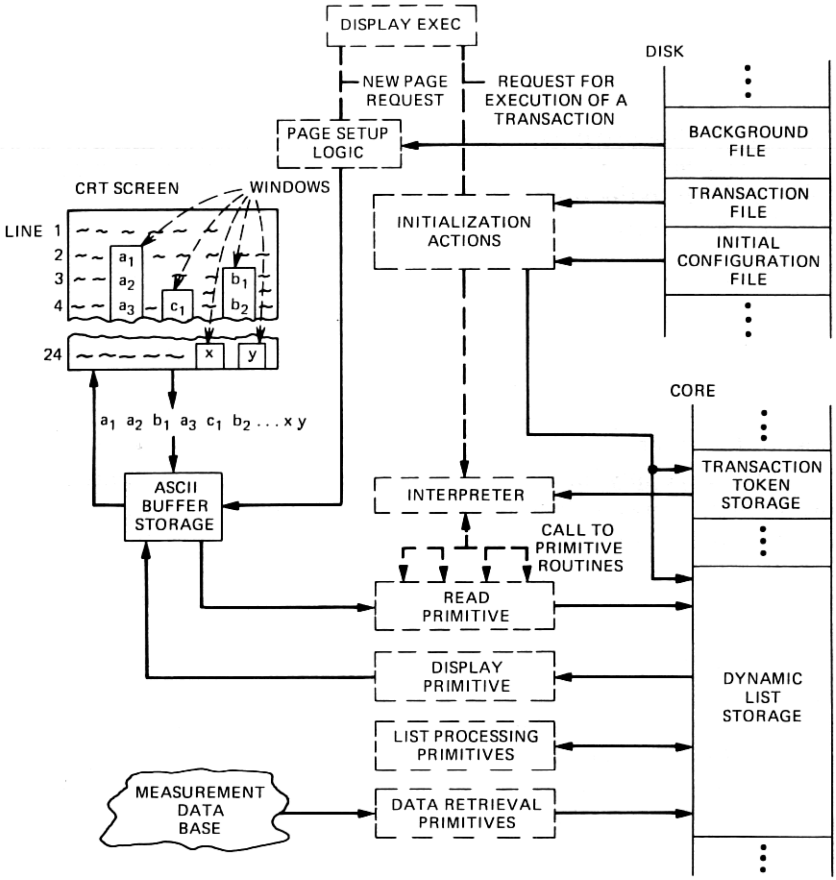

Fig. 8—CRT page processing.

The starting point for this sequence is a transmission from the CRT device. As shown in Fig. 8, the information in the unprotected window areas is transmitted as a serial string of ASCII characters resulting from a left-to-right, top-to-bottom scan of the unprotected areas by the device. These transmitted characters are placed into a dedicated buffer by the I/O system software and the executive function of the display system software receives an entry at the end of the transmission.

On receiving the stimulus from the I/O system software, the display system software checks if another page has been requested by consulting the last two strings of ASCII characters (x and y in Fig. 8). If a new page has been requested, the CRT background for the new page is fetched from filing structure on disk and placed into the ASCII buffer for transmission to the CRT by I/O system software.

After transmission of the background information to the CRT has been completed, the execution of the transaction is instituted through the following actions:

(*i*) Retrieve the initialization data from disk and establish the starting configuration for list storage.
(*ii*) Retrieve the transaction information from disk and set up the token pointer to the start of the transaction data.
(*iii*) Pass control to the interpreter function.

The interpreter starts a scan of the tokens in a transaction. The tokens are identified as either operators or arguments. The arguments are placed on a pushdown stack in preparation for encountering an operator. (Postfix notation is used in describing a processing statement and, therefore, all the arguments for an operation are encountered before the operator.) When an operator is encountered, the argument stack is "popped" to supply the arguments for the operation and an appropriate primitive routine is called by the interpreter. The primitive can be one of four basic types:

(*i*) Data retrieval, which provides access to the measurement data stored in the No. 4 ESS and, as a function of their arguments, generates lists of various sizes which are dynamically allocated space in list storage.
(*ii*) Processing, which provides the capability for data analysis and control of processing actions.
(*iii*) Input, called Read, which constructs lists of window elements which have been transmitted from the CRT and reside in the ASCII buffer. This operation will also convert the ASCII information into an appropriate form for list storage. As shown in Fig. 8, the window information contained in the ASCII buffer is in a scrambled form and the unscrambling algorithm uses the window geometry tables mentioned above.
(*iv*) Output, called Display, which fetches lists from the list storage area and, after doing an appropriate transformation into ASCII representation, places the information into the ASCII buffer for transmission to the CRT. Since windows on the CRT screen are filled one at a time, the output primitive incorporates cursor-positioning characters into the ASCII buffer along with the window information.

The end of a transaction is indicated by an encoded EXIT statement in the token string. When the interpreter encounters the EXIT statement, it returns control back to the executive function which then has the ASCII buffer transmitted to CRT. The executive function now will honor a request from another channel for the execution of a transaction.

When another transmission from the CRT is received, the I/O system software informs the executive function of the display software system.

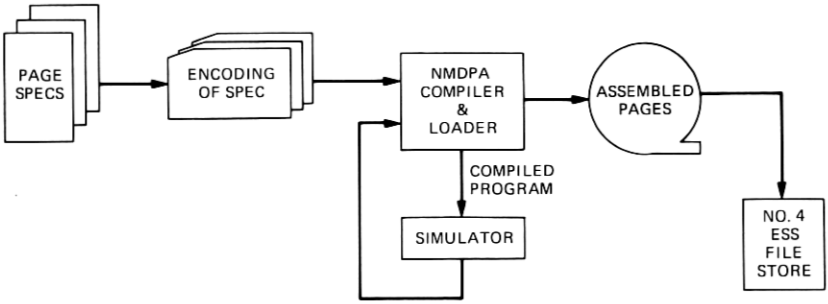

Fig. 9—Generation of CRT display pages.

If the transmitted data do not contain a request for a new page, the execution of the transaction is reestablished. In this case, the background is not transmitted to the CRT and the interpreter starts its scan of the tokens in the transaction at a point programmed to handle the interactions with the user.

4.3.3 Generation of page specifications

As shown in Fig. 9, the information for CRT displays that resides in the No. 4 ESS memory is generated by means of a network management display page assembler, NMDPA. This program, which is executed in a general-purpose computation center, operates on an input specification of the following form:

(*i*) Background text and window geometry for a CRT page is described by an 80- by 24-character array (similar to its presentation on the 80- by 24-character screen of the CRT).

(*ii*) The transaction is described by high-level language statements.

(*iii*) The initial lists for a transaction are described by a set of text strings.

The NMDPA program performs extensive format and syntax checks on this input data and provides an assembled version of the input information. A collection of assembled pages can be provided as input for a loader run which produces a tape containing the image of the filing structure discussed above. The display page specifications on this tape can be read into the No. 4 ESS memory by the standard tape input facilities of the processor.

V. MEASUREMENT DATA FOR MACHINE AND NETWORK ADMINISTRATION

The administration of a No. 4 ESS and the long-distance telephone network connecting it to other toll offices depends heavily on the measurement data made available to the operating company personnel by the network management and traffic and plant measurement systems. These systems establish a data base of over 200 k call-store and file-store words and provide a retrieval system for the timely reporting of this data.

5.1 Groups of measurement data

Conceptually, the entire data base can be divided into four broad categories: engineering data, machine performance data, network performance and control data, and division of revenue data.

5.1.1 Engineering data

Peg count, usage, and overflow measurements are provided on each of the facilities that must be engineered for the No. 4 ESS office. Every message trunk subgroup has peg counters which register incoming seizures, outgoing seizures, and overflows, and an occupancy counter to provide a usage measurement representing both service and maintenance usage.

Each of the sets of memory registers in the No. 4 ESS, such as Call Registers (CRs) and Output Message Registers (OMRs), and many of the software queues, such as the MF transmitter queue, have peg count and usage measurements generated for them. For those software facilities for which attempts can exceed available resources, overflow peg counts are provided. Facilities that can be idled by call abandonment have associated abandon peg counts.

5.1.2 Machine performance data

Measurements are available for the machine administrator and for the plant manager that detail the ineffective attempts in the office—both equipment-related, such as permanent signal time-outs, partial dial time-outs, and partial dial abandons, and traffic-related, such as no circuits and vacant codes. Measurements of the occurrences and duration of phases and interrupts and of the maintenance usage for both processor and peripheral equipment serve as indicators of the quality of service provided to customers.

5.1.3 Network performance and control data

Peg count data is accumulated for each Numbering Plan Area (NPA) code and for each office code in the home NPA and in a maximum of six foreign NPA's that are specified by the network manager. The peg counts reflect the number of calls that failed to be switched through the office, the number of calls that were successfully forwarded out of the office, and of the latter, the number of calls that failed to receive answer. As discussed previously, these data provide the information for determining codes which are hard to reach.

Counts of calls that are affected by the network management automatic and manual controls on a trunk subgroup and maintained on a trunk subgroup basis. This allows for the monitoring and evaluation of these control actions as well as providing data which can be used to adjust the normal engineering data for trunk subgroups.

5.1.4 Traffic separations data

The "from-to" relationships of traffic flowing through the No. 4 ESS can be analyzed with the traffic separations data. Each incoming trunk subgroup in the office is assigned to 1 of 32 incoming categories and each destination code is assigned to 1 of 64 outgoing categories. For each call switched through the No. 4 ESS, peg count and usage measurements are registered in a cell of the 32 by 64 matrix of traffic separations data. A primary use of this data is to develop the separations factors which are used each month to divide the interstate revenues based on the plant investment attributable to interstate usage.

5.2 Organization of the data base

Both 5-minute and 15-minute data are made available in the network management and traffic and plant measurements data bases.

The 5-minute data base in generated primarily for the use of the network manager. Every 5 minutes the data are collected and stored on file store in one of four 5-minute blocks so that the current quarter hour's data is always available for reference.

The 15-minute measurements encompass over 8000 counters required in all offices plus several thousand counters which are dependent on the office size. Associated with each measurement is a minimum of two and in several cases three distinct areas of memory: a required call store counter, an accumulating register in call store or file store, and four required file store quarter-hour holding registers. Each call store counter contains either a peg count, which is incremented each time an event occurs, or an occupancy count, which is incremented each time a facility is seized and decremented when it is released. Accumulating registers contain cumulative totals of occupancy counts over the 15-minute interval.

5.3 Data collection and storage

The starting point for the collection of data is the registering of event occurrences in call store counters by programs throughout the No. 4 ESS software system, as shown in Fig. 10. Occupancy counters are maintained

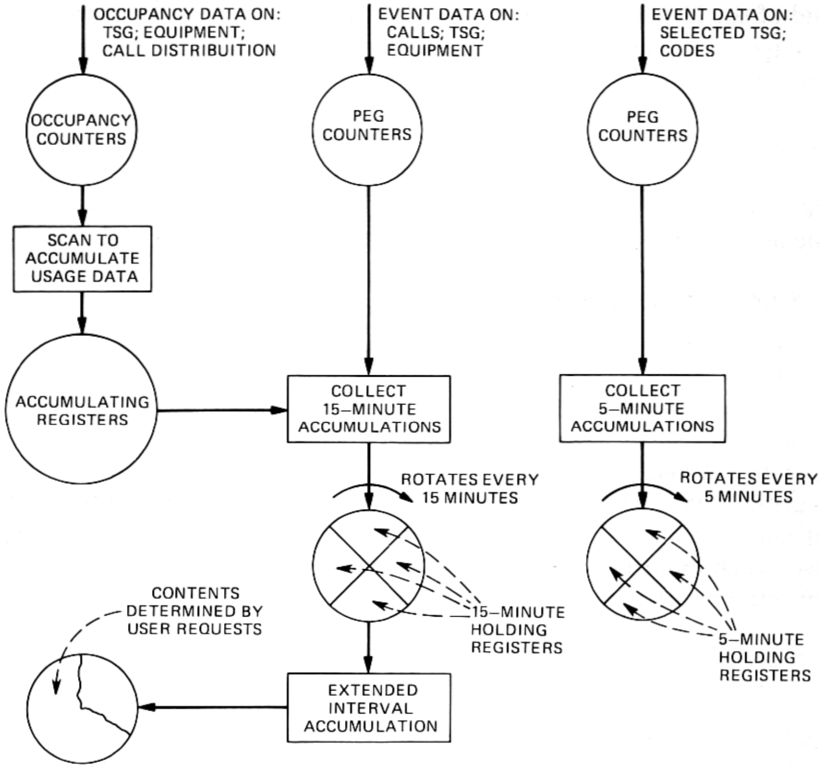

Fig. 10—Measurement data base.

to reflect the current number of facilities busy at any point in time. To obtain usage measurements, the occupancy counts are scanned at an interval that approximates the holding time of the facilities. Thus the occupancy counts for memory registers such as CRs or OMRs are scanned every 10 seconds. The value of the count at the time of the scan is added to an accumulating register which accumulates over a 15-minute interval. Counters for trunk subgroups are scanned every 180 seconds and similarly are stored in accumulating registers for 15 minutes.

Every 15 minutes the data in the counter blocks and the accumulating register blocks is collected and written into the holding registers on file store containing the oldest quarter-hour data. The counters and accumulating registers are recycled to zero to begin accumulating the data for the next quarter hour. The entire collection requires less than 1 minute to complete. In a similar fashion, 5-minute data, which only include peg counts, is transferred from call store to file store.

The primary source of measurement data is the last hour's peg counts and usage measurements stored in the four 15-minute holding register blocks on file store. Measurements from this data base can be accumu-

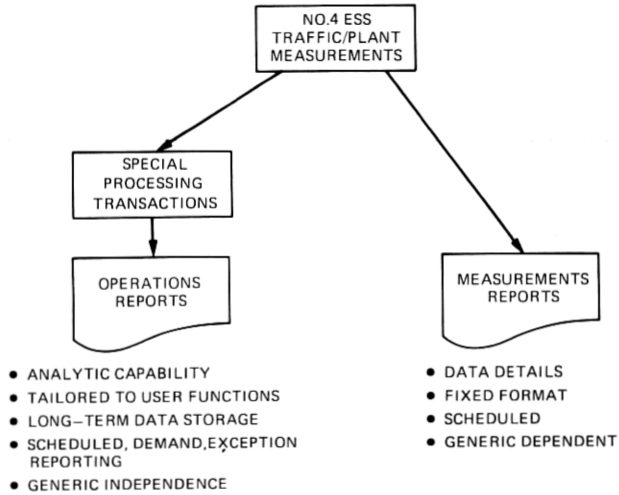

Fig. 11—Types of reports.

lated for extended intervals. The measurements to be summed into extended interval registers are specified by office personnel with teletypewriter input messages.

VI. ADMINISTRATION AND ENGINEERING REPORTS

The No. 4 ESS provides on-site reports, containing processed data presented in various formats, to meet the needs of machine administration and network engineering. The quality and reliability of the traffic data in these reports benefits from utilizing the same overall system, with its equipment reliability and validity check features, that is used in the switching of calls. As shown in Fig. 11, two types of reports are generated: operations reports and measurement reports. Operations reports are formatted to meet specific user requirements concerning the performance of central office equipment. Measurement reports contain selected subsets of essentially "raw" data needed for monitoring trunk subgroups and performing other special studies.

6.1 Operations reports

The traffic engineering and administration reports generated by the No. 4 ESS list critical machine items (such as MF transmitters), related measurements, and other data in formats specifically tailored to user functions. Provisions are made for the printing of these reports on demand, on a scheduled basis, or automatically when manually preset thresholds are exceeded. When reports are demanded, they are immediately printed and contain data for specified past measurement intervals, while scheduled reports are printed immediately after measurement intervals for which they were requested. Listings can be obtained on demand for all scheduling details.

Format flexibility is provided for all reports to allow the addition and deletion of measurements and the rearrangement of data on the report. This flexibility is independent of the No. 4 ESS generic and, therefore, allows changes to be made more frequently than generic updates.

6.1.1 Machine Service Report (MSR)

The MSR is a comprehensive service report aimed at reflecting the quality of service being provided by the No. 4 ESS machine by focusing on the numbers and kinds of ineffective attempts and other service irregularities. The report provides a measure of the service level and a means of comparing that service level with other switching machines providing the same function. The report is partitioned to provide a brief, comprehensive presentation of the most sensitive service-level indicators in addition to providing details on the components of these indicators. Figure 12 shows the overview report which can be supplemented by reports on components, such as VACANT CODE. Hourly and daily MSR reports serve as the source of information for administering the office to meet service objectives. A monthly MSR serves to document whether or not an acceptable level of service is being achieved.

6.1.2 Machine Load and Service Summary (MLSS)

The MLSS is used by the machine administrator and network engineer to monitor the traffic-load performance of the machine and interrelationships of traffic-sensitive equipment such as MF receivers and MF transmitters. Data, such as per-call holding times and maintenance usage on equipment, is provided to aid in identifying trends which could indicate potential problems. With this information, timely action can be taken to avoid service degradation. The MLSS serves as a common meeting ground for discussions between administration and engineering personnel.

The MLSS contains an hour's worth of data, beginning on the hour of half hour, and contains the data upon which the Load Distribution Report and Load Service Report (discussed below) are based. An MLSS can be scheduled for periodic printing and can also be printed on demand for any data-hour out of the preceding 24-hour period.

6.1.3 Load Distributing Report (LDR)

The LDR is provided for use by the machine administrator and network engineer to determine the busy hour and distribution of load throughout a day for a category of traffic-sensitive equipment. Presently,

```
33REPT:
                          MACHINE SERVICE REPORT

                                   MSR

                                 PART 1


REPORTING OFFICE:  KSCY MO 09 04T DATE:  05/31/76  TIME:  0030  PAGE 1
OF 3 REPORT TYPE:  HOURLY                              SCHEDULED REPORT
REPORT PERIOD:  MAY 1976


                                                       COUNT *   PERCENT
CONDENSED INEFFECTIVE MACHINE ATTEMPT (IMA) REPORT
                                                       IS-FS

    TOTAL NO CIRCUIT (NC)                              166        0.80

         NC-INTERTOLL (NC-IT)                          166        0.80
         NC-TOLL COMPIETING (NC-TC)                      0        0.00
         NC-OTHER                                        0        0.00
         NC-NETWORK MANAGEMENT (NM) BLOCKED
         CALLS                                           0        0.00
              NM BLOCKED AFTER AUTOMATIC OUT OF
              CHAIN (AOOC) ROUTING                       0        0.00

    TOTAL VACANT CODE (VC)                               0        0.00

    TOTAL REORDER (RO)                                  23        0.11
         EQUIPMENT CAUSED REORDER (REO)                 13        0.06
         NON-EQUIPMENT CAUSED REORDER (NRO)             10        0.04

    TOTAL IMA                                          189        0.92

    TOTAL ADJUSTED IMA                                 189        0.92

ANNOUNCEMENT GROUP OVERFLOW REPORT
                                                                 IS-FS

    GROUP A OVERFLOW                                     0        0.00

    GROUP B OVERFLOW                                     0        0.00

ABANDON CALL REPORT
                                                                 IS-FS

    TOTAL ABANDONS                                       2        0.00

05/31/76  00:33:07 #450
```

Fig. 12—Operations report output.

eight such categories have been identified. In general, the LDR contains data for 28 hourly intervals, overlapping on the hour and half hour, covering a span of 14½ hours. The busiest hour is flagged on the basis of a key measurement such as usage. An exception to the above is the LDR reflecting the real-time utilization of the processor which is based on 15-minute time intervals of data with no overlap and within a span of 14 hours. For this item, the busiest 15-minute interval is flagged.

Daily LDRs can be scheduled for selected traffic-sensitive equipment categories and are primarily used for studies of traffic on special days such as Mother's Day and Christmas. Five-day average (weekly) LDRs can also be scheduled for selected traffic-sensitive equipment categories. These reports employ the daily LDR format described above and assume a consistent span of contiguous hours throughout the week. Weekly LDRs are primarily used to determine busy hours for scheduling MLSSs and Load Service Reports.

6.1.4 Load Service Report (LSR)

The LSR is a specialized engineering report provided for traffic-sensitive central-office equipment. An LSR contains a yearly summary of busy-hour traffic data for a traffic-sensitive equipment category (similar to the LDR) in a form usable by the network engineer. For a defined busy hour, the LSR contains a descending listing of the 15 highest days within the past year, the average of the ten highest days, averages of each of the 3 highest months (busy season), an average of the 3 highest months (average busy season), and averages of each of the remaining 9 months.

Capability is provided for scheduling the printing of two sets of LSRs as a function of selected days of data. One set of LSRs summarizes only weekdays (Monday through Friday) and the other set considers only Sundays. For a set of LSR, capability is provided for the selective assignment of each traffic-sensitive equipment category to two different busy hours, e.g., a morning and afternoon busy hour. Provisions exist for the LSR to be based on data accumulated for the previous year with the capability of zeroing the historical data base when required. Capability also exists for manually changing the designated busy hours to accommodate any significant shifts in traffic concentrations.

Load Service Reports are also generated and printed based on peak intervals of traffic as opposed to a predetermined busy-hour basis, for the output message registers (OMRs) and the processor's real-time capacity. For each of these components, the peak interval of daily measurements is determined automatically and the associated data is processed.

6.1.5 Exception reports

Exception reporting is planned to alert the machine administrator in real time to individual abnormal conditions affecting service. Exception reporting occurs when manually preset thresholds are exceeded, and provides automatic monitoring of the performance of the switching machine during periods not covered by scheduled reports.

Machine status exception reports result from any of three conditions that can cause traffic data to vary significantly from what is normally expected:

(*i*) Equipment malfunctioning
(*ii*) Inordinate variations in offered load from the network
(*iii*) Data inaccuracies

Machine-status exception reports are based on measurement intervals used in the generation of central-office equipment reports. Exception reporting on the processor's real-time utilization is based on 15-minute measurement intervals, since this is the interval used for engineering purposes. Exception reporting on all other exception measurements is based on measurement intervals of 1 hour. Besides containing exception notification of "key" measurements used to detect abnormal conditions, machine-status exception reports contain "slave" measurements which comprise a predetermined grouping of related measurements to provide additional understanding regarding possible causes of exceptions.

Machine-status exceptions result in automatic flagging of exception measurements, the underlying data, and all other measurements based on the underlying data. For example, an exception report on holding time has directly related data of usage and peg count and all three measurements are flagged. Flagging of related measurements on the engineering reports is necessary since the network engineer does not monitor exception reports and should be aware of questionable data.

6.1.6 Reliability

Provisions are available for the machine administrator to manually delete invalid data from the LSR historical data base. Questionable data cannot be automatically suppressed since human judgment is needed to verify that the data is unreasonable. A grace period is provided during which the machine administrator can delete invalid data before it becomes permanently imbedded in the data base.

Software and hardware are provided for purposes of data validity and reliability. Data base redundancy is provided for both the LSR and previous 24-hour historical data bases. Periodic audit checks of the data are made to ensure high reliability and data integrity. A tape backup system is provided for the LSR data base to give the machine administrator the capability of initiating an automatic overwrite of bad data. Should the tape backup fail, periodic hard copy outputs provide the network engineer sufficient information to manually construct data for engineering purposes. Data base protection mechanisms are provided to prevent the destruction of historical data by system users. The use of software "keys" and "policing of communications" are employed along with the use of a manual key to prevent storage wipe-out.

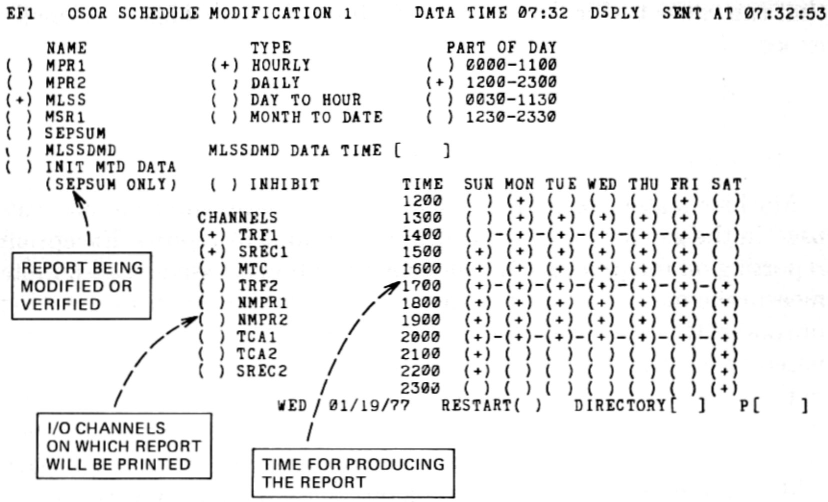

Fig. 13—Scheduling for operations report.

6.1.7 Report handling

The input of all related report messages regarding scheduling, demanding, editing, and setting of thresholds are performed by the machine administrator at the machine administration center (MAC). A report can be broadcast to as many as five teletypewriter channels. Should the printing device be inoperative, the reports are directed to backup printers to ensure data collection. Reports are printed in 8½ by 11 inch page increments on sprocket-fed fanfold paper to facilitate handling and storage.

6.1.8 Scheduling and generation

Operations reports can be scheduled for output to the various work centers in the No. 4 ESS. The scheduling and generation for these reports uses the capabilities of the network management display system discussed in Section 4.3.

Office personnel establish schedules for the printing of operations reports by interacting on CRT pages as illustrated in Fig. 13. After requesting this page on a display terminal, the user modifies the scheduling information through an interactive sequence. After the user specifies the report, the system responds with the scheduling information currently active. In Fig. 13, the system has responded with the MLSS schedule. The user can modify the scheduling for the MLSS by changing the channel and/or time indications and then activating this information.

Operations reports are generated with the processing capabilities

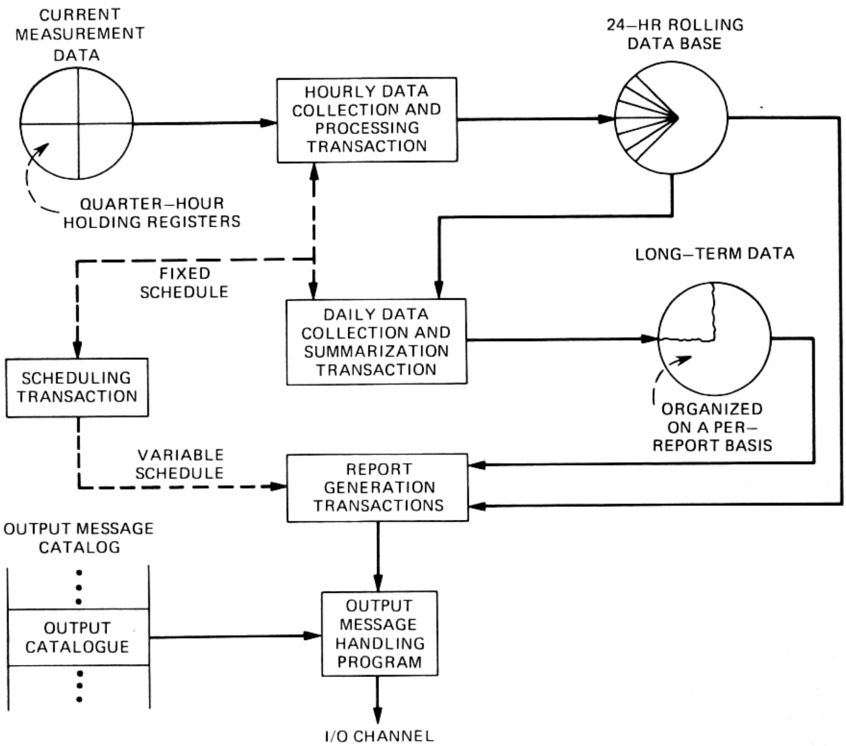

Fig. 14—Data flow for operations reports.

described in Section 4.3 for CRT displays. Transactions, as described in Section 4.3.1, are encoded to:

(*i*) Establish and administer a data base tailored to the information needs of the operations reports.

(*ii*) Process information for output on a report.

(*iii*) Schedule the administration of the data base and the printing of reports.

The general flow of information for the reports is shown in Fig. 14. Transactions are routinely scheduled which generate a compacted data base to service the information needs of the operations reports. Hourly and daily reports on the performance of the No. 4 ESS access the processed data which is maintained in a rolling 24-hour data file. Data from this file is summarized to establish historical data files for monthly and yearly performance reports.

In producing hard-copy reports, the output lists of processed data, which are formed by a transaction, are merged with the background form in the processor before the ASCII characters are transmitted through an I/O channel to a teletypewriter.

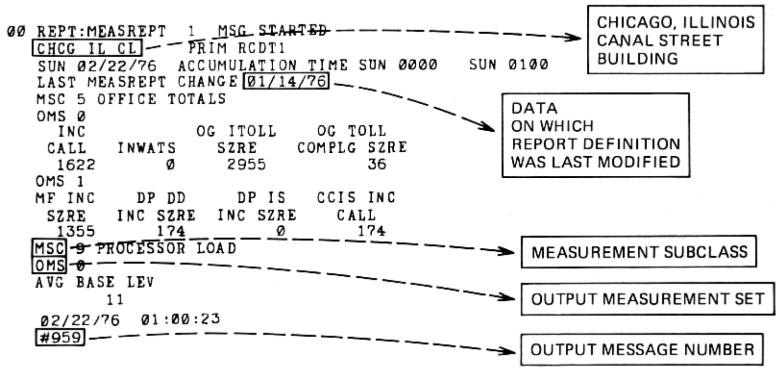

Fig. 15—Measurement report output.

6.2 Measurement reports

In addition to operations reports, the No. 4 ESS generates measurement reports that provide the machine administrator and other users with subsets of essentially "raw" data needed on a scheduled basis to monitor trunk subgroups, perform special studies, and monitor malfunctioning processor and peripheral equipment.

Twenty-four measurement reports can be produced by this system to meet user needs. The user defines a measurement report by specifying with input messages the report number and the primary input channel over which all subsequent changes to the report must be made. The accumulation interval of the report, which can range from 15 minutes to a full week, is input as well as one or more output times. The user can direct the output of each report to a maximum of five teletypewriter output channels as well as specify one output report to be written onto tape. The report output contains the specific sets of measurements requested on input messages. Measurement report 1 in Fig. 15 was defined using the RCDT1 channel as the primary channel (PRIM). It is a report produced by the Chicago 7 office in the Canal Street building in Chicago, Illinois on February 22, 1976, at 1 A.M. This is a 1-hour report whose definition was last changed on January 14, 1976. It contains several sets of related measurements which are identified in terms of measurement subclasses (MSC) and output measurement sets (OMS).

The input messages used to define a measurement report are processed by the measurement reporting system (Fig. 16). The definitions are stored on file store and remain in effect until changed by subsequent input messages. These definitions are backed up on magnetic tape. The definition of a measurement report reflects the need for additional memory if the report is required to output data accumulated over an interval longer than 1 hour. A sufficient number of extended interval

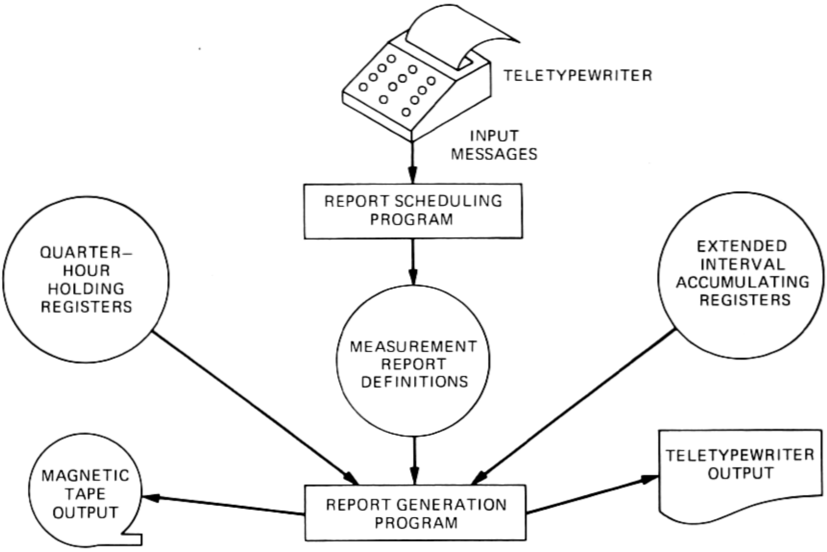

Fig. 16—Measurement reporting.

accumulating registers are allocated to the report to accumulate the measurements. If the report has an accumulation interval of an hour or less, it can be output entirely from the quarter-hour holding registers which are available to all reports.

Each quarter hour after the data is collected and stored on file store in one of four sets of quarter-hour holding registers, the extended interval accumulation program checks the measurement report definition on file store to determine if any reports are currently accumulating data for a period longer than 1 hour and adds an additional hour's data into the extended-interval accumulating registers for those reports. All accumulations are stored back on file store and then output processing begins.

The report generation program first checks the measurement report definitions to determine if any reports are scheduled for output. The reports are output in numerical order with one report being printed on all its assigned output channels before processing of the next report begins. If a report has extended-interval accumulating registers, the data are output entirely from the registers. If not, the data are gathered from one or more quarter-hour accumulating registers, added together if necessary, and output.

VII. CONCLUSION

The service afforded by a No. 4 ESS relies upon effective engineering and administration in handling predictable traffic loads as well as unusual traffic conditions. In this paper, the features designed to provide these capabilities have been discussed.

For the unusual traffic conditions, automatic actions to control traffic overloads have been made more effective by controlling traffic parcels which have been determined to have low completion probabilities. To complement the automatic control system through manual control actions, a real-time surveillance system has been organized to provide problem detection by reporting exceptions and to allow problem investigation using interactive CRT displays of real-time performance data.

For engineering and administrative personnel the No. 4 ESS is generating reports on its switching performance. A comprehensive set of measurement data is gathered, screened, and processed to provide information in a directly usable form.

VIII. ACKNOWLEDGMENTS

The authors wish to acknowledge the efforts numerous colleagues in the system planning and development of the traffic administration and network management features. Their contributions were equal to those of the authors. In particular we would like to acknowledge the contributions of M. J. DiCarlo-Cottone in the development of the network management features and R. I. Rosenthal for the system planning associated with the network management surveillance arrangements.

REFERENCES

1. Y. Nakagome and H. Mori, "Flexible Routing in the Global Communication Network," Seventh International Teletraffic Cong., Stockholm, June 1973, Paper 426.
2. J. H. Weber, "Some Traffic Characteristics of Communications Networks and Automatic Alternate Routing," B.S.T.J., *41,* No. 2 (March 1962), pp. 769–796.
3. P. J. Burke, "Automatic Overload Controls in a Circuit-Switched Communications Network," Proc. National Electronics Conf., December 1968, pp. 667–672.
4. W. B. Macurdy, "Network Management in the United States—Problems and Progress," Seventh International Teletraffic Cong., Stockholm, June 1973, Paper 621.
5. E. E. Lewis and E. E. Schwenyfeger, "Automatically Controlling Heavy Telephone Traffic," *Bell Lab. Rec., 44,* No. 4 (April 1966), pp. 135–138.
6. L. A. Gimpelson, "Network Management: Design and Control of Communications Networks," *Elec. Commun.,* IT&T Corp., *49,* No. 1 (1974), pp. 4–22.
7. C. A. Dahlbom, "Common Channel Signaling—a New Flexible Interoffice Signaling Technique," International Switching Symp. Record, 1972.
8. R. L. Franks and R. W. Rishel, "Optimum Network Call Carrying Capacity," B.S.T.J., *52,* No. 7 (September 1973), pp. 1195–1214.