ENGINEERING AND ADMINISTRATIVE DATA ACQUISITION SYSTEM (EADAS) DESCRIPTION

CONTENTS PAGE

1. GENERAL

1.01 Engineering and Administration Data Acquisition System (EADAS) provides an electronic, software controlled means of collecting traffic data. EADAS, which is a basic part of the Total Network Data System (TNDS), provides for the acquisition of all forms of basic traffic data. (See Fig 1.) Network Management (EADAS/NM) and the Individual Circuit Usage Recorder (EADAS/ICUR) are part of the TNDS. EADAS/NM is covered in Sections 190-540-060, 190-540-100, and 190-540-200. EADAS/ICUR is covered in this section and in Section 190-510-062.

1.02 This section is reissued to include information pertaining to network design and network administration and to delete portions of the previous issue pertaining exclusively to system maintenance. This section is a general description. There are separate documents supporting network maintenance, network design and network administration.

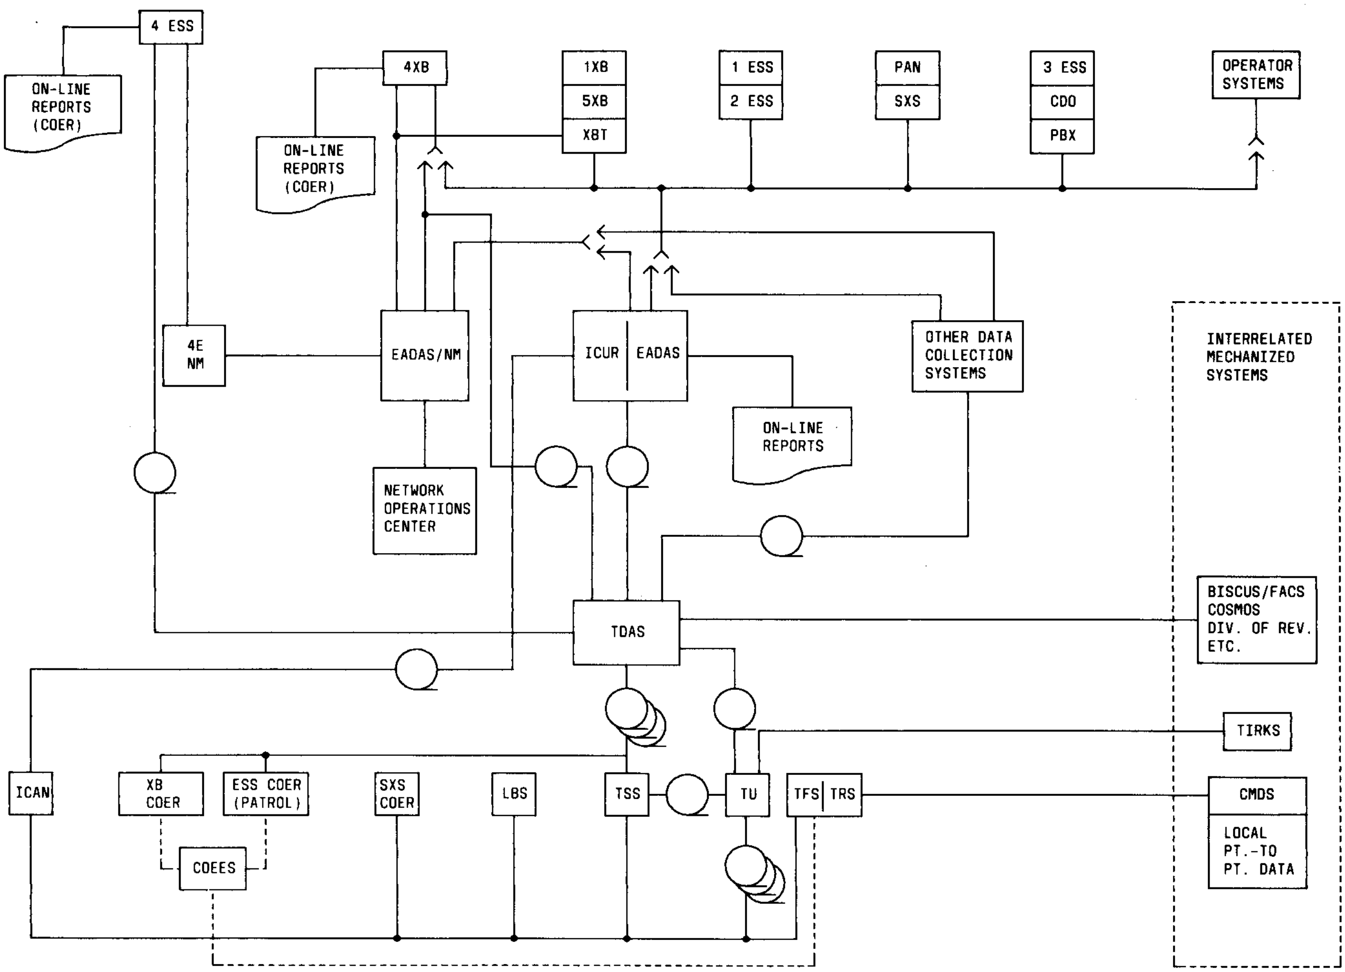

Fig. 1—Total Network Data System

1.03 Fig. 1 depicts the EADAS central control unit (CCU) receiving data from EADAS data collection devices in various types of central offices. The system is composed of traffic measuring and indicating devices, data converters, data accumulators, an EADAS CCU, and a general purpose computer. Data converters are devices that transmit near real-time data to the CCU on a continuous basis. Data accumulators are devices that collect and store data and then transmit it to the CCU on either a scheduled or polled basis. The downstream general purpose computer provides data to the data management system which in turn provides the raw data, properly formatted and for the measurement intervals requested, to other downstream programs.

1.04 Fig. 2 illustrates the overall EADAS portion of the TNDS. The data converters or data accumulators are located in the central offices served by the system. The EADAS CCU is located at a site which is generally centralized in relationship to the offices it serves. Converters and accumulators in the outlying central offices are connected to the CCU via dedicated data links or dialed-up data links. Magnetic tape, in a sense, connects the CCU to the large scale general purpose computer where downstream batch processing is performed.

1.05 Input data (ground closures) points come from central office traffic measuring devices (traffic usage recorders and dial tone speed machines) and traffic indicating leads (peg count, overflow, etc). The EADAS traffic data converter (ETDC) sends data to EADAS; however, existing Traffic Data Recorder System 1A (TDRS-1A) converters, the pollable data terminals (PDT-1A), Electronic Switching System (ESS), or other commercial terminals which meet the defined interface can also send data to EADAS. The ETDC immediately generates distinctive binary identities (addresses) for each traffic indication and transmits them to the EADAS CCU via dedicated data links. The data link required is a 4-wire dedicated data link suitable for transmitting data at the rate of 1200 baud.

1.06 The EADAS CCU collects and accumulates the incoming traffic data as it is transmitted via data links from all associated central offices. Data, as received, is temporarily stored in core memory. Every few seconds, data is summarized by adding the new data in core memory to existing totals already stored on disk. After each data collection interval, surveillance checking is performed. Calculations on certain user defined groups are compared with predetermined levels of traffic (thresholds) which have been stored in memory to indicate abnormal threshold conditions. The reports are sent out via data links to network terminals at an administrator's or maintenance location. Every collection interval (30 minutes), the accumulated register totals in disk memory are placed in the proper format for subsequent recording on 9-track magnetic tape. Register totals are written on tape for the time periods specified by the user's schedule. This data is subsequently provided to other downstream programs for further processing.

1.07 ***General Purpose Computer:*** A general purpose downstream computer is used on a batch basis. Magnetic tape produced by EADAS is periodically (daily or more often if required by the quantity of data recorded) removed and sent to the general purpose computer site for subsequent processing. At locally required intervals, the accumulated EADAS tapes are processed in a general purpose computer using the individual circuit analysis (ICAN) program for EADAS with ICUR and the Traffic Data Administration System (TDAS) program. Further analysis and reporting is then accomplished via the specialized downstream software systems such as Trunk Servicing System (TSS), Trunk Forecasting System (TFS), Central Office Equipment Reports (COER), Program for Administrative Reports On Line (PATROL), and Load Balance System (LBS). The individual circuit analysis (ICAN) program provides support and analysis reports for EADAS/ICUR Systems. Individual circuit usage tapes are generated by EADAS/ICUR Systems. These tapes contain copies of the ICUR data base and the usage for individual circuits in the electromechanical offices served by EADAS/ICUR. With this information ICAN:

- Reports the collection schedules for ICUR
- Validates the ICUR data base and reports the anomalies between ICUR and the TDAS data base.
- Analyzes the individual circuit usage for central office equipment and reports apparent irregularities in the use of the equipment.

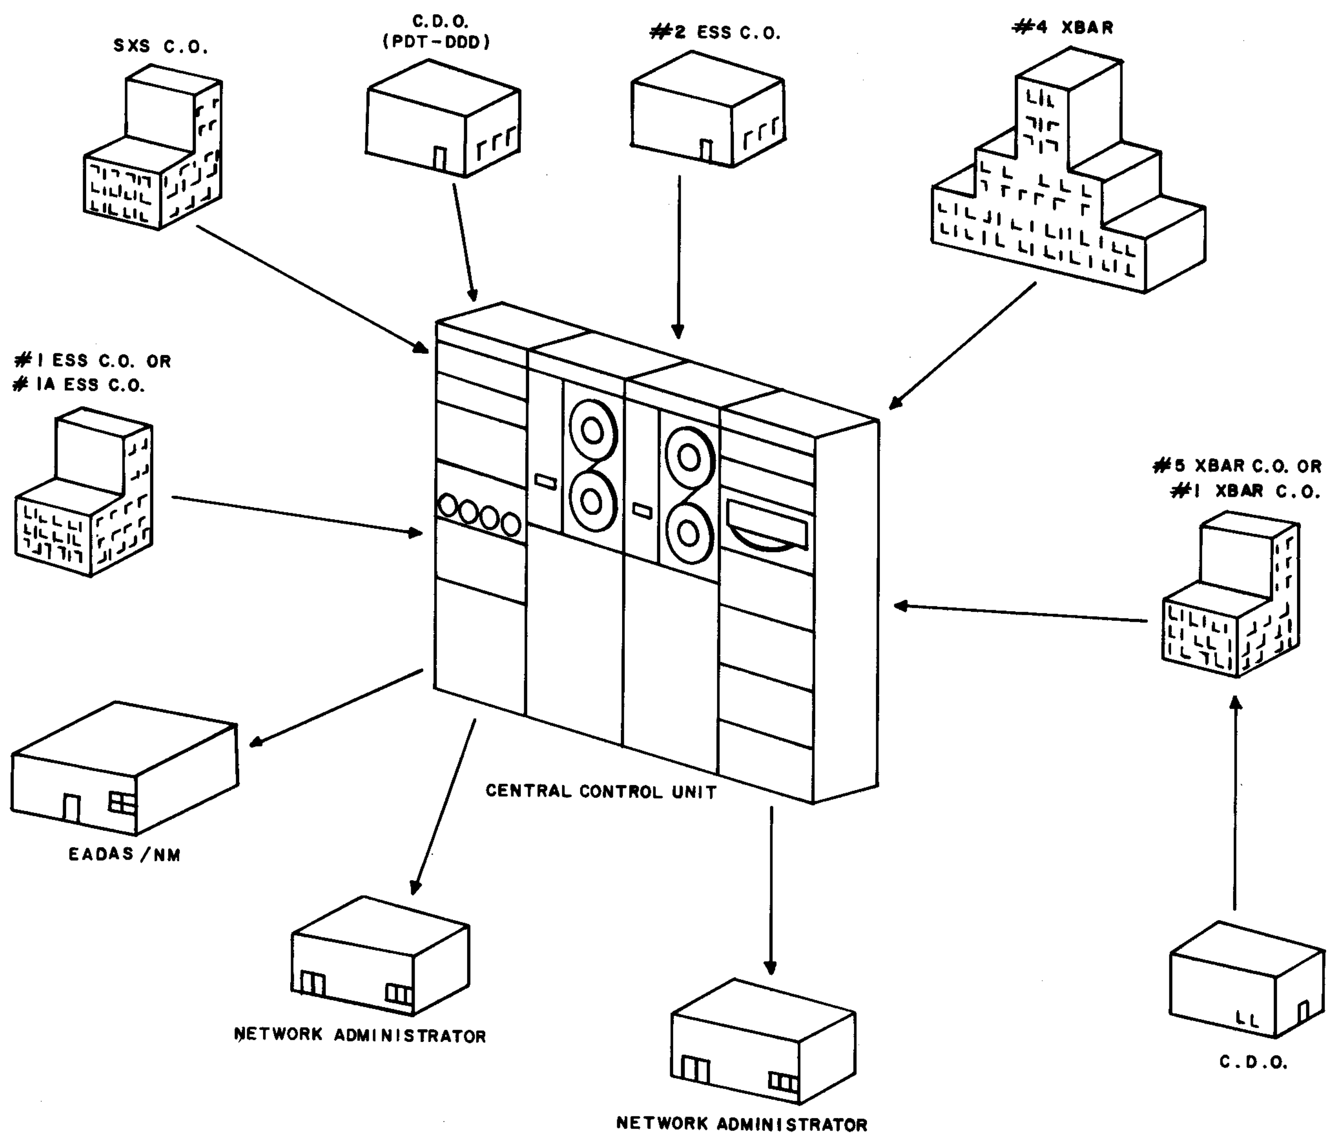

Fig. 2—Overview of EADAS

The TDAS program merges the data from various data acquisition systems and performs the following functions:

(a) The establishhment of schedules or data collection

(b) Maintenance of assignment records for all data collection devices

(c) The acceptance of measurement data for any time interval

(d) The reporting of measurement data to downstream processes via a set of standard interfaces, or optionally, in printed form for manual analysis

(e) Performance and quality control reports specifically designed to permit effective management of the data collection effort

(f) Adjustment and validation of measurement data.

TDAS provides a single system for data collection support services required to make measurement data usable by other downstream programs. The following is a description of these programs.

1.08 ***Central Office Equipment Reports (COER):*** This large scale computer software package accepts central office engineering data formatted by TDAS or its equivalent, subjects this data to a series of validation tests, and produces final summarized reports designed to meet both administrative and engineering requirements. The programs which comprise the No. 5 Crossbar Engineering and Administration Data System software package are designed to operate in a standard data mode.

1.09 ***Program for Administrative Reports on Line (PATROL):*** The PATROL system for ESS is a time-shared computer program, designed to provide the network administrator with data concerning service and traffic levels in an office. ESS block "H" data is sorted by TDAS and transmitted to the PATROL system located in New Jersey. Transmission is accomplished using a high-speed data operation, called T-TRAN. Block "H" PATROL data is stored in history files arranged to allow administrators to obtain daily, weekly, and monthly reports. Using the daily data, PATROL also furnishes machine load and service summaries for engineering.

1.10 ***Load Balance System (LBS):*** The LBS is a software package which uses load balance data, properly formatted by TDAS, to produce reports necessary to correctly load and to balance switching machines. The LBS provides index summaries, data summaries reflecting current loading group levels, and weekly historical trends. It also provides assignment lists indicating load and balance in the form of + or -CCS per loading group and customer line usage summaries.

1.11 ***Trunk Servicing System (TSS):*** The TSS is designed to provide the message trunk information necessary to assure efficient, economical use of an operating company trunk network. The TSS receives trunk measurement data properly formatted by TDAS, validates the data for the proper quantities and quality, converts the measurements to offered and overflow CCS, calculates the required trunks for direct and network trunk groups, provides exception reports to assist the Trunk Administrator in the detection of trunking problems, and prepares the traffic trunk base for use by the Trunk Forecasting System.

1.12 ***Trunk Forecasting System (TFS):*** The TFS is designed to assist the trunk forecaster in the determination of future message trunk requirements. The TFS uses the traffic trunk base established by the TSS and provides forecasts for multiple future years and, if desired, subdivision within those years.

1.13 In summary, a complete and responsive traffic data system is provided in the best way through the full array of modules (EADAS, TDAS, and downstream). EADAS plays a vital role in the total network data system concept; it is not designed, however, to fulfill the total needs for complete analysis and reporting of traffic data. This point is mentioned in order to help each company plan for the modernization of its entire traffic data gathering and processing system. Planning for EADAS should be considered as an integral part of this mechanization plan, but modernization of the total data system should be considered the overall objective.

2. EADAS SYSTEM DESCRIPTION

GENERAL

2.01 EADAS is a near real-time data collection and surveillance system which provides an electronic, sofware controlled means of collecting traffic data. The EADAS CCU is a DEC* PDP*-11/40 processor. The processor conrols the operation of the EADAS equipment including (1) EADAS CCU equipment, (2) remotely located data collection devices, and (3) remotely located network administrator and maintenance terminals.

*Registered trademark of Digital Equipment Corporation.

2.02 EADAS offers options to the user for collection of different types of data for use suitable to a particular type of equipment and/or situation. These options are:

- EADAS with EADAS/Network Management (EADAS/NM) interface
- EADAS with Individual Circuit Usage Records (ICUR) feature

- EADAS with Network Operations Report Generating (NORGEN) feature.

2.03 Expanded calculations for near real-time data analysis and report generating by an EADAS System is obtained by the NORGEN feature. This option permits development and implementation of report generating programs in a new high level program language. A standard set of programs provides the reporting needs common to most operating telephone companies and permits the development of programs by telephone personnel.

SYSTEM FEATURES

2.04 EADAS is capable of collecting and summarizing peg count and usage data and several additional functions as follows:

- Provides key real-time traffic data validity checks and calculations to network administrators on scheduled exception basis
- Reports results of calculations to network administrators on an exception basis
- Collects data from collection devices such as ETDC, PDT, TDC, ESS offices, etc
- Permits real-time status inquiries about the entities being measured
- Creates a data base for network management use
- Provides peg count and usage concentration to satisfy the needs of small central offices
- Stores data on magnetic tape for downstream TNDS processes.

2.05 EADAS, in conjunction with the TDAS and other downstream BIS programs, provides load balance, engineering, trunk servicing, and forecasting data for the operating telephone companies, as well as immediate key real-time data. The CCU collects and summarizes peg count and usage data, records the results on magnetic tape for subsequent downstream processing, and makes certain real-time calculations and presents exceptions to the appropriate network administrator and maintenance personel.

SYSTEM HARDWARE

2.06 The EADAS System consists of five basic groups of hardware. These are data collection devices, interface circuitry, CCU, storage devices, and input/output devices. The EADAS System configuration is shown in Fig. 3.

A. Data Collection Devices

2.07 There are different types of data collection devices. The type used depends on the type of office the data is being collected in and the type of data being collected. The data collection devices are as follows:

EADAS Traffic Data Converter (ETDC)

Equipment Description

2.08 The ETDC is a highly reliable, flexible, and modularly expandable unit for data collection. The ETDC utilizes integrated circuits on interchangeable input circuit cards, control cards, and data set cards to collect, encode, and transmit peg count and usage data to the CCU.

2.09 The ETDC is a relay rack mounted, 27-inch high unit composed primarily of solid state devices. The ETDCs can be located in most any on-premise equipment bay. The ETDC is designed to control associated dial tone speed, sender attachment delay recorder, and TUR devices. TUR scan controls are provided for the 2A, 3A, 3B, and 4A TUR circuits with detector test controls provided for the 4A TUR. TURs, DTS and SADR machines will require certain modifications when connecting to the ETDC.

2.10 The ETDC consists of three interconnected units: scanner, control circuit, and buffer/output circuit. The scanner has the capacity for 32 input circuit cards of 32 inputs each, or a maximum of 1024 inputs. Of these, 989 may be used for peg count and usage data as discussed above.

2.11 The ETDC is provided as the interface for associating electromechanical central office traffic load and service measuring devices via dedicated data links. The ETDC is compatible with all types of central office equipment that operates a 14-type register by placing ground on the load or service indicating leads when a traffic event occurs. An ETDC can function as a single unit

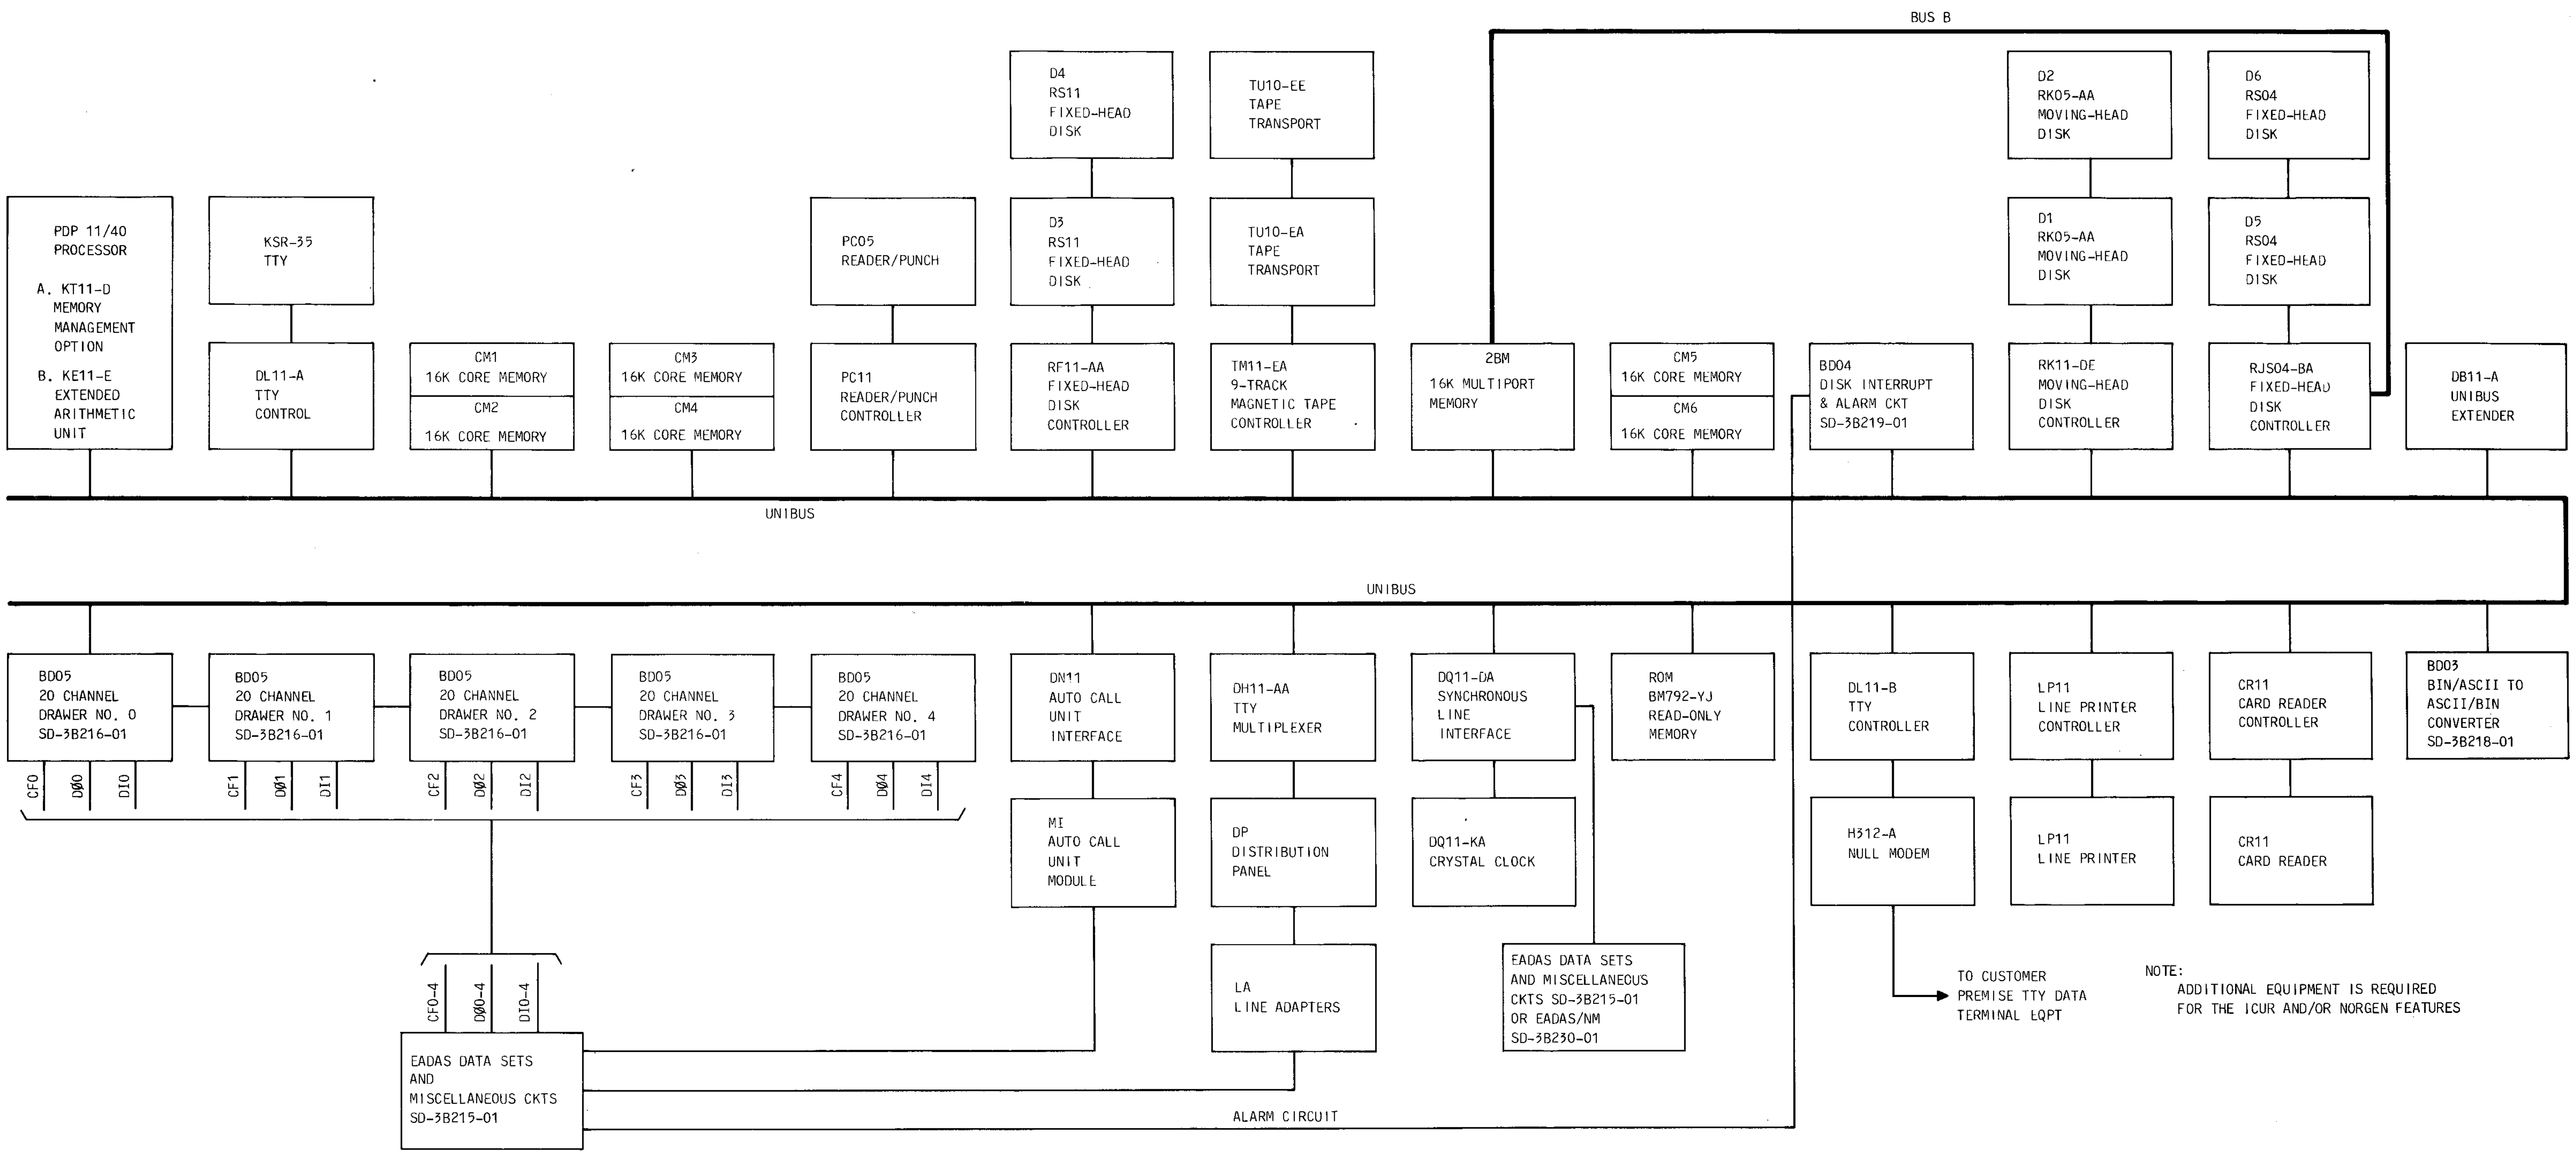

Fig. 3—Typical EADAS CCU Configuration and UNIBUS Structure

(home ETDC) serving one central office location via one dedicated data link to the EADAS CCU, or up to six ETDCs (one home ETDC and five remote ETDCs) can be concentrated to serve up to six different office locations. With the concentrated arrangement, up to six data sets are provided with each set occupying a single plug-in slot at the home ETDC. One data set is required at the home ETDC for connection to the CCU, whereas an additional set is required for each remote ETDC that is concentrated to the home ETDC. This feature allows full utilization of lengthy data links between ETDCs and the CCU when the ETDCs are not fully occupied.

2.12 Plug-in control cards contain common control circuitry. Control and display cards contain alarms, status lamps, switches, and circuitry for manually initiating maintenance routines. (These can also be remotely activated via CCU.) These enable the user to make all inputs appear busy or idle. Similarly, TUR and dial tone speed machines can be locally as well as remotely controlled. There are several types of input cards which are described as follows:

(a) ***Unscaled*** input cards are designed for use with peg count leads, output (register) leads from TUR frames, dial tone speed machine output leads, and any other type lead on which a single score is desired for all indications of 25 milliseconds and longer. It produces a data word (score) when a ground is detected only if the lead was not grounded on the previous scan. This "last look feature" assures that only one score is produced for each ground indication regardless of the ground duration. These 31 cards can be installed in any card position 0 through 30.

(b) ***Scaled*** input cards are designed for use with high volume input leads. This card, employs the "last look feature" to prevent multiple scoring of prolonged grounds. Unlike the unscaled card, this card produces only one data word for every ten valid ground pulses detected on a given input lead.

(c) ***100-second*** usage cards are designed for the measurement of usage. These cards scan their input leads (32 input leads per card) at 100-second intervals, synchronized with the CCU, thus producing CCS measurement. Each ground indication encountered during the scan produces a data word containing the lead identity and if the ground is still present on the next scan, another data word is produced.

(d) ***Concentrator*** input cards are only required in a home ETDC which is serving as a concentrator for associated remote ETDCs. One concentrator input card is required for each associated remote ETDC.

(e) ***Discrete*** input cards are similar to 100-second usage input cards except they scan their 32 input leads at approximately 10-second unsynchronized intervals, producing a data word for each input ground detected. A maximum of two discrete cards per channel can be utilized. Their sole purpose is to provide discrete data to an EADAS/NM via EADAS. Discrete data is an indication of "on" or "off" conditions in the office. For example, the Dial Tone Speed machine.

(f) ICUR: The ETDC in EADAS System requires the use of input data cards to collect, encode, and transmit peg count and usage data from TURs to the CCU. In the EADAS/ICUR System the data from the TUR is sent directly from the TUR detectors to the ICUR applique circuit on the ETDC thus eliminating the need for many of the input cards. The overall result is a savings due to fewer input cards that need be purchased. Converting an EADAS ETDC to ICUR operation is done by adding up to six new cards. Space has been provided on the ETDC for these additional cards. One card for each TUR associated with the ETDC (maximum of four), and two control cards are required. These cards may be ordered with an initial system or individually on a conversion.

Operation

2.13 Inputs to the ETDC are in the form of contact closures to ground (ability to operate a 14-type traffic register). Each input is scanned at regular intervals (approximately every 20 milliseconds) for the presence of a ground input. When ground is detected, the output is a binary number representing the address (1 out of 989) of the active input. The inputs, once scanned and processed, are encoded and immediately transmitted over a dedicated data link to the CCU.

2.14 ***Maintenance and Testing:*** Maintenance was a primary concern in the design of the ETDC. Hence, several significant design approaches and features have been incorporated into the ETDC to simplify the trouble detection and fault location.

2.15 Three maintenance commands are provided: inhibit, busy, and restore. The inhibit command causes all inputs to an ETDC to appear idle, the busy command causes all inputs to appear busy, and the restore command restores the ETDC to normal operation. By determining what data has been transmitted by the ETDC during the test cycle, troubles may be immediately detected, and with the aid of the trouble locating Bell System Practice (Section 190-510-213), fault isolation may be determined within a few circuit packs. As an additional aid, alarms and status lamps are provided to indicate various off-normal conditions.

Polliable Data Terminal (PDT-1A)

2.16 The Pollable Data Terminal (PDT-1A) is a small solid state electronic data collection device designed for use in central offices which are too small or too remotely located to justify installation of an ETDC and its associated data link. The PDT-1A will accumulate counts from up to 250 peg count and TUR register leads for either 30 or 60 minute intervals, and store the results in local memory. The terminal can then be polled by the EADAS CCU over the switched network to collect data stored in the local memory.

Equipment Description

2.17 The PDT-1A is a modular unit which contains a data set and a maximum capacity for 250 registers. The PDT-1A consists of control and counting circuits which accept peg count and usage data from TURs, totalizers, traffic register circuits, or other circuits that generate a -48 Vdc to ground contact closure. Power is obtained from -48 volt central office battery. The PDT requires approximately 13 inches of space on a standard 23-inch frame.

Operation

2.18 The PDT-1A can accumulate and store data under three different modes of operation:

- ***Collect and Transfer Mode:*** This mode will generally be used to collect data over extended periods of the day.
- ***Selected Data Mode:*** If daily measurements for one or two preselected hours are sufficient, the selected data mode can be used to divert some of the switched network usage involved in data collection to off busy hours.
- ***Peak Data Mode:*** The peak data mode, like the selected data mode, may also be used to minimize CCU polling operations. This is an optional feature of the PDT-1A terminal. The basic intent of the peak data mode is to obtain a single set of measurements for that interval of the day during which overall traffic volumes were highest.

2.19 Basic measurement intervals of either 30 or 60 minutes may be chosen by means of a manual switch setting on the unit. The interval thus selected ***cannot*** be overridden by commands from the EADAS CCU. The PDT-1A has two memories for storing traffic data received on up to 250 register leads. At any given time, one memory is designated ***active*** and the other is designated ***passive.*** The function of these two memories vary depending upon the data collection mode in effect.

Application of Traffic Data Recorder System (TDRS) Type Converters

Peg Count Converter (PCC) and Traffic Usage Recorder Converter (TURC)

2.20 Each CCU will be capable of receiving data from any TDR-type converter. If these devices are to be left in service, one dedicated data link will be required per converter. Concentration in this situation is not possible.

Outside Supplier Data Accumulators

2.21 EADAS is compatible with scanner/accumulator type data terminals manufactured by outside vendors if the interface requirements are met.

Polled ESS Interface

2.22 A polled, 1200-baud binary interface is available only for No. 1 ESS. The interface is available for No. 1 ESS with 1E(B3)4 and later generics. The interface consists of ESS hardware and software changes.

Monitored ESS Interface

2.23 The 110-baud data collection interface for No. 1, 1A, and 2 ESS, and the 1200-baud for No. 1A ESS, is arranged for EADAS to monitor the data link which is bridged on the ESS network administration teletype circuit. When the ESS prints a data block (ie, C, H, or W) according to its traffic schedule, the EADAS senses the presence of the data and records it on disk. EADAS cannot control the acquisition of the data; network management (5-minute) data or features are not available with the monitored interface.

B. Interface Circuitry

2.24 A large portion of the inputs to the CCU from the input devices are through a 20-channel drawer interface. A total of 100 channels (ie, 5 drawers) may be provided. This interface provides the connection to the central office data collection devices. These channels are wired on dedicated facilities to offices equipped with ETDCs, TDRS converters or outside vendor devices. For PDTs, a channel(s) is equipped with an automatic call unit for polling. A dedicated synchronous line interface may also be provided for connection to EADAS/NM. Interface connections (up to 16) may be provided for remote network terminals. Systems not equipped for NORGEN are limited to 16 terminals.

C. CCU For Systems Without ICUR and NORGEN

PDP-11/40 Processor

2.25 Operation of the data collection system is controlled by the PDP-11/40 computer and its associated software (programs). This computer is a 16-bit word machine which is characterized by a unique common bus system termed the UNIBUS.* This bus interconnects the processor with the core memory and all the peripherals, producing a high degree of flexibility and efficiency by permitting the processor to employ the same set of signals to communicate with both core memory and peripheral equipment. This allows peripheral devices such as the input channel interface registers to be addressable in a manner identical to that of core memory.

*Registered trademark of Digital Equipment Corporation.

2.26 The CCU provides approximately 64K words of core memory in its minimum configuration which accommodates up to 72 input channels. 80K words of core are required for a full system of 100 input channels. If the ultimate system configuration will encompass more than 72 inputs, it is recommended that sufficient core memory be ordered initially to handle the entire projected system size.

2.27 The PDP-11/40 processor performs all arithmetic and logical operations required in the system. It also acts as the arbitration unit for UNIBUS control by regulating bus requests and transferring control of the bus to the requesting device with the highest priority.

UNIBUS Structure

2.28 The UNIBUS is a common, high-speed data path that interconnects the processor and all devices within the CCU. It is a 120-conductor ribbon cable (white in color) that weaves through the devices comprising the CCU. The UNIBUS uses 56 lines for information and the remaining lines grounded to provide noise immunity. Fifty-one of the signal leads are bidirectional which include, for example, the address lines (18), data lines (16), control lines (2), and the lines dedicated to monitoring the power supplies. The remaining five lines are unidirectional.

2.29 All devices attach to the UNIBUS by paralleling off the desired signal and control lines as the lines are passed serially through the device enroute to the UNIBUS terminator (located on the end of the UNIBUS section). Fig. 3 shows a typical EADAS CCU UNIBUS structure.

2.30 The allowable length of the UNIBUS and the number of devices which are attached to it are limited by a PDP-11 system requirement of 50 feet and 20 unit loads. A UNIBUS unit load is defined as one UNIBUS receiver and two UNIBUS drivers which is the approximate equivalent of one device (ie, RK11 disk controller, 8K of MF11 core memory). If additional length or loading is required, as it is for EADAS, a UNIBUS extender may be added. It allows an additional 50 feet and 19 unit loads to be added to the UNIBUS structure. The UNIBUS characteristics for a typical 100-channel EADAS System are as follows:

Main UNIBUS:

KD11-A PDP-11/40 Processor
DL11-A TTY Interface
80K Core Memory
PC11 High-Speed Paper Tape Reader
BD04 Disk Interrupt and Alarm
RK11/RK05 Moving-Head Disk
RF11/RS11 Fixed-Head Disk
TM11 9-Track Magtape Controller
DB11-A UNIBUS Extender
KE11-E Extended Instruction Set
KT11-D Memory Management Control
MA11-FA Multiport Memory
RJS04-BA Fixed-Head Disk

UNIBUS Extension:

DB11-A UNIBUS Extender
BD03 ASCII/BIN To BIN/ASCII Converter
DH11-AA 16-Line Asynchronous Multiplexer
DN11 Automatic Calling Unit
5 DB05 Channel Interface Drawers
LP11-FA Line Printer Interface
Read Only Memory
CR11 Card Reader
DL11-B TTY Interface
DQ11-DA Synchronous Line Interface

D. Storage Devices

Core Memory

2.31 The memory provided with EADAS is magnetic read/write core memory. A basic 50-channel system is equipped with 64K words expanding to 80K words for 100 channels.

Memory Management Control

2.32 This control unit provides the hardware facilities for the management of the CCU core memory. Its basic purpose within EADAS is to allow additional words of core memory to be added to the standard 28K word PDP-11/40 system. The core memory size is greater than the standard memory size, therefore the memory management control is required as a standard device for the EADAS application.

Fixed-Head Disk

2.33 The fixed-head disk (RF11) system consisting of a disk controller and a disk drive is a fast random-access bulk-storage system. An RF11 provides 262,144 17-bit words (16 data bits and 1 parity bit) of storage. Up to 2 disk drives may be provided.

RK05/RK11 Moving Head Disk

2.34 The RK05/RK11 disk consists of an RK05 disk drive and an RK11 disk controller. The disk cartridge used in this unit holds over 1.2 million words. An RK11-D includes a control unit and the first disk drive. Up to 2 disks may be provided. The RK05/RK11 is used only with non-NORGEN Systems.

Magnetic Tape System

2.35 The magnetic tape system is used for writing, reading, and storing large volumes of data and programs in a serial manner. The system reads and writes in a compatible format for downstream processing. The magnetic tape unit accommodates 10 1/2-inch tape reels that contain up to 2400 feet of tape. Each reel can contain over 180 million bits of data stored on 9 tracks. Tape transports are either 800BPI (TU10) or 800/1600 BPI (TE16). Two transports may be provided.

E. Input/Output Devices

Modified KSR-35 Teletypewriter

2.36 The modified KSR-35 teletypewriter is modified by Digital Equipment Corporation for compatibility with the EADAS CCU. The operations performed at the EADAS CCU via the teletypewriter are as follows:

- Schedule changes
- Define calculations (non-NORGEN)
- Add and delete system functions
- Define equipped channels
- Program inputs, outputs.

LA36-CA DECwriter

2.37 A LA36-CA DECwriter may be used instead of a modified KSR-35 teletypewriter to perform the same functions. It is used as the input/output terminal during the running of the off-line DEC diagnostic programs for on-line EADAS system tests, and for interaction with EADAS and NORGEN software.

2.38 The LA36 DECwriter prints from a set of 64 characters at speeds up to 30 characters per second.

LP11-FA High-Speed Line Printer (Used in Early Installation)

2.39 The LP11-FA high-speed line printer is an 80-column, 64-character model. The printer is an impact type using a revolving character drum and a hammer per column. Fanfold paper 9 1/2 inches wide may be used with adjustment for pin-feed tractors. The print rate is 600 lines per minute.

LP11-VA High-Speed Line Printer

2.40 The LP11-VA high-speed line printer is a 132-column, 64-character model. The printer is an impact type printer using a revolving character drum and one hammer per two columns. Forms with up to six parts may be used for multiple copies. Included with the printer is a control unit for interfacing to the computer.

PC11 High-Speed Paper Tape Reader/Punch

2.41 The high-speed reader and punch is capable of reading eight-hole unoiled perforated paper tape at 300 characters per second and punching tape at 50 characters per second. The system consists of a paper-tape reader/punch and control.

F. Additional CCU Requirements for Systems With ICUR

PDP-11/40 Processor

2.42 The ICUR option provides an additional 32K words of memory expanding the EADAS/ICUR system to a minimum of 96K words and to a maximum of 112K words. 16K of the 32K additional words of core memory required for ICUR must be of the multiport memory type.

2.43 The second 9-track magnetic tape drive is mandatory for an ICUR installation. The additional magnetic tape recorder requires an additional cabinet in the lineup. The addition of a punched card reader, a read only printer, two high density fixed-head disk units, and a moving-head disk are also required.

Fixed-Head Disk and Controller

2.44 The RJS04 includes a controller and an RS04 fixed-head disk drive with a storage capacity of 512K 16-bit words. The RS04/RJS04 fixed-head disk and controller are used in the EADAS/ICUR System. The RJS04 is expandable by adding up to two RS04 drives per controller.

Card Reader

2.45 The CR11 card reader is used to input circuit grouping information for ICUR. The card reader reads EIA standard 80-column punched data cards at 300 cards per minute.

Multiport Memory

2.46 The MA11-FA multiport memory allows the computer memory and data to be shared on the UNIBUS in the EADAS/ICUR System. In the basic EADAS System configuration, the memory and data appear only on UNIBUS A. UNIBUS B is required for load balancing of the additional 32K of core memory required for the ICUR feature. The multiport memory contains 16K of core memory located on UNIBUS B and a memory control circuit. The other 16K of memory is located on UNIBUS A.

G. NORGEN Feature

2.47 When an EADAS System is equipped with the NORGEN feature, additional equipment is required. The system must have KE11-F floating point option added to the processor, an RP05 (which replaces the RK05/RK11 disks and controller) moving-head disk, MA11-FE multiport memory, and an additional cabinet with an expander drawer in the lineup. If more than 16 network terminals are used, additional multiplexers may be provided to furnish up to 48 terminals.

H. EADAS/NM Synchronous Line Interface

2.48 The DQ11-AA synchronous line interface is a high-speed, double buffered interface unit used to connect the EADAS System with EADAS/NM. This allows EADAS to provide remote concentrated data for processing by EADAS/NM. An interface module is located in each system . When the systems are located more than 50 cable feet apart, 209A data sets (or equivalent) and data facilities are required.

3. SYSTEM SOFTWARE

A. System Operation

3.01 The EADAS mission is accomplished by a series of interrelated task programs, each of which performs (or helps perform) a system function. These functions include:

- Collecting and summarizing data for up to 100,000 registers on the fixed head disk plus an additional 100,000 on the moving head disk spread over a maximum of 100 channels from a variety of sources on both an event-by-event and accumulated total basis

- Periodically performing several thousand arithmetic and logic calculations on the summarized data (non-NORGEN System)

- Reporting results which exceed predefined limits on up to 16 (48 with NORGEN optinal equipment) remotely-located network terminals (NT sites and optionally at the central site) (non-NORGEN System)

- Recording the summarized data on a magnetic tape on a scheduled basis for downstream processing

- Issuing formatted reports about each entity on a scheduled basis to the NT and central sites

- Retaining for a period of time the results of calculations performed in the system

- Accepting demand requests from all NTs for supplementary data such as additional results and raw register totals

- Performing other ancillary functions such as administering user-defined schedules and data tables, performing remote detection tests on input channels, etc.

3.02 EADAS programs will be required to pass summarized register totals on a frequent basis to the Network Management (NM) System and to accept and execute network control information from the NM System.

3.03 The relationship between these functions is best understood by looking at an overview of the system. Traffic data, transmitted over the input channels, is received by interface circuits which temporarily store single-event data in buffers associated with each channel. As the core buffers fill, the data is placed on disk as summed totals in areas dedicated to each channel. This continues for a period of 30 minutes and the data just collected becomes passive (ie, new data is no longer added to it). The previous 30 minute data collected is now erased and the storage area on disk becomes active (ie, new data coming in is accumulated). This active to passive change is called the "swap". At this point, other functions such as performing calculations, printing results, and writing magnetic tape are activated to operate on the passive data. Also at this time, those channels which transmit register totals are polled. Meanwhile, the single event data is being received, buffered, and placed on the active area of disk. When 30 minutes have elapsed, the new data is frozen, the active and passive areas of the disks are swapped, the various tasks are activated again, and the cycle repeats itself.

B. Program Organization

3.04 The EADAS program operates on-line in real time in a multiprogramming environment to perform its many tasks. It is structured so that several partially completed tasks may be processing concurrently in a main repeating loop called base level. In this base level loop, only one task is actually executing at any instant in time, but over a period of time all tasks are performed, each one alternating between using the processor and waiting for input/output requests to finish. This structure has the effect of dividing each task into sections where only one section of each task executes in a given base level loop. Thus, each task will require many sections, ie, base level loops, to complete.

3.05 Interrupts, handled by interrupt service programs, are used to temporarily suspend the execution of the base level loop. Interrupts occur when an input/output (I/O) action or file operation has been completed and on a timed basis for scanning input channels. Knowledge that the I/O action has occurred is communicated to the concerned base level program. The base level loop is then continued at the place where the interrupt occurred. The next time the concerned base level program is executed, it can continue on with the section of code that interprets the effects of the I/O action.

3.06 In addition to base level and interrupt service programs, there is a third class of programs in EADAS called service routines. These are pieces of code, needed in common by many tasks, which do not need a real-time break. Examples of this type of service routine include printing CCU terminal messages and entering disk requests, program timing, using ASCII-to-binary and binary-to-ASCII converters.

3.07 Sequencing of base level programs is controlled by an executive program. This program causes each base level program to be called in a predefined serial order. After the last program in the sequence is executed, the executive restarts the cycle by calling the first program, thus forming a closed loop.

3.08 Most of the functions in EADAS which process collected data are periodic in nature, usually occurring once per system period. These programs are initiated by a time monitor program which operates as part of the disk interrupt. The time monitor does this by comparing a system software clock to a generic time schedule and then placing the first progress mark of programs scheduled to be run in their slot in the executive table. This causes the first segment of the program to execute within the next base level loop.

3.09 Many of the functions in EADAS do not take a full system period to complete. In fact, most are quite short, taking only a few minutes to execute and then remaining in a dormant state for the rest of the swap period. Also, some sets of functions have a natural execution sequence among themselves, ie, they are naturally executed in mutually exclusive time intervals. Other periodic functions have no restriction as to the time during a swap period that they execute and, consequently, may be in mutually exclusive execution intervals. In addition to period functions, certain CCU terminal requests are also mutually exclusive. These facts are used to reduce the amount of core occupied by programs. This is done by having programs which execute in mutually exclusive time intervals share the same core area on an overlay basis. Thus, all programs of a set will reside on disk while only the currently executing program will be in core. When it finishes, the next sequential program will be read into core from disk and subsequently executed.

3.10 A common paging program is used to bring the programs from disk into core and start their execution. The same paging program is called when the time monitor program initiates a sequence and when a program to handle a CCU terminal request is overlayed in core. All programs in a mutually exclusive set will also use the same executive table progress mark slot.

3.11 EADAS provides four distinct types of communication with its user via the network terminal (NT), the CCU terminal, and the line printer. Two of these consist of output initiated by EADAS to produce routine reports and EADAS error messages. The error messages are printed automatically on the CCU terminal.

3.12 System error messages can be issued by any task program in the system and are printed on the CCU terminal. They are divided into four categories: ***information, action, warning,*** and ***fatal.*** Some messages can call for human intervention and will discontinue a function until that action is taken. For example, if the magnetic tape is off-line at the beginning of a swap period, the program will issue an error message and cancel magnetic tape writing until an operator places the tape on-line. The other types of communication are of an interactive nature and are initiated by the user. These are:

- Modes for defining installation dependent parameters needed by the system,

- Demand requests for additional information on collected data or for verification of parameter definitions, and

- Execution of NORGEN commands.

Five modes are implemented to specify such things as channel definitions, hourly report formats, (non NORGEN only), and collection schedules. These modes may be entered only at the CCU terminal and are conducted in a semiconversational manner.

3.13 Demand requests may be made from any of the NTs or from the CCU terminal. Replies to requests made from an NT are printed only at the NT. These inquiries may request almost any information stored in EADAS whether or not the information pertains to an entity normally associated with the requesting NT. If an inquiry breaks the printing of reports, these reports are lost in non-NORGEN Systems. In addition, the current message being printed and several other queued messages could also be lost. Input from an NT cannot alter information in the non-NORGEN System.

3.14 Demand requests at the CCU terminal may be about any information stored in EADAS. Depending on the size and the internal way in which it is constructed, the response may appear on the line printer or the NT. If any inquiry interrupts routine printing on the line printer, this printing will not continue after the inquiry is answered in non-NORGEN Systems. Certain commands, such as change time of day, can also alter information in the system. NORGEN, discussed in paragraphs 3.43 through 3.64 replaces the calculation and hour report mode operations and is more flexible with interactive operations from the NT and CCU terminals.

3.15 Since the man/machine interface in EADAS is so diverse and involves a multiplicity of devices, many highly interrelated task programs are needed to accomplish this function. These task programs are partitioned by device and operate both at interrupt level and base level. The interrupt programs interface with the hardware, while the base level programs interpret input messages and assemble replies which are passed to the interface programs. In some cases, several base level programs are executed in order to print a single line of output. Fig. 4 gives a general overview of the complex program structure needed to handle the man/machine interface.

C. Disk Memory Organization

3.16 The basic disk memory system used in EADAS consists of a fast access disk (fixed-head) having 128 tracks of 2048 words for a total storage capacity of approximately 260,000 words. In addition, there are two disk cartridges (moving-head) each having a capacity of 406 tracks of 3072 words for a total memory capacity of around 1.2 million words for non-NORGEN. The data is subdivided into six sections:

- Program and system parameters
- A raw register accumulating area
- User defined calculation definitions
- Hourly accumulation areas for raw registers
- User defined hourly report formats
- Long term exception report calculation results.

The six major sections of data have fixed sizes and locations on the disk.

Registers

3.17 Two identically formatted sets of 50 contiguous tracks on the fixed-head disk are used to store register totals. One set, known as "passive", is used to hold the data totals for the previous swap period. This includes both event-by-event data collected during the previous swap period and accumulated totals received at the beginning of the current swap period but representing the previous swap period. The other set, known as "active", is used to accumulate the data for the current swap period including both partial event-by-event totals and intermediate accumulated data totals needed for network management. The active/passive areas are swapped just prior to the beginning of a swap period by zeroing the then passive set of tracks and interchanging the starting addresses of the two sets of tracks.

3.18 Within a set of tracks, the data is organized by dedicating one half track to each one thousand register group or fraction thereof ("software channel") from an input source. For event-by-event data sources (where the maximum input is 992 or less), the entire source corresponds to a single software channel. For accumulated data sources, a single input source can correspond to up to five software channels (eg, ESS). The numbering of software channels is in one-to-one correspondence

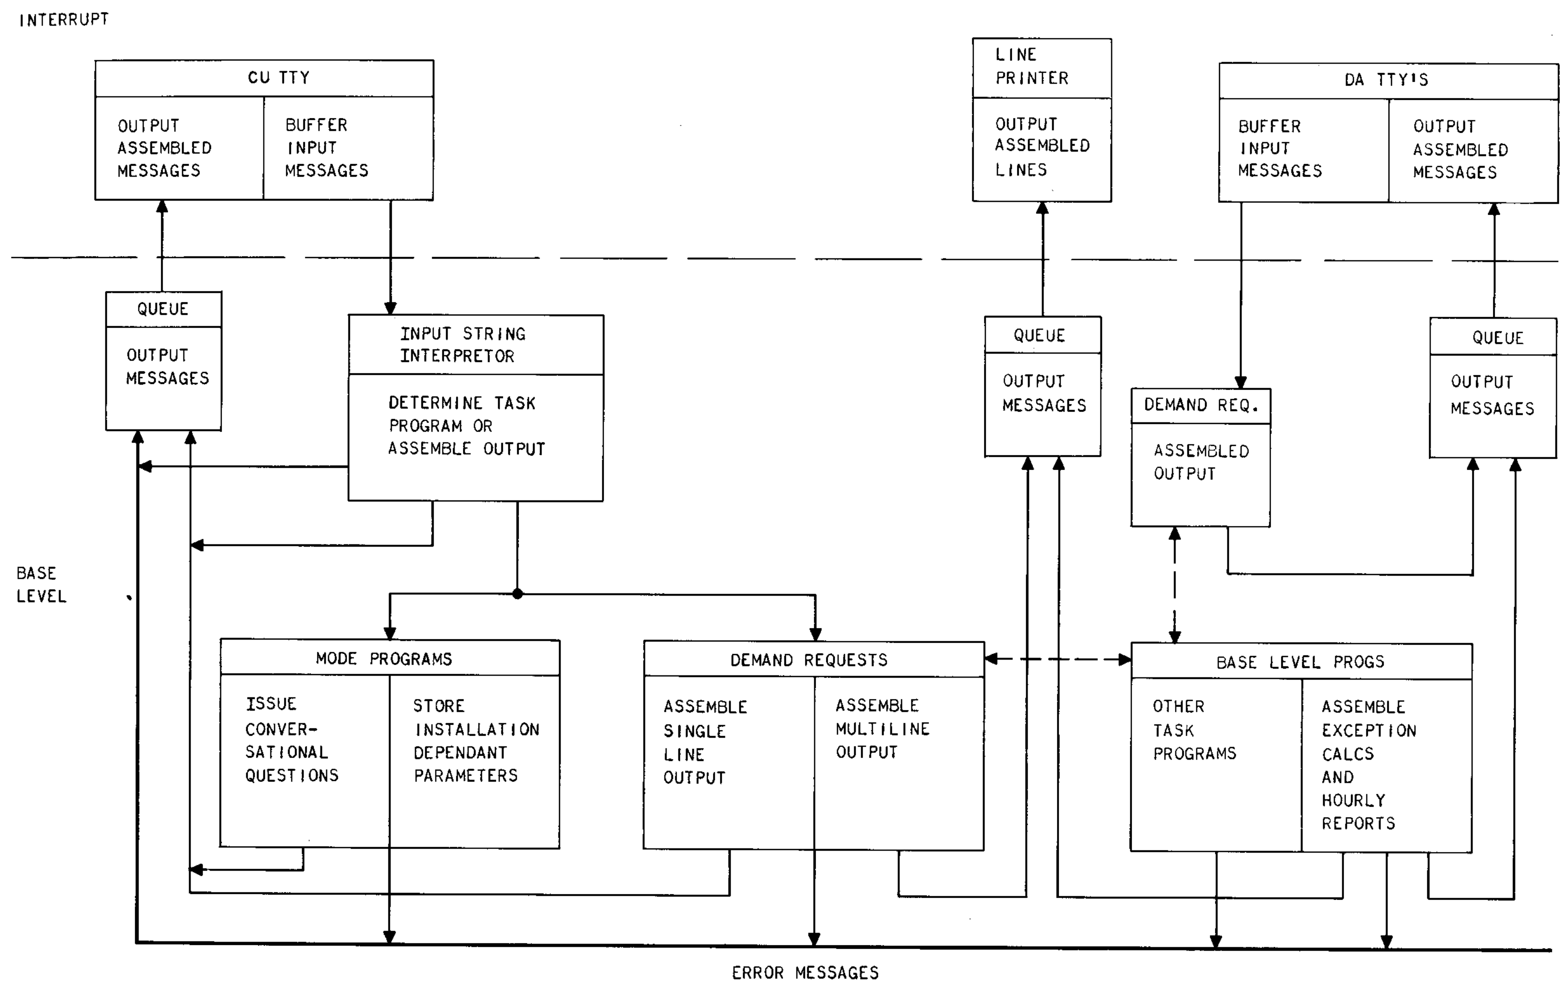

Fig. 4—EADAS Man/Machine Program Flow

with the contiguous track addresses in a set, ie, sotware channels zero and one are stored on the first track, channels two and three on the next, channels 98 and 99 on the 50th.

3.19 Another set of register storage exists on the moving head disk. Here 400 blocks, each of which holds 250 registers, are available for some accumulated data register sources. The channel headers for these exist on the fixed head disk and a map in the header identifies which register blocks belong to that header. These registers exist on a passive basis only. There are ESS channel headers available for this type of register storage.

3.20 Software channel numbers are assigned by the program when the input sources are defined and are for internal program use only. The assignment separates the input sources into two mutually exclusive groups in which event-by-event channls are numbered contiguously in ascending order starting with zero, and accumulated total channels which are fixed head disk register storage are numbered contiguously in descending order starting with 99. The accumulated data channels using the moving head disk software channel storage are numbered from 100 up to 255 in ascending order. Assignment in this manner ensures the most efficient processing inputs and also expedites system growth since the software channel numbers and the channel interface slot number to which an input source is physically connected do not have to be the same.

3.21 Each software channel storage area is divided into two sections, a 1000-word block for storing totals and a 24-word header. The header is used to store information related to the input source which is needed for various system functions. This includes such items as source identity, collection schedules, and scaling factors. For event-by-event sources, the word assignments in the 1000-word block bear a one-to-one ordered relationship to the assignments of inputs on the data collection device. For accumulated sources, the same order relationship holds within each of the associated software channels. For example, an accumulated data source with 1800 contiguous inputs would have two software channels assigned. The first would hold inputs 1 to 1000 in order; the second would have the remaining registers in words 1 to 800 in order. Hence within each set of tracks, a disk word ("software register") is dedicated for every input in the system.

3.22 In addition to the active and passive sets of raw register totals kept on the fixed head disk, a third set of "accumulated" raw register totals for these fixed head disk channels is kept on the moving head disk. This third set, formatted identically to the other two, holds the summed hourly totals. It is used both for making hourly report calculations and for writing magnetic tape if the swap period is different from the magnetic tape writing interval.

Program

3.23 The remaining area on the fixed-head disk is used to store a copy of the object code of the EADAS program. This copy is used both in paging during system operation and for primary backup in case of system failures. In addition to the program image, a bootstrap loader and a system backup module are also stored. The bootstrap loader is a small program which is used to load the initial core image during system startup. It must be physically located at the beginning of track zero in order to be manually accessible. The system backup module is another small program which is used to create and access secondary system backup information on magnetic tape.

3.24 The program image is divided into a fixed area and a variable area. The fixed area consists of parameters and nonpaged programs. It is stored on the disk as an exact core image of the lowest number addresses in core. Parameters consist of vector addresses needed to handle interrupts and special processor conditions and miscellaneous installation dependent information which is in core rather than on disk. Examples of such information include availability of the line printer, a table of the current status of each input channel, and supplementary data for using channels and calculation definition blocks. The variable area contains all paged programs. Except for the initial image of paged programs, all paged programs are stored in arbitrary order on the disk. A table within the common paging program maintains the location and size of paged programs on the disk.

Calculation Definitions (Non-NORGEN)

3.25 A modular highly-structured set of tables which contain the information needed to define exception report calculations and hourly-report calculations are stored on the moving-head disk. The modular pieces, known as blocks, contain 512 words (1/6 track) and can hold up to 17 (on the average) calculation definitions about a single software channel. A multiple number of blocks, typically between one and three, may be linked together for a single channel.

3.26 In all, a maximum of 400 blocks (up to 6800 calculations) are provided. A chained list, which accords easy administrative and growth practices, is used to associate software channels with their blocks. The first block for each channel is found via an index table kept in core. Additional channel-related blocks, if any, are pointed to by linkage words stored on each block. Each block is essentialy self-contained and encompasses nearly all of the data needed to define, perform, output, and administer each of its calculations. This includes an alphanumeric identification, multiple thresholds, intermediate and final results, and exception printing formats for each calculation. The actual definition is stored as a variable length string of register numbers, constants, and arithmetic operations which is interpretively processed to perform a calculation. Thus, the user has complete flexibility in defining each calculation and the output generated by it.

Long Term Data (Non-NORGEN)

3.27 The major portion of the moving-head disk is used to store exception report calculation results over an extended period of time. This area is arranged to hold results of ninety-six swap intervals for a fully-equipped system (100 channels). The period of time which these 96 intervals cover varies since the user defines not only the length of a swap interval but also the continuous hours per day to store long term data. For example, if the user specifies the swap interval as 30 minutes and wishes to save results for the entire day, then at any time, the data base will contain calculated results for the previous 48 hours.

3.28 Internally, the long term data is stored in a rotating buffer with 96 slots. Each slot is 34 sections in length and holds all the results for one swap period. The results from six calculation definition blocks are packed into one disk section (1/12 track). Along with the results, pass/fail indications are kept on a per calculation basis. As a check when retrieving data on demand, the time of the results is stored on each sector.

D. Interrupt Programs

3.29 A multilevel (seven) interrupt system is provided in the PDP-11/40 processor to allow entry to programs immediately on demand. EADAS uses four of the seven available interrupt levels. Interrupts have a two-dimensional priority structure such that the highest level interrupt demand is executed. In addition, for simultaneously appearing requests within a single interrupt level, the most important device (determined by physical closeness to the processor) is serviced. When an interrupt signal is generated, an interrupt service program (ISP) is entered at a fixed programmable address corresponding to that device. The ISP will execute at the processor level corresponding to the status word assigned to the vector. When the service program is finished, control is returned to the program that was interrupted. Fig. 5 depicts the interrupts assigned to each priority level and their relative importance within each level.

3.30 Most demand interrupts signal the completion of an input/output action. The action itself was performed by a direct memory access (DMA) transfer without processor supervision. For example, to perform a disk read, a program merely informs the disk controller of the location of the designed data and the location of the core buffer into which to put the data, and then the program relinquishes control to the executive. After the disk has positioned itself, the actual data transfer is done via a DMA request while the processor is executing other codes. After the last word is transferred, an interrupt signal is sent to the processor causing the disk ISP to be executed. This program will communicate to the requesting program that the disk read has been completed. Interrupts are also generated on a fixed-time basis to perform frequently occurring periodic functions. For example, scanning for register data on the input channel interfaces is done on a fixed periodic basis using a timed interrupt.

Input Channel Scanning

3.31 The program (SCAN) that scans input data channels is divided into two sections one of which handles event-by-event data, the other accumulated totals. Since channels of the former-type are more real-time critical, the system software is structured to handle these as efficiently as possible.

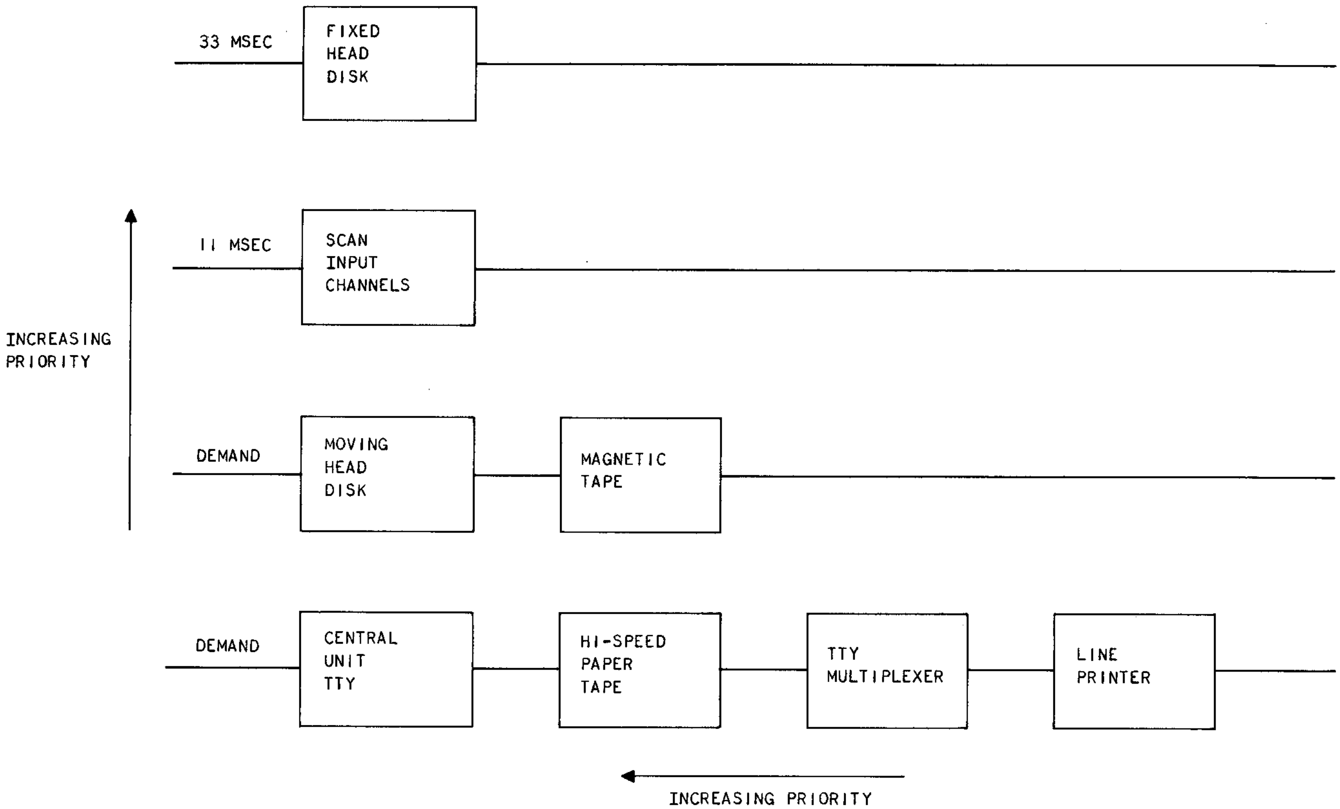

Fig. 5—EADAS Interrupts

Event-By-Event Data

3.32 Event-by-event data is transmitted to EADAS at the rate of 80 words per second per channel and is buffered in core temporarily. The areas on disk corresponding to each channel are continuously read, one after the other, into core, updated with the new data and rewritten onto the disk. In order to reduce the amount of core memory needed, an algorithmic scheme was developed for buffering the input data. Briefly, this scheme, known as "triangular buffering", eliminates core memory which would otherwise be empty when the arrival and processing of input is uniform. It does this by dynamically altering the buffer slots assigned to each input channel over time. Thus, while the total buffer space always remain full, the input channels which were recently updated will have only a few allotted slots and the channels to be updated next will have the most allotted slots. The fundamental assumption of the triangular buffering scheme is that the input scanning and disk updating process is uniform. This means that a fixed integer number of scans of the input channels must be made before each source is updated.

3.33 The fixed-head disk used for storing register totals has a rotational speed of 33.3 msec. However, a word of input is transmitted over each input channel every 12.5 msec. Therefore, in order to synchronize the scanning of inputs with the disk, a fixed-time interrupt of 11 msec was chosen for scanning. To guard against drifting of the interrupt, generation of the scan interrupt pulse is physically tied to the disk timing gap. The scan interrupt is offset by msec from the disk timing gap so that the program that updates the disk can execute partially (and empty some of the buffer) before the scan program places more data into the buffer. Every 11 msec, the program reads each interface and places the contents into the core buffer slot determined by the triangular buffering algorithm. Every twelfth scan, after all channels have been scanned, the scan program initiates an update program.

Accumulated Totals

3.34 Devices containing their own local memory for accumulating traffic data can be polled by EADAS at regular intervals. Presently, a number of such devices exist including offices with outside vendor accumulating devices, ESS offices, and pollable data terminals. Each register total transmitted from these devices is identified by its position in the data stream. The data, therefore, will be transmitted in small groups or blocks in order to avoid loss or retransmission of large blocks of data, should synchronism be lost. Each block will contain an identifier and up to 250 register totals.

3.35 The number of accumulated sources polled simultaneously is a function of the amount of core memory allotted. This is a function of the percent of accumulated devices in the system and the importance associated with accumulated data. Between 0 and 72 buffers will be provided for accepting simultaneous transmission of accumulated data. Each of the buffers will be around 300 words, 250 for a block of register totals and 50 for control information.

3.36 Since each device has unique features, each must be treated differently. These differences include: the format of data transmitted, the data rate, the polling frequency, and the polling technique. The program (POLACM) which does the polling and processing is divided into two sections, base level and interrupt level. The interrupt program interfaces the software with the channel data links. The base level program deals with the differences in channels. It decides the channel and function to be worked on next, passes output messages to be transmitted to the interrupt program, and processes data previously accepted by the interrupt program.

3.37 The interrupt level portion of the accumulated data program executes during the 66 msec timed interrupt after the event-by-event data channels have been updated. It looks at each active input buffer, determines the channel using that buffer, and the function currently being worked on. For example, if the channel is currently transmitting totals, it scans the input interface for data and buffers it. If transmission is an ASCII, the interrupt program assembles the characters into binary numbers. When a block has completed transmission, the interrupt program flags the base level program to process the data. A reverse process is followed when the buffer is being used to transmit a polling request to a device.

3.38 The fixed-head disk is used primarily for accumulating register totals. The process of scanning inputs and updating the corresponding software registers must be done synchronously. Since input scanning is an interrupt process, updating must also be initiated by an intercept synchronized to the scan interrupt. Since it is initiated by an interrupt, updating cannot wait for disk accesses and hence cannot use the normal method for obtaining disk data. Therefore, a structure was developed within the framework of EADAS which maintains the proper time sequence between scanning and updating the disk and which ensures that the correct software registers are in core for the update program. This structure involves both software within the disk interrupt service program and a minor variation to the standard disk interrupt arrangement.

Service Routines

3.39 There are many well defined subjects which are needed in common by many task programs. Some of these, such as using the hardware ASCII-to-binary converter, can be accomplished immediately and involve nothing more than a straightforward subroutine call. Others require a real-time break and are considerably more involved. These are generally accomplished through a series of progress marks either in common or within individual task programs.

3.40 The disk monitor program (DMONTR) is used by task programs when requesting disk accesses. It has multiple entry points, one for each type of data on the disk. When a task program needs data, it calls DMONTR at the appropriate entry point with an indication of the specific data to be accessed. Other informtion pertinent to the request is passed via a table in the task program.

3.41 A disk request can be in any of four distinct stages, each of which is characterized by a progress mark in the task program.

- Progess mark 1—A call to the appropriate subroutine. This is needed for retries when the queue is full.

- Progress mark 2—Request is in queue or being performed. This can be an immediate return to the executive, or timing can be performed as a defense against system malfunction.

- Progress mark 3—An error occurred on the disk request. The task program, not the monitor, now decides the action to be taken.

- Progress mark 4—Disk access was completed unsuccessfully. Task program continues.

This combination of progress marks, together with the calls to DMONTR, affords the task programs considerable flexibility when using the disk, while still placing all of the common jobs in a single program.

E. Base Level Programs Polling Accumulated Data Devices

3.42 The base level program (POLACM) that initiates the polling of accumulated devices is activated every system period. This program looks at each accumulated data channel and determines from an internal schedule if any data is to be received. If no data is received, the channel is marked completed. If there is data to be received, a core buffer is assigned to the channel. The program then determines the size and number of blocks to be polled and formulates a polling message in the core buffer. After the message has been sent and a block of data has been received by the interrupt program, the base level program checks the data for error and, if none, decides where to store the data. If an error is determined, retransmission is requested when possible. If not possible, the data is marked as bad. This procedure continues until all blocks for a device have been transmitted. At this time, the device is marked completed and the buffer is released. Each polled device should have a buffer allocated whenever polling is initiated. Input buffers can be polled in parallel for as many devices as there are. There are two exceptions to these descriptions. First, in EADASs which are passing data to a network management system, POLACM will be executed every five minutes. The number of blocks received at these five-minute intervals will usually be much less than at the swap period. When the swap period and a five-minute interval occur simultaneously, polling for all network management is completed before any nonnetwork management data for that channel is received. Second, the swapping of active/passive memories in some devices is controlled by EADAS. When entered on the system period, POLACM will swap the memories in a device before polling for data.

F. NORGEN

3.43 NORGEN is a complete multiuser time-sharing system written to share the computer with EADAS. In conjunction with the NORGEN applications programs, it replaces the user defined calculation and hour report programs.

NORGEN Characteristics

3.44 NORGEN was developed to provide the operating telephone company (OTC) with an improved EADAS to handle the expanding requirements for more surveillance and near-real-time reporting of network maintenance and network administration data. NORGEN provides the OTC with:

- Expanded and more flexible calculation capability with sophisticated logic for better exception reports

- Improved hourly reports and 24-hour data accumulation for daily, weekly and monthly reports

- Up to 48 remote terminals that can selectively receive reports and interact with the system to change report schedules, thresholds, etc

- Standard programs designed for applications to various types of central office switching equipment

- A high level programming language.

In general, NORGEN provides the OTC with increased near-real-time surveillance and reporting capabilities more consistent with the functions of Network Operations Administration.

3.45 NORGEN provides the facility to perform and report general logic and arithmetic calculations on EADAS register data. The data is stored by EADAS on the moving head disk, where it can be read by any program running under NORGEN. There is enough data storage space for the EADAS data (about eight million registers) so that EADAS will not overwrite old data for many system periods. As a result, close synchronization between data collection and calculation is not necessary, and NORGEN programs can process data as they get to it (so long as, on the average, they process it as fast as it comes in). This looseness in timing allows honoring interactive user demands in spite of the real-time nature of the data being processed.

Software

3.46 NORGEN is a separate subsystem of software in EADAS operating somewhat independently of the data collection process. This separation of NORGEN and EADAS data collection functions was developed for two primary purposes:

(a) To make the EADAS report generating potential as flexible and general as possible

(b) To protect the basic data acquisition functions against interference from processes that generate reports.

The NORGEN subsystem consists of a main generic program which controls the operations on:

- Data bases created by the user or supplied with the system
- Programs for generating reports
- Network terminal input and ouput plus line printer output.

3.47 The NORGEN Operating System (NOS) is a completely general operating system. Under NOS, the user can create "K" (a high-level language roughly equal in richness to ALGOL or extended FORTRAN) programs, which can be run to process whatever data EADAS has collected, or for any other purpose. The NOS can be used in either of two ways:

(a) The user obtains the applications programs which have been written in K to process various switching machine types and generate reports. The user must learn to use and administer the data base manipulation programs which come with the applications programs to set up the input parameters to enable each office to be processed properly. Thus the user obtains reports for each office.

(b) The user creates K programs to process the EADAS data when it becomes available. To do this, the user must learn to program in K and some of the structure of the EADAS data storage on disk. NOS may also be used to augment the applications program.

3.48 The main features of the software in NOS are:

(a) Complete spooling (storage) of user input and output; then user can type input as far ahead of the system as desired, and programs need not pause waiting for terminal ouput to finish.

(b) A simple and effective scheduling algorithm which allows highly interactive users (eg, those doing editing) fast response even when users requiring more processor time are competing for service.

(c) User controls over the CCU or network terminal:

- Interactive and noninteractive modes
- Abort job control (CTRL/A)
- Abort printout control (CTRL/B)
- Reprint printout (CTRL/C).

(d) A simple and powerful file system.

(e) Ability to take input commands from a text file, instead of from the terminal. The user, with the aid of a text editor, can make any sufficiently noninteractive series of commands into a single command, by using a single special command and its arguments.

(f) Distribution of outputs to any terminal or set of terminals, or to file, without program modification.

(g) Protection against unauthorized user activity:

- Write protection for each set of files associated with a user identity code
- Log on passwords, execute permissions, output permissions.

(h) A large set of user commands.

3.49 The commands available in NOS include:

(a) A text editor

(b) A K compiler

(c) A K interpreter, which executes compiled K programs

(d) File system commands to copy, delete, rename, etc, files

(e) Backup to tape commands to save, restore, and verify NORGEN tapes

(f) Commands to change the input source or output destination.

3.50 The user can create programs in K using the text editor, compile them, and run them under NOS. Depending on how the commands are strung together, they can run on demand or on a schedule. Subject to total time availability, this can be done independent of every other user.

3.51 The NORGEN applications programs, which contain a comprehensive set of calculations for various office types (and a large set of programs to administer parameters describing how they should run for each particular office) were all written in K, and could have been written by the user.

3.52 In the main path, NOS appears to be a single-user machine, where inputs and outputs never hold up the computation. The primary reason for coding an operating system is to create this programming environment. The main path includes programs which swap in and out to perform foreground functions for the particular user; parsing commands, setting up spooled output, printing error diagnostics, in addition to the obvious command loads in response to a user request. These are all logically in the computation path for the user. This path is occupied by only one user at a time as determined by the scheduler in path three. The user currently in core may change every two seconds or so. This is transparent to the main path software and is what makes a time-sharing system.

3.53 The main path executes in the base level loop of the EADAS System monitor, which is just a ring structure of programs. The NOS attempts to keep the "base level loop" going by returning to it from time to time even if NOS could continue processing. Every command must be sure to call the operating system frequently enough to keep this base-level loop going. There are two null calls in NOS that cause no operation, but cause a cycle through the base level loop and give NOS a swap opportunity.

3.54 The input-output path copies text input from user terminals to disk and copies main path text output from disk to the user terminals. This path exists in multiple copies, one for each terminal, and is interrupt-driven. The input path is started by the reception of a character from a terminal and continues when another character or disk completion occurs as appropriate. After the output path has been started, it continues to be driven by disk or terminal-driver interrupts until the linked test disk blocks have been sent to the terminal command. This path is implemented through low-core subroutines named NORGIN, NOROUT and NORPUT.

3.55 The scheduler path implements the communication between the first two user paths and also schedules the computer time between users. This path executes from time to time at one of the checkpoints in the main path, but it is logically transparent to the main path. Every two seconds, the main path (at a checkpoint) will call SWAP, a subroutine which executes one iteration of this path. The SWAP loads in system programs which check the I/O status of every terminal. Every terminal has an 8-word status area allocated to it in low core shared by paths two and three. If a terminal has finished its allocated output from disk, it is given more if there is any (a call to NORPUT). The management of the spoiled I/O disk area is done at this point also. Then the scheduling program is loaded and a user swap may occur. Finally, SWAP returns to the main path at whatever checkpoint is implemented by nonresident programs named S#SWAP and S#SCHED. They are nonresident because they are executed infrequently, and there is a better use for the memory that they would occupy (eg, for I/O buffers).

3.56 The device drivers, except for the tape driver used by NOS are interrupt driven for both requests and completions.

3.57 NORGEN core resident programs are located as follows:

In low core (accessed by interrupt-level programs):

- Terminal status words (8W/terminal, 64 terminals)—512 words I/O path
- Programs—1024 words.

Anywhere in memory:

- Moving head disk accessible buffers for the I/O path—298 words
- "Minibuffers" for the I/O path—512 to 4096 words (optional)
- High core moving head disk accessible—command and buffer area; 10000 to 14000 words
- Resident NOS and file system—2048 words.

Like all of EADAS, NOS does its own memory management register manipulations.

Applications Programs

3.58 Provided with NORGEN is a series of applications programs. Each program is designed to produce exception and periodic reports for specific central office switching equipment types. The equipment types included are:

- No. 1 crossbar,
- No. 5 crossbar,
- Step-by-step,
- Crossbar tandem,
- No. 1 electronic switching system,
- No. 2 electronic switching system,

and a separate package for the pollable data terminal 1A (PDT 1A).

NORGEN Data Base

3.59 NORGEN operation is based on a system of catalogued storage locations on disks called files. All information used by NORGEN is obtained from files.

3.60 Files are either system files (provided with the EADAS/NORGEN generic) or user files and must be created for each office. The user files which must be created constitute the data base(s) of offices for the generation of reports.

Operation

3.61 NORGEN has two general categories of operation:

(a) An interactive process where users supply, modify, and verify information in the system.

(b) An automatic process conducted each system period when the applications programs are executed. This is called report generation.

Interactive Process

3.62 Interactive processing is used whenever NORGEN is required to perform a specific operation from a terminal. A network or CCU terminal operation uses this process for entering any of the commands to perform such tasks as:

- Create and/or modify data bases
- Request demand reports
- Execute programs.

Report Generation

3.63 After an office data base is created, the office network data will be automatically reported according to the assignments made in the data base. In this process, NORGEN executes the specific applications program packages for the particular equipment type offices initialized.

3.64 Each system period, NORGEN first makes a copy of the raw data just collected (excluding ICUR load balance data). It then begins to process this data according to the sequence of offices initialized. An office is selected, its equipment type determined and the applications program package for that equipment type begins execution. The applications programs perform calculations for reports, store the results, (ie, 24-hour data accumulation, exception results, etc), and generate reports to the appropriate terminal(s).

G. ICUR Description and Operation

Flow of Data in EADAS/ICUR

3.65 EADAS without ICUR interfaces with the 4A traffic usage recorder (TUR) by cabling leads from the TUR grouping terminals to the ETDC registers. With the ICUR option, the output of the TUR detectors will go directly into an ICUR applique on the ETDC (Fig. 6). Twenty-six leads from sequence circuits on the TUR brought through the ETDC will indicate which crosspoint is being scanned. The six detector leads indicate the usage status, busy or idle, for each of the six circuits brought to that crosspoint (Fig. 6). Usage data is individually obtained for all 3600 inputs (switch, contact, horizontal, vertical [SCHVs]). This data is transmitted to the CCU and accumulated in the ICUR Subsystem.

3.66 With this option, peg count data continues to be collected in the same manner as collected in basic EADAS. That is, peg count leads are cabled to ETDC inputs and the real-time counts are transmitted to the CCU where they are accumulated in the EADAS Subsystem.

3.67 The CCU sorts incoming data, based on word construction, and directs peg count and usage to the EADAS and ICUR Subsystems, respectively.

3.68 The ICUR Subsystem performs the following functions:

(a) Accumulates all usage on an individual circuit basis.

(b) Counts and accumulates TUR transitions on an individual circuit basis. (TUR transitions are changes of state, from busy to idle, or vice versa. Transitions are detected by a "last-look" comparison feature of the software which operates on every scan point on each TUR cycle.)

(c) Segregates LB data from non-LB data to facilitate separate processing.

(d) Groups all usage data according to a software assignment record called the circuit grouping map (CGM). The CGM is maintained by entering punched update cards at the CCU.

(e) Forwards grouped non-LB usage data to EADAS where they are accumulated in data collection devices (DCDs) on the software channels.

(f) Accumulates grouped LB data in special LB DCDs.

(g) Writes scheduled accumulations of grouped LB data directly on the Traffic Data Administration System (TDAS) magnetic tape.

(h) Writes scheduled accumulations of ***all*** individual circuit usage and TUR transitions (changes in usage state—idle to busy to idle), and certain mapping and scheduling data, on the ICAN tape.

3.69 Once EADAS has combined the grouped non-LB usage data and the peg count-type data, it writes the accumulated data on magnetic tape. This magnetic tape then becomes available for downstream processing by TDAS or other downstream equivalents (Fig. 7). This tape will have the same format as the one provided by basic EADAS. The individual circuit usage and other data is written on another magnetic tape for downstream processing by ICAN. The ICAN program will utilize this data to identify trunk and equipment problems and to provide essential administrative reports used in maintenance of the CGM.

3.70 The ICUR Subsystem automatically generates several types of output reports relating to system operation. For example, it provides near real-time reports regarding rejected CGM update cards and unexpected usage on unassigned or unequipped TUR inputs. These reports provide information on certain classes of mapping errors. Another report, pending status flag removals, lists CGM updates which take place dynamically. In addition to these automatic reports, the user can request at the CCU listings from the CGM for either entire TURs, specified data collection devices DCDs (circuit groups), or specified SCHVs. Certain TUR malfunctions (for example; detector failures, crossed verticals, and skipped crosspoints) are reported in system error messages on the receive only teletypewriter (TTY). Rejected update card messages will appear at the control unit terminal reports of pending status flag removals,

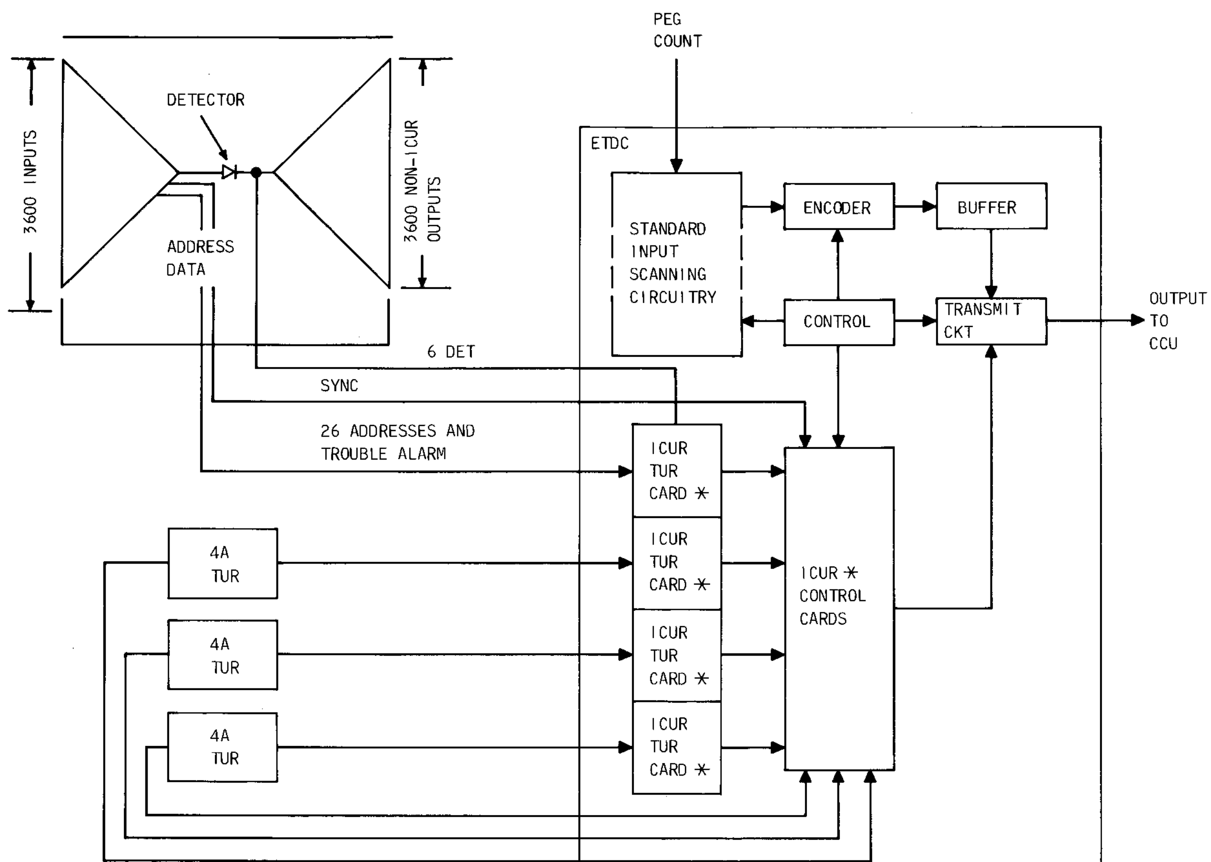

Fig. 6—ETDC With ICUR Option

unequipped/unassigned usage reports; TUR error messages will appear on the receive only TTY; and CGM listings will appear on the high-speed line printer.

Individual Circuit Analysis Program

3.71 The downstream ICAN program will read and process the data on the ICAN tape produced by the EADAS/ICUR System. ICAN accesses the TDAS-common update data base and in that way certain circuit group identification data for the ICAN data base will be obtained directly from common update.

3.72 The two major functions performed by ICAN are ***administration analysis*** and ***usage analysis.*** Administration analysis supports the EADAS/ICUR System by helping to ensure the integrity of the CGM. Usage analysis is essentially a maintenance function. It identifies abnormal load patterns observed on individual circuits (eg, "killer trunks," always busy trunks, and always idle trunks). Various periodic and demand reports are issued by ICAN in the performance of these two functions.

4. NETWORK OPERATIONS RESPONSIBILITIES

4.01 EADAS requires planning for its operation long before the actual ordering of the equipment. The Network Design Engineer, Network Administrator and maintenance personnel are key members of the coordination and planning of the system.

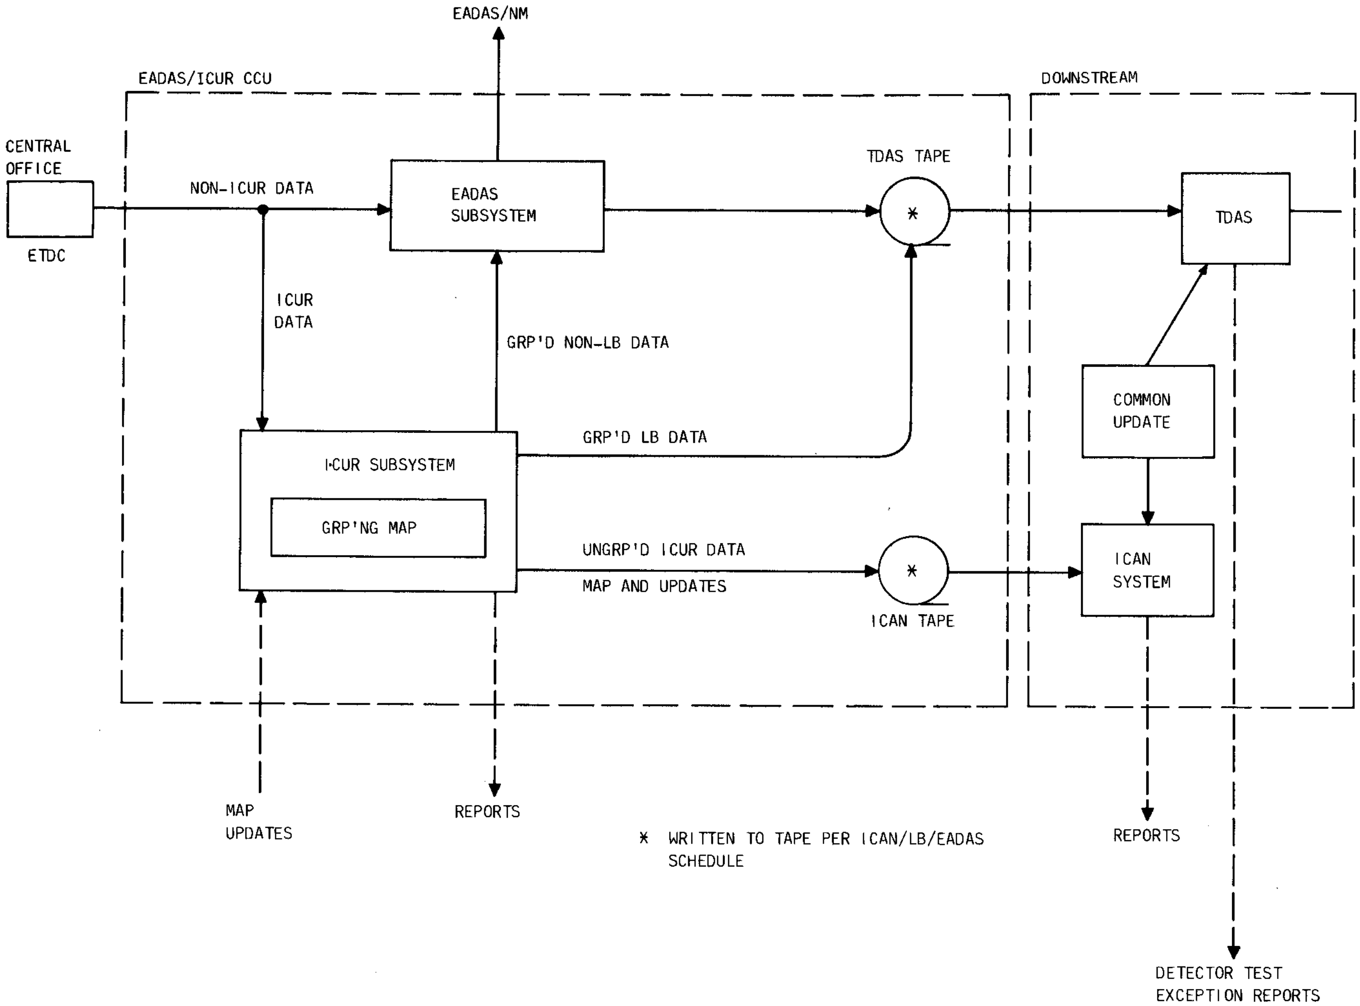

Fig. 7—Simplified Diagram of EADAS/ICUR System

A. Network Design

4.02 The engineer responsible for ordering EADAS must coordinate the individual end office equipment jobs as well as the CCU ordering. Types of offices (ie, No. 5 crossbar, No. 1 ESS etc) and their requirements determine quantities and capacities of the CCU. The following list of references is to aid the Network Design Engineer:

190-510-061	EADAS Determination of Quantities
190-510-062	EADAS/ICUR
190-510-100	EADAS General Description
190-510-202	EADAS Traffic Data Converter Description
190-510-203	EADAS Pollable Data Terminal Description
190-510-416	EADAS Dial Administration TTY Operation
190-510-420	EADAS Surveillance-General

B. Network Administration

4.03 The Network Administrator has prime responsibility for the operation of EADAS. This requires the administrator to:

- Coordinate with engineering on the configuration of the system

- Organize the administration of the CCU operations (ie, tapes, schedules, etc)

- Organize the administration of the ICUR and NORGEN data losses to maintain their accuracy

- Establish preventative and corrective maintenance procedures for the system.

4.04 The administrator is referred to the following documents for further detail:

190-510-202	EADAS Traffic Data Convertor Description
190-510-203	Pollable Data Terminal No. 1A-Description
190-510-416	EADAS Dial Administration TTY Operation
190-510-420	EADAS Surveillance—General
190-510-475	EADAS Dial Administrator's Training Guide
190-510-476	EADAS Instructor's Guide—Dial Administration Training
190-510-406	EADAS—Implementation Coordination
190-510-410	EADAS—CCU Operations
190-510-411	EADAS—CCU Position Practices
190-510-412	EADAS—CCU Position Practices
190-510-413	EADAS—CCU Position Practices
190-510-440	EADAS—Operational Maintenance
190-510-470	EADAS—CCU Training Material
190-510-471	EADAS—CCU Training Material
190-510-472	EADAS—CCU Training Material
190-510-473	EADAS—CCU Training Material
190-510-474	EADAS—CCU Training Material
190-510-501	EADAS/ICUR—General Description
190-510-504	EADAS/ICUR—Ordering & Coordination
190-510-520	EADAS/ICUR—System Definitions.

C. Network Maintenance

4.05 Maintenance has responsibility for acceptance testing at job installation and diagnostic testing for both preventative and corrective maintenance of the system. Practices for further description of the maintenance on EADAS are:

190-510-202	EADAS Traffic Data Converter Description
190-510-203	Pollable Data Terminal No. 1A Description
190-510-204	EADAS Central Control Unit Operating Procedures
190-510-205	EADAS Central Control Unit—Trouble Sectionalizing
190-510-211	EADAS Traffic Data Converter Remote Tests
190-510-212	EADAS Traffic Data Converter Tests of ETDC at ETDC Location
190-510-213	EADAS Traffic Data Converter Trouble Locating Procedures
190-510-214	EADAS Traffic Data Converter Connection and Verification Tests
190-510-215	EADAS Traffic Data Converter Connection Verification Tests Using EADAS Test Set
190-510-216	Traffic Data Converter Circuits Channel Definition and Validation
190-510-221	Pollable Data Terminal No. 1A—Remote Tests
190-510-222	Pollable Data Terminal No. 1A—Tests for PDT at PDT Location
190-510-223	Pollable Data Terminal No. 1A—Trouble Locating Procedures.