British Telecommunications Engineering

VOL 3 PART 4 JANUARY 1985

System X

The Journal of
The Institution of British Telecommunications Engineers

BRITISH TELECOMMUNICATIONS ENGINEERING

VOL 3 PART 4 JANUARY 1985

Contents

Published in April, July, October and January by *British Telecommunications Engineering Journal*, 2–12 Gresham Street, London EC2V 7AG.
(Formerly *The Post Office Electrical Engineers' Journal* Vols. 1–74: April 1908–January 1982)
Price: £1·00 (£1·50 including postage and packaging). Orders by post only.

Annual subscription (including postage and packaging): home and overseas £6·00.
(Canada and the USA $10·00).

Price to British Telecom and British Post Office staff: 51p.

 ISSN 0262-401X

For System X, full support from Plessey

Plessey is leading development work for British Telecom's System X. To ensure that System X lives up to its promise in service, and its potential for the 21st century, the Plessey commitment is total and absolute. Plessey provides fast, reliable access to its team of support engineers selected for their skills in the designing and commissioning of System X. Plessey expertise is always yours for the asking – whenever, wherever that you decide that you need it.

Plessey Major Systems Limited, Edge Lane, Liverpool.

PLESSEY

British Telecommunications Engineering

Telecommunications—Meeting the Challenge

J. F. BOAG, DIP.E.E., C.ENG., F.I.E.E., F.B.I.M.,† and P. B. FRAME, C.ENG., F.I.E.E.*

INTRODUCTION

Major technical articles on System X were published in earlier issues of this *Journal* between January 1979 and April 1981. During this period, development of the system was still proceeding, but it was coincident with the first exchanges becoming available for public service. In July 1980, a 200 erlang tandem switching unit opened for service at Baynard House in London, followed by a local switching unit at Woodbridge, Suffolk, in July 1981, having a capacity of 930 customer connections. Over the following months, other installations entered service within the UK network, and all have provided valuable early experience in operating and maintaining digital common-control exchanges in a widely diverse network.

The continued development has resulted in some significant changes in a number of areas, and mainly reflect the advances in technology that have taken place during this time. The first exchanges to take advantage of these more advanced designs were accepted for service in the Autumn of 1984 at Coventry and Leeds.

These events were significant milestones in the development of the UK network, since they heralded the start of a major modernisation programme based on System X digital exchanges.

The purpose of this issue of the *Journal* is twofold: to detail the major changes in the design of System X that have occurred since the turn of the decade, and to describe the influence that it has had on future UK network strategy, on operational and maintenance policy and on procedures.

THE WAY AHEAD

Major organisational changes recently introduced in British Telecom (BT) are to some extent the consequence of the changing circumstances for the provision of telecommunications services in the UK, where the Government has liberalised the terminal market and introduced competition into the network.

Two of the business divisions in BT which have been established are charged with specific responsibility for the inland network. Responsibility for the trunk network is vested in National Networks (NN), which embraces both traditional telephony and specialised services such as KiloStream data. The local exchange network is the responsibility of Local Communications Services (LCS).

† Chief Executive Trunk Services, British Telecom National Networks
* Director Local Exchange Systems, British Telecom Local Communications Services

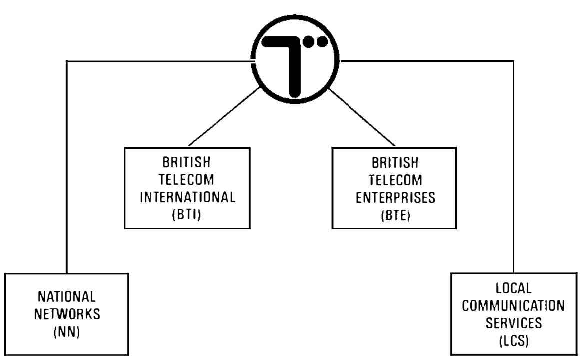

FIG. 1—Business structure of British Telecom

The business structure is illustrated in Fig. 1: British Telecom International (BTI) has responsibility for the international network interface, and British Telecom Enterprises (BTE) for consumer products and value-added services.

BUSINESS OBJECTIVES

BT's main objectives can be summarised as follows:

(*a*) to sustain and improve the quality and range of services available to customers, and
(*b*) to exploit new capabilities to support emerging market opportunities.

These objectives will be achieved by:

(*c*) accelerating the modernisation of the network, and
(*d*) minimising capital and running costs, consistent with quality customer service.

Over the next four years, a digital trunk network will be established so that all trunk traffic can be loaded onto it not later than 1990, and thus enable closure of the existing analogue switching centres (see Fig. 2).

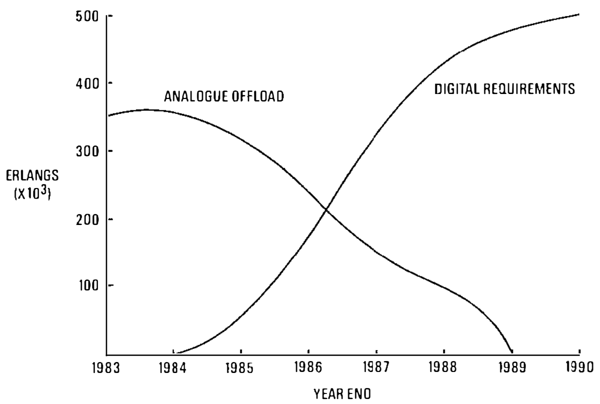

Present	Future
Analogue network (limited digital penetration)	60 digital main switching units
360 group switching centres	9 digital derived services switching centres
36 district switching centres	400 000 E switching capacity
11 main switching centres	32 000 2 Mbit/s modules
400 000 trunk circuits	850 000 trunk circuits
3·6 billion trunk calls	5·2 billion trunk calls

FIG. 2—Trunk switching modernisation

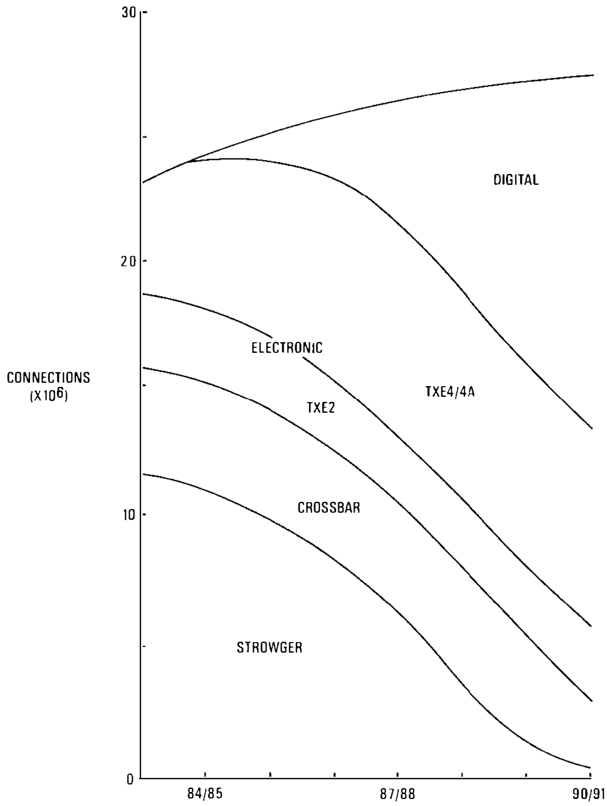

FIG. 3—Local exchange capacity

The local exchange network modernisation programme matches the trunk network plan, and some 80% of the total exchange connection capacity will be provided by either System X or TXE4/4A exchanges by 1990 (see Fig. 3).

Particular attention is being given to meeting the needs of the UK's business customers, and current plans are designed to ensure that digital switching and transmission capacity is available to serve at least 30 of our major centres by early 1986.

THE ROLE OF NETWORK MODERNISATION

An effective and responsive telecommunications network is essential if customers' expectations are to be met, particularly in relation to the wider market for information technology. Here, speed in the provision of service, acceptable cost, quality of service, and the evolution of a range of services and facilities that meet present and future needs, together with a responsive interface between the customer and BT, are of particular importance.

QUALITY OF SERVICE

Although much has been achieved through the use of more modern plant and better monitoring systems, the current network technology still leaves room for improvement in terms of calls with poor transmission, high noise level and long set-up times. Inter-processor signalling over a fully digital switching and transmission network, coupled with keyphones using MF signalling, will very substantially improve the service as perceived by the customer.

The availability of fast common-channel signalling will greatly reduce call set-up times across the network, and will be an important element in BT's ability to provide an integrated services digital network (ISDN).

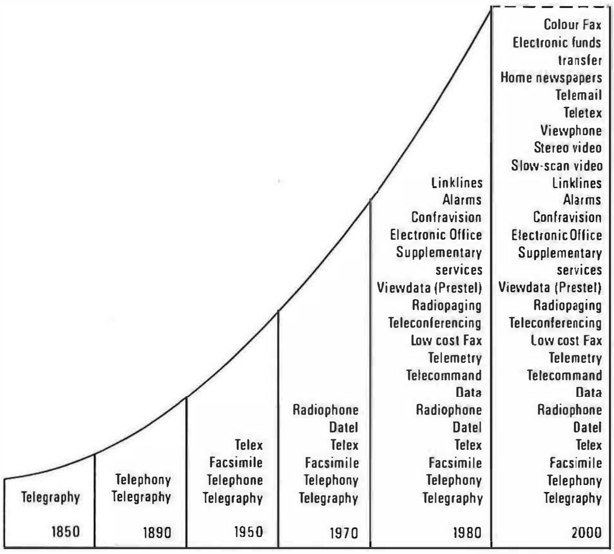

FIG. 4—Growth of switched and private services

FACILITIES AND SERVICES OFFERED

A fully digital network has the inherent capacity to provide a very extensive and diverse range of switched and private services (see Fig. 4). The ability to handle these services is essential to the future livelihood of any telecommunications business.

Up to now, requirements for additional services have generally been met by the creation of new physically-independent networks such as Telex and the packet switching network (Packet SwitchStream). A digital network offers the potential of a single flexible carrier able to handle all forms of information flow, allowing speedy provision and expansion of embryonic services while minimising the commercial risk inherent in the setting up of small separate networks.

The services and facilities that can be offered to customers can be broken down into two broad categories—voice and non-voice—each of which are capable of further subdivision. The degree of sophistication of service offered closely relates to the penetration of digital switching within the network, as a sample of the voice facilities that become possible clearly shows (see Fig. 5).

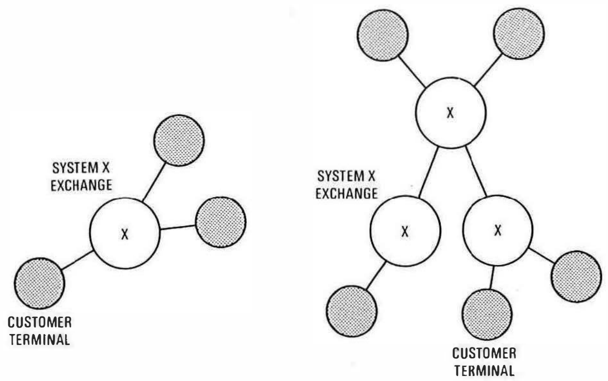

Facilities Available from Home Exchange	Facilities Available from End-to-End Digital Connection Between Two or More Exchanges
Charge advice	Ring back when free
Reminder call	Conference call
Call diversion	Auto credit call
Three-way call	Auto free phone
Code calling	Auto reverse charge call
Repeat last call	Remote control of call diversion
Call barring	Remote call forwarding
Call waiting	Universal access number
Selective accounting	
Diary service	

FIG. 5—Potential for services offered by a digital switching and transmission network

INTEGRATED SERVICES DIGITAL NETWORK

Non-voice services are expected to grow rapidly, but full exploitation of this market can be achieved only by extending the digital path, available between exchanges, through to customers' premises. This concept, on which the ISDN is based, opens up a wide range of possibilities (see Fig. 6).

BT's plans for a pilot ISDN service are about to come to fruition. The service will be centred on a switching node in London and will soon be extended to Birmingham and Manchester with line extensions to remote locations, where customers will be multiplexed onto a high capacity bearer. Although intended as a pilot service with defined coverage, it is capable of immediate and rapid expansion (dependent

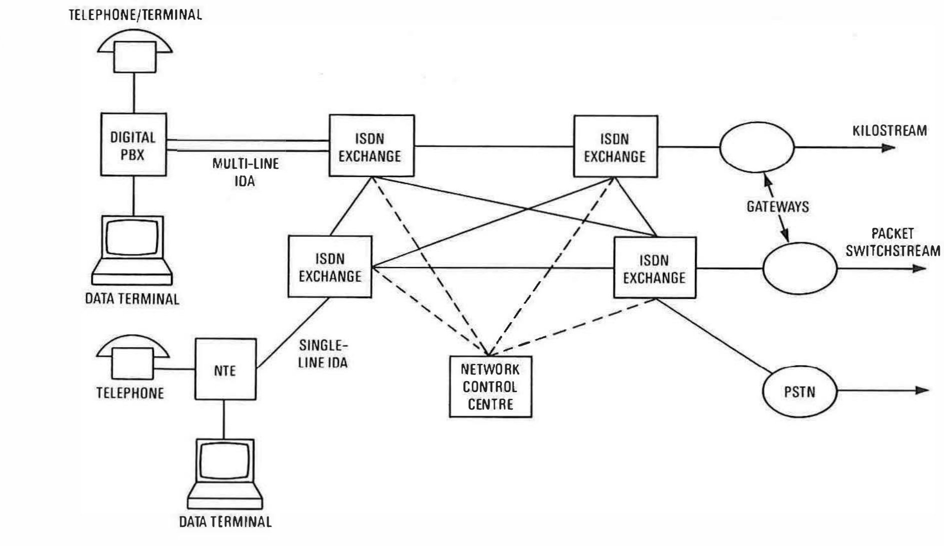

IDA : Intregrated digital access
NTE : Network terminating equipment
PSTN : Public switched telephone network
ISDN : Integrated services digital network
PBX : Private branch exchange

FIG. 6—Integrated services digital network

only on demand) to take in all System X exchanges.

The availability of a 64 kbit/s channel, customer to customer, opens up an enormous potential for new services; the examples given below are just a few that could be available at the outset.

Circuit-switched data (range of speeds up to 64 kbit/s)
Facsimile service at 64 kbit/s
Telex access at 2400 bit/s
Access to Prestel at 8 and 64 kbit/s
Private-circuit circuits
Access to packet-switching services
Slow-scan television
Digital telephony service
Electronic banking
Electronic funds transfer

It is probable that wideband applications such as local area networks (LANs) will soon require interconnection in one form or another. A high degree of standardisation in the interface between networks will need to be established to achieve the full potential of interconnection.

An important interface is seen to be with the digital PABX and, with the current increase in the rate of provision of such units, a significant demand for wideband local networks may arise within major business conurbations.

Agreed CCITT† Recommendations for interworking between the PABX and the exchange have not yet been established. In the absence of such standards, and until they are available, BT will establish support for an interim signalling system so that full customer–customer digital working can be introduced into the ISDN plan.

SERVICE BENEFITS

A digital network offers also the ability to provide more responsive management of service affecting faults. Automation of back-up transmission plant—the service protection network—and the development of network administration centres will allow quick reaction to loss of plant capacity.

The concept of an automated digital distribution frame (DDF) is an example which will allow rapid rerouteing of traffic by the diversion of blocks of digital bit rate capacity around fault locations. DDFs will be provided as flexibility points at various multiplexing levels, and will offer rerouteing at a number of standard bit rate interfaces.

Flexibility within the design of System X for automatic alternative routeing and a fully-interconnected trunk network will also ensure that traffic can be rerouted should a route fail.

† CCITT—International Telegraph and Telephone Consultative Committee

As well as plans for improved maintenance systems and equipment repair centres having modern sophisticated testing equipment, it will be possible to provide the effective rapid response to repairing faults that customers expect to be offered.

CONCLUSION

System X is at the focal point of plans geared to meet the future demand for telecommunication facilities. Its modular concept and the family of systems available, from small rural exchanges to large international gateway units, offer an economic and efficient solution to the problems of modernising a large, diverse and complex network. With the establishment of operational and maintenance procedures to manage and control the evolving network, the accompanying articles in this special issue of the *Journal* describe how BT is poised to commence the major build up of a System X modernisation programme.

ACKNOWLEDGEMENTS

The authors wish to thank Alan Bealby and David Holmes for their assistance in the preparation of this article and their work in co-ordinating the preparation of this special issue.

Biographies

John Boag joined the then Post Office in 1948 as a Youth-in-Training at Aberdeen. During the 1950s, he worked on the first transatlantic and Pacific submarine cables. In 1967, as a Senior Executive Engineer, he was seconded to COMSAT, in Washington DC, for two years to work on various INTELSAT satellite systems. He was promoted to Assistant Staff Engineer and was responsible for microwave radio-relay construction and then for network planning, a task which he continued as a Head of Division. In 1979, he was appointed General Manager of British Telconsult, and shortly afterwards as Director of the Overseas Liaison and Consultancy Department. He is currently Chief Executive of National Networks' Trunk Services. He is also Chairman of the Institution of British Telecommunications Engineers and Chairman of the Board of Editors of this *Journal.*

Brian Frame joined the then Post Office as a Youth-in-Training in the Circuit Laboratory in 1947. In the 1950s, he worked in the team developing subscriber trunk dialling and the transit network and afterwards spent some years in the development of analogue and digital switching and signalling systems. This was followed by a spell as a member of the Engineering Promotions Board before he took charge of the development of customer apparatus and coinbox equipment. In 1975, he was appointed Head of the New Exchange System Development Division with responsibility for developing national and international switching and signalling systems, and in 1980, as Deputy Director, he added maintenance and planning for switching systems to his responsibilities. He is currently Director of Local Exchange Systems with responsibility for the modernisation of local exchange switching systems. He is also a member of the Board of Editors of this *Journal.*

Evolution of System X—A Review

R. W. BRANDER, B.SC., C.ENG., F.I.E.E., F.INST.P.†, and P. J. BURVILLE, B.SC., PH.D., C.ENG., F.I.E.E., M.B.C.S.*

UDC 621.395.34

This article reviews some of the factors that have influenced the evolution of System X and shows how the original design of the system, based on a modular concept, has enabled advances in technology and changes in requirements to be incorporated easily into the system.

INTRODUCTION

This article gives a brief update on the development of System X exchanges and on their introduction into the British Telecom (BT) network. As has been recorded in previous articles[1] in the *Journal*, the strategic objective in developing the family of System X exchanges was to make available a range of units using digital techniques and common-channel signalling under stored-program control. Given the inevitable uncertainties about the future requirements for new services and facilities, the system design has been realised with the in-built capability to enable it to evolve to meet a variety of potential needs.

The basic feature that makes possible this evolutionary capability is the modular design of both the hardware and software elements of the system. These elements, or *subsystems* as they are normally called, are the building blocks which form the system architecture and, with proper control over the interface definitions, enable changes to be made to the separate parts of the system.

Clearly, there are limits to the evolutionary potential of a switching family, such as System X, and even to the desirability of incorporating unlimited options because of the consequent cost penalties. For example, the system is based on a real-time rather than a packet-switching concept. Also, the framework of the subsystems is pre-emptive in establishing boundaries which inevitably act as an impedance to changes to that framework.

This limitation, however, is not as restrictive as it may at first appear since, although the system is designed around a central processing utility, extensive use is made of distributed processing capability. The modular approach also offers the opportunity to exploit new technologies. This can be seen quite dramatically in the processor area, where the physical size has been reduced by a factor of 13 and its processing power increased fourfold from the original design[2].

In addition to the evolutionary capability, other aspects such as maintainability and expandability have been designed into the system. The philosophy on maintenance enables the exchange to provide diagnostic information which can be acted upon during operation by the system itself in managing its resources, and by the administration running the system.

The system design allows the exchange units to be extended, without any interruption to service, to provide increased capacity as well as the addition of new facilities to meet changing requirements. Whilst additional facilities are normally provided by loading a new software build, the ability to make hardware and firmware changes, should this prove necessary, is also built into the design. This feature obviously enables new technologies to be exploited in order to take advantage of improvements in design as well as in component performance and cost.

Finally, in this introductory section, it should be noted that the arrangements under which the system is being developed changed during 1982. Originally, with BT funding, BT, General Electric Company plc (GEC), Standard Telephone and Cables plc (STC) and Plessey Telecommunications Ltd. were four equal partners sharing the development work although, quite obviously, BT as the UK network operator, was primarily responsible for stating the system performance requirements. Under the new arrangements BT is still providing the funding and statement of requirements, but Plessey Major Systems Ltd. (PMSL) is the prime development contractor with GEC as the major subcontractor, whilst STC has withdrawn from the development.

CURRENT FAMILY

As forerunners to the start of the main installation of System X exchanges, which is a major element in the BT network modernisation programme, there have been System X exchanges operating in the public switched telephone network for several years. Since the showing of a local exchange at the Geneva Telecom 79 exhibition, both local and trunk units have been brought into service and interconnected by using the CCITT† No. 7 common-channel signalling system. Operations and maintenance centres (previously known as *local administration centres*[3]) have also been brought into service. These centres provide the main management tool for the exchange, one centre serving a number of exchanges.

In order to gain early in-service experience of some of the features of the second generation of exchanges, which features the new processor referred to above, a trunk unit was brought into public service in Coventry at the end of 1983. During 1984, both local and trunk second-generation exchanges were opened for service. These were single-cluster units, and will be extended to multi-cluster processor operation as the demand for increased capacity emerges in the future. The first full multi-cluster units are expected to be connected into the network later this year. A summary of these products and their target capabilities is given in Table 1.

EVOLUTIONARY FEATURES

System X has been the subject of evolution both in terms of its design and the way in which its development is being managed. A brief reference to the latter has been made above and more will be said later, but first the technical design is considered.

† Director, System Evolution and Standards, British Telecom Development and Procurement
* System Evolution and Standards Department, British Telecom Development and Procurement

† CCITT—International Telegraph and Telephone Consultative Committee

TABLE 1
System X Operating Objectives

	Termination Capacity	Switch Capacity (switched erlangs)	Processing Capacity (BHCA)
Concentrator:			
Telephony only	2 000	160	6 500
{20% Telephony {80% ISDN	2 000	160	8 000
Trunk exchange:			
Single-cluster processor	14 400	3 500	135 000
Multi-cluster processor	60 000	18 500	500 000
Local exchange: (including DPLE)			
Single-cluster processor	16 000	3 400	135 000‡
Multi-cluster processor	60 000	15 000	500 000‡

DPLE: Digital principal local exchange
BHCA: Busy hour call attempts
Note: these figures exclude overload capabilities
‡ For DPLEs, the effective BHCA depends on the amount of double processing as a result of network considerations. Assuming 20% double processing, the effective DPLE BHCA will be 128 000 for a single-cluster processor and 450 000 for a multi-cluster processor

Within the subsystems, which are the building blocks of the exchange, many advances have been made. Whilst the boundaries between the subsystems have remained basically sacrosanct, some modification has been allowed in order to better exploit developments and to accommodate expected trends. The architecture of the system has not changed, but there has been convergence of design, in some instances, of certain parts of the system, in the different exchange types.

Thus, so far in the life of System X, where there has been evidence of a tendency towards diversification to exploit new technologies, this has been more than balanced by the move towards common designs for different applications. An example of this convergence is the use of the enhanced processor, which serves the whole range of exchange products, rather than two quite different designs being used which was the case in the first generation of exchanges.

There have been several sources of trends which have led to the controlled change to interfaces both within the system (between subsystems) and within the network. The range of these sources includes changes in the network itself (for example, the use of digital principal local exchanges), the provision of new facilities (see Fig. 1), and improvements to the system design as described above. There appears to be little evidence to suggest that the future will be any more stable. With the development of the integrated services digital network (ISDN)[4] and further exploitation of the common-channel signalling system, the ability to adapt System X will doubtless be utilised.

One example of the development currently being scheduled into the programme is provision of a wideband digital switching facility, which will give a switched-circuit capability considerably in excess of the current 64 kbit/s.

ENHANCEMENT

In the current phase in the life of System X, facility enhancement is achieved by providing new software builds with only very minor hardware changes. These software upgrades are made on-line by carrying through the information on the current state of the exchange in terms of calls in progress etc. from the old to the new build.

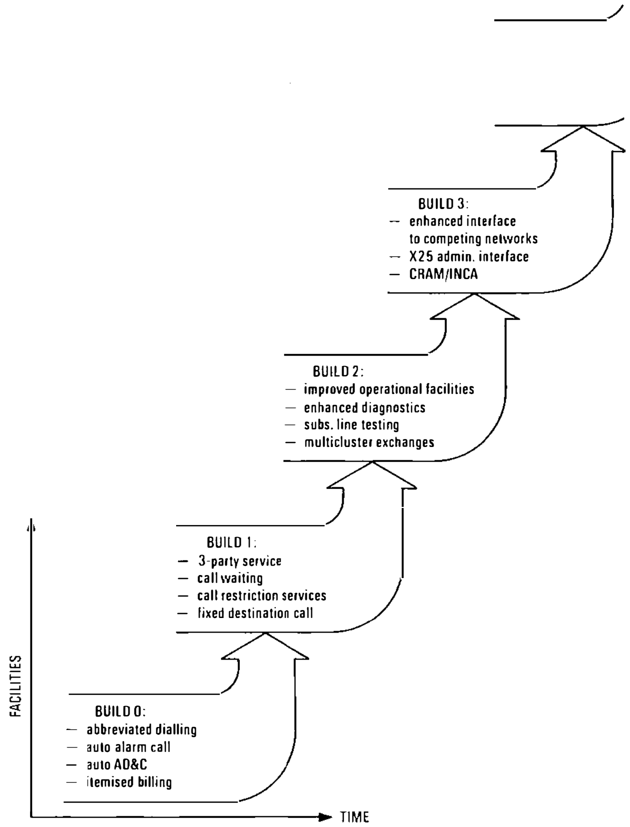

CRAM : Call revenue apportionment measurement
INCA : Internetwork call accounting

FIG. 1—Facility upgrades for System X

The facility upgrading sequence between builds for a local exchange is illustrated in Fig. 1. It will be seen that the facilities relate to both those offered to the customers and to the adminstration running the network.

In the Fig. 1, the original build '0' is upgraded through various builds with new or improved facilities being added on each occasion. The three builds shown as examples can be regarded as the drops which have gone into the System X exchanges over the past three years. It will be noted that the eight star services were added in two phases[5].

SUMMARY AND FUTURE PROSPECTS

This article has indicated the way in which System X has been evolving in the recent past and the forces which have led to these changes. The expected requirements of the future have been anticipated in the sense that 'hooks' are being designed into the system and these will ease the introduction of new facilities. Obviously, not all future requirements can be catered for in advance, but the framework for change is well established.

In addition to what can be referred to as the micro-forces influencing the system design, such as network requirements, facilities and exploitation of improved components, there are also macro-forces whose effect could be far more radical and have an impact on the system architecture in the longer term. Amongst those which spring to mind are technology, competition, social conditions and politics.

Some changes have already been described which have been the direct result of technological improvements, but radical evolution might be required for these macro-changes. An obvious example is the development of an optical switch which could offer the possibility of using optical techniques right through substantial parts of the exchange. The modularity of the exchange should facilitate the incorporation of optical techniques as they become cost effective. Indeed, the use of optical transmission links is an option which has already been developed[6].

Direct competition between BT and others to provide customers with telecommunications services, at some time in the future, could undoubtedly lead to further exploitation of the network at local level. Also, a possible factor which might require restructuring of the large central switching unit is its vulnerability to attack and disruption by minority elements in the community.

Factors of this nature could be accommodated by having an 'organic' or 'distributed' network evolution in which primary switching is carried out at the periphery of the network at say the cabinet level. Thus, the local community, within which much of its own traffic is actually confined, would be interconnected by a local switch and not make demands on the central switch or the associated transmission plant for such locally targetted calls.

Also, given that telephones increasingly will have in-built intelligence, the loss of access to certain central support facilities (such as short-code dialling) may be no great loss to the customers. However, it is recognised that increasingly sophisticated services, covering a wider range of features including banking, will be demanded of the future central exchanges.

Given these points, the prospect exists for simplifying central switching units and, in particular, reducing the size of the software needed to run the systems. The latter feature is attractive both from the cost of provision point of view as well as the reduction of the maintenance and enhancement burden. System X has this option built into its modular structure.

On the subject of the management of the development, the major change has been PMSL taking over the role of prime contractor to BT for the current phase of System X design, thus giving the manufacturing organisations, PMSL and GEC, greater responsibility for the development. This has made taking initiatives and introducing new ideas much easier, but without in any way reducing or weakening the requirement to meet the original design objectives.

With this new level of responsibility, the contractors are in a better position, when designing and manufacturing the products, to respond to the world market for public switching systems, while at the same time meeting BT's requirements.

It would, of course, be quite easy to develop a scenario in which the trend for switching units was rather different to that given above. However, the main point to be made is that switching-system, or indeed transmission-system, design is not set in concrete, and the network designers—or rather network evolvers—of the future will have to respond to forces which at the moment can only be the subject of conjecture regarding their nature and strength. The design of the System X exchange is well placed to meet the wide range of demands which might be placed on it in the evolving network.

References

[1] MARTIN, J. System X. *Post Off. Electr. Eng. J.*, Jan. 1979, **71**, p. 221.

[2] TROUGHTON, D. J. *et al.* System X: The Processor Utility. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 226.

[3] WALIJI, A. A. Architecture of System X. Part 4—Local Administration Centre. *Post Off. Electr. Engs. J.*, Apr. 1980, **72**, p. 36.

[4] PRICE, C. D. E., and BOULTER, R. A. Integrated Services Digital Network. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 311.

[5] STREET, F. P. To the customer's advantage. *Br. Telecom J.*, 1981, **2**(3), p. 4.

[6] HARRISON, F. G., and LE GOOD, R. K. In-Station Cabling with Optical Fibres. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 259.

Biographies

Bob Brander spent 15 years working on semiconductor materials and device research at GEC's Hirst Research Centre at Wembley before joining British Telecom Research Laboratories in 1972. Initially, he was responsible for the fabrication of high-reliability transistors for undersea telephone cable systems, and subsequently for device research for advanced systems such as optical-fibre links. In 1981, he was appointed Deputy Director of Research for Advanced Technology, with responsibility for micro-electronics, optical communications and component reliability. In May 1984, he was appointed Director of System Evolution and Standards, and is responsible for research and development activities relating to the design, operation and management of the rapidly expanding telecommunications network, and for the interface standards and protocols necessary to ensure satisfactory interworking.

Peter Burville joined BT as a Youth-in-Training and spent some time in the Canturbury and Southend-on-Sea Areas before moving to the Engineer-in-Chief's Office in 1960. After working on submarine cables, computer projects, network planning and teletraffic studies, he joined the System X development team in 1978. He is currently Head of the Switching Development (Management) Division of the System Evolution and Standards Department.

System X: The Processor Utility

D. J. TROUGHTON, B.SC., M.PHIL., C.ENG., M.I.E.E.†, A. P. LUMB, B.SC., A.I.M.E.E.†, R. C. BELTON, C.ENG., M.I.E.E., M.B.C.S.*, I. D. GALLAGHER, B.ENG.*, M. D. BEXON, B.SC(ENG)., M.SC., A.C.G.I., C.ENG., M.I.E.E.*, S. R. MOOR, B.SC., A.M.I.E.E.*, and S. C. J. STEGMAN, A.M.I.E.E.*

UDC 621.395.34 : 681.31

This article describes the second-generation processor utility developed for the System X stored-program control switching units. It deals with the main features of hardware and software design and operation and includes operating system, man–machine interface and reliability mechanisms.

INTRODUCTION

A previous article in this *Journal*[1] described the System X processors, the Release 1 Processor Utilities (known as *POPUS 1* and *PPU*)‡, and mentioned the necessity for the processor utility (PU) to be developed in conjunction with that of the overall system. This article describes the reasons for the development of the new processor, and discusses its design. Fig. 1 shows the PU in a TEP 1(H) rack.

The reasons for the ongoing development were twofold:

(*a*) The store size and the instruction rate had to be increased as the total system development gained momentum, as the amount of software required became known and as the instructions to be processed per call were identified.

(*b*) The size and cost of the PU could be reduced by the use of new technology and revised design rules.

† GEC Telecommunications Ltd.

* System Evolution and Standards Department, British Telecom Development and Procurement

‡ A glossary of terms and abbreviations used is given as an appendix to this article

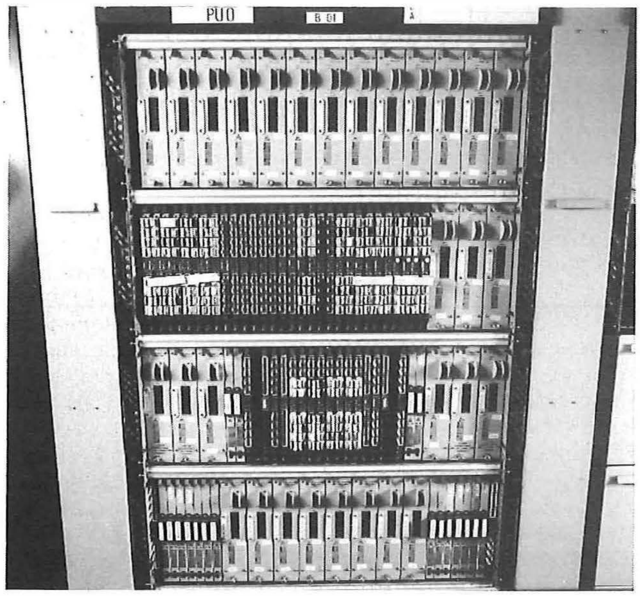

FIG. 1—The processor utility

Constraints: Only one major constraint existed—that relating to the instruction set. A particular requirement was to ensure that the new machine would be capable of supporting the existing System X applications software suites. These had taken a great deal of time and effort to produce, were well understood, and were supported by a large software development facility. The decision was therefore taken that the processor arising from the ongoing development should provide the required software facilities. The high-level language required the instruction set developed for POPUS 1. The new machine provided this, together with three new instructions to assist the high-level language compiler and software designers. It would also be able to process, with minimal changes, the applications software produced for use with POPUS 1.

New technology and rule changes: The new technology used included the improved integrated circuits that became available; for example, the large dynamic random-access memory (RAM) chips and magnetic bubble memory parts. The rule changes were mainly of a detailed nature, such as track spacing and new edge connectors. Two fundamental changes were agreed which did most to produce the new machines; these were the adoption of fan cooling, and the use of printed-wiring boards with up to eight layers.

The most important requirement was reliability. The machine has to work for many years continuously without faults occuring that cause complete switching-system failure. This requirement is therefore reflected throughout the design.

This article gives some details of the self-checking and report-generating features, including the restart phase, that allow for the automatic reconfiguration of processes or resources within the PU to enable it to provide continuous service during fault conditions.

SYSTEM OVERVIEW

Hardware Configuration

The processor is configured so that up to four central processor units (CPUs) can be coupled together to form a multi-processor. Each CPU is free to be allocated any processing activity within the exchange, without a specific function being given to any particular one. This is achieved by giving each CPU shared access, on an equal priority basis, to all storage and peripherals in a tightly coupled configuration, known as a *cluster*.

A number of separate multi-processing clusters can be linked together to achieve extra power for large exchanges. In this configuration, each cluster is independent of the others, except for a limited amount of intercommunication

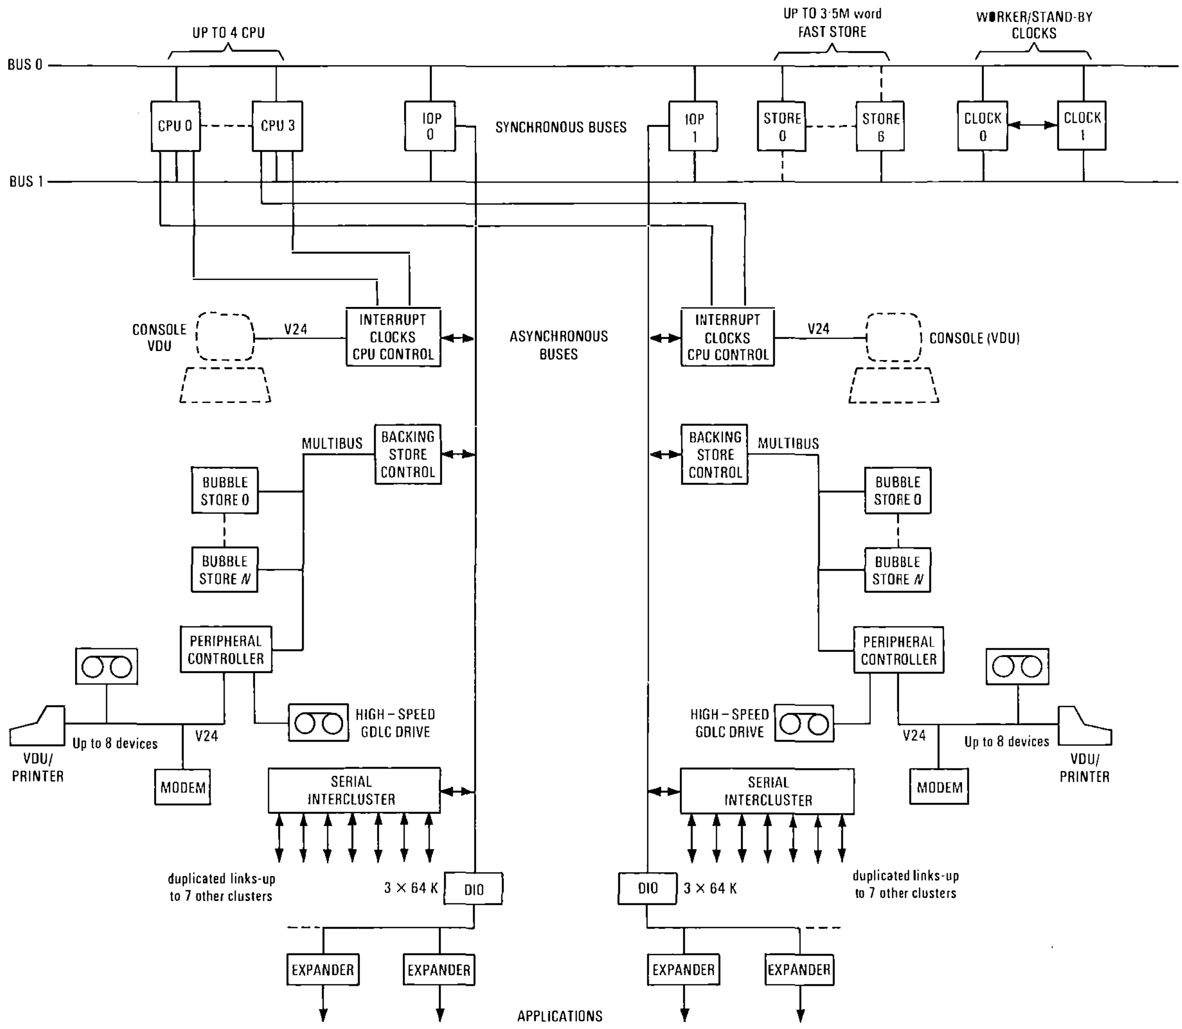

FIG. 2—Basic modules of a processor cluster

between the executives or between the application software suites. This intercommunication is kept to a minimum by a particular cluster being dedicated to a specific part of the exchange.

The basic modules of a processor cluster are shown in Fig. 2.

The overriding principle in the design of the hardware is reliability and every functional unit is duplicated as a minimum. In normal operation, all security units are in use on a load-sharing basis; the exceptions to this are the interrupt unit and the synchronous clock units, which are operated in a WORKER/STAND-BY mode with a routine change-over to minimise the effect of latent faults. The processor is dimensioned so that it can carry the normal exchange load with any faulty security unit out of service. A rigorous boundary is maintained between security units to prevent faults or errors being spread. Error propagation is reduced by employing error detection on information passed over the boundary between security units. If any part of a security unit is found to be faulty, the entire unit is taken out of service. It can then be powered down in order to allow faulty slide-in-units (SIUs) to be replaced.

Software Structure

The software is divided into the operating system and the application software, which is further partitioned into a number of functional subsystems such as call processing, call accounting or message transmission.

The operating system can be viewed at two different levels: the interface with the application software; and the interface with the processor hardware and internal workings of the system.

The most important role of the operating system is the allocation of processor resources for use by the application software. For example, this includes the allocation of CPU time to a program and the allocation of space in main memory for its code and data. Other examples are the allocation of input/output (IO) devices for man–machine communication, backing-store space for information storage and retrieval, and secure data areas for long-term preservation of information.

The various functions of both the operating system and the application suite are partitioned into separate units of software called *processes*. The process boundaries are primarily selected on a functional basis, but they also ensure that communication between processes is kept to a minimum. Communication between processes is by means of tasks routed by the operating system or by the use of shared store.

A process normally has a particular function or activity to perform within the exchange, such as call control, switching or maintenance. The process has its own private data area but, in some circumstances, data can be shared with another process. Each process has a unique identity and has associated with it a priority and a state.

Structurally, a process is a collection of procedures, which specifies sequences of operations to be carried out one after another. As far as the operating system is concerned, a process is represented by its process descriptor block, which is the means by which it can schedule and invoke the execution of procedures within the process.

Each process usually alternates between periods of independent activity, that is, executing code on a processor, and communication, that is, passing tasks to another process.

Because processes have to share time on the same processor or processors, one process has to be protected from another. Each process has to be regarded as a secure and independent unit so that another process cannot access its private data or start executing its code. This protection, in the form of virtual memory, is built into the processor hardware and the operating system.

PROCESSORS AND MAIN MEMORY

At the centre of the processor hardware are up to four identical CPUs sharing a pair of busses for access to 3·5 Mwords of main memory. The CPUs and memories are clocked synchronously from a single, secured clock source to permit a simple high-speed bus protocol.

The average instruction execution rate for the complete processor is about four million instructions per second.

Synchronous Bus

The two synchronous busses provide high-speed communication between the CPUs and memory cards (see Fig. 2). Under normal operating conditions, the load is shared by both busses, but if a fault occurs all activity can be switched to the fault-free bus.

Each bus consists of 60 signals, of which 22 are address bits and 16 are data bits, and is implemented by using tri-state devices driving unterminated etched tracks in an 8-layer backplane 610 mm long.

The basic clock frequency is 8 MHz, so that any transfer of data on the bus is allocated a 125 ns time-slot. Alternate 125 ns time-slots are referred to as A and B phases; an A phase followed by a B phase represents one 250 ns CPU cycle. Fig. 3 illustrates the basic bus protocol. During the CPU cycle that specifies a memory access, the A phase is used to check that the memory card is not busy and the B phase to arbitrate with other CPUs for the bus. If the memory is free and the CPU wins the bus, the memory address is driven onto the bus in the next A phase. The memory card staticises the address at the end of that phase and then returns the data in the next but one B phase.

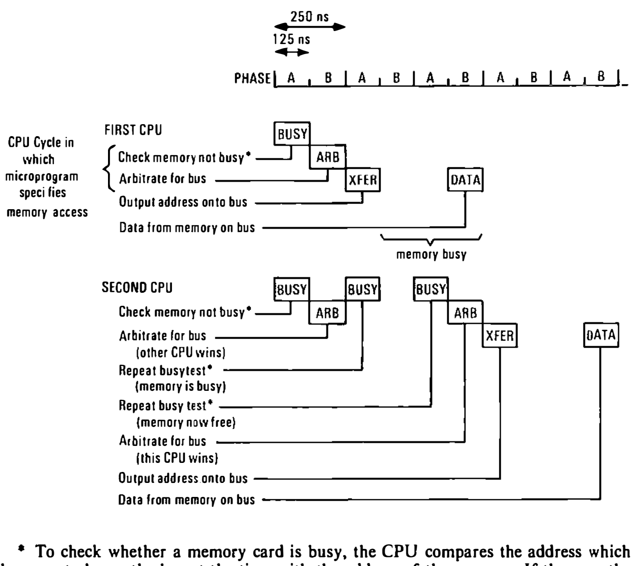

* To check whether a memory card is busy, the CPU compares the address which happens to be on the bus at the time with the address of the memory. If they are the same, then there is a transfer to the memory in progress and it would be busy if the CPU attempted to access it. If they are different, the memory is free.

FIG. 3—Synchronous bus protocol

During the access, checks are made on address and data parity, command validity and to ensure that only one device is driving the bus at any time. Any errors in the bus, or in the contents of the memory are signalled back to the CPU at the end of the access.

Synchronous Clock

Synchronism between devices connected to the synchronous bus is maintained by a single, secured clock source provided by two clock cards, which operate as a worker/stand-by pair. The three signals generated by the clock are:

(*a*) a 125 ns clock signal,
(*b*) a 250 ns phase signal, and
(*c*) a refresh signal to instruct the memory cards to refresh their dynamic RAMs (one row every 15 μs).

In addition, a register on each clock card is used to control the allocation of memory cards between the two busses, and their IN- or OUT-OF-SERVICE state. These registers may be read from or written to by a CPU using the synchronous bus. Each clock card monitors its own outputs and the clocks decide between themselves which should be the worker and which the stand-by. CPUs and memories automatically select the output from the worker clock. The error checking logic in each clock is periodically routined by the operating system using the synchronous bus to access a clock and inject faults into it.

Memory

Up to seven memory cards, each having a capacity of 512 k × 16 bit words, making 3·5 Mword in all, can be connected to the synchronous bus. The memory devices used are 64 or 256 kbit dynamic RAMs, which are all refreshed simultaneously under the control of the clock card.

The memory is configured to allow double-word reads and single-word writes. For double-word reads, data from an even address is returned down the data bus and data from the following odd address is returned down the address bus. Thus, performance is improved by reducing the bus occupancy needed for instruction fetching.

Each memory card has an interface to both busses, although only one bus at a time is looked at under the control of a signal from the clock card. In the event of a bus fault, all the memory cards are allocated to the other bus, although normally the load is evenly spread to minimise bus contention.

Central Processor Unit (CPU)

Fig. 4 is a simplified block diagram of the CPU showing its main functions. Each CPU occupies four cards, and up to four CPUs may be equipped.

At the centre of the CPU is a 24 bit arithmetic logic unit (ALU) whose operands are provided by the A and B busses. The main sources for the A bus are the 32 registers in the register files, a programmable ROM (PROM) containing various masks and constants called ***wired-in*** data, and a register containing the current instruction. The main sources for the B bus are the 32 registers, and a register containing the latest data retrieved from memory. The data on the B bus can be shifted up to 16 bits in a barrel shifter and then have various bits masked out before being input into the ALU. The ALU output onto the output bus (O bus) can be written to the file registers or into the 20 bit virtual-address registers.

The processor uses a virtual-memory scheme based on 256-word pages. Therefore, for each memory access, a translation must be carried out to convert the virtual-page address (held in the address registers) to the physical address to be output to the memory. The look-up tables for this

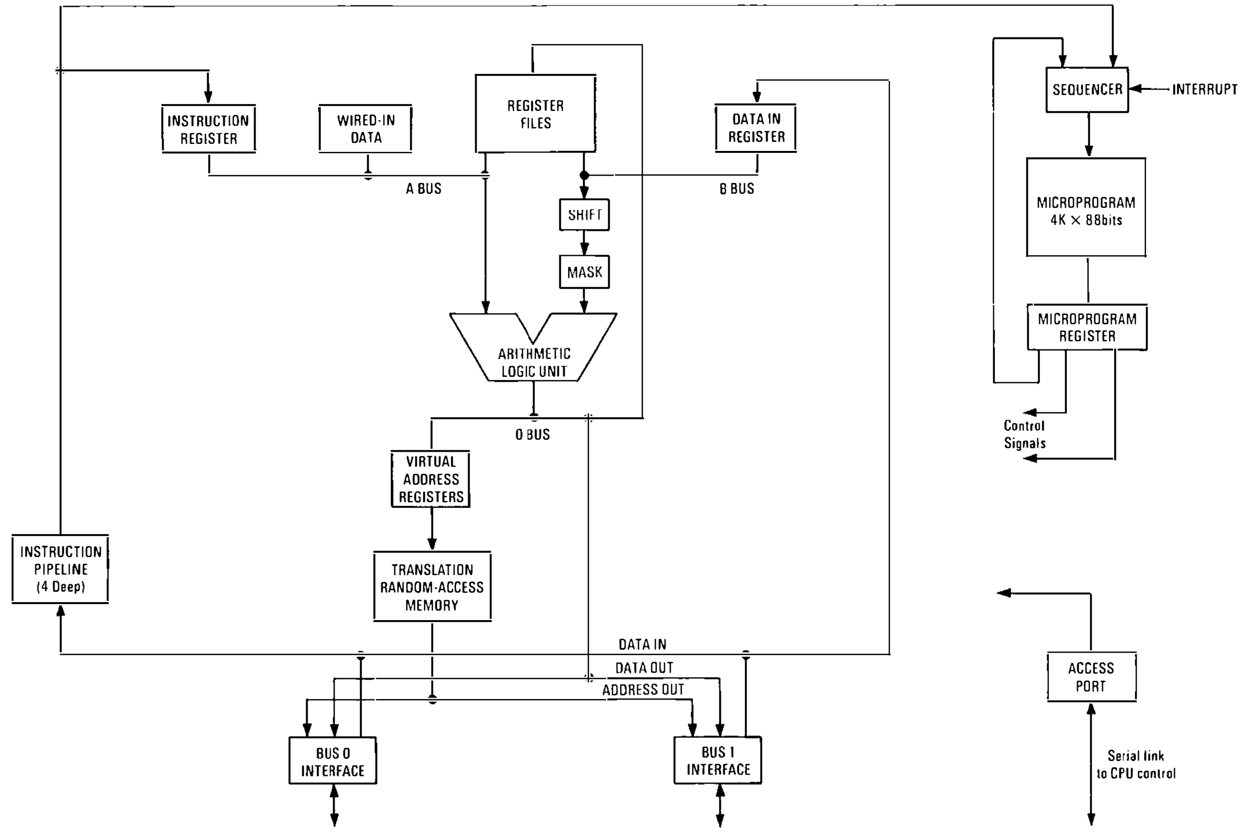

FIG. 4—Simplified block diagram of the CPU

translation are held in the main memory. To prevent these tables having to be read for every memory access, a translation RAM in each CPU is used to hold translations for recently used pages. Up to 1024 virtual-to-physical translations may be held at a time.

When instructions are fetched from memory, they are queued in a 4-deep instruction pipeline which they advance through as the instructions ahead are executed. The next instruction is clocked into the instruction register and used at the same time by the sequencer to derive the start address of the microprogram required to execute it. The first micro-instruction is then clocked into the micro-instruction register.

The microprogram is 88 bits wide, and provides control signals such as register file addresses, ALU functions, and requests for memory accesses. One micro-instruction is executed every CPU cycle of 250 ns, and 4 K of micro-instructions are available to implement the instruction set, access port commands, and the process allocator (PA) part of the operating system.

The times for some typical CPU functions are as follows:

Register-to-register operation	1 cycle
0–16 bit barrel shift	1 cycle
Address calculation (base + offset)	1 cycle
Address calculation (base + offset + index)	2 cycles

The normal flow of instruction execution may be changed by the arrival of an external interrupt which causes the sequencer to jump to the interrupt handling section of the PA microprogram.

Faults internal to the CPU are detected by parity error or by time-outs. When a fault occurs, the CPU stops and opens its access port; this allows other CPUs to read its internal registers, carry out tests and, if required, return it to service.

INPUT/OUTPUT PROCESSOR (IOP)

The input/output processor (IOP) interfaces the CPUs to the asynchronous modules in the processor. The IOP is accessed from a CPU in the same way as a memory except that after accessing an IOP, the CPU enters a WAIT state. The CPU is restarted by a command from the IOP upon completion of the asynchronous access. The IOP can make read, write or read/clear accesses to memories on the synchronous bus; this enables the IOP to provide a direct memory access (DMA) facility for the backing store. The IOP has 8 kwords of ROM. This contains program and data which has to be secured against corruption; for example, cluster restart procedures.

The IOP is normally scanning for requests either from the synchronous bus or from the backing-store controller (BSC). These requests are treated on a first-come first-served basis; but if simultaneous access occurs, the access from the synchronous bus is given the highest priority. The IOP is allocated to one of the synchronous busses by a duplicated signal which comes from the synchronous clock card. The IOP responds only to accesses along the correct bus. Upon receipt of a *synchronous bus request*, the IOP initially performs a parity check from the received address. If this check is satisfactory, then address and data are loaded into a first-in-first-out (FIFO) register. This FIFO is 8 words deep, which allows it to store up to two accesses from each CPU in the system. The address data and command information are removed from the FIFO by the asynchronous bus control logic, which performs the asynchronous access down the IOP bus and initiates the IOP write back along the synchronous bus when a response has been received.

The protocol on the asynchronous bus takes the form of a *request/proceed* handshake. The IOP puts an address (plus data in a write access) onto the bus and then sends a dedicated *request* signal to the module being accessed. When the module has performed the function it was instructed to carry out, it sends a *proceed* signal back to the IOP and puts data onto the bus if a read access has been requested. On receipt of the *proceed* signal, the IOP accepts the data and removes its *request* signal. The access finishes when the addressed module, seeing the *request* signal going away, removes its *proceed* signal.

The response from the asynchronous module is parity checked, and subject to a 20 μs time-out. If the parity of the received data is incorrect or if the module does not reply in time, a fault response to the CPU is generated and fault data is returned. Asynchronous modules on the IOP bus can return a fault proceed and fault data if they have detected a fault. If this happens, the IOP generates the fault response to the CPU and forwards the received fault data unchanged. Fault data is, of course, still subject to the normal parity checks in the IOP.

INTERRUPT UNIT AND CPU CONTROL UNIT

Interrupt Unit

The interrupt unit contains the logic required to inform quickly the operating system of external events in situations where the IO polling techniques are not considered to be appropriate. The interrupt unit contains a register (known as the *LCPU* register) that holds the identity of the CPU currently running the lowest-priority process, so that interrupts can be steered into the multi-processor in such a way as to cause minimum disruption to the overall processing throughput. The LCPU register can be loaded via IO command and is regularly changed by the operating system as processes are rescheduled.

The interrupts fall into two categories:

(*a*) *immediate interrupts*, which are communicated to the LCPU with the minimum of delay, and

(*b*) *deferred interrupts*, which do not cause immediate interruption to processing, but which are stored until an immediate interrupt occurs (within 10 ms maximum).

Upon receipt of an immediate interrupt, the interrupt unit sends a signal to the lowest priority CPU and sets a bit in a status word to record the interrupt that has occurred. Upon the receipt of the *interrupt* signal, the CPU will TRAP on execution of the current instruction and this activates the PA, which interrogates the interrupt unit through the IOP to see which interrupts are set.

To ensure the security of the PU, the interrupt unit is duplicated; each duplicate is accessed via a different IOP and each has its own power supply. Each hardware module that can generate interrupts sends two independent signals, one to each. The duplicates are operated in a WORKER/STANDBY mode and both receive the interrupts, but only the worker generates interrupts to CPUs.

The interrupt unit also contains the system clocks. Three types of clock are provided:

(*a*) the real-time clock, which generates an immediate interrupt every 10 ms,

(*b*) the digital clock, which records the time of day, date and the year, and

(*c*) the pseudo-random clock, which generates an interrupt with variable mean time.

The digital and pseudo-random clocks are implemented as routines on a microprocessor within the interrupt unit. All these clocks derive their timing from a triplicated 2·048 Mbit/s signal which the processor receives from the synchronisation utility. A two-out-of-three majority vote is taken on the three clocks.

CPU Control Unit

The CPU control unit allows access to the CPUs via the access port to enable interrogation and control. The CPU control unit can be accessed either by the IO from a CPU, thus allowing one CPU to interrogate or control another, or from a visual display unit (VDU)/console to allow an operator to control the CPUs.

The access port consists of a 2 Mbit/s serial datalink. All the functions of this unit are controlled by a microprocessor which performs checks of validity on data received from a CPU and which formats the data for presentation to the VDU.

The CPU control unit also contains the hardware alarm collection logic; this sets a flag in an IO location, which can be polled by the operating system, under certain alarm conditions; for example, power supply unit failures.

BACKING STORE AND PERIPHERAL CONTROLLER

Backing Store

The backing store contains all the code and data for the system. This includes the information needed to control a telephone exchange, and the data which must not be lost if the power supply fails; for example, subscribers' call records. If a failure occurs in the processor, the PU can be restarted automatically, by reloading the software from the backing store, and be functional within seconds. The backing store is duplicated for reliability.

The backing-store information is stored in magnetic bubble devices mounted on bubble memory units (BMUs). Each BMU has 8 × 1 Mbit devices, providing 512 kword of secure storage. The backing store controller (BSC) controls up to nine BMUs, including a spare which can be automatically brought into service if a failure occurs. This further increases the reliability of the backing store. The BSC has an Intel 8086 microprocessor to control its main function, which is to transfer information between the main store, the BMUs and the peripheral controller (PC). To achieve this as quickly as possible, the BSC has a DMA unit which can transfer data between its own buffers and the main store or between the main store and the PC buffers. To do this it has to be able to control the asynchronous highway and the Multibus® highway and can be master on both. Since the only other master on the asynchronous highway is the IOP, the arbitration logic is simple, but on the Multibus® all the BMUs are also masters, so more complex arbitration has to be used.

The BMU also has a DMA facility to control the flow of information to and from the bubble devices. Information is normally transferred on a page basis, and each page of information can be secured in the backing-store. The following precautions are taken to provide security of information:

(*a*) error correction and detection is carried out within the BMU by using microprocessor logic,

(*b*) the page identifier within the BMU is compared with the page identity requested,

(*c*) cyclic redundancy checksums (CRCs) are used in the BSC to secure the path between the BMU and BSC,

(*d*) the file and process identity is compared within the BSC and with that requested by the storage allocator (SA) process, and

(*e*) a parity check is added by the BSC to secure the asynchronous highway and main store.

Multibus® is a trade mark of the Intel Corporation (UK) Ltd.

These precautions are taken with every page of data in the backing store, but some sensitive data will require further security and this is provided within the software.

When the backing store is running normally, the BSC continuously reads a portion of the main store to look for requests from the SA process. A single-word pointer is used to optimise performance and to indicate that a request is outstanding, and where in the main store the details of the request are stored. Each BMU is divided into eight sectors, with each sector having a single-word pointer allocated to it. If there is data to be transferred to or from a BMU, then the BSC collects the additional information from the main store and sets up a request for the appropriate BMU. The request may be to transfer a page of data from the main store to a BMU, in which case the BSC transfers the page to its own buffers before sending the request to the BMU. When the BMU successfully completes the transfer, the BSC sends a completion message to the SA. Alternatively, to transfer data from the BMU to the main store, the BSC gives the BMU details of the area of the BSC to which the page should be transferred and, when the BMU has finished, transfers the page to the main store and sends the completion message.

When dealing with the PC, the BSC gets its tasks from the PC buffer store; this is arranged so that both the PC and the BSC can read from and write to it. Any transfers of data are done by the BSC between the main store and the shared RAM without the data first being buffered in the BSC.

Peripheral Controller (PC)

So that peripheral equipment (for example VDUs) may be connected to the PU, a PC is provided and this terminates the Multibus® within each backing store. The PC takes advantage of the BSC's ability to gain control of the asynchronous bus from the IOP, and transfers large amounts of data directly from the main memory to the PC's internal buffer (see Fig. 5). Additionally, it is well positioned to allow direct transfer of the contents of the bubble memory to and from a fast external cartridge unit.

The PC is equipped with nine peripheral ports, eight V24 interfaces and one V11 interface known as *GDLC*. The V24 ports, under the control of a microprocessor system, can service keyboards, cartridges and modems at a variety of speeds between 110 baud and 9·6 kbaud. In addition, the GDLC port drives a balanced line link that runs at 880 kbaud and is connected to a high-speed cartridge unit.

The transfer of data between V24 peripherals and software takes place via a set of input and output buffers that are located in the main memory and pointed to from the control page. The PC uses control-page messages and interrupt signals to request DMA transfer of data to its own buffer from the output buffers and vice versa when data is input. Similarly, data can be archived by direct transfer from the backing store to the GDLC port as a facility whilst the PU is ON-LINE.

The GDLC port also allows the complete exchange software to be loaded in less than 15 minutes; for example, for a new exchange installation. To achieve this without the operating system, the PC firmware controls the flow of data from the high-speed cartridge unit to the Multibus®, and requests the BSC to transfer it into the backing store.

The PC has a self-test facility that ensures its hardware is operating correctly before it is put into service. If during normal operation a fault is detected at a port, a message is conveyed to the software. A simple error code is written into the control page and greater detail is presented in the fault page, which is also in the main memory. For some peripheral faults such as bad parity from cartridges, the PC requests a repeat attempt which, if successful, generates a transient fault report. In the case of serious faults from which recovery is not possible, the PC withdraws itself from service and generates fault interrupts.By using these mechanisms, the recovery process can take the appropriate action.

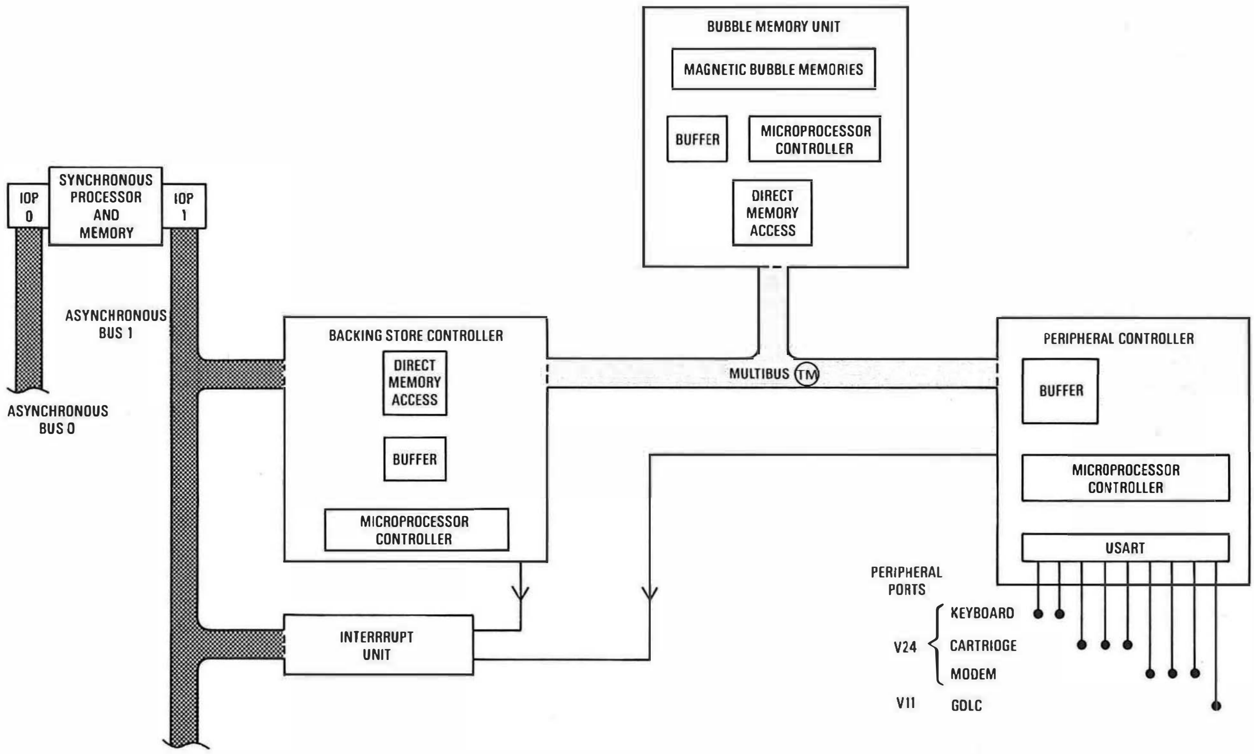

USART: universal synchronous/asynchronous transmitter/receiver

FIG. 5—Backing store and peripheral controller

APPLICATIONS AND INTER-CLUSTER INPUT/OUTPUT

Direct Input/Output (DIO)

The PU provides parallel connections to the rest of the exchange by expanding and extending the asynchronous busses through the DIO. A portion of the PU's physical address space is allocated to IO, and applications hardware units (AHUs) can be accessed by using the DIO simply to read and write as if to memory; this is known as *memory mapped* IO.

Fig. 6 shows a simple block diagram of the IO, which is duplicated on each asynchronous bus. The applications interface consists of a balanced line highway originating and terminating on an IO expander, 'daisy chaining' up to eight eight AHUs on the way. Each loop may be up to 150 m long. For security, AHUs are dual ported with connections to the IO expanders in each DIO duplicate. This ensures that faults, maintenance action and IO additions do not interrupt service. An IO expander is not physically part of the PU, but is a separate shelf located alongside the AHUs and up to 150 m distant from the IOP/IO interface of the PU. Three IO expanders can be linked by using the balanced-line IO expansion interface from the PU.

Each IO expander occupies a 64 kword† segment of the IO address space; 62 applications each of 1 kword are supported by an IO expander; the remaining 2 kword are reserved for testing purposes.

An access through the IO to the AHUs takes place under the control of the IOP. Asynchronous bus protocols extend out to the AHUs and include the fault detection mechanisms of the bus such as parity checking. If any fault is detected by the DIO, then the current access is frozen to allow the IOP to time-out. If an AHU detects a fault, then both time-outs and fault proceeds returned to the IOP may be used; this causes the CPU to TRAP.

A set of test registers is positioned to locate faults throughout the IO and extending to the far termination of the applications interface. Each register has two addresses, one in the top and one in the bottom 1 kword of each segment. Three tests can be conducted with each register to establish

(*a*) that there is a working IO path to and from the register,
(*b*) that the databus is fault free, and
(*c*) that the IOP parity checking mechanism is functioning correctly.

These simple tests enable the recovery process to locate faults within the PU's IO or to the AHUs, and ensure that the appropriate resources are withdrawn from service.

† An alternative implementation is half size and utilises only 32 kword out of each segment.

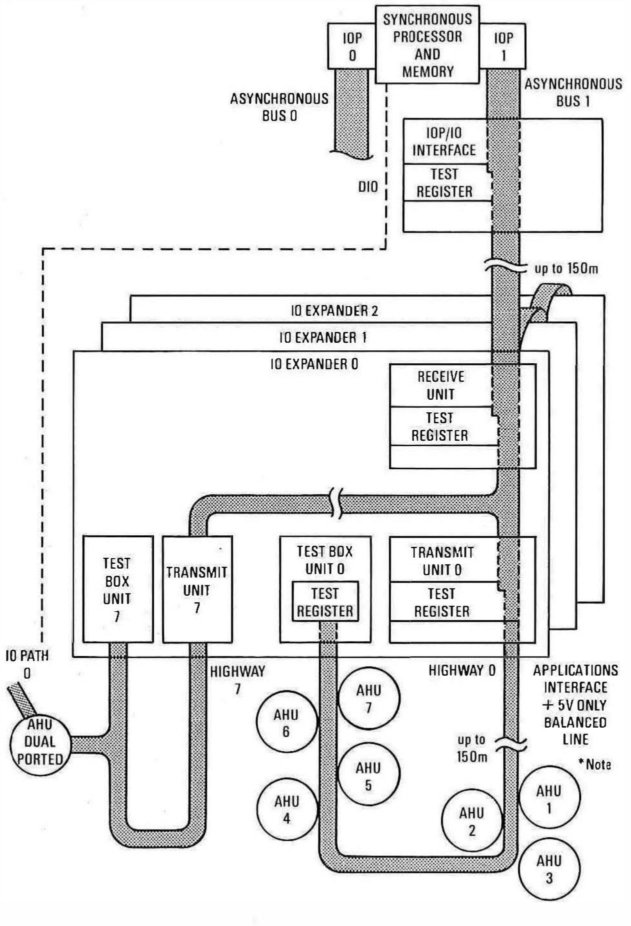

Note: There is no AHU 0 on highway 0 nor AHU 7 on highway 7

FIG. 6—Block diagram of the input/output

Inter-Cluster Communication

The inter-cluster communication controller logic (ICCL) is provided to facilitate the passing of tasks between processes residing on separate clusters under the control of the PA. An essential feature of this communication is that any additional processing by the PA must be minimised and that it must not introduce undue task delays. To achieve this, a simple PA/ICCL interface using barrel queues and fast serial links at 1·3 Mbit/s directly connecting all clusters is provided.

An outline diagram of the ICCL is presented in Fig. 7.

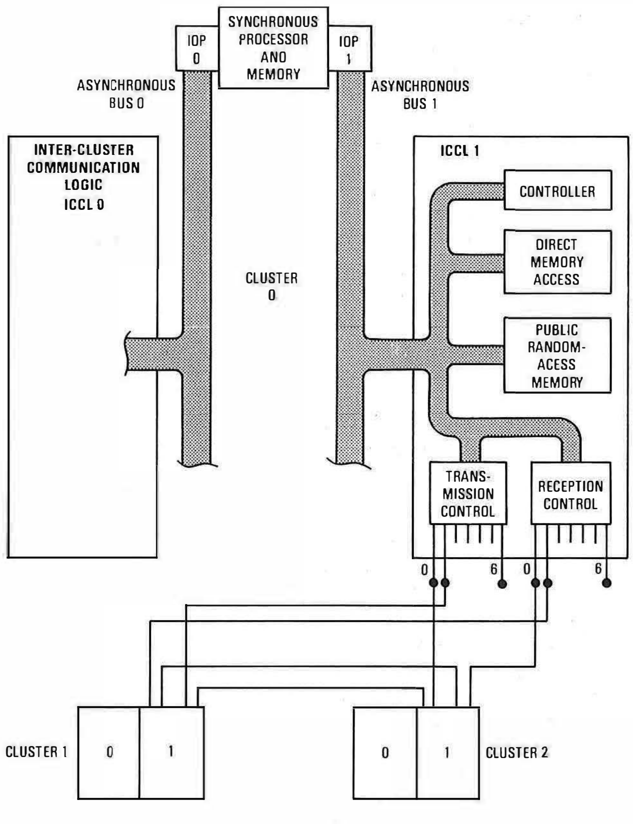

Note: All links are duplicated in plane 0 (not shown)

FIG. 7—Inter-cluster communication

An ICCL simply consists of a microprocessor-based system controlling a DMA, two semi-custom serial-to-parallel converters and a large task buffer, all linked by a bus system. A single ICCL occupies one SIU and can provide up to seven links, enough for the largest configuration of eight clusters. For security, two ICCLs are provided, one being connected to each asynchronous bus. Under normal conditions, both ICCLs may be carrying tasks, some links being workers and some stand-bys.

The task buffer area is divided into eight queues: one for all incoming tasks, the rest allocated one to each outgoing link. The queues are controlled by using head pointers and tail pointers. The tail is changed when task(s) are added and the head when they are removed. The ICCL scans these pointers to see if the PA has inserted any tasks for transmission. Once a task needs to be sent, the ICCL establishes the link to the destination cluster. A cyclic redundancy check (CRC) is added to the task, which is then labelled with flags to identify the start and the end of the message. The CRC conforms to CCITT No. 7 signalling system and provides a fast and secure way of identifying transmission errors. A task received without errors is loaded into the incoming queue to await collection by the receiving PA at the next 10 ms interrupt. The task takes on average slightly longer than 5 ms to travel between PAs; this is due almost entirely to the 10 ms periodicity of the poll, and therefore inter-cluster tasks are transmitted almost as quickly as the intra-cluster tasks.

Various fault and overload handling techniques are employed by the ICCL and the recovery process:

(*a*) Repeat transmission of a task is used if the CRC fails; after two consecutive failures, the link is taken out of service. A count of CRC failures/link is also kept so that action can be taken if repeats become excessive.

(*b*) Under heavy load conditions, difficulty may be experienced in setting up links between the ICCLs. This causes tasks to be held in queues until the links are established. However, the queues are dimensioned so that a fault condition is indicated if they ever become full.

(*c*) A cluster undergoing a restart or complete failure is identified at the local ICCL by a stay-alive timer. The failure of the timer causes the ICCL to stop accepting tasks and results in links timing out and the remote clusters quickly detecting the failure.

These and other fault conditions are indicated to the PA in the queue pointers and to the recovery process in various fault counters as well. In addition, the ICCL has a full self-test capability which runs at restart, and the outcome is reported to the operating system in its status words.

REAL-TIME OPERATING SYSTEM

Processor Allocation

Processor allocation is concerned with the scheduling of process activity. Each cluster has four CPUs and no process can run simultaneously on more than one CPU. The selection of which process to run, and on which CPU, is performed within the PA, which is implemented in each CPU microprogram. This selection is achieved by reference to data held in the main memory.

To enable processes to be stopped and restarted, a place for storing information during periods of inactivity must be provided; this area is termed a *process descriptor*. In addition to preserving the necessary register contents, it also contains other information about the process which the operating system uses for other functions such as process communication.

A process reschedule on a CPU entails the writing of the current register contents to the appropriate process descriptor and then reading the registers of the next process to be executed from its process descriptor.

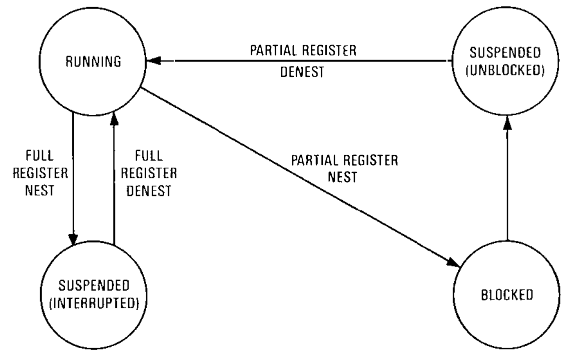

FIG. 8—Process states

For the purposes of rescheduling, there are three states in which processes can exist as depicted in Fig. 8:

(*a*) RUNNING the process currently occupies a CPU and its registers are held by that CPU.

(*b*) SUSPENDED (unblocked or interrupted) the process requires CPU time and is waiting for a CPU on which it can run to become free. The registers are held in the process descriptor.

(*c*) BLOCKED (that is, not competing for CPU time) the process is not scheduled until an external event causes it to move into the SUSPENDED state.

Scheduling is concerned firstly with the selection of a process to run from amongst those suspended, and secondly in finding a free CPU or, alternatively, one which can be freed to run the process selected (that is, pre-emption of an active process).

The selection of which of the suspended processes to run next is based on two factors. The governing factor is the priority of the suspended processes; the other factor is the order in which they are suspended. The secondary factor is relevant only when more than one process on the same priority level is suspended.

Priority on System X is fixed, and any number of processes may share the same priority level. This is advantageous where process replication is used. Replication means that the same code but with different working data can be scheduled as different processes. This has the effect of increasing the available CPU power to a particular function. An example of process replication is the call-control process.

Processes are selected to run by accessing a queue in the form of a linked list, the elements in the queue corresponding to the processes in a SUSPENDED state. The queue is maintained for each priority level, and any of the PAs can access the queues. When a process enters the SUSPENDED state, it is added to the queue. If this newly suspended process is of higher priority than one of those currently active, then that process has to be displaced as quickly as possible.

In order to alert the PA on the CPU running the lowest priority process, the PA that has decided that pre-emption is required, accesses the interrupt hardware to cause an interrupt on the appropriate CPU. The occurrence of this pre-emptive or suspended queue interrupt is then handled in the normal way by the interrupt handling part of the PA, and the higher priority process on whose behalf the interrupt was created will run.

Inter-Process Communication

Although any one process performs its function autonomously, its activity is driven by external requests, which could be in the form of an interrupt, or a task from another process.

Process communication can be achieved by using areas of memory or by making use of registers in the CPU. Whichever mechanism is used, there must be some buffering of information if the processes are to run truly asynchronously.

The mechanism of task passing is implemented also in the CPU microprogram as part of the PA. Task passing is by reference to a task index translation table (TITT), one being provided for every process. A TITT is used by the PA to determine the destination process and task priority. In setting up a task the process also gives the appropriate task index that is used to index the process's TITT.

A process has the choice of sending either short or long tasks. A long task is essentially a short task with an associated data area.

Short tasks are contained in 12-word blocks, although only eight words are available for data. A pool of free blocks is held in the form of a linked list and, when a process wishes to send a task, a free block is taken and linked into the task queue of the destination process by the PA. Tasks are linked into the process's queue at a position determined by the task priority. The act of placing a task in a process queue causes it to move out of the BLOCKED state into the SUSPENDED state. When the process is scheduled to run and the transmitted data has been processed, then the task block can be returned to the free pool.

Within a cluster, the transfer of data from long-task areas is achieved not by actually moving the data, but simply by passing a pointer to the long-task area involved from the sending process to the destination process. In this way, the overhead of copying large amounts of data from one process to another is avoided. However, if a task crosses an inter-cluster boundary, then movement of data over the inter-cluster link must take place.

An advantage of the task-priority scheme is that it enables the process to be selective about its choice of work. For example, it can voluntarily enter a BLOCKED state until a task of a particular priority is received. This is particularly useful when it is not possible for the process to continue doing other work, until it has received a particular message.

Storage Management

Storage management controls the access and the use of the main store and the backing store, and the movement of code and data between the two. This is one of the functions of the storage allocator (SA) process.

Allocation and de-allocation of memory is facilitated by the use of virtual addressing. It is a function of the compiler to translate the identifiers of a given process into a virtual address. The allocation of the virtual address determines only the relative positions of code or data items without reference to physical locations; the translation into a physical address is performed at run time. The store is organised into pages, and groups of pages into files. Protected store access is also provided by separation of the files into the following types:

(*a*) execute-only,
(*b*) read-only, and
(*c*) read/write.

To prevent mutual interference between processes each must access only the store that has been allocated to it, except under certain conditions where data is shared by two or more processes. This means that a process must have access to a range of addresses that overlap the address available to other processes only where such sharing is allowed.

A feature of the virtual-addressing scheme is that the total virtual address space occupied by all the processes may exceed the space physically available in main memory. This means that a mechanism for paging is required.

Paging Mechanism

The store is treated as if it were partitioned into a number of pages each of 256 words. Once pages have been brought into store from the backing store, they remain there unless selected for over-writing. The selection of the page is not arbitrary since this is usually carried out according to a keeping priority. The different pages of a process have different keeping priorities, reflecting their relative importance. These priorities are not necessarily fixed; for example, a page which has been written to ranks higher in keeping priority than a page not written to, since the overhead of moving out of store is much greater (an extra backing store transfer is required). Also, if a process is currently running on a CPU then its pages are obviously not selected for overwriting.

In addition to these dynamically changing priorities, further static keeping priorities are applied. The software is partitioned into three different categories of page as shown in Fig. 9:

(*a*) those permanently resident for time-critical activities (permanent pages),
(*b*) those usually resident, but overwritten when demand for pages is high; for example, under fault conditions (reserved pages), and
(*c*) those overwritable for non time-critical activities (overwritable pages).

The time-critical activities are parts of the operating system and the call handling parts of the application suite. The non-critical activities include functions such as man–machine communication and maintenance control.

Use of the Backing Store

The backing store is used for several purposes:

(*a*) for holding code and data which are fetched into the main store at restart,
(*b*) for holding the initial versions of data in a secure manner, and
(*c*) for holding data in the overwritable category when it is different from the initial version.

Timing

The timing functions within the operating system provide

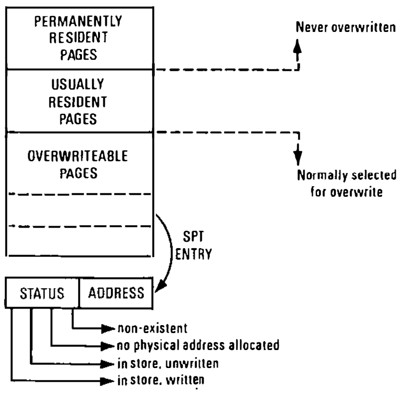

FIG. 9—System page table areas

three basic facilities: absolute time, relative time, and periodic timers.

These facilities are for general use by the application processes (APs) and for control purposes; for example time-outs on accesses to AHUs or on messages to other processes.

The basic source of timing to the exchanges is provided by the synchronisation utility, which drives a real-time clock system in each cluster. This generates an interrupt every 10 ms. This interrupt is received by the PA and an interrupt task is returned to the timing process for it to maintain the timing facilities. The only exception to this is periodic timing, which is handled directly by the PA. It generates periodic tasks and scans the inter-cluster link for incoming tasks.

A software calendar is provided on each cluster for use by any process within the cluster. Two digital clocks are provided on the base cluster to maintain the accuracy of the calendar. These maintain time whilst the cluster undergoes a restart.

There is a need to co-ordinate the maintenance of time between clusters and, to ensure consistency, any changes concerning the absolute time are handled on a system basis by the synchronisation and control of timing (SCOT) process, whereas the other timing facilities can be provided on each cluster by the timing process.

APs can also make use of absolute time by a facility called the *initiates function*. This allows a process to request the operating system to send it a task at a particular time in the future.

Relative time is used for the time-out mechanism. APs can request the operating system to return a task to that process when a specified time interval has elapsed.

The periodic timing facility is handled by both the timing process and the PA. The timing process is responsible for starting and stopping periodic timing, and the PA is responsible for sending periodic tasks to all those APs which have periodic timing started. The periodicity is a multiple of the basic real-time clock period of 10 ms, from the minimum 10 ms up to a maximum of 2·55 s. Normally, the periodic tasks continue indefinitely and no action is required by the processor apart from it receiving the tasks.

It is also possible to arrange for periodic tasks sent to different processes to have a fixed phase relationship. This can help to distribute the load on the system so that, for example, two 20 ms processes run on alternate 10 ms interrupts rather than both together.

Load Control

The purpose of load control is twofold. Firstly, it is designed to maximise the amount of call processing whilst not exceeding the capability of the processor and, secondly, it prevents abnormal overloads causing the exchange to fail.

The first is achieved by the rejection of a proportion of newly arrived calls. Processing demand is monitored by the call processing subsystem (CPS) measuring the number of calls accepted and rejected and by the PU measuring the amount of spare processing power it has available. This information is then processed according to a pre-defined algorithm which generates new limits for the application subsystems, such as CPS, to determine the amount of work that can be handled. These activities occur on a cyclical basis and provide graduated control of the amount of work the system is handling.

The second of these functions is needed to deal with a failure within the processor or problems within the network. Detection of overload is done by monitoring both the process task queues and the total number of tasks within the cluster.

The function of detecting a potential overload is time-critical as the quicker a cluster can be saved from entering overload, the more secure the exchange. The function of load management is not as time-critical, so to prevent its work from delaying the action taken in overload conditions, this function is implemented as a separate lower-priority process called the *load monitor process* (LMP). The process dealing with the overload condition is called the *overload control process* (OCP).

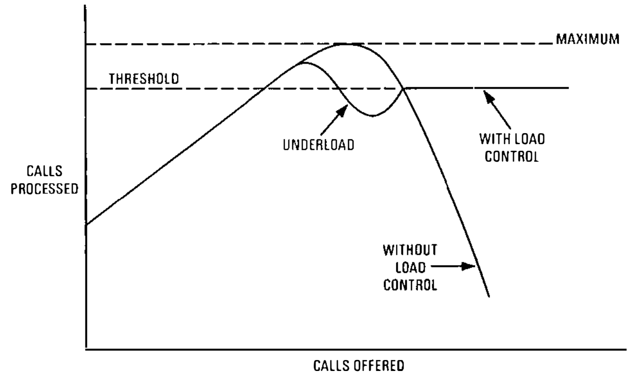

FIG. 10—Load control

For the purposes of the control of the level of work in the application subsystem, the subsystems are classified as *type 1*, *type 2* or *type 3 targets* of load control depending on how finely they can control the work they perform.

A type 1 target is capable of measuring its workload directly in terms of the work it does; for example, calls in set-up. This is a fine control that is achieved by reducing the number of simultaneous call set-ups.

Type 2 targets are subsystems such as the maintenance control subsystem[2] (MCS) and the management statistics subsystem (MSS), where certain functions can be abandoned according to a workload limit which it receives from the operating system. The lower the workload limit the more functions have to be abandoned; this is a coarse control.

Type 3 targets are those subsystems which can operate only in a completely ON or OFF mode. The OFF state merely means that the subsystem processes the absolute minimum amount of work.

The amount of spare processing power available is continuously monitored. At the end of each load monitoring period, if the power available exceeds a certain value, then the system is deemed to be underloaded and the spare power is distributed between type 1 and type 2 targets which may be rejecting work. This is done by having their workload limits increased. If the power available is below a certain value, the system is deemed to be overloaded and the workload limit of the various type 1 and type 2 targets is decreased.

To detect traffic overload, length thresholds are pre-set on each process task queue and, if these limits are exceeded, the operating system reduces the workload of subsystems which may have contributed to the overload. A matrix is used to define for each process whose queue reaches its specified limit a list of targets which may be responsible. It also takes into account the number of tasks by which the queue length must fall before the overload condition can be said to be ended. This avoids oscillation between overload and non-overload (see Fig. 10).

FAULT RECOVERY

Recovery from Software Faults

The basic mechanism of software error recovery is to invoke the appropriate level of *roll-back*; that is, to reload into the main store from the backing store the affected portions of the system software and to restart the affected processes.

To minimise roll-back activity, an algorithm is used to control the escalation of roll-back from one level to another. The algorithm is implemented in the roll-back master process which exists only on the base cluster, and the roll-back process which exists on each cluster.

There are five levels of roll-back:

Level 1 roll-back of the read/write data of a single process,

Level 2 roll-back of all the application process (AP) read/write data on that cluster,

Level 3 roll-back of all files on the affected cluster. This includes the operating system and is termed a *cluster restart*,

Level 4 cluster restart of every cluster in the system. This is identical to level 3 on a single cluster system,

Level 5 system restoration; this is initiated when a system restart fails to execute correctly or when an escalation of system restarts takes place. This results in a complete initialisation of all software and hardware in addition to a cluster restart.

The roll-back algorithm has three mechanisms to control the escalation of roll-back. They are process control, cluster control and system control.

Process Control

In process control, two leaky bucket counts are held against each process. Each count has its own preset threshold. One set of counts holds roll-back requests made against each process and the other set holds roll-back requests made by each process. When a roll-back request is made, the requests-by count of the requesting process and the requests-against counts of the requested process are incremented.

If a count reaches a threshold, the process concerned suffers a roll-back, upon which both counts for that process are set to zero. An ignore period is then started, and during this period all roll-back requests made by or against the process are ignored. The only exception to this is if a TRAP occurs; then the process must be rolled back immediately. A TRAP during the ignore period causes escalation to the next level.

Cluster Control

In cluster control, there is a leaky bucket count for each cluster. The count is incremented by one for each level 1 process roll-back and by a pre-defined value for level 2. The count is used for escalation from levels 1 to 2 and levels 2 to 3. After a cluster has undergone a level 3 roll-back, an ignore period is started against the cluster. During this period any requests by processes on that cluster, or for that cluster, are ignored.

System Control

There is only one leaky bucket count for the system. This count is incremented by a pre-defined value when a level 3 roll-back occurs and a pre-defined value when a level 4 roll-back occurs. The count is used for escalation from levels 3 to 4 and levels 4 to 5.

All of the above counts are decremented periodically if not zero.

Recovery from Hardware Faults

As mentioned earlier, the processor is required to run for hundreds of years with only a small probability of total failure. The failure rate of individual integrated circuits and the quantity of them in the processor means that it is inevitable that hardware failures are going to occur much more frequently than has been allowed for in the total system failure rate. It is the job of the hardware fault recovery process to ensure that the system can recover from these faults so that processing can continue with minimum disruption. This process is known as *recovery*.

The process of recovery from hardware faults falls into several distinct stages: fault detection, isolation, reconfiguration and reporting. Initially, a fault has to be detected and recognised as a hardware fault. This is accomplished by the use of the fault detection logic built into the hardware. Also, the software can detect hardware faults by using time-outs and by checking data for consistency.

In general, detection of a fault causes the recovery process to be run; however, some faults may only be transient, that is to say there is no permanent fault in the hardware but for some reason a single operation, for example, a store access by a CPU, has failed, even though subsequent repetition of the operation succeeds. Modern memory systems using densely-packed RAMs are particularly susceptible to this type of fault. If a transient fault occurs, it is inefficient to have to run the recovery process to sort out the problem; so for most faults the failed operation is re-attempted by the hardware/firmware in the system. If this second attempt succeeds, then no further recovery action is necessary except to record the fact that a transient has occurred. Counts of transient faults are kept, and if a resource generates too many in a given time, it is taken out of service.

At the other end of the spectrum, certain faults can have such a disruptive effect on processing that it is impossible to continue running without an immediate reconfiguration of the hardware. An example of this is the failure of a memory card containing operating system data. Under these circumstances, an immediate cluster restart is initiated. The cluster restart, which is identical to the level 3 roll-back, is a fundamental part of the recovery from these types of hardware fault. At the commencement of a cluster restart, all CPUs stop and open their access port. Backing stores also stop because the operating system ceases to access the BSCs and a stay-alive timer inside them times out. Within a few seconds, hardware inside the CPUs causes them to close their access port and to attempt to start processing again. Normal component tolerances ensure that one CPU will start first. This CPU takes charge of the restart operation and, initially, it prevents the other CPUs from starting to process until a stable configuration has been established.

When the CPU begins the restart, it assumes nothing about the state of the processor. Firstly, it tries to find an IOP in which it can access the ROM containing the restart software, known as *kernel fetch*. This is executed to enable the most important parts of the operating system to be loaded into memory. Before any hardware is used during the restart, test routines are run to check that it is working correctly; any hardware that fails these tests is left out of the new configuration. Once the kernel of the operating system is loaded and running, the remaining CPUs are restarted and the remainder of the operating system together with the application software is loaded from the backing store. This is performed as one large data transfer known as *bulk retrieval*. If at any time during the restart a hardware fault occurs, the whole procedure is started again from the beginning. The hardware ensures that the second restart is initiated by a different CPU to prevent a faulty CPU permanently failing to restart.

Within the two extremes of transient fault and cluster restart, there are a whole range of fault types for which the recovery process is required to reconfigure the processor into a working state with the minimum possible amount of disruption to processing. There are two aspects of this. Firstly, there is the hardware reconfiguration. All the essential resources in the system are secured by having spares.

These are provided either by load sharing or by worker/stand-by arrangements. Hardware reconfiguration just makes use of the spare or stand-by resources. The rigid boundaries that are maintained around the security units ensure that it is always possible to isolate a faulty module from the working hardware. Reconfiguration in this way is always possible under single fault conditions. Secondly, after the reconfiguration of the hardware, it is necessary to reconfigure the use of that hardware by the software.

Some hardware failures interfere with the normal running of software; for example, if a CPU fails, a process will be interrupted or its data corrupted. In these cases, the software fault recovery mechanisms also have to be invoked.

In most cases, when a fault has been located to a hardware resource, that resource is immediately reconfigured out of the system. However, some faults can occur on the interface between two resources or there may be misleading fault information. In these cases, the initial reconfiguration may fail to isolate the fault. If this happens, a sequence known as *trial reconfiguration* is invoked. This procedure involves taking resources out of the configuration one at a time until a configuration is found that will run without faults. It is possible that this procedure may take out of service a resource which does not contain the faulty hardware; however, the overriding principle is to obtain a configuration which can operate securely and continue to process traffic until repairs can be performed. To enable such a fault to be repaired, it is possible manually to force the recovery process to try different configurations.

The trial configuration algorithm uses only resources that were previously in service before the fault occurred. If a fault occurs while a resource is out of service, for example, during maintenance operations, it is possible that there may not be a viable set of resources left to enable the system to run. In this situation, as a last resort, the recovery process attempts to use resources which were previously out of service. This procedure, known as *restoration*, may seem extreme, but as stated above, the overriding principle is to obtain a working configuration: consequently the processor will never stop attempting to do this.

Having reconfigured the hardware into a viable configuration, the recovery process informs the maintenance staff of what actions have been taken so that the faulty equipment can be repaired as quickly as possible. The repair normally takes the form of powering down the faulty resource, replacing the appropriate SIUs and powering up again. After the repair, the recovery process performs a series of tests on the resource before allowing it to be returned to service. The recovery process is involved in reconfiguring the system to use the resource if these tests are passed.

MAN–MACHINE INTERFACE

Communication between application processes (APs) and both local and remote man–machine peripherals is provided by the man–machine interface (MMI), which also enables administration data to be sent to a remote administration centre. ASCII (ISO 7-bit code, British Standard BS4730) and binary characters can be accommodated: ASCII for the man–machine language (MML) transactions via keyboard peripherals and binary for administration data and some applications using magnetic tape cartridges.

Input/Output Facilities

For input of man–machine transactions via keyboard devices, the MMI supports the CCITT MML for which commands are of the form

COMMANDCODE: PARAMETERS, PARAMETERS;

The MMI checks the syntax and generates any necessary prompts and error messages before passing the input to the destination AP.

Output is directed to a channel specified by the AP. The MMI translates the channel number to a device number, a stand-by device number and an optional copy-channel number. Thus, output can appear on one or more devices, and can be diverted to stand-by devices in cases of device failure.

The MMI supports the ISO-extended system for labelling and formatting (high density) cartridges and can read and write multi-file volumes and multi-volume files. The MMI keeps a list of cartridges known to the system and can prompt for a particular cartridge, for example, the next one in a multi-volume set, when required. Fixed and variable-length blocks with single or multiple records can be handled.

Access security is provided by means of user-codes and passwords; also, command password levels (CPL) are used to restrict access to individual man–machine transactions. When a user successfully logs-on to the system, a CPL is calculated from the user-code/password and port-of-entry. This CPL is compared to another CPL stored with each command code and the transaction is allowed to proceed only if these correspond. Thus, different users can have different levels of access to command codes on the system.

Remote Communication

Man–machine IO, together with other data from the APs can be passed over a link to a remote site to enable man–machine transactions to be remotely controlled, and administration data, such as billing records, to be sent from the exchange to an administration centre.

The MMI performs the link-management function of link set-up, verification and flow control, and ensures the integrity of the data across the link by re-transmitting messages when necessary.

An administration data back-up store (ADBUS) in the form of one or more cartridges is provided to store data destined for the remote link during periods of link failure. The MMI can, according to the type of data, decide to divert data to ADBUS if the link has failed, divert to ADBUS irrespective of the link state, or refuse the data if the link has failed. The ADBUS can be divided into up to four separate streams, each being either a single cartridge or duplicated pair, with stand-by cartridges allocated for use when the working cartridges fail or become full. In situations where the remote link is not provided, then all administration data can be sent to the ADBUS and the cartridges later shipped to a data processing centre.

MMI Processes

Man–Machine-Manager (MMM)

The man–machine-manager (MMM) resides on the base cluster only and controls the man–machine processes on all clusters. Its functions are:

(*a*) handling inter-cluster MMI transfers and the remote MMI protocols,

(*b*) providing access security by validating user-codes/passwords,

(*c*) providing transaction security by validating commandcodes, controlling transactions and maintaining transaction serial numbers, and

(*d*) ensuring that man–machine IO is logged.

Character Generation Process (CGP)

The character generation process (CGP) resides on all clusters equipped with man–machine peripherals. It converts the raw data and output instructions from the APs into ASCII characters for output.

Device Control Input Analysis (DICK)

Device control input analysis (DICK) is provided on all clusters with man–machine peripherals. It has several functional areas:

(*a*) *DEVCON* (Device Control) handles input/output via the peripheral controller and deals with device failures.

(*b*) *IAN* (Input Analysis)

(*i*) converts ASCII input into the data format required by the destination AP, and

(*ii*) generates syntax error messages and prompts.

(*c*) *CARTMAN* (Cartridge Manager)

(*i*) controls output to magnetic tape cartridges,

(*ii*) maintains duplicated and stand-by cartridges for the ADBUS and generates operator prompts when necessary, and

(*iii*) handles ISO-extended cartridge labelling.

(*d*) *KEYMAN (Keyboard Manager)*

(*i*) controls output to VDUs and printers, and

(*ii*) diverts output to stand-by devices and copies output to other devices when required.

Interface Package Process (IPP)

The interface package process (IPP) resides on the base cluster and other clusters where required. Its functions are:

(*a*) management of the remote communication link,
(*b*) control of data transfers across the link, and
(*c*) diversion to the administration data back-up store.

The MMM and DICK are *multi-threading* and use job control blocks to schedule individual jobs. The IPP is run periodically and the CGP can queue 16 jobs with three running simultaneously.

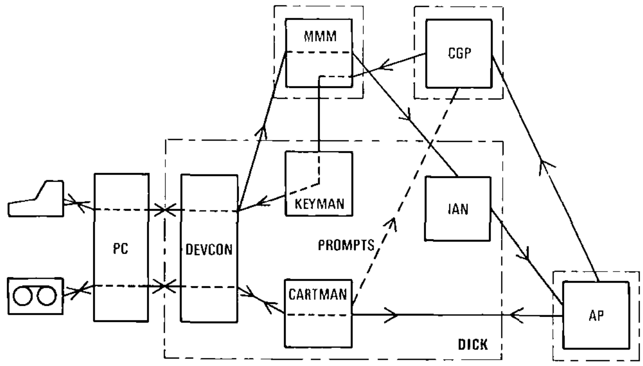

FIG. 11—Keyboard and cartridge input/output

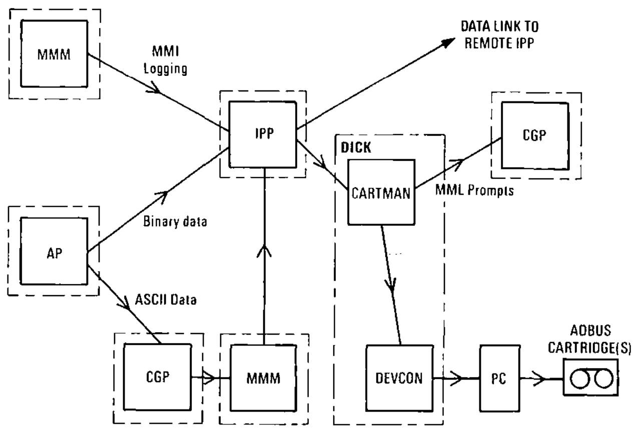

FIG. 12—Administration data and ADBUS

Man–Machine Paths

The basic communication paths through the MMI processes for IO are shown in Fig. 11. Keyboard input is routed by the DEVCON to the MMM, which validates password and command codes before passing it to the IAN for syntax analysis and transmission to the destination AP. Keyboard output is handed by the AP to the CGP, which generates the characters and hands them via the MMM to the KEYMAN, which routes them via the DEVCON to the appropriate device(s). The MMM is responsible for logging all input and certain types of output.

Binary cartridge input and output is passed between the AP and the PC via the DEVCON and the CARTMAN. ASCII cartridge input and output is also routed via the MMM, IAN or CGP in a similar manner to ASCII keyboard input and output.

The routes for administration data are shown in Fig. 12. The APs (and the MMM for logging) send binary data direct to the IPP, which can transmit it over the link or divert it to the ADBUS via the CARTMAN and the DEVCON. ASCII data can also be sent over the link or diverted to the ADBUS; in this case the route is via the CGP and the MMM to the IPP.

ADMINISTRATIVE FUNCTIONS

Traffic Monitoring

The purpose of the traffic monitoring function is to provide a measure of processor performance and efficiency. The data obtained can be used as an aid to system dimensioning and to verify teletraffic models or to indicate whether a given processing configuration is adequate for the presented load. Periodically, the load monitor process (LMP) gathers statistics on the number of rejected calls. These are added together and the total sent to the MSS.

Other measurements are carried out by the processor utility monitor (PUM) process, which makes use of the pseudo-random clock interrupts to start the measurement cycle and to determine whether to take a sample; this ensures that the measurements have no fixed relationship with the periodic clock interrupts.

Measurements can be taken of the following parameters:

(*a*) process state occupancies,
(*b*) task queue lengths,
(*c*) inter-cluster link overload count,
(*d*) backing store transfers,
(*e*) backing store use, and
(*f*) storage requests.

All of these measurements can be requested by administrative staff via the MSS. The data collected during the measurement is sent back to the MSS.

Archiving

Archiving provides a means of copying the system code and data by transferring data from the backing store to cartridge tape whilst the system is running normally. It is possible either to select a complete system archive or to archive specific files. In the case of the complete system archive, contents of the tape may be subsequently reloaded into the backing store. The partial data archives are not reloadable, but they can be used for carrying out code and data audits elsewhere.

Archiving facilities are provided on each cluster, with overall control from the base cluster. Before the archive can start, the backing store from which it is to take place is assigned a special state. In this state, no changes to initial versions can occur, although updates are still performed on the other backing store during the archive operation. It is for this reason that both backing stores must always be in normal use for an archive to start.

The archiving functions are implemented in three separate processes; archive master, archive slave and high-density cartridge (HDC) controller. The archive master process runs only on the base cluster and co-ordinates the activities of all clusters involved in the archives; archive slave is present on each cluster and controls those parts of the archiving operation which are specific to that cluster; and the HDC controller is responsible for all aspects of management and use of the HDCs.

A complete archive can be achieved in a matter of a few minutes.

Software Enhancement

The purpose of the software enhancement facility, which resides in the operating system, is to provide a secure means of loading new software into an existing installation without interruption to service and to subject the software to a trial before it is finally accepted into the system.

This new software could be in the form of a part process, a new process or any number of processes.

The purpose of the trial is to ensure that the new software is tested before being finally accepted. This is achieved by allowing the modified system to run for a pre-specified period, during which time the old system is preserved and can be returned in the event of the new system being unacceptable. Should the system under test fail at any point, then the necessary recovery action is automatically carried out.

Simultaneous modification can be made to any number of processes on any number of clusters. To facilitate this there is a controlling process, known as *loader master*, on the base cluster and a secondary process, known as *loader*, on each cluster. The loader is responsible for carrying out the input and syntax analysis of a load file.

So that the modified system can be set up while preserving the current system, both systems are held in the backing store. During normal operation of a cluster, two backing stores are in use, with the initial versions of the files duplicated on each. A modified system is set up by using the loader to modify one copy while the other copy runs the system.

At the start of a trial, all modified processes must be restarted with modified versions of their files. This requires a cluster restart to ensure that all modified files are rolled back. Similarly, to end or scrap a trial, a cluster restart is again required.

A trial is successfully completed when the trial period has expired and no failures have occurred. The operator must then finally accept the system by use of the appropriate command which causes the new software to overwrite the old system.

Hardware Upgrades

It is possible to sub-equip some of the resources on the processor. For example, where a particular installation requires less store or processing power, then fewer storage modules or CPUs can be provided initially and added to later if and when required. This also applies to the backing store, which is modular.

The most significant upgrade is the addition of another cluster to an in-service exchange. As far as the processing subsystem is concerned, this can be considered as a hardware enhancement, but the system would treat it as a software enhancement also, since additional application code and data is being added. The basic procedure is described in outline below.

The new cluster is tested initially in the STAND-DOWN mode, firstly, by commissioning the cluster on site by using a suite of test programs that test the operation of each hardware module. The second part of this test is to load the operating system onto the new cluster as if it were a base cluster. This includes the MCS and any necessary data to allow a complete test of that cluster.

The links are now cabled between the new cluster and the existing ones and the new cluster is run as part of the system, but under probation. This requires that the master processes loaded onto the new cluster as part of the first stage must be prevented from running; that is, the new cluster must behave as a true slave cluster.

The final stage is to carry out a complete reload of the system on the new cluster. This contains the new application software suite and removes any base-only processes. The software enhancement facility is then used to make the necessary changes to the existing clusters and run a system trial. Once the trial has been successfully completed, the cluster is ready to accept traffic.

CONCLUSIONS

The machine has, during its trialling period, undergone and completed satisfactorily a very extensive series of hardware and software tests. It was noteworthy that the hardware which uses current technology and System X tolerancing design rules has proved to be particularly resilient to the quite severe environmental conditions to which it was subjected.

ACKNOWLEDGEMENTS

Acknowledgement is made to GEC Telecommunications Ltd., who developed this processor under a BT/System X contract, for permission to publish the information contained within this article.

References

[1] BELTON, R. C. System X: Subsystems Part 3—Processor Utility Subsystem. *Post Off. Electr. Engs. J.*, Oct. 1980, **73**, p. 158.

[2] BATY, R. M., and SANDUM, K. N. System X: Maintenance Control Subsystem. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 276.

Biographies

Derek Troughton joined GEC as a student in 1967 and gained a degree in Electronic Engineering at Lanchester Polytechnic in 1971. He worked in the Transmission Division on microwave radio systems before returning to Lanchester Polytechnic in 1972 to study encoding techniques for digital transmission systems. In 1975, he joined the Telephone Switching Group at GEC as part of the System X processor team, working on fault-recovery software. He ran the operating system development for 5 years until 1983 when he was given the complete responsibility for processor development. He is currently Head of Development, Processing and Signalling Systems.

Peter Lumb gained an honours degree in Physics at the University of Aston in Birmingham before joining GEC in 1974. He worked as a member of the POPUS 1 processor design team until 1980 when he was appointed Group Leader with responsibility for the design of the hardware for the new processor. He has recently gained responsibility for the operating system development.

Rex Belton joined the Post Office in 1943 as a Youth-in-Training. He has worked on various aspects of research and development of telegraphs, Telex power plant, data, facsimilie, computer designs, computer systems procurement and provision, computer applications including planning , maintenance and early knowledge-based systems, packet switching and multi-service networks. At present he is Head of Section in BT's System Evolution and Standards Department responsible for System X processor development.

Ian Gallagher joined the Post Office as a student in 1969 and obtained a degree in Electronic Engineering from Sheffield university in 1973. He was involved with the development of TXE4A before transferring to the System X processor development group in 1978 where he was closely concerned with the design and testing of the CPU and memory area of the processor. He is currently Head of Group investigating multi-service networks.

Mike Bexon gained a degree in Electrical Engineering at Imperial College, London before joining BT (then the Post Office) in 1970 as an Executive Engineer. He worked on pulse-code modulation switching trials until 1972 when he was sponsored to study Telecommunications Systems at the University of Essex. On his return to BT, he worked on the line scanning and cyclic store areas of the TXE4A exchange. In 1978, he joined the System X development team working on the specification and testing of the man–machine interface and loading/editing areas of the processor utility subsystem.

Steve Moor joined BT in 1974 having obtained an honours degree in Electronics at Sussex University. He worked with FDM transmission equipment and was part of the 60 MHz project team. In 1979, on promotion to Executive Engineer, he joined the System X processor team where he has specialised on processor hardware and software to test hardware.

Steve Stegman joined BT as an Assistant Executive Engineer in 1971 after having successfully completed an Electrical Engineering course at Liverpool Polytechnic. He worked in the Telecommunications Development Department on pulse-code modulation system signalling units during which time he designed an automatic processor-based tester for commissioning them. In 1979, he was promoted to Executive Engineer and joined the System X processor team where he has specialised in the hardware design aspects of the processor.

APPENDIX
Glossary of Terms

ADBUS Administration data back-up store.
AHU Applications hardware unit.
ALU Arithmetic logic unit.
ASCII American standard code for information interchange.
Barrel Queues A first-in-first-out buffer of finite length with input and output cylically following each other.
BMU Bubble memory unit.
BSC Backing store controller.
CARTMAN Cartridge manager.
CCITT International Telegraph and Telephone Consultative Committee.
CGP Character generation process.
Control Page A special page in the main memory that contains PC commands, responses and buffer information.
CPL Command password levels.
CPS Call processing subsystem.
CPU Central processor unit.
CRC Cyclic redundancy check—a binary code generated from a block of data used for error checking.
Daisy Chaining A highway that loops in and out of its connections.
DEVCON Device control.
DICK Device control input analysis.
DIO Direct input/output.
DMA Direct memory access, without CPU interaction.
FIFO First-in first-out.
Firmware The name given to the software that controls the microprocessors which form part of the ancilliary functions of the PU. It is generally held in ROM.
GDLC GEC data link control—a version of high-level data link control (HDLC).
HDC High-density cartridge.
IAN Input analysis.
ICCL Inter-cluster communication control logic.
Initial Version The secure copies of the exchange code and data as distinct from the working versions.
IO Input/output.
IOP Input/output processor.
IPP Interface package process.
ISO International Standards Organisation.
KEYMAN Keyboard manager.
LCPU CPU running the lowest priority process.
LMP Load monitor process.
MCS Maintenance control subsystem.
Micro-instructions Instructions internal to the CPU that are sequenced to produce the CPU's instruction set and other high-order operations.
Microprogram A program of micro-instructions.
MMI Man–machine interface.
MML Man–machine language.
MMM Man–machine-manager.
MSS Management statistics subsystem.
ON-LINE The normal state of the PU when running an exchange.
PA Process allocator.
PC Peripheral controller.
Poll A regular interrogation to check if an expected event has occurred.
POPUS 1 Early System X multi-processor utility.
PPU Early System X main/stand-by processor utility.
PU Processor utility.
PUM Processor utility monitor.
RAM Ramdom-access memory.
ROM Read-only memory.
SA Storage allocator.
SCOT Synchronisation and control of time.
SIU Slide-in-unit.
Stay Alive Timer A timer used to ensure that another entity is still functioning correctly.
Task In the System X context, a task refers to a message passed between processes.
Time-out A secure mechanism to ensure that a function does not take too long.
TITT Task index translation table.
TRAP The CPU stops running and executes a fault handling process.
VDU Visual display unit.

System X: Digital Subscriber Switching Subsystem

R. C. WARD, M.A., M.SC., C.ENG., M.I.E.E.†

UDC 621.395.34

The digital subscriber switching subsystem of a System X local exchange is used to concentrate traffic from individual analogue and digital subscribers' lines into the central digital switch. This article describes the architecture of the subsystem developed as part of the System X enhancement programme. Particular attention is given to the critical line-interface area.

INTRODUCTION

The concept of a subscriber switching subsystem and an early architecture have been described hitherto in this *Journal*[1]. However, as part of the System X system enhancement programme (SEP), the design has been modified to provide a product suitable for widespread introduction into the UK network. The digital subscriber switching subsystem (DSSS) is capable of supporting both analogue and digital subscribers' lines, and is a constituent of both the local exchange and the digital principal local exchange. A local exchange comprises a number of DSSS concentrator modules, which may be either co-located or remotely connected via 2 Mbit/s pulse-code modulation (PCM) digital links to a parent exchange[2].

Up to 2048 analogue subscriber lines can be supported on a single concentrator module. The traffic generated is concentrated onto a maximum of eight 2 Mbit/s PCM systems linking the concentration module to the digital switching subsystem (DSS), with the speech data conveyed in 64 kbit/s channels across the switch.

† Local Exchange Systems Department, British Telecom Local Communications Services

A feature of the design is the extensive use of large-scale integrated (LSI) and very-large-scale integrated (VLSI) custom circuits in the analogue subscriber line circuit to ensure competitiveness in an area of the exchange traditionally associated with the majority of the cost. As well as analogue subscriber lines, both single- and multi-line integrated digital access (IDA) are supported at transmission rates of 80 kbit/s and 2 Mbit/s to the digital subscriber. The DSSS uses the digital access signalling system (DASS) to provide 2 Mbit/s interfaces for both integrated services private branch exchanges (ISPBXs) and remote multiplexers (RMUX). Application of the DSSS in an integrated services digital network (ISDN) is described more fully in another article in this issue of the *Journal*[3].

SUBSYSTEM ARCHITECTURE

A simplified diagram of the DSSS architecture is shown in Fig. 1. The DSSS consists of central software resident on the host local exchange processor (PU) and other software resident on a duplicated microprocessor, known as the *module controller* (MC), within each concentrator module. The central software is responsible only for maintenance co-

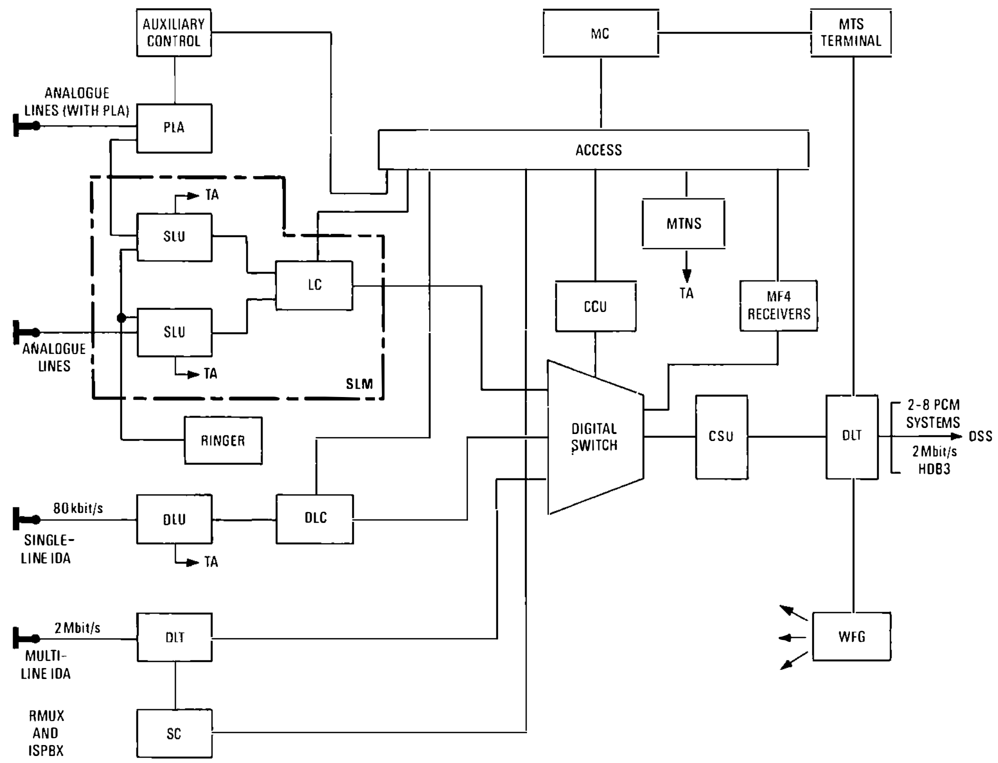

CCU: Call connection unit
CSU: Channel supervisory unit
DLC: Digital line controller
DLT: Digital line termination
DLU: Digital line unit
DSS: Digital switching subsystem
IDA: Integrated digital access
ISPBX: Integrated services private branch exchange
LC: Line controller
MC: Module contoller
MTNS: Mini test network subsystem
MTS: Message transmission subsystem
PCM: Pulse-code modulation
PLA: Per line auxiliary
RMUX: Remote concentrator
SC: Signal converter
SLM: Subscriber line module
SLU: Subscriber line unit
TA: Test access
WFG: Waveform generator

FIG. 1—Digital subscriber switching subsystem (DSSS) architecture

ordination and some auxiliary functions. The main call-handling function is performed by the MC. Time-slot 16 (TS16) in two of the 2 Mbit/s PCM systems linking the DSSS to the DSS is dedicated to carrying signalling information at 64 kbit/s between the MC and the PU by using the message transmission subsystem (MTS). Local generation of clock waveforms is performed by the waveform generator (WFG), which is synchronised to 2 Mbit/s at the digital line termination (DLT).

The digital switch provides for the connection of up to 2048 bi-directional 64 kbit/s time-slots, presented on up to sixty-four 2 Mbit/s digital highways, to any of up to 240 time-slots in eight 2 Mbit/s PCM systems linking the DSSS to the DSS. Paths across the switchblock are set up in accordance with instructions passed from the MC to the call connection unit (CCU). A medium-speed input/output interface, known as *ACCESS*, is provided between the MC and all peripheral control units such as the CCU.

The interface to the subscriber's line is provided by a subscriber line module (SLM) comprising subscriber line units (SLUs) and a line controller (LC). Line supervision and loop-disconnect digit reception are performed by the LC.

Control messages pass between the LC and MC via ACCESS. Support for multi-frequency (MF4) signalling is provided by switching a path across the digital switch to an MF4 receiver on receipt of a *seize* condition from a subscriber requiring this facility. Subsequent digit information is passed direct to the MC. Subscriber speech information is converted to 64 kbit/s within the SLM and multiplexed onto 2 Mbit/s digital highways connected to the digital switch.

Per line auxiliaries (PLAs), under the control of auxiliary controllers, can be associated with an analogue subscriber line to provide P-wire, 50 Hz subscriber private meter (SPM) and high resistance (30 kΩ) loop detection facilities. A variant of the SLU is capable of providing a 16 kHz SPM facility. This is a new standard for the UK, the use of which will become increasingly economic with the projected sharp increase in demand for SPM signalling to support call-charge display and call-logging equipment.

For single-line IDA, a digital interface is provided by digital line units (DLUs) in conjunction with a digital line controller (DLC). A customer interface of (64 + 8 + 8) kbit/s comprising primary (B), secondary (B′) and signalling (D) channels is supported by using a time-compression multiplex (TCM) burst-mode type transmission system with WAL2 encoding of the line signals. For multi-line IDA, a 2 Mbit/s HDB3-encoded interface is provided by a DLT. Signalling is provided in TS16 by using DASS. A signal converter (SC) performs similar functions to the DLC and provides the interface between the 64 kbit/s DASS data injected/extracted from TS16 by the DLT and ACCESS. Both the DLC and SC provide for line supervision and signalling functions in association with the MC.

Tones required during the handling of calls by the DSSS are injected digitally in the channel supervisory unit (CSU). It should be noted that the DSSS has a limited call-handling capability under isolation conditions and it is only then that the full repertoire of tones is used. Under normal conditions only *proceed* and *special proceed* indications are provided to a local subscriber.

Ring tone is provided to the distant subscriber via the appropriate PCM channel. Other tones are provided centrally via the DSS. Ringing current is supplied from a central source with cadencing performed within the SLM.

With the widespread introduction of remote concentrator units (RCUs), a message-based version of the test network subsystem (TNS), known as *mini test network subsystem* (MTNS), has been developed to overcome the problems of remote measurements over long metallic connections. The DSSS has the capability of supporting the MTNS as an ACCESS user. Messages are passed via the MC and central DSSS maintenance software to the central TNS software resident on the PU. It is thus possible to perform accurately remote test measurements on the full range of subscriber lines. Test access (TA) highways from the SLUs and DLUs are either connected to MTNS or extended over metallic pairs to the TNS if the distance to the DSSS is small.

CALL HANDLING

The DSSS can handle either switched-traffic or private-circuit connections for both analogue and digital terminations. The design of the system is such that to reduce the load on the PU, a substantial amount of call set-up, supervision and clear-down signal processing is performed within the DSSS by the MC. The peripheral intelligent control units such as the LC, DLC and SC similarly reduce the processing load on the MC by interpreting low-level line conditions such as *loop* and *break* into higher-level messages such as *seize*, *digit* etc. Such messages are passed via the MC to the call processing subsystem (CPS), which comprises software resident on the PU. The CPS provides call set-up functions such as routeing translations and call-charging supervision.

As well as handling up to 6500 busy-hour call attempts (BHCA) under normal conditions with a 40% margin for overload, each DSSS concentrator module is required to maintain a satisfactory performance despite being grossly overloaded. As the DSSS provides the interface to the subscriber, it has been necessary to ensure that the design provides considerable protection against abuse. Under gross overload, the capability exists to ensure that new seizures do not propagate beyond the peripheral controllers (LCs) until the load on either the PU (short term) or MC (longer term) has subsided. To provide compatibility with existing customer expectations, delay working is implemented by periodic repetition of *seize* messages by the LC until the MC is able to resume normal processing and establish call set-up. Thus the DSSS is capable of continuing to provide normal processing capability with overload in excess of 65 000 BHCA.

For ISDN subscribers, the MC supports the call-handling functions carried in Level 3 of the DASS protocols. Both telephony and ISDN digital calls are supported. For digital calls, call progress information is provided by appropriate messages rather than tones. The statistics of overload and customer behaviour are significantly different. Loss working is acceptable because network terminating equipment (NTE) can make repeat-call attempts. What must be guarded against is very rapid repeat-call attempts being made across the network during congestion. The design of the DSSS provides for the rate restriction of initial service requests (ISRs) within the peripheral controllers (DLCs, SCs). Thus, the network is protected against both a faulty NTE running amok and malicious attempts to send large numbers of ISRs into the network.

An additional feature for ISDN customers is the provision of closed user group (CUG) supplementary services. These provide considerable security for customers requiring secure data networks. The DSSS provides the interface between the simple *local identifier* available to the customer and the complex *network interlock code* required by the network. The DSSS also ensures end-to-end compatibility of the data service selected by the customer on all digital calls, otherwise the call is rejected.

ADMINISTRATIVE FACILITIES

As well as providing the subscriber interface and call-handling functions, the DSSS must be maintainable and perform certain actions in support of charging and statistical functions.

Maintenance

The DSSS provides maintenance of its own hardware under the control of the maintenance control subsystem (MCS) by using the defined MCS-to-subsystem interfaces and protocols. Some maintenance activities such as testing of the subscriber line interface are carried out by using the facilities of the TNS. The DSSS provides the following facilities:

(*a*) *On-Line Update of Resources* This enables the DSSS functions to be taken into and out of service.

(*b*) *Parent–Dependant Handling of Resources* This ensures that on-line update of resources occurs in a manner consistent with the maintenance state of resources higher in the resource hierarchy.

(*c*) *Diagnostics* This enables tests to be performed on an out-of-service resource.

(*d*) *Routining* This enables tests to be performed on an in-service resource.

(*e*) *System Recovery*

These facilities are required so that faulty units can be detected, reported and subsequently replaced. They are also required to enable recovery from faults and to shed traffic away from faulty equipment so that reliability requirements are achieved. Wherever possible, the DSSS provides fault resolution to a single slide-in-unit. If, in a particular case, this is impracticable, a means exists to enable appropriate rapid manual location and repair of the faulty unit.

Call Charging

Call charging is principally a software function performed by the call accounting subsystem (CAS), which is resident on the PU. However, if a subscriber requires SPM facilities or if a payphone is to be supported, the DSSS is directly involved in providing the line supervision and signalling related to call-charge information. Charge-rate tables compatible with the CAS are downloaded to the MC, which is then responsible for the subsequent generation of meter pulses and supervision of coin-value signalling via the appropriate peripheral control units.

Management Statistics

A feature of System X is the extensive statistical information available from the management statistics subsystem (MSS). The DSSS is involved in providing data to the MSS with regard to the following types of statistic:

(*a*) per-call details collected on a 1-in-*N* basis for both analogue and ISDN telephone calls (for example, call status, time to dial tone and MF4 receiver hold time);

(*b*) per-call details for all ISDN digital calls based on instructions from the MSS (for example, channel used (B or B′), whether the CUG facility is invoked and the clearing reason); and

(*c*) on-demand statistics on a per-concentrator basis relating to resource congestion and traffic levels (for example, MF4 congestion, number of PCM channels busy, number and listing of parked subscribers).

ANALOGUE SUBSCRIBER LINE MODULE

A critical area of any digital exchange is its performance in the line interface area. The DSSS economically meets the diverse requirements of the local line interface and yet retains evolutionary flexibility to take full advantage of the rapidly changing technology available in this area. The following types of line are supported:

(*a*) *Ordinary Subscriber* Both loop-disconnect signalling (7–22 pulse/s) and MF4 signalling are supported. Malicious call identification (MCI) can be provided by either earth or recall signalling.

(*b*) *Coinbox Subscriber* Two types are supported:

(*i*) pay-on-answer coinbox—utilising line reversals (coin-slot control) and 5 kΩ loop signals (coin-pulse detection); and

(*ii*) prepayment coinbox—treated as an ordinary subscriber with SPM.

(*c*) *Private Branch Exchange (PBX) Subscriber* Both manual (PMBX) and automatic (PABX) subscribers are supported. PABXs can be either earth or loop calling.

(*d*) *Outgoing Junctions* For direct-dialling-in (DDI) PBX lines and isolation-working junctions, the DSSS provides loop-disconnect digit pulsing at 10 or 20 pulses/s.

An LC and a number of SLUs form an SLM. A number of SLMs are grouped together on a line shelf. By choice of a standard interface at the plug-in connector on the cables connecting the line shelves to the core of the DSSS, earlier and future types of SLM can be accommodated. Earlier SLMs comprised 32 subscribers' lines per SLM. Current designs comprise 64 subscribers' lines per SLM. This now gives a maximum packing density of 192 subscribers' lines per shelf, an improvement of 100% on earlier designs. Such a dramatic improvement in just two years is due to the rapidly increasing use of LSI and VLSI custom circuits to perform the subscriber interface function.

At the start of the development of the current generation of SLM, British Telecom (BT) recognised the state-of-the-art nature and high risk in this area of digital exchange development and initiated a number of competing compatible SLM developments. As a result of this decision, both Plessey and GEC have now successfully produced SLM designs for use in System X SEP exchanges. With the present designs, functions are usually divided between a high-voltage analogue subscriber line interface circuit (SLIC) and a low-voltage digital subscriber line audio processing circuit (SLAC).

Fig. 2 shows a block diagram of the SLIC functions. The SLIC is responsible for providing line-feed, detection of loop-disconnect signalling conditions, 2-to-4 wire conversion, ringing, test access and over-voltage protection against extraneous voltages on the local line plant (for example, lightning and mains cross). The present technology in the DSSS provides the first three functions by using an LSI custom circuit. A feature of System X is the provision of a constant-current line feed. This has two advantages: it reduces power dissipation and better controls the impedance of the local line as seen by the exchange.

As the voltage fed to line is now a function of loop resistance, it is possible to use this as a measure of the attenuation and to use automatic gain control (AGC) to provide compensation. With the use of switched-mode power supplies, battery boost is possible without excessive dissipation on short lines. As a consequence, higher loop resistances can be supported with an improved standard of transmission

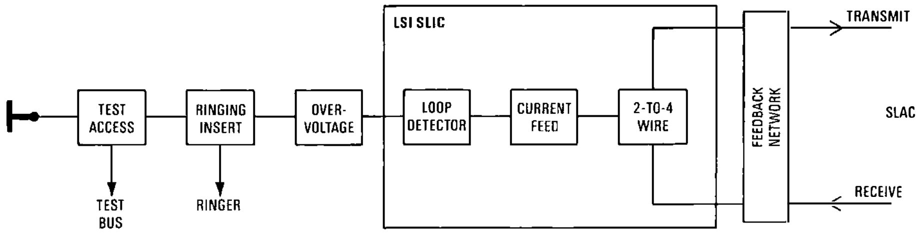

FIG. 2—Subscriber line interface circuit (SLIC)

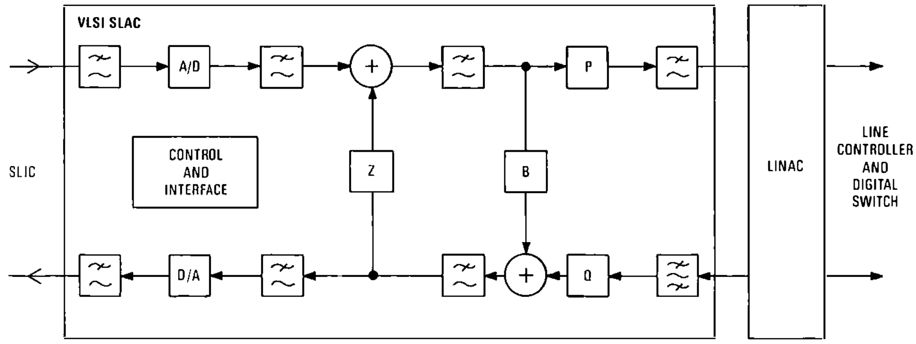

FIG. 3—Subscriber line audio processing circuit (SLAC)

performance. Typically, the current is limited to 40 mA with at least 25 mA provided to a loop resistance of 1800 Ω. This gives the local line planner two distinct economic advantages. Remotely-located customers can now be served without the additional cost of providing loop extenders. Cable of a lighter gauge than hitherto can be provided, with greater opportunity to use cheaper aluminium cables of lower conductivity.

In contrast to the conventional use of a transformer, 2-to-4 wire conversion is implemented with operational amplifiers (Op-Amps). The outgoing signal on the 2-wire side is fed into the 4-wire side via an Op-Amp. The incoming signal on the 4-wire side is connected to the 2-wire side via the Op-Amps used as current generators. A simple balance network attenuates some of the signal now fed back to the transmit path. However, it is the digital filter in the SLAC that cancels out the reflected signal to enable the SLU to function adequately.

At this stage of SLIC evolution, ringing, test access and over-voltage protection are still implemented with discrete components. The use of unbalanced ringing in the UK results in higher voltages than integrated SLIC substrates can at present withstand. Consequently, a simple miniature relay is still an economic solution for this function. Similarly, relays are still used for test access to both exchange and line because of the inability of semiconductor technology to meet the enduring requirement for clean metallic access.

The SLAC is responsible for analogue-to-digital (A/D) conversion (and vice versa), for multiplexing the resultant signal onto a 2 Mbit/s digital highway and for controlling the transmit and receive levels so that the UK transmission plan is met. The production of an A-law encoded PCM signal from an analogue input is already well documented[4]. What is now of particular interest is the implementation of such functions on VLSI custom circuits by using digital signal processing techniques. This has been achieved with a version of the current generation of DSSS line card. Fig. 3 shows a simplified diagram of the processes involved. Whereas Nyquist's criterion dictates that sampling of the analogue input signal shall be performed at least at twice the highest frequency of the analogue input signal (8 kHz) to avoid aliasing distortion, considerable simplification of the filters required is possible if a much higher sampling rate is chosen.

In the send direction, the analogue signal is band limited by a simple resistor-capacitor (RC) filter before conversion to digital form. Because of the very high sampling frequency chosen, the following decimation filters can be implemented by using digital filtering techniques. Fixed time-slot assignment provides the multiplex function onto a 2 Mbit/s digital highway. Linear to A-law conversion is performed separately by means of a LINAC (a custom integrated circuit designed for this purpose).

By choice of suitable coefficients, the digital feedback filter, B, can be used to simulate a balance network. Different line characteristics can thus be accommodated by a simple download of appropriate parameters. Similarly, the digital feedback filter can be set to accommodate a fixed value of nominal line impedance. Relative gain in both transmit and receive directions can be controlled by choice of coefficients for the digital filters P and Q. With information fed back from the SLIC, AGC is achieved without manual intervention. This allows 0–15 dB local lines to be accommodated. In the case of some types of long line, manual intervention may be necessary to account for different gauge cable.

By use of the technology described above, it is now possible to meet the diverse requirements identified earlier for the majority of ordinary and PBX lines by using a single SLU variant with downloaded control of gain and balance parameters.

CONCLUSION

This article has described the digital subscriber switching subsystem which has been developed for the System X SEP programme. As well as an outline description of the architecture, particular attention has been given to the design decisions which have had to be taken to meet the requirements of the UK local line network. Emphasis on the subscriber line interface has been given because of the state-of-the-art technology used. With the large number of lines in the network, it is the author's belief that the economic solution of the problems encountered in this area will underwrite the success of any modern exchange system.

ACKNOWLEDGEMENTS

The author wishes to thank his colleagues in BT and both Plessey Major Systems Ltd. and GEC Telecommunications Ltd. for their contribution to the development described and for the many long and stimulating discussions which have taken place over the past two years to ensure that the DSSS fully meets BT's requirements.

References

1 FOXTON, M. G. System X: Subsystems. Part 5—The Subscribers' Switching Subsystem. *Post. Off. Electr. Eng. J.*, Jan. 1981, **73**, p. 216.

2 RABINDRAKUMAR, K. The Function and Use of Remote Concentrator Units. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 245.

3 PRICE, C. D. E., and BOULTER, R. A. Integrated Services Digital Network. *ibid.*, Jan. 1985, **3**, p. 311.

4 CARLSON, A. B. Communication Systems. 2nd Edition. McGraw-Hill, 1975. p. 294–330.

Biography

Roger Ward is an Executive Engineer in BT's Local Exchange Systems Department. He has responsibility for the definition and evaluation of the subscriber line interface of digital exchange systems obtained by competitive procurement. He joined BT in 1974 after graduating from Trinity College, Cambridge, with an honours degree in Electrical Engineering. His initial work was with the operational area of BT International; during this time he returned to study and obtained an M.Sc. in Telecommunications Systems from Essex University in 1978. In 1979, he joined the System X Development Department where he led the joint BT/contractor team responsible for the system design of the DSSS during the SEP programme.

The Function and Use of Remote Concentrator Units

K. RABINDRAKUMAR, M.SC., C.ENG., M.I.E.E.†

UDC 621.395.34

This article describes the use of the System X digital subscriber switching subsystem as a remote concentrator unit.

INTRODUCTION

The System X digital subscriber switching subsystem (DSSS) is described elsewhere in this *Journal**. This article outlines the function and use of the DSSS as a System X remote concentrator unit (RCU).

An RCU is a System X concentrator (DSSS) remotely located from a parent exchange and connected to the parent System X local exchange by a number of 30-channel 2 Mbit/s pulse-code modulation (PCM) digital paths. The digital switching subsystem (DSS) and the exchange processor (PU) are a part of the parent exchange. One or more RCUs can be located at a remote concentrator centre (RCC). An RCC with two or more RCUs is called a *remote multi-concentrator unit* (RMCU).

† Local Exchange Systems Department, British Telecom Local Communications Services

* WARD, R. C. System X: Digital Subscriber Switching Subsystem. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 241.

SYSTEM STRUCTURE AND CAPACITY

A block diagram of an RCC is shown in Fig. 1. A single RCU can serve up to 2048 connections and has a bothway traffic capacity of around 170 erlangs. Up to eight 2 Mbit/s digital paths can be provided to the parent exchange, depending on the traffic requirements of the RCU.

There is no technical limitation on the maximum number of connections or concentrators that can be located at an RCC within the connection, traffic and call-handling capacity limit of the parent local exchange. However, practical considerations such as relative economics, availability of 2 Mbit/s digital paths between the RCUs and the parent, and a slight reduction in security in some applications are amongst the factors to be taken into account when the optimum size of an RCC is decided.

The transmission of information between the RCU and its parent is over duplicated common-channel signalling links (MTS-G) over time-slot 16 of two of the 2 Mbit/s digital paths between the RCU and its parent exchange.

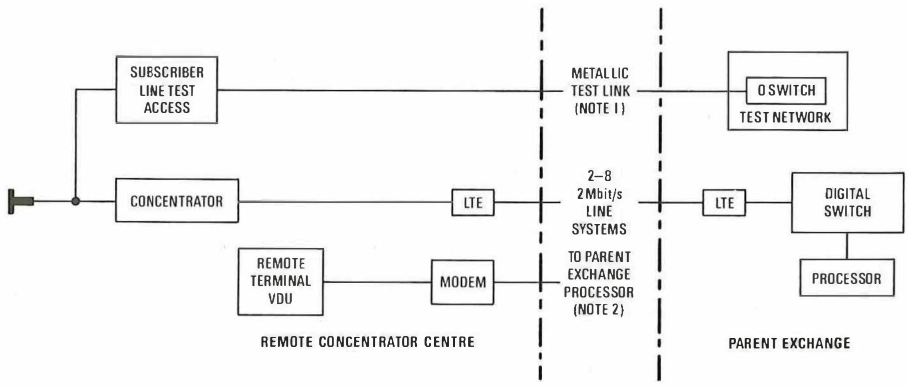

(*a*) Test network at parent exchange

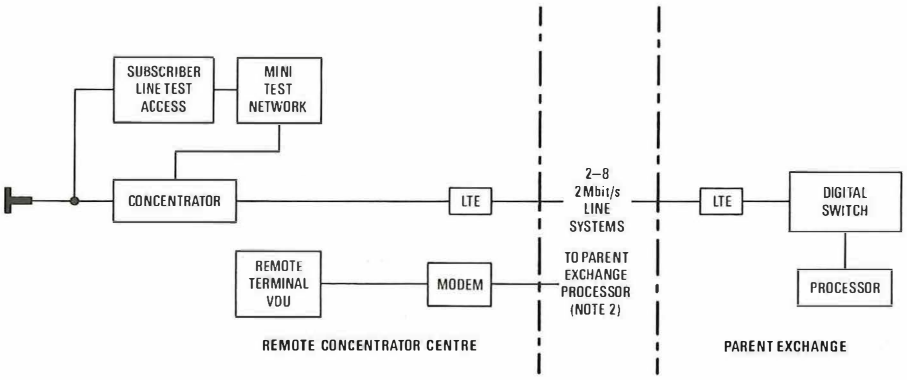

(*b*) Mini test network at remote concentrator centre

LTE: Line terminating equipment VDU: Visual display unit

Notes: 1 Up to four links are required per concentrator
2 One link per site

FIG. 1—Remote concentrator centre

The testing of subscriber lines can be carried out either directly from the parent exchange or, where the resistance of the metallic test links exceeds the maximum permitted value, by the use of a mini test network provided at the RCU. Where required, the mini test network is provided on each RCU.

SYSTEM FEATURES AND FACILITIES

An RCU provides the same full range of System X customer and administration facilities available to a concentrator co-located with its parent exchange. However, the RCU is almost entirely dependent on the parent exchange processor and digital switch for the facilities it provides.

An RCU is isolated when it cannot communicate with the parent processor. This may be because of the loss of all digital paths, the loss of both common-channel signalling links to the parent exchange, or the failure of the parent-exchange processor.

Under isolation conditions, subscribers connected to the RCU can still make own-concentrator calls (at present, there is no facility for inter-concentrator calls under isolation conditions). All call charging, customer supplementary facilities (star services such as abbreviated dialling) and automatic announcements cease on isolation.

The RCU can, during isolation, route emergency (999) and, if required, operator assistance (100) calls over a limited number of analogue junctions to operators.

Studies into the security of network access with the deployment of RCUs have been carried out that take into account the failure rates of exchange equipment and transmission plant. The studies show that, although the nature of the failures may change, there is no significant increase in the risk of overall loss of service by the deployment of RCUs.

To reduce the risk of the RCC being isolated by cable failure, the digital paths to the parent exchange are routed, where possible, via different cables and different ductways. Even if diverse routeing is not provided, the risk of isolation is still low and comparable to a remote exchange served by a group of junctions, all of which are routed via a single ductway.

A remotely located concentrator has some additional facilities for handling alarms (for example, environmental alarms) and has an optional facility to enable preference working to be invoked during isolation conditions.

A System X local exchange can serve as a parent for RCUs in charge groups remote to the charge group containing the parent exchange. The limit on the total number of charge groups supported for originating subscribers, including RCUs, is 16.

RCU APPLICATIONS

RCUs can be used for a number of potential applications, and some of the applications for which early System X RCUs are likely to be used are:

(*a*) as a means of rapidly providing digital exchange facilities, including an integrated services digital network (ISDN) capability, to target areas of commercial potential (for example, business centres);

(*b*) as an additional unit to cater for growth to avoid installation of more analogue equipment;

(*c*) as an intermediate complete replacement of a local exchange unit which will later grow into a System X local exchange;

(*d*) as an 'overlay' unit in an exchange which is to remain analogue for a period of time;

(*e*) as a means of turning round an exchange to a digital local exchange within the existing accommodation;

(*f*) as a complete replacement of an analogue local exchange; and

(*g*) as a means of providing digital facilities early to remote or isolated communities.

The overlay arrangement (*d*) is being used to provide System X facilities, where required, in advance of complete exchange replacement by transferring subscribers' lines.

Normally, RCUs used in applications (*a*) to (*e*) above would later be absorbed into a co-located System X local exchange.

SUBSCRIBERS' NUMBERING

Each RCU is normally allocated a discrete number block sufficient to cater for the forecast requirements.

In director areas, the RCU will normally be given a dedicated all-figure number (AFN) code that is spare or has been allocated to cater for growth at the unit where the RCU is being installed.

An RCU can share an AFN code with the analogue unit. However, it will then be necessary for all local exchanges and switching units with direct routes to the analogue unit to examine the first numerical digit to determine routeing. This may not be possible with some existing switching systems.

In non-director area applications, the RCU is normally identified uniquely by the 'DE' or 'DEF' digits of the national number (see Fig. 2). To ensure the maintenance of transmission standards, it must be possible unambiguously to identify the RCU at the analogue group switching centre (GSC) in order to route incoming analogue main network traffic destined for the RCU to the parent digital local exchange.

Where RCUs are used in charge groups remote from the parent exchange, the number range of the RCU will lie within that of the numbering group containing the RCU.

The replacement of GSCs with digital main switching units (DMSUs) will remove some of the existing constraints on RCU numbering. However, RCUs will continue to be allocated discrete number blocks, particularly in cases where future re-parenting is planned.

TRAFFIC ROUTEING

Outgoing traffic from RCUs will be treated in the same way as the traffic from the customers on the parent local exchange, irrespective of whether the RCU is in the same charge group or in a remote charge group to the parent exchange.

To ensure the maintenance of the existing transmission standards, the link between the terminal (incoming) analogue main network unit (that is, a GSC) and the parent local exchange, which will normally be via the DMSU or the digital principal local exchange (DPLE), will be provided on digital plant.

The main network switching centre (MNSC) serving the RCU home charge group will have a direct route to the RCU parent exchange or to the DMSU or DPLE of the parent exchange. Where the parent exchange supports a number of RCUs, then more than one route may be required (see Fig. 3).

Code 1/999 traffic will normally appear at the auto-manual centre, directory-enquiry bureau or repair service centre appropriate to the parent exchange, although special arrangements may be provided to route the traffic to appropriate alternative centres.

RE-PARENTING RCUs

Ultimately, several RCUs would be integrated into a local exchange unit with the installation of an exchange core (primarily, the exchange processor and digital switch). In these cases, the RCUs will be re-parented on to a co-located host processor.

Fig. 4 indicates an outline of the different stages of re-parenting.

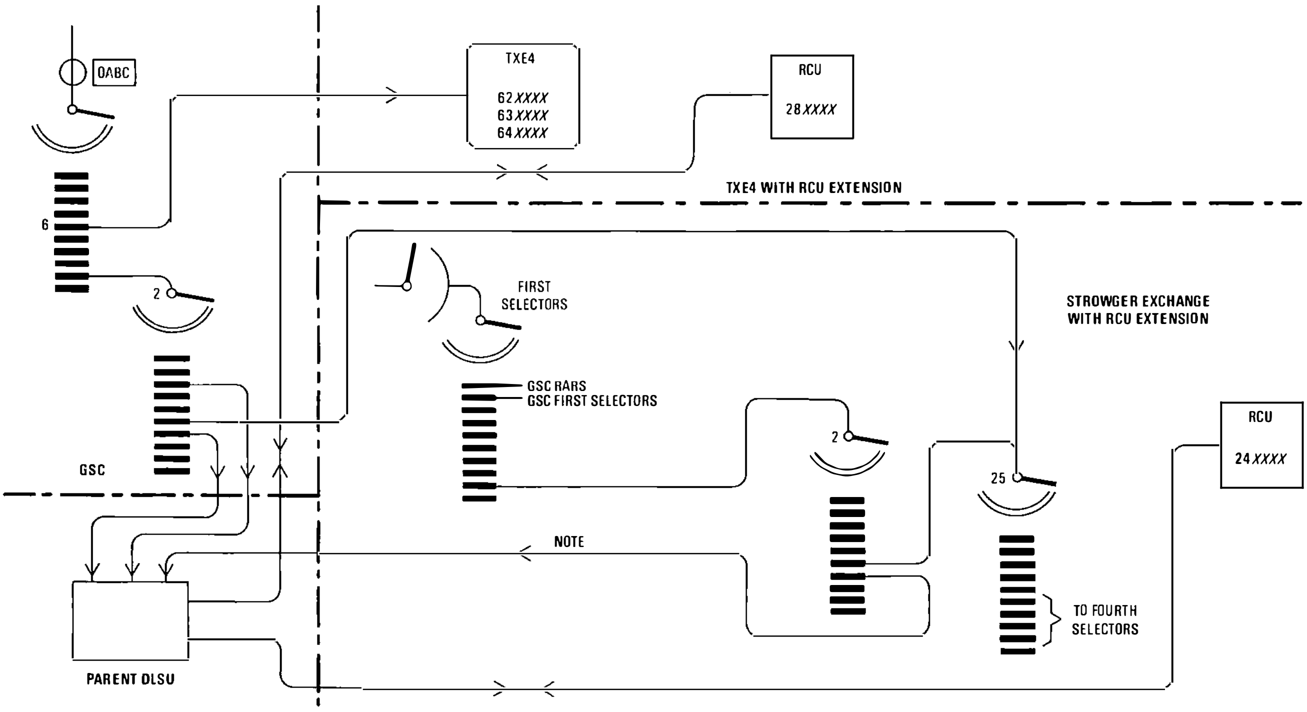

DLSU: Digital local switching unit GSC: Group switching centre RARS: Register access relay-set RCU: Remote concentrator unit

Note: If the GSC was a TXK1 exchange, the '24' route could be taken to the GSC rather than directly to the parent DLSU

FIG. 2—Provision of an RCU at a Strowger and a TXE4 exchange

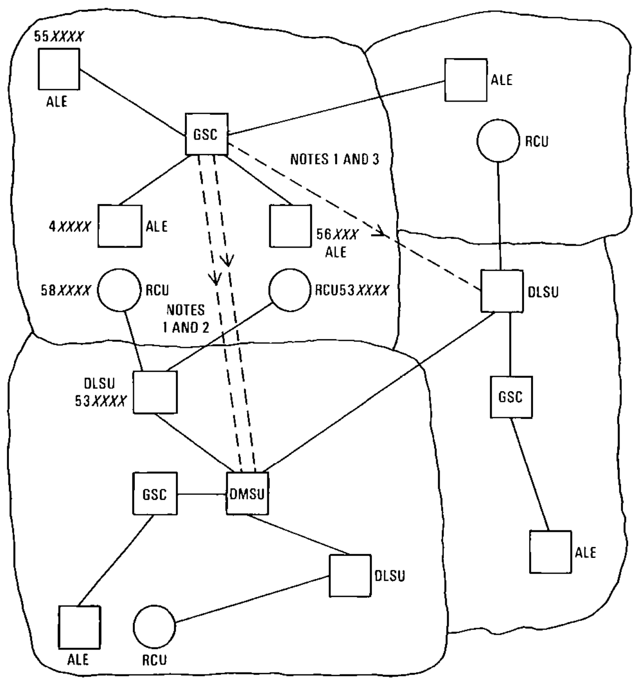

ALE: Analogue local exchange
DPLE: Digital principal local exchange
DLSU: Digital local switching unit
DMSU: Digital main switching unit
GSC: Group switching centre
RCU: Remote concentrator unit

Notes: 1 Direct routes are required for incoming traffic from the analogue network
2 Only one route is required if the GSC is a TXK1 exchange. This route may be taken direct to the parent DLSU instead
3 This route may be provided to the DPLE rather than direct to the parent DLSU

FIG. 3—Provision of RCUs in remote charge groups

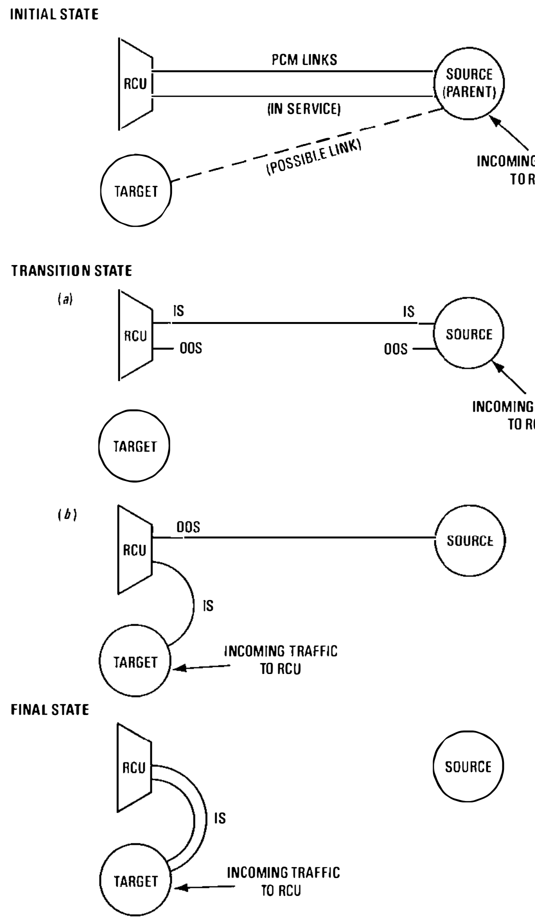

IS: In service
OOS: Out of service
PCM: Pulse-code modulation
RCU: Remote concentrator unit

FIG. 4—Stages in re-parenting an RCU

The salient aspects of re-parenting RCUs are as follows:

(*a*) The subscribers connected to the RCUs are transferred in bulk to the target exchange.

(*b*) At the time of re-parenting, simultaneous routeing changes are carried out at distant exchanges (for example, the GSC, the parent exchange and all exchanges with direct routes to the parent exchange). As indicated previously, this means that complete number blocks have to be allocated to the RCUs to facilitate re-parenting of RCUs.

(*c*) Several concentrators can be re-parented at the same time, assuming that the number block allocations have been planned accordingly.

(*d*) An off-line computer facility, known as the *data compiler and decompiler* (DCD), is used to decompile the exchange data of the source and target exchanges, and to extract the data relating to the RCUs to be re-parented. The final exchange data for the source and target exchanges are prepared by merging the prepared data changes and compiling the new source and target data.

(*e*) The exchange data will need to be 'frozen' during the above process. Any data changes will have to be input subsequently by man–machine commands.

(*f*) The new system images are then loaded into the source and target processors.

(*g*) For the subscribers being transferred, calls in progress will be lost during the re-parenting process. Hence, re-parenting will normally be carried out during a period of light traffic. Calls in set-up will be lost for a very short period during the re-start process in both the source and target exchanges. However, the period of service disruption is expected to be short.

ECONOMIC BENEFITS

The results of cost studies covering a range of exchange sizes, customer calling rates and network configurations indicate that a modernisation strategy using System X RCUs gives significant economic benefits over a digital implementation strategy based solely on System X processor exchanges.

Economic factors relating to the use of RCUs are as follows:

Cost Savings

There are significant savings in processor and switching costs at the RCC, and some cost savings relating to other subsystems that will need to be provided if a digital local exchange is installed at the RCC.

Additional Costs

Against the savings given above, some additional costs are incurred:

(*a*) Additional transmission costs arise because 2 Mbit/s digital paths to the parent exchange have to be provided. The 2 Mbit/s paths can, however, be re-used for junction traffic if the RCUs are absorbed by a co-located local exchange at a later date.

(*b*) Some increase in cost at the RCC results from the provision of a mini test network (if needed), additional alarm facilities and, where provided, analogue junctions.

(*c*) There is an increase in the processor and switching costs at the parent exchange, but as these are marginal additional costs, they would be substantially less than the equivalent cost savings at the RCC. The extra switching and processor capacity can be re-used to cater for growth or for turn-around of analogue equipment if the RCUs are re-parented at a later date on a co-located local exchange.

(*d*) Trunking changes and re-routeing work at distant exchanges incur some costs.

OTHER BENEFITS

In addition to the cost benefits given above, the selective use of RCCs has a number of other advantages:

(*a*) The early provision of digital exchange facilities to a larger number of customers over a wider catchment area can be made.

(*b*) A quicker response in meeting service demands for advanced facilities at a particular location can be achieved (for example, the provision of digital exchange facilities at business centres).

(*c*) ISDN facilities can be provided more economically by deploying RCUs, especially in the early years of low penetration.

(*d*) Maximum use can be made of the limited number of installations of System X local exchange processors, particularly in the early part of the System X implementation programme.

(*e*) Accommodation requirements for exchange and power equipment are reduced, hence obviating the need for building extensions.

(*f*) Initial effort required for installation and commissioning is reduced.

SUMMARY

This article has outlined the functions, applications and some of the planning implications of the use of System X RCUs in the growing digital network. RCUs will have an important role in the modernisation of the existing analogue network, and will enable the benefits of digital switching to be offered over a larger part of the British Telecom (BT) network, especially in the early stages of the implementation of digital local exchanges.

ACKNOWLEDGEMENTS

The author wishes to thank his colleagues in BT Local Communications Services on whose work this article is mainly based.

Biography

K. Rabindrakumar is Head of the System X Local Exchange Planning and Design Group in BT's Local Exchange Systems Department. He joined BT in 1970 and worked on exchange planning and equipment provisioning standards for TXE4 and subsequently TXE4A electronic exchanges. In 1978, he joined the System X Launch Division and was responsible, initially, for the hardware design and dimensioning aspects and, later, for the software and data generation procedures for Release 1 System X exchanges. In 1980, he was seconded to the Hong Kong Telephone Company as a Senior Engineer on digital switching and was responsible for the preparation of the technical specifications for an international tender for digital local and tandem exchanges and an international toll exchange, and for the subsequent evaluation of the tenders. He returned to BT in 1983 and now heads a group responsible for System X digital concentrators, customer facilities and interworking in the Local Exchange Evaluation and Planning Support Division.

Power for System X

W. P. OWEN, B.SC.†

UDC 621.395.34 :621.311.6

This article briefly descibes the concept of the power equipment rack and shows how it has been applied to serve System X installations in the form of the Power System 2020.

INTRODUCTION

During the early development of System X, it was recognised that much shorter planning and installation time-scales were necessary and, hence, that similar criteria needed to be applied to the provision of DC power supplies. This gave British Telecom's (BT's) Power Division the opportunity to appraise thoroughly the state-of-the-art in power-conversion and battery technology. It was realised that the technology was becoming available to develop a new generation of power systems in which both rectifiers and batteries could be mounted within System X equipment practice and installed *en suite* with switching equipment. This could replace the traditional centralised power plant, and provide economy in power accommodation and distribution busbars, which had always been prone to working-party faults. Most important of all, these power systems could be provided and installed well within the radically shortened time-scales for manufacturing and installing switching equipment.

Development of the new rack-mounted power system, known as the *Power System 2020* (PS2020), started in 1981, with a target date for its introduction into service in early-1984. A performance specification was drawn up and development contracts were placed with three manufacturers. At the same time, discussions went ahead with a major UK battery manufacturer to develop a battery suitable for rack mounting. The battery was to have a much increased energy density ratio and operate on an oxygen recombination principle, limiting the products of electrolysis given off to the rack environment. Prototype racks and batteries were produced during 1982, and exhaustive testing took place during the following years. One of the prototypes was exhibited at the 1982 Washington Intelec Conference, where it proved to be the first application of a complete rack-mounted power system in the world.

DESCRIPTION

The PS2020 is based on a TEP-1(H) rack, which has six shelves: the lower three accommodate batteries and the top shelf houses rectifiers. The other two shelves contain a DC connection panel, power alarm unit (PAU), and the incoming AC supply fuses and associated isolator. Fig. 1 shows a production model PER being installed at Guildford ATE.

Up to three batteries (one per shelf) can be accommodated on the rack. The batteries consist of a number of monoblocks containing either two or three cells. There is no free electrolyte within the cells, this being fully absorbed into a micro-

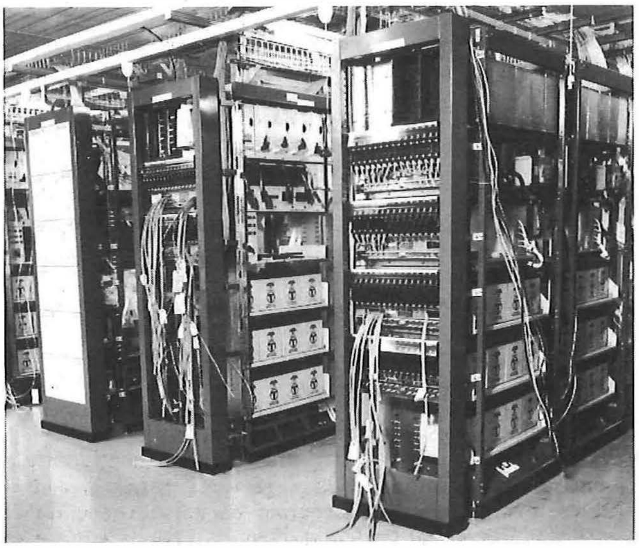

FIG. 1—Power equipment racks at Guildford ATE

porous glass-fibre separator surrounding the plates. Each battery contains either 8 × 3 cell monoblocks or 7 × 3 cell plus 1 × 2 cell monoblocks where a lower supply voltage is required. The major advantages of the new sealed cells are:

(*a*) they are fully-sealed under normal operating conditions,
(*b*) they are virtually maintenance free,
(*c*) they have a high energy density, coupled with low weight, and
(*d*) their construction makes them easier to handle and install.

Each battery is floated at a constant potential equivalent to 2·27 V per cell, and an eight-monoblock 24-cell battery is capable of supplying 2 kW for one hour when discharged to 1·92 V per cell.

The Rectifier No. 160* used in the rack is of the switch-mode type, which enables a significant reduction in size and weight over more traditional types of rectifiers. The main points that make these types of rectifiers attractive to use in rack-mounted power systems are:

(*a*) voltage conversion and isolation takes place at high frequency, enabling the size and weight of the transformer to be reduced;
(*b*) energy storage is at relatively high voltage, thereby reducing the size of capacitors required;
(*c*) high-frequency rectification reduces the size of the filter components; and
(*d*) output regulation is achieved by using pulse-width

† Energy, Transport and Accommodation Department, British Telecom Local Communications Services

* WITTEY, B. A., and HENDERSON, J. P. Modern DC Power Systems: The Power Equipment Rack. *Br. Telecommun. Eng.*, Oct. 1983, **2**. p. 186.

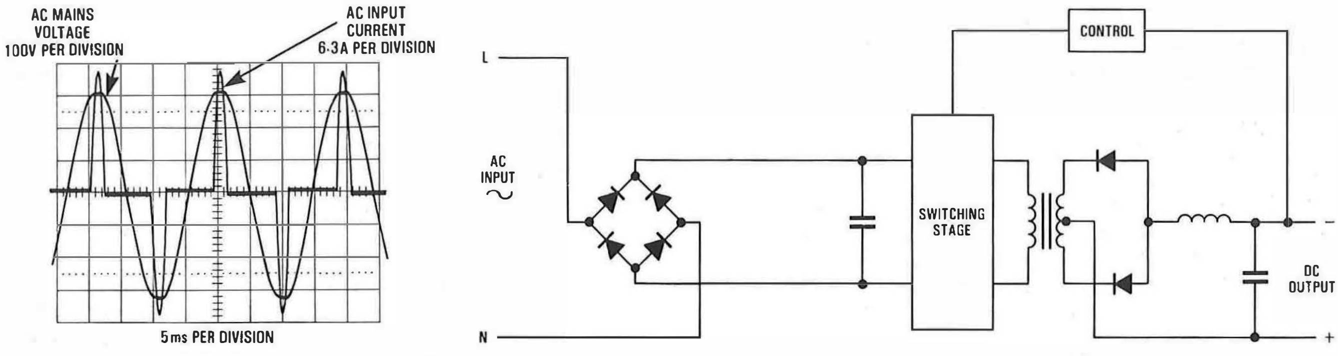

Note: The oscillogram shows AC input waveforms for a conventional switch-mode power supply delivering 1600 W

FIG. 2—Switch-mode power supply and typical input waveforms

modulation techniques, which minimise the power loss in the regulator and the size of the heat sink.

Fig. 2 shows a diagram of a switch-mode power supply and typical input waveforms. Each rectifier has a single-phase AC input with an output of 28 A at a nominal 54·9 V, although this can be modified for applications where a lower output voltage is required. Five rectifiers are mounted in the rack to give an output of 6 kW, one rectifier acting as an on-line spare.

A main end panel, as used at the end of a standard System X equipment suite, is used in conjunction with the PS2020, and is shown in Fig. 3. This end panel contains up to five 15-way fuse panels, which serve as distribution points for 10 mm^2 cables that supply power to the equipment shelves.

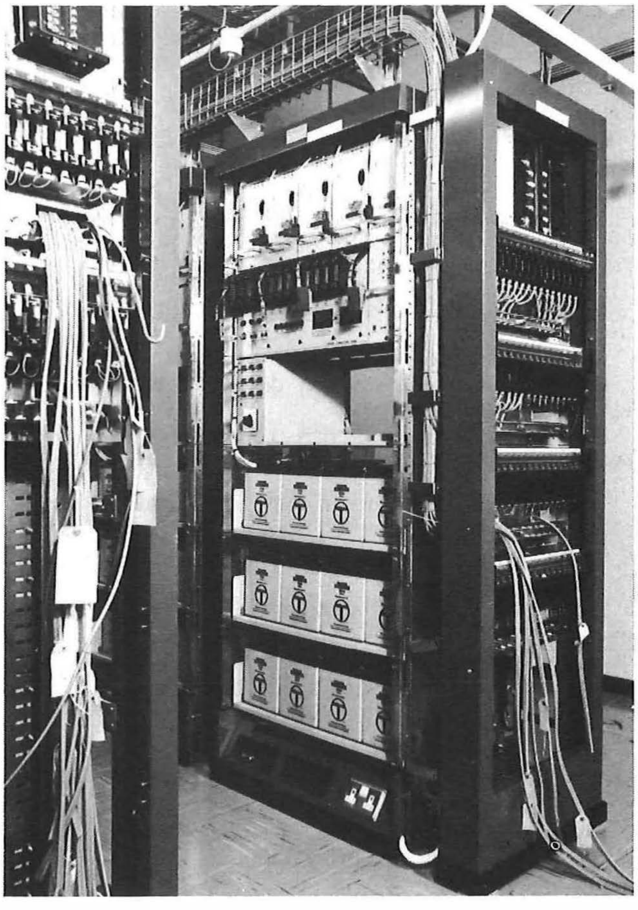

FIG. 3—Main end panel

In addition, a suite distribution connection panel provides a cabling interface betwwen the PS2020 and the fuse panels.

USE

Clearly, the rack-mounted power concept has a number of advantages over traditional centralised power plant; however, it must also be cost effective. It can be shown that, for a 'green-field' site, where power plant with spare capacity does not already exist, there are substantial economic benefits in using the PS2020 rather than traditional power plants; these range from savings of over 50% for an exchange load of 12 kW, to 25% for an exchange load of 80 kW or more.

A substantial capital investment in power plant exists in the network at present, but most of these plants are obsolescent and over 10 years old. The cost of maintaining them is relatively high and it will become progressively more difficult to keep them in service; hence, providing PS2020 when the exchange system is being modernised with System X has obvious benefits.

The question of whether or not to use the PS2020 to replace the comparatively smaller number of relatively modern power plants is perhaps more difficult to answer. Spare capacity often exists because of the relative inflexibility of centralised power plants and because exchange loads have to be planned for 10 or 20 years ahead. Therefore, the decision on whether to use the PS2020 at these sites or continue with the existing power plant can be taken only after a full appraisal of all the factors involved when modernisation of the exchange takes place.

PLANNING AND INSTALLATION

Because the PS2020 is installed *en suite* with System X equipment, power planning is closely related to the exchange-design function, and power racks and their associated distribution arrangements can be regarded as a separate power-equipment subsystem (PES).

An important planning activity is dimensioning; that is, determining the number of racks, rectifiers, batteries and fuse panels required for a particular installation. This can be tackled only after the basic design of the exchange, in terms of the numbers of racks and their power requirement, has been completed. The application of some simple rules, such as limiting the number of PS2020s to no more than two per suite and allowing cross-feeding to adjacent suites, then permits preparation of the exchange floor layout.

Some System X equipment, mainly in the signalling interworking subsystem and in line terminating equipment, requires a supply to narrower voltage limits than that mentioned earlier.

A narrow-voltage-limit (NVL) supply is provided from a PS2020 having its output voltage set at a nominal 52·65 V. The problem of ensuring interchangeability is solved by making all Rectifiers No. 160 with a dual-voltage output. The output voltage of a particular rack/rectifier combination is then determined by the presence or otherwise of a pair of links contained in the DC output plug to each Rectifier No. 160. The links, when inserted, short out resistors in the rectifier control circuit; this alters its output voltage and makes it fully interchangeable with wide-voltage-limit (WVL) systems. A PS2020 is therefore engineered for either narrow or wide voltage limits during manufacture, although the change in voltage can be made on site if necessary. When it is operating within the narrow-voltage limits, a 23-cell battery is required as opposed to the normal 24-cell version.

Installation involves erecting racks in position at the end of equipment suites, fitting the required number of rectifiers and batteries, and making the AC input connections. Once the main end panel and its associated hardware have been installed, the DC distribution to racks and shelves can be completed.

COMMISSIONING

The PS2020 requires minimal on-site testing which should help to minimise delay and reduce on-site costs, and prove a considerable advantage at the commissioning stage.

The test equipment required to carry out commissioning checks consists of a DC load frame, a 500 V ohmmeter, a digital multirange meter and a line loop impedance tester. Tests fall into the following broad categories:

(*a*) checking connections to ensure that the main AC and DC connections are secure and correctly terminated;

(*b*) checking insulation and earth-loop impedance;

(*c*) checking the output of each rectifier in turn to ensure that it is operating and delivering its output within limits;

(*d*) checking the alarms to ensure that the three light-emitting diode (LED) indicators on each rectifier panel are operating correctly. Two of these indicate that the rectifier input and output are healthy, whilst the third indicates a current-limit condition and is purely for information. In addition, there are high- and low-voltage prompt alarm conditions to be checked on the power alarm unit; and

(*e*) ensuring that the batteries are fitted correctly and are capable of supplying the load.

MAINTENANCE

Routine maintenance of the PS2020 is confined to an annual battery-performance check and a quarterly test of the alarms. All the functional units in the system are designed for easy replacement on a pull-and-replace basis, the repair of faulty rectifiers being undertaken by BT at designated electronic repair centres. The use of multiple rectifiers coupled with a low mean-time-to-replace (MTTR) of five hours, ensures high system reliability. Studies show that the design aim of 1000 years mean-time-between-failure (MTBF) for a DC power system using PS2020 will easily be met.

A microprocessor-based supervisory and alarm unit, known as the *power management scheme* (PMS), has been developed, and this will enable operation and maintenance centres to supervise the power system along with the rest of the exchange. It is capable of automatically monitoring the state of the power system at regular intervals and of initiating alarms on a more selective basis than at present. It will also be possible to check completely automatically battery-discharge capacities, and interrogate the equipment to determine the power status and battery reserves remaining under discharge conditions.

The PMS, which is to be standard equipment on the PS2020, is now becoming available.

THE FUTURE

The Power System 2020 described in this article is the beginning of a family of similar systems, using the same basic building blocks. The concept is finding a variety of new applications in the markets for both the customer and public switched telephone networks. Systems are now in use in transmission stations, UAXs and customers' premises.

The concept of rack-mounted power enables advantage to be taken of any new developments in power conversion that could arise in the future. It is also a further stage in the gradual move towards distributed power with a possible reduction in the number of conversion stages involved. System X has the most modern power conversion system available in the world today and one that is capable of evolving in the future along with System X itself.

Biography

Paul Owen is Head of the DC Operational Support Group in BT's Energy, Transport and Accommodation Department. In 1967, he gained an honours degree in electrical engineering at Salford University and joined BT in 1970. After working on power and accommodation services planning in BT's North West Regional Office, he transferred to his present position in 1982. He is currently involved in co-ordinating the introduction of the PS2020 into service with System X.

The Evolution of Cooling Techniques for Digital Exchanges

P. HOWELL, B.SC., D.M.S.†

UDC 621.395.722 : 621.374

The matching of exchange equipment performance to a practical operating environment is a key task for the equipment designer and the building-services engineer. This article examines how the environmental requirements for British Telecom's digital exchanges were developed and how these were translated into building-services hardware.

INTRODUCTION

At the start of the System X development programme, typical average power densities (W/m^2) for the exchange types then in service[1] were:

electromechanical (Strowger) step-by-step:	50 W/m^2
electromechanical crossbar:	70 W/m^2
small semi-electronic (TXE2):	80 W/m^2
large semi-electronic (TXE4):	180 W/m^2

Estimates of the likely power densities in System X exchanges covered a range of values with an upper limit of around 600 W/m^2. While such power densities were common in computer centres, they were too high for traditional exchange cooling practices. Of particular concern were the 4000 or so small rural exchanges where the available space and the quality of the AC power supply would have made the provision of reliable, active cooling both difficult and costly.

ENVIRONMENTAL REQUIREMENTS

A consequence of these anticipated power densities was the requirement for digital exchanges to survive at elevated temperatures. A power density of 400 W/m^2 was considered to be a likely power density for a digital exchange at this time. Under conditions of no cooling, the ambient temperature in such an exchange would reach 55°C in about five hours. This was considered possible should there be a prolonged failure of the AC supply at a high-power-density digital exchange which was backed up by a large central battery inherited from the exchange system it had replaced (that is, larger than strictly necessary). The exchange, therefore, had to survive a 5 hour excursion of room ambient between 40°C and 55°C with the power applied to the system.

To reconcile this and other requirements related to export requirements, the following environmental envelope was chosen:

Temperature range:	+5°C to 40°C
Normal occupied room temperature:	24°C
Relative Humidity:	10% to 70%
Absolute Humidity:	2 g/m^3 to 20 g/m^3
Short-Term Temperature Range:	−5°C to +55°C
Short-Term Humidity:	5% to 90%
Short-Term Absolute Moisture:	1 g/m^3 to 25 g/m^3

† Energy, Transport and Accommodation Department, British Telecom Local Communications Services

The environmental systems were then planned using this envelope around the structural limitations imposed by the existing buildings. Six fundamental tasks were undertaken to translate the requirements into building-services hardware. These were:

(*a*) the identification of specific cooling requirements as a function of increasing power density,

(*b*) the selection of possible cooling systems to meet the requirements of (*a*) above,

(*c*) the objective assessment of the reliability of the chosen systems,

(*d*) the creation of an environmental simulation model to test cooling design and reliability,

(*e*) the development of cooling plant to meet a particular set of requirements, and

(*f*) the creation of a reporting and information system to assess system performance.

Table 1 shows how tasks (*a*) and (*b*) were resolved and implemented. The Table applies to both switching (TEP-1H) and transmission (TEP-1E) rack practice[2].

The TEP-1H rack height was limited to 2·2 m to prevent excessive heat dissipation and to retain sufficient headroom to promote good air mixing. This height also ensured that all buildings would be capable of accommodating the system.

The assessment of reliability (task (*c*)) was a particularly intractable problem. Prior to the System X era there was little incentive to monitor the performance of cooling plant. No historical performance data existed and reliability data for cooling plant was scarce. Through the use of sampling, interviews and research, the basic reliability parameters were established[3]. Table 2 shows the 10-year minimum-reliability targets of the cooling system as a function of switched traffic.

The development of a computer program to model temperature rise after a cooling-plant failure has proved to be a powerful design tool. From data relating to the power density, building construction, air infiltration and external ambient conditions, a thermal footprint of the equipment room is estimated by using the model. Fig. 1 shows a typical output from the model. By varying the input parameters, the effects of structural changes to the building can be determined; in addition, the minimum DC battery reserve which would keep the exchange thermally safe can be estimated.

At an early stage of the investigations, a need for reliable low-energy low-maintenance cooling plant was identified. As suitable commercial plant was not available, a small number of manufacturers were commissioned to develop a range of standard cooling units. The units all use the same control system and employ fresh-air cooling whenever possible. A full description of these units can be found in Reference 4.

TABLE 1
Preferred Cooling Systems for Dedicated TEP-1H and TEP-1E Installations

Power Density	Preferred Cooling System	Alternative System
Below 120 W/m²	Fan powered ventilation (no refrigeration) using a single-speed fan-filter unit to bring in fresh air at a rate of 3 air changes per hour. Fan controlled by thermostat set at 24°C Unducted (free blow) air distribution Exhaust by natural leakage Reliability requirement is met automatically	Natural ventilation may be used instead of fan-powered ventilation when ambient pollution is low and dust problems are rare (average ambient particle deposition below 50 mg/m² per day). Actual deposition should be assessed by local trial. Install two grills 450 mm from floor level on one external wall, and another two grilles at high level on the opposite wall Each grille to be protected by insect screen and to have true free area of 0·01 m²
Above 120 W/m² Room load less than 20 kW	If room is permanently staffed: Refrigerated cooling incorporating fresh-air cooling where possible: 24°C room temperature Unducted (free-blow) air distribution Exhaust by pressure-relief ventilators (fresh-air mode) High reliability requirement If room is not permanently staffed: Modular fan-powered ventilation units using variable-speed fans, each controlled by a separate thermostat set at 24°C. Design based on 0·2 m³/s per kW of cooling load Unducted (free-blow) air distribution Exhaust by pressure-relief ventilators High reliability requirement	Ventilated ceiling where physical constraints make unducted distribution impractical Exceptionally, ducts and diffusers air distribution
Above 120 W/m² Below 260 W/m² Room load greater than 20 kW	Refrigerated cooling incorporating fresh-air cooling where possible: 24°C room temperature Unducted (free-blow) air distribution Exhaust by pressure-relief ventilators (fresh-air mode) High reliability requirement	Ventilated ceiling where physical constraints make unducted distribution impractical Exceptionally, ducts and diffusers air distribution
Above 260 W/m² Room load greater than 20 kW	Refrigerated cooling incorporating fresh-air cooling where possible: 24°C room temperature Unducted (free-blow) air distribution Exhaust by pressure-relief ventilators (fresh-air mode) High reliability requirement	

TABLE 2
Reliability Targets for TEP-1H Installations

Total Traffic Carrying Capacity (erlangs) of System X Unit	Cooling Plant MTBSF Target*
Less than or equal to 500	200 years
Greater than 500 but less than 1000	400 years
Greater than 1000 but less than 20 000	2000 years

* The cooling plant will be deemed to have failed if the temperature of the apparatus room exceeds 40°C or falls below 5°C
MTBSF: mean time between system failures

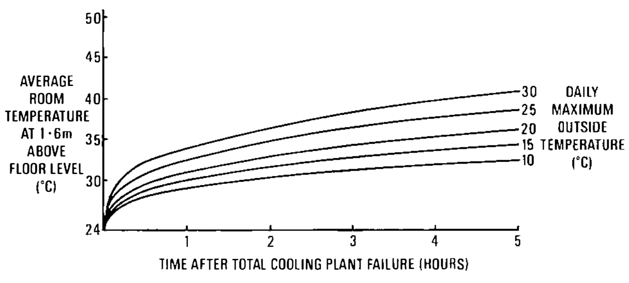

FIG. 1—Rise in room temperature following cooling-plant failure (200 W/m² equipment)

PRESENT SITUATION

As the design and development of System X has progressed, so too has the quality of the information available to the building-services engineer within BT increased. Over this period, the advances in switching technology have moved rapidly, and these are reflected in the cooling requirements. Fig. 2 shows the power density of the latest version of System X[5]. To date, the upper power-density limit has been set at 400 W/m² for present technology. There is now 99%

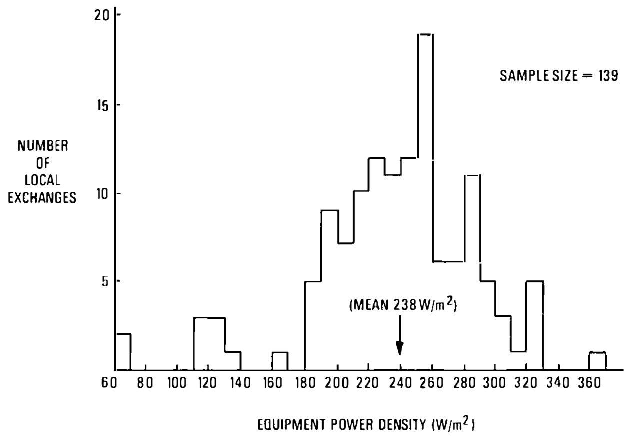

FIG. 2—Power density of System X local exchanges

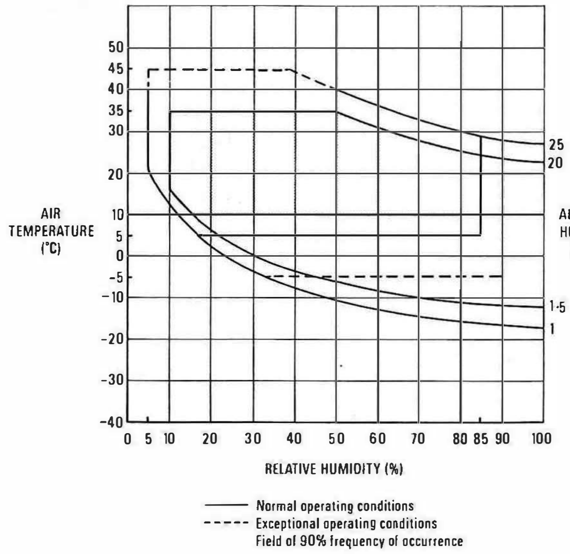

FIG. 3—Climatogram for operation in indoor rooms

FIG. 4—Standard ISO container fitted with high-reliability cooling plant

confidence that exchange power densities will not exceed 360 W/m². The use of power equipment racks[6] (PERs), with their 1 hour reserve automatically ensures the thermal safety of the largest exchanges. A high-power-density exchange with a large central battery is now a 'special' rather than a 'likely' situation, and suitable precautions can be taken to improve the security of the AC supply.

Recent temperature tests on a System X exchange demonstrated a traffic switching capability at room ambient temperatures in excess of 50°C. Because the exchange is only required to survive at these temperatures, the thermal robustness of the design and the match between exchange performance and thermal environment were clearly demonstrated.

On a wider environmental front, BT co-operates with other European administrations through the CEPT†. Since October 1984, BT has used the proposed CEPT climatogram as its prime standard for telecommunication equipment environments (see Fig. 3). The adoption of this standard will allow BT greater freedom in exchange purchase with reciprocal competitive benefits for System X in Europe.

FUTURE

There seems little doubt that, whereas total exchange power consumption will decrease, rack power densities will increase because the exchange will no longer be spread out over a large area. This will require the components of particular racks to have great thermal robustness. Additionally, traditional exchange functions may be housed in a variety of enclosures remote from the exchange. While these will have low absolute powers, they too could generate high power densities. Hence, the main direction of cooling development will lie in identifying single low-cost high-reliability cooling methods of the following types:

Natural and forced-convection cooling
Fresh-air cooling in very small buildings
Passive cooling
Raised modular floor air supply from uninterruptable power supply backed fans
Direct rack cooling
Ground cooling for street-located equipment.

Fig. 4 shows a simulated 2000-line remote concentrator unit fitted with high-reliability cooling plant housed in a standard ISO container. Intended for mobile/emergency/service continuity during equipment turnrounds, the power dissipation is typically 3–4 kW, and gives a power density of around 400 W/m².

The use of simulation models will become increasingly important, and these will be based on fundamental research now being undertaken into building heat-transfer mechanisms under conditions of cooling failure.

Future exchanges will probably make use of forced convection, high-power-density racks and a far greater use of fresh-air cooling. Refrigerative cooling will still be used, but, perhaps, for peak-lopping rather than for load matching. The cooling techniques developed for the present design of System X exchanges will satisfy most of the applications which will arise during the next 10 years.

† CEPT—Conference of European Post and Telecommunications Administrations

References

[1] SPURGIN, D. A. Low Energy Cooling for Digital Telecommunications Systems. INTELEC 1979.

[2] BRITISH TELECOM Cooling for TEP-1H and TEP-1E Installations. ISIS Guide: BES/VAC/B182.

[3] BRITISH TELECOM Reliability of Cooling Plant for TXD Exchanges. ISIS Guide: BES/VAC/B157.

[4] BROWN, W. G. Equipment Cooling for Modernisation. *Br. Telecommun. Eng.*, Jan. 1984, **2**, p. 246.

[5] Thermal Control Working Group (BT, PTL, GEC), Discussion Paper SX/TCWG/DP(84)01, Issue 03, GEC Telecommunications Ltd., Coventry.

[6] OWEN, W. P. Power for System X. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 249.

Biography

Peter Howell is Head of Building Services Forward Study Group. He joined the Building Services Division of BT in 1974 after graduating from Middlesex Polytechnic. Since then he has been involved with lifting equipment testing, BES maintenance and the cooling of telecommunications systems. He is currently the BT representative on the CEPT Environmental Standards Working Group.

System X: Common-Channel Signalling—Progress on Installation and Testing

D. C. BUTTERWORTH, B.SC., M.B.C.S.†, and A. WILLIAMS, B.SC.*

UDC 621.395.34

This article addresses the operational use of CCITT‡ No. 7 signalling in the UK network, the facilities provided to assist in its use, the proving of the system and the current in-service experience. Operational requirements for the future indicate the necessity of connecting System X exchanges to other exchange types in the network and other service providers at the common-channel signalling level; potential problems in this environment are briefly considered.

INTRODUCTION

The common-channel signalling system adopted for System X is based on the CCITT Signalling System No. 7[1, 2]. The general principles and advantages of common-channel signalling, and its implementation in the message transmission subsystem (MTS) of System X, have been discussed in previous articles in the *Journal*[3, 4].

This article looks at the operational use of CCITT No. 7 signalling in British Telecom's (BT's) network, the facilities provided to assist in its use, the proving of the signalling system and the current in-service experience. In addition, the use of CCITT No. 7 signalling between a number of different systems in the UK network is imminent and the problems that need to be solved to permit such interconnections are outlined.

BRITISH TELECOM POLICY

It is BT's policy to introduce a fully-digital network as soon as possible. This is demonstrated by the very rapid introduction of pulse-code modulation (PCM) and optical-fibre systems, together with a high capital investment programme for the rapid installation of System X exchanges over the next few years.

Associated with this policy is the decision that all System X exchanges will be interconnected from the outset with CCITT No. 7 common-channel signalling.

Network Architecture

The principal options for interconnecting the various types of network nodes are shown in Fig. 1. All 60 digital main switching units (DMSUs)[5] in BT's main network will be fully interconnected by secure duplicated signalling links. Digital principal local exchanges (DPLEs)[6] will be connected by duplicated links to at least two DMSUs. Local exchanges (LEs) will be connected by unsecured signalling links to at least two DMSU/DPLEs. Direct links between LEs, where they exist, will also be unsecured.

Security for signalling at the LE is achieved by the use of alternative routeing, with the DMSU or DPLE acting as a signal transfer point (STP).

Use of CCITT Signalling System No. 7

The main use of CCITT No. 7 signalling in the early network will be for telephony. The telephony user part (TUP) of the

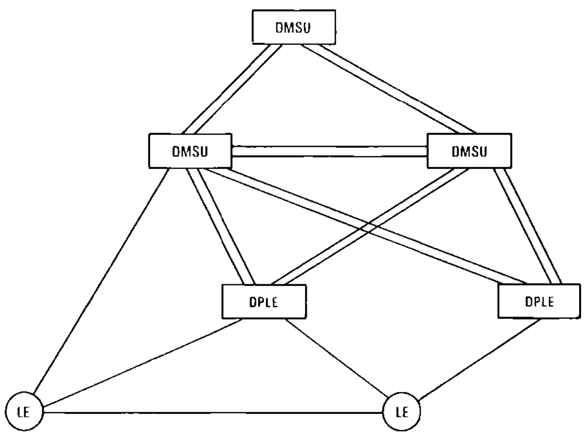

DMSU : Digital main switching unit
DPLE : Digital principal local exchange
LE : Local exchange

Note: This diagram shows the typical connections that can be expected in the UK network. Particular configurations in director areas may cause the provision of additional signalling paths

FIG. 1—Provision of signalling links between network nodes

System X implementation is based on the CCITT TUP Recommendations, but modified to meet additional BT requirements. Included in these requirements is the ability to offer service to integrated services digital network (ISDN) customers. An ISDN service can therefore be offered by System X prior to more definitive Recommendations for an ISDN user part being available from the CCITT.

A subset of CCITT No. 7 signalling is also used for signalling between the remote concentrator units[7] and the main exchange. A concentrator is connected to the main exchange by a minimum of two PCM systems; time-slot 16 is used to provide the 64 kbit/s signalling link. The use of PCM systems in this way permits the remote operation of concentrator units without the necessity for changes to the design of the LE.

A further instance of the use of CCITT No. 7 signalling is in the interconnection of System X exchanges and the operations and maintenance centres (OMCs)[8]. These links are used to convey the man–machine language (MML) commands to the exchange, and for reporting faults from the exchange to the OMC. This use of CCITT No. 7 signalling for large exchanges is temporary, the long-term objective being to use links conforming to CCITT Recommendation X25 to connect support systems to the exchanges; small exchanges, however, may continue to use CCITT No. 7 signalling.

† System Evolution and Standards Department, British Telecom Development and Procurement
* Local Exchange Systems Department, British Telecom Local Communications Services

‡ CCITT—International Telegraph and Telephone Consultative Committee

International Links

The extension to British Telecom International's (BTI's) Ericsson AXE10 switching system at Keybridge House, London includes the provision of CCITT No. 7 signalling. All links between Keybridge House and System X exchanges will use CCITT No. 7 signalling. BTI sees a rapid expansion in the use of CCITT No. 7 signalling to foreign administrations as soon as the Keybridge House unit is ready for service.

OPERATIONAL SUPPORT FOR CCITT No. 7 SIGNALLING

As well as providing the ability to route signalling messages through the signalling network, the MTS must also provide facilities to manage the signalling network whilst in service. Three aspects of this support are outlined below.

Man–Machine Language

In conjunction with the maintenance control subsystem (MCS)[9], the MTS provides a wide range of MML commands for use by maintenance personnel. These commands allow the interrogation of data that determines the current state of the signalling network; for example, which links are in service, which route from a number of alternatives is being used etc. Commands are also available to change the configuration of the signalling network; that is, to change the operational state of a link, force a particular alternative route into use, etc. Further commands allow additional links, link sets, routes and route sets to be added or deleted from the exchange to facilitate extensions and enhancements.

Statistics

Statistics are required to dimension the network so that, for example, additional links can be provided if necessary, and to observe the quality of the signalling links. System X, therefore, provides measurements on the volume of messages of various types (that is, incoming, outgoing and STP), and also provides statistics on message discard due to queue-full conditions, total number of retransmitted messages, frequency of retransmission requests and the number of signal units received in error (this includes fill-in signal units (FISUs) that are not retransmitted). This provision of statistics is in accordance with CCITT Recommendation Q791, which has just been published.

The capacity of a 64 kbit/s signalling link is very large; for example, there is an expected occupancy of only 20% for the signalling traffic associated with 1000 erlangs of telephony traffic. The quality of PCM systems is high, hence the number of errors/retransmissions will be low. CCITT No. 7 signalling networks should, therefore, exhibit high quality and large capacity; few dimensioning problems are thus envisaged.

Fault Handling and Diagnostics

MTS software automatically detects and reconfigures the signalling network in the event of hardware or signalling-link failures. Fault reports are made, via the MCS, to personnel at the OMC and/or locally as required. Fault detection mechanisms include error-rate monitors, parity and checksums on data transfers, and time-outs for lack of acknowledgement of messages by both hardware and software.

In the particular case of the hardware timing out as a result of the software not being available within 100 ms, due, say, to processor failure or a major restart, the hardware will autonomously change the line state on all signalling links to PROCESSOR OUTAGE. When this condition is received at the distant exchange, these signalling links are marked as faulty and new telephony traffic is diverted from the affected route. This action prevents the build up of a backlog of messages. When the exchange recovers, the PROCESSOR OUTAGE condition is removed and normal traffic handling is resumed by the distant exchanges.

Diagnostic routines are provided to self-diagnose the signalling terminals by looping them back, either to themselves or to each other. A further feature of these routines is the ability to check the identity and link number of the distant exchange so that staff can confirm that the network configuration matches the data in the exchange.

This confirmation of network configuration also takes place before a link is permitted to handle user traffic. It is achieved by the interchange of link test messages, and the link may not be used until a successful acknowledgement is received. This particular use of link test messages is a BT requirement additional to the CCITT Recommendations; these require only that link test messages are exchanged periodically, but with no necessity to be successful before user traffic is allowed access.

PROVING OF SYSTEM X No. 7 SIGNALLING

The proving of CCITT No. 7 signalling on System X has taken place in a number of stages and has had to address several different areas. The CCITT No. 7 signalling Recommendations were ratified in the form of a Yellow Book published as a result of the 1980 CCITT Plenary Assembly. It is, therefore, a relatively new and untried set of Recommendations, which has required proving to confirm that the signalling system defined is coherent (that is, it specifies a workable signalling system), is unambiguous in that it can be implemented by separate development teams and, further, that these products can then be interconnected successfully.

The three major stages of this proving to date are as follows:

Testing of the MTS

The MTS is required to work to all System X exchange systems and, because of the rapid change in processor technology[10], it was required in three versions: Release 1 PPU for OMC use, POPUS 1 for existing trunk exchanges and R2PU for new trunk/local exchanges. The hardware and the software source code for all three of these variants are identical, but all three required separate testing to prove conformance to the MTS requirements.

Because the MTS provides the signalling between systems, it was considered vital during subsystem testing to provide all three processor types and to interconnect them. This greatly reduced the time required at subsequent stages of overall system testing. After completion of subsystem testing, the MTS was incorporated into the various systems to allow system and network testing to be undertaken, with the final BT tests and product acceptance taking place at the field support unit (FSU) at British Telecom Research Laboratories (BTRL) at Martlesham Heath.

Proving CCITT No. 7 Signalling Recommendations

With the publication of the CCITT Recommendations in 1980, it was proposed that, rather than establish an international trial under the aegis of the CCITT, individual manufacturers and administrations should set up their own trials to prove out the Recommendations.

Standard Telephone and Cables plc (STC) (then a participating company in the joint System X development) arranged to conduct an interworking trial of different implementations of the Recommendations between a System X model at their New Southgate site and an International Telegraph and Telephone Corporation (ITT) System 1240 at the Bell Telephone Manufacturing Company (BTMC) plant in Antwerp, Belgium[11].

When STC withdrew from the System X development in

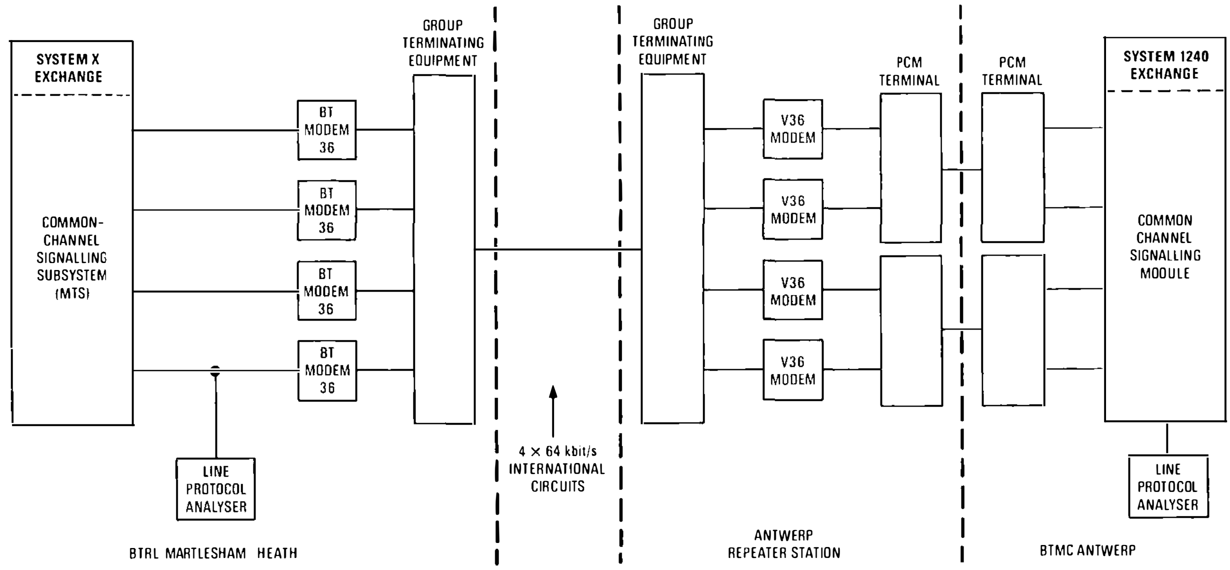

FIG. 2—Circuit arrangements of the CCITT No. 7 international trial

October 1982, the trial was considered to be of such importance that BT took over the organisation and running of the trial using a System X model at BTRL, but retained the use of the BTMC facilities already agreed.

The trial, managed by a four-party committee comprising representatives from BT, RTT (the Belgian Administration), BTMC and Plessey Major Systems Ltd. (PMSL) (the System X prime contractor), commenced in October 1983 (see Fig. 2). Four international 64 kbit/s analogue circuits were provided by BTI and the free provision of these otherwise revenue-earning links limited the trial duration to some 9 months. In this time it was possible to test and evaluate only aspects of levels 1, 2 and 3 of the protocols; that is, those concerned with message transmission and signalling network management—the message transfer part (MTP) of the Recommendations.

The principal outcome of this first international trial of two different implementations of CCITT No. 7 signalling on proprietary exchange systems (as distinct from laboratory models) has been the general sufficiency of the Recommendations—at least at the MTP level—and the conformance of the System X implementation to the requirements of this very complex signalling system[12]. A number of considerations have been forwarded to CCITT Study Group XI/2 for improvements to the system. These are mainly concerned with flag generation, timers and timing tolerances; without some tightening of the Recommendations in these areas, it will be difficult to ensure that different implementations will interwork.

Another outcome of BT's involvement in this trial has been the need for BT to examine critically the underlying philosophy and methodology of interworking trials. This aspect is examined in more detail later in this article.

UK Field Trial

As a further extension to the testing undertaken during development, a field trial is planned using all the System X exchange types to be found in the UK network; Fig. 3 shows the arrangement. The test facilities and systems at Martlesham Heath are also to be connected into the network, and much of the traffic and measurements will be produced at that site.

The object of the trial is to confirm the network performance of common-channel signalling in a wide range of operational conditions, and to validate all the operational procedures inherent in running such a network. The exchanges in the trial network are being used prior to in-service use, and the opportunity will be taken to inject fault, overload and other conditions into the network which would not be possible once the exchanges are in service. This trial, which started in October 1984, should provide BT with a high degree of confidence in the operation and stability of a common-channel network.

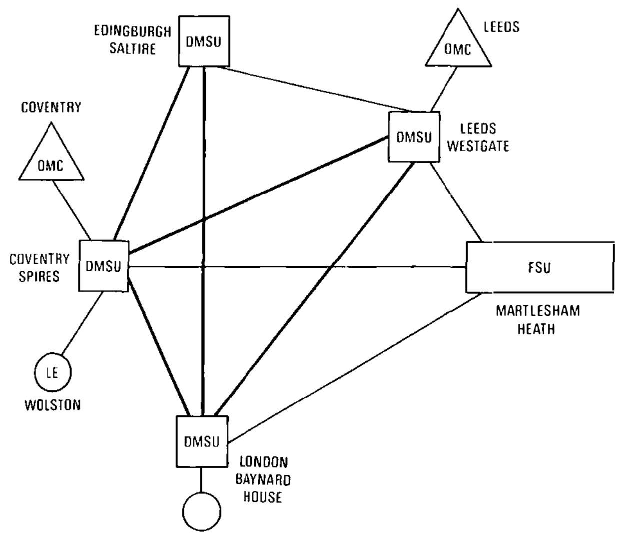

DMSU : Digital main switching unit
FSU : Field support unit
LE : Local exchange
OMC : Operations and Maintenance Centre

Note: All exchanges have additional connections to the existing network

FIG. 3—Interconnection of System X exchanges for the UK field trial

IN-SERVICE EXPERIENCE

Early versions of System X trunk and local exchanges at Cambridge and Arrington have been operating common-channel signalling since 1982.

The equipment used does not conform to the current CCITT No. 7 signalling Recommendations since it was produced before 1980. However, it has validated the principle of common-channel signalling on System X and has provided the confidence to proceed with the immediate installation of the MTS on all System X exchanges, so that they are fully connected by using only common-channel signalling. System X exchanges will not be interconnected by existing DC/AC/MF signalling techniques.

A further demonstration of the common-channel signalling on System X is provided by the successful cut-over to service of a System X local exchange supplied by the General Electric Company plc (GEC) to Cable and Wireless plc in St. Vincent, in the Caribbean. This system uses CCITT No. 7 signalling between the main unit and eight subscriber concentrators; seven of these are remote and one is co-sited with the main unit at Arnos Vale. See Table 1.

TABLE 1
St. Vincent System X Common-Channel Signalling Network

Phase	Concentrators	Subscribers	PCM Links to Arnos Vale
1	Arnos Vale	700	3
1	Mesopotamia	600	2
1	Kingstown (1) (2)	2500	9
2	New Montrose	500	2
2	Camden Park	400	2
2	Villa	300	2
2	Prospect	500	2

All phase 1 units were cut-over on 29 September 1984, and it is planned that phase 2 units will be in service before the end of 1984. The system has given excellent service since cut-over, and the CCITT No. 7 signalling system has provided trouble-free operation.

INTERCONNECT TO OTHER SYSTEMS

In parallel with the use of CCITT No. 7 signalling between System X exchanges, it is also BT's policy to provide interconnection from System X to other systems by using CCITT No. 7 signalling.

Some of these systems represent enhancements to existing UK systems such as TXE4 and UXD5; others include systems such as those supplied to BTI and Telecom Eirean (Ericsson AXE10).

In the near future, BT will be incorporating other stored-program control (SPC) systems into its network apart from System X and, with the liberalisation policies to be adopted, will be connecting to other networks by using CCITT No. 7 signalling.

With such a large number of different systems being interconnected by a complex common-channel signalling system, important questions are raised concerning the proving and testing of such systems prior to operational use. Two possible approaches are available to undertake such testing.

(*a*) To test specifically the combinations of systems that are to be interconnected. This will certainly prove that the immediate interconnect is valid: it does not prove that the systems are suitable to connect to any other system. Neither does it necessarily prove that the systems will still work together if either, or both, are enhanced.

(*b*) To prove conformance of each system to the specification of the CCITT No. 7 signalling system. This approach requires the development of a specialist test system, which must itself be validated before being used as a test environment. The advantages of this method are that systems could be tested once, and that there would be a high degree of confidence in the ability to interconnect such systems once they had been so tested. However, the problems of producing a test environment to the standard required, and validating it, should not be underestimated.

Both these methods of proving interconnect are being actively considered at the time of writing, and it is quite probable that a mixture of both techniques will be used.

CONCLUSION

This article has sought to provide some appreciation of the wide scope and range of activities that have been necessary to provide a CCITT No. 7 signalling system for the BT network. It demonstrates a very firm commitment by BT to common-channel signalling in the UK network.

The provision of common-channel signalling, together with a fully digital network comprising modern SPC exchanges, places BT in an excellent position to satisfy customer requirements now and in the future.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the help and assistance given by many of their colleagues in the preparation of this article. PMSL and GEC provided much assistance during the running of the UK–Belgium trial; without this, and the generous provision of international circuits by BTI, the trial would not have been so successful.

References

[1] CCITT Signalling System No. 7. CCITT Yellow Book. Vol. VI, Fascicle 6, Recommendations Q701–Q741.

[2] MOORE, B. W. Progress Towards Common-Channel Signalling. *Post Off. Electr. Engs. J.*, Oct. 1981, **74**, p. 210.

[3] JONES, W. G. T., KIRTLAND, J. P., and MOORE, B. W. Principles of System X. *ibid.*, July 1979, **72**, p. 78.

[4] PRICHARD, T. W. Common-Channel Signalling and the Message Transmission Subsystem. *ibid.*, Oct. 1980, **73**, p. 165.

[5] GARBUTT, B. N. Digital Restructuring of the British Telecom Network. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 300.

[6] ROLFE, M. The Role of the Digital Principal Local Exchange. *ibid.*, Jan. 1985, **3**, p. 304.

[7] RABINDRAKUMAR, K. The Function and Use of Remote Concentrator Units. *ibid.*, Jan. 1985, **3**, p. 245.

[8] STRICKLAND, L. F., and HEWITT, M. A. New Operation and Maintenance Centres for Second Generation System X Exchanges. *ibid.*, Jan. 1985, **3**, p. 286.

[9] DANN, A. T. Maintenance Control Subsystem. *Post Off. Electr. Engs. J.*, Apr. 1981, **74**, p. 45.

[10] TROUGHTON, D. J., *et al.* System X: The Processor Utility. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 226.

[11] FISHER, D. G., *et al.* Experience in the Implementation of the CCITT No. 7 Signalling System. Fourth World Telecommunication Forum: FORUM 83, Geneva, Oct. 1983.

[12] THOMAS, D. L., HAERENS, F., BEKAERT, R., and MANTERFIELD, R. J. Testing Results from the Implementation of the CCITT No. 7 Signalling System. International Switching and Signalling Symposium: ISS 84, Florence.

Biographies

David Butterworth joined BT as a Youth-in-Training in 1953 in the Leeds Telephone Area. Later, while with the Central Training School, Stone, he graduated with a B.Sc. in Mathematics and Statistics from the University of Birmingham. He moved to Telecommunications Headquarters in 1967 where he was involved principally with software engineering and management aspects of computer systems applied to telephone exchange maintenance. In 1979, he transferred to the System X Development Department, with responsibilities in project planning, management and system economics. At the end of 1982, he moved to BTRL where he had responsibility for four of the System X subsystems, including the MTS. He is currently a Head of Section working in the Network Evolution Division.

Anthony Williams joined BT in 1967 as a BT Student and graduated from Bristol University with an honours degree in Electrical Engineering in 1971. As an Exective Engineer he worked on the TXE4 switching system until 1975, when he became involved with System X. This work embraced system design, reliability and maintenance, maintenance software for the MCS and, more recently, common-channel signalling. As an extension of this work he was responsible for running the trial with Belgium. He has now been appointed to the Local Exchange Systems Works Support Division and is responsible for data builds for in-service exchanges.

In-Station Cabling with Optical Fibres

F. G. HARRISON, B.SC., C.ENG., M.I.E.E., and R. K. LE GOOD, C.ENG., M.I.E.R.E.†

UDC 621.391.63 : 666.189.2

This article briefly outlines the potential applications of optical-fibre cables for in-station cabling. The advantages of optical fibres over the conventional metallic cables are discussed in the context of a System X exchange and a transmission repeater station. The appropriate optical-fibre technology for these applications is also considered.

INTRODUCTION

The installation of optical-fibre transmission systems in British Telecom's (BT's) main and junction networks has proceeded rapidly over the past few years[1]. Initial systems have provided BT with considerable experience of the design, installation and operation of optical-fibre systems. The benefits, when compared with conventional copper cable, of increased repeater spacing, reduced cable size and large bandwidth have all been demonstrated. The technology has now been developed to an advanced state, and this has given high confidence in the performance of optical-fibre transmission systems and opened up opportunities for using optical fibres in other areas. Additionally, a maturing optical-components industry is now able to offer low-cost devices and connectors. These factors encourage the use of optical fibres for in-station and in-suite cabling applications, where conventional metallic technology is now beginning to show technical limitations.

ADVANTAGES OF OPTICAL FIBRES

As electronic circuitry continues to shrink in physical size, owing to ever-higher levels of integration becoming available, the sheer bulk of the interconnecting cabling and of the multiway connectors begins to dominate the packing density and physical realisation of equipment. With metallic cable, using either symmetric or coaxial pairs, reducing conductor size and separation tends to increase attenuation and worsen the crosstalk performance. Optical fibres, with their inherently small size, immunity to interference, and low attenuation, offer solutions to many of the equipment-cabling problems. In addition to their reduced physical size, optical fibres can provide a very large bandwidth.

In large system installations, such as telephone exchanges and transmission repeater stations, the inter-equipment cabling can represent a considerable capital investment. While the practical limitations of metallic cables are now being reached, optical fibres can offer a highly versatile and upgradeable cabling system for forseeable equipment enhancement.

APPLICATIONS OF OPTICAL-FIBRE CABLING

Two specific application areas are dealt with in more detail, with the emphasis on the cabling of digital circuits.

System X Digital Exchange

In simple terms, a System X exchange consists of a digital switch, containing both time and space switching elements and operating on standard 2 Mbit/s line interfaces. The switch is controlled by a computer, which may comprise several physically separate processors. Within the switch, signals are carried from stage to stage by special high-performance metallic cables, and strict limitations are imposed on the cabling distances between switching stages. The use of very-large-scale integration (VLSI) in the switch modules would enable a substantial reduction in the size of the switch, but then the physical volume of the metallic cables would dominate the design. Consideration has already been given to the use of optical fibres as a solution to this problem, and a joint development between BT and Plessey has produced transmitter and receiver hardware for such an optical cabling system. To date, no clear cost advantage in the use of optical fibres has arisen, and so the current switch design relies on metallic cables. However, the arrival of a new generation of more compact switch modules may well tip the balance in favour of optical-fibre cables.

Transmission Repeater Station

A typical transmission repeater station in the BT network houses line and radio transmission system terminals, multiplex equipment and service-protection switching apparatus. The equipment is mounted in racks, and digital signals at bit rates of 2, 8, 34 and 140 Mbit/s are carried between these racks by means of coaxial cables, with cabling distances of up to several-hundred metres. Multiplex equipment, which combines several digital tributaries onto a single path, is now becoming so compact that the physical size of the cabling causes interconnection difficulties. As an example, a fully equipped rack of multiplexers could require some 600–800 coaxial cables. In addition, line transmission systems with operating rates of 565 Mbit/s will soon be in service, and the distance over which cables carrying such high-speed signals can be used within a station will be severely restricted if normal internal coaxial cables are used.

Optical-fibre cabling can offer a solution to both of these problems by virtue of the advantages already discussed. A detailed study into the use of optical fibres for application within repeater stations is being conducted and some of the initial findings concerning the appropriate optical technology are discussed later in this article.

The applications discussed so far have been for point-to-point traffic signals. A more general use of fibre cabling can be envisaged for signalling and control circuits, because a ring or local area network arrangement could be used. Such a configuration gives flexibility for growth and the ability to incorporate remotely located equipment into a system.

OPTICAL-FIBRE TECHNOLOGY

The applications discussed place stringent requirements on the various components of an optical-fibre cabling system. In order to realise the benefits of optical fibres, the cable, connectors and associated opto-electronic devices should equal or better the attributes of the existing metallic cable systems. Important parameters include the physical size of the cable and devices, the maximum cabling distance and the maximum bit rate. The environmental performance and the comparative cost of the different cabling technologies are also pertinent.

The fundamental element of an optical-fibre cabling system is the fibre itself, and the choice of fibre type affects many of the characteristics of a link. Main and junction optical-fibre transmission systems are based on small-core multimode fibre or singlemode fibre; this is necessary to achieve very long repeater section spans[2]. However, an in-

† Trunk Services Planning and Works, British Telecom National Networks

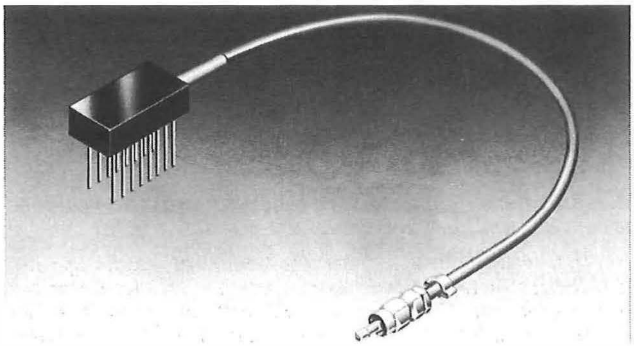

FIG. 1—Plessey LAN LINK opto-electronic module
Photograph courtesy of Plessey Optoelectronics and Microwave Ltd.

station optical-fibre system with a maximum reach of about 1 km can make use of large-core silica fibres or plastic-clad silica (PCS) fibres. Large-core fibres considerably ease the handling, connection and termination requirements of a cable, and provide a cost advantage over small-core fibres. For the interconnection of high-bit-rate equipment, small-core fibre becomes more desirable because of its high bandwidth.

Optical connectors form an important part of an optical-fibre cabling system. High-performance optical connectors are currently used in optical-fibre transmission systems in order to obtain the required transmission parameters and to provide adequate reliability. Large-core fibres allow simple optical connector mechanisms and termination techniques to be used, since the fibre alignment and end-surface finish are not critical. The use of plastic component parts in connectors has the potential for making them cheaper. Small-core fibres demand more accurate alignment in order to achieve a low insertion loss, and probably require the now commonly used adhesive-and-polish termination technique to achieve satisfactory finish to the end of the fibre. Whatever connector is selected, it must be robust enough to allow for handling during installation and maintenance and to operate under the normal station environmental conditions.

The choice of line driver and receiver components is dictated by the bit-rate requirement. For low-speed applications, a light-emitting diode (LED) source is ideal and can even be suitable for rates up to 280 Mbit/s. For 565 Mbit/s, however, a semiconductor laser is necessary to avoid dispersion penalties arising from the wide linewidth of the LED source. Receivers based on avalanche photodiodes or PIN diodes are suitable for all of the envisaged bit rates. It is anticipated that in-station cabling systems will operate in the first wavelength window of the fibre (850–900 nm) as components for this wavelength range are now well established and are currently cheaper than those operating at 1300 or 1550 nm.

This brief outline of optical technology suggests a potential division in the realisations of in-station cabling according to the application and bit rate. For example, a cabling system operating at up to, say, 50 Mbit/s could employ large-core fibres, simple plastic connectors and an LED transmitter. The Plessey LAN LINK[3], produced under a joint development with BT for System X applications, is an example of how such a system can be implemented (see Fig. 1). For higher bit rates, the use of small-core fibre, precision connectors and a laser source may be required, depending upon the exact system performance desired.

IMPLEMENTATION OF OPTICAL-FIBRE CABLING

Clearly, optical-fibre in-station cabling offers a number of advantages over currently used metallic cables in several applications, and the optical-components technology has developed to a stage where optical-fibre cabling is now viable. However, two barriers are hindering the introduction of optical fibres on a wide scale.

The first of these is an economic barrier. There is, at present, no clear cost advantage in the use of optical fibres, and the inertia of the existing metallic technology is difficult to overcome. The introduction of a new technology will result initially in extra costs for specialised manufacturing equipment and test apparatus, and for the training of staff in new practices required for both the installation and maintenance of optical-fibre cables.

The second barrier is one of compatability. In a large and evolving system installation, such as a repeater station, a considerable quantity of equipment is already provided with electrical interfaces. The introduction of an optical-fibre cabling scheme would either restrict the interconnection of the different types of interface, or necessitate retrospective modification of existing equipment to provide optical-fibre interfaces. It is also desirable that an optical-fibre interface which is to be used extensively for interconnecting parts of the network should gain international recognition as a standard.

These are short-term difficulties, but it is anticipated that, as the technical advantages of optical fibres increase and the cost decreases, the inertia currently retarding their introduction will be overcome.

CONCLUSION

This article has shown that optical fibres offer considerable potential for use as an in-station cabling system for telecommunication applications. The advantages of optical fibres over metallic cables include reduced size, high bandwidth, long cabling lengths and interference immunity. These attributes are becoming enhanced as equipment packing density and complexity continue to increase. The optical technology now available offers options for the ways in which links can be implemented, and the optimum design for repeater station applications is now under study. Similar work on System X development has produced technically viable hardware for use in new-generation exchange equipment. The introduction of optical fibres is now nigh, pending a more favourable cost advantage in changing to fibre. Undoubtedly, in the near future, optical fibres will play a major role as inter-suite and in-suite transmission media.

References

1 MARTIN-ROYLE, R. D., and BENNETT, G. H. Optical-Fibre Transmission Systems in the British Telecom Network: An Overview. *Br. Telecommun. Eng.*, Jan. 1983, **1**, p. 190.

2 MUIR, A. W. Optical-Fibre Transmission Systems in the British Telecom Network: System Design. *ibid.*, Oct. 1983, **2**, p. 180.

3 SUMERLING, G. W. New Developments in Fibre Optic Data Links. IEE Conference Publication No. 230, p. 19, 1983.

Biographies

Fred Harrison is an Executive Engineer working on the development of optical-fibre transmission systems in the Trunk Services Transmission Systems Engineering Division of BT National Networks. He joined BT in 1974 as a sponsored university student, and in 1978 received an honours degree in Electronic Engineering from Southampton University. He has spent six years on various aspects of optical-fibre systems, including in-house development, equipment specification and technical support. Recently, he has been involved in 565 Mbit/s optical system design.

Keith Le Good is head of the Proprietary Optical Fibre Systems Group in the Trunk Services Transmission Systems Engineering Division of BT National Networks. He joined BT from industry in 1963 by way of open competition. During his career in BT he has worked on the development of video systems and equipment, digital data multiplexers and transmission test equipment. He currently leads a group responsible for engineering the proprietary optical line systems, 565 Mbit/s optical systems and optical components.

System X: Build Control

Part 1—Product, Hardware and Software Build

M. J. FOX, B.SC., C.ENG., M.I.E.E.,‡ and M. H. STOREY, C.ENG., M.I.E.R.E., M.B.C.S.†

UDC 621.395.34

Build control is a procedure to ensure that changes to a working telephone exchange are implemented without risk to the overall integrity of the system. Part 1 of this article, describes the procedures used to control the overall system and the hardware and software used to make up the system. Part 2 will descibe the techniques adopted for the management of the exchange data.

INTRODUCTION

With any item of equipment that is required to interwork either internally, because of the subsystem design, or externally to other equipment, there is a need to control what is happening at all the various interfaces for the life of the equipment. This is especially true with a telephone exchange such as System X, where the problem is compounded because a mixture of hardware, firmware and software is used.

The interworking with other equipment is controlled by adhering to well defined specifications for the interfaces; for example, the specification for CCITT* No. 7 common-channel signalling and the CCITT Recommendation G703 for pulse-code modulation line conditions.

Within the telephone exchange itself, changes can occur for a number of reasons; for example:

(*a*) the introduction of improved facilities,
(*b*) changes imposed on manufacture by components becoming obsolete,
(*c*) maintenance problems arising because of faulty design,
(*d*) as a result of cost-reduction exercises by the manufacturer, and
(*e*) manufacture of equipment which is still undergoing development.

‡ Trunk Services Planning and Works Department, British Telecom National Networks
† Trunk Services Operations and Maintenance Department, British Telecom National Networks
* CCITT—International Telegraph and Telephone Consultative Committee

Before any changes can be accepted onto a working exchange, it has to be proven that the introduction of the change from the agreed standard will not cause problems to the existing equipment, and that the exchange will continue to function properly after a modification is made. In addition, there should also be cost benefits accruing from the change either in terms of increased revenue from new facilities or reduced maintenance costs. The process of ensuring that this happens is called *build control.* Four different areas of build control are encountered with System X.

(*a*) *Product Build* This defines the overall system.
(*b*) *Hardware Build* This defines the hardware used to make up the system.
(*c*) *Software Build* This defines the exchange processor operating system and the application processes.
(*d*) *Data Build* This defines the exchange-dependent data[1].

The structure of the build for System X in terms of its software and hardware constituant parts is outlined in Fig. 1[2].

The levels in the build definition are defined in terms of three-letter codes, as shown in Table 1.

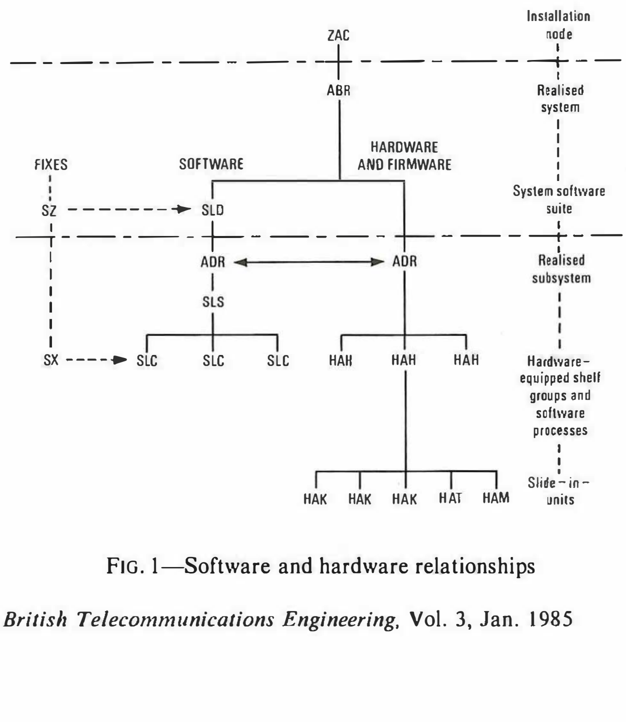

FIG. 1—Software and hardware relationships

TABLE 1

System Build Definitions

ZAC :	All the constituent parts of a particular installation
ABR :	A generic system design including hardware, software and firmware.
SLD :	A generic software suite. Fixes are added to this as SZ files
SLS :	Supplementary software at subsystem level.
ADR :	A realised sub-system, either software or hardware and software.
SLC :	Processes of a realised software sub-system. Fixes are added to this as SX files. Several software processes may form a total subsystem software module.
HAH :	Modules of a realised subsystem; that is, shelf groups.
HAK :	Submodules of hardware; that is, slide-in-units.
HAT :	Submodules of hardware; that is, slide-in-units which include firmware.
HAL :	Submodules of hardware; that is, power supplies.
SX :	Fixes at process level
SZ :	Fixes at system level

PRODUCT BUILD

The System X development programme will move through a number of phases. The earliest product build, known as *Release 1*, has provided both the developers and British Telecom (BT) with useful information and experience, and has formed the basis for the design of the product that is going into general service within the BT network. From this basic design, an overall development strategy, known as the *system enhancement programme* (SEP), has been evolved. The SEP is in a number of phases, SEP1, SEP2, SEP3, etc., the later phases being mainly software enhancements to increase the range of facilities offered by the exchanges.

At the inception of SEP1, some of the exchange subsystems were more stable than others in their design. Various realisations were planned for both trunk and local exchanges in order to obtain early operational experience. These became known as the *LE1* and *LE2* builds for local exchanges and the *TE1* and *TE2* builds for trunk exchanges. An ongoing requirement placed on those developing the exchanges is that each realisation must be capable eventually of being raised to the full requirements of the SEP system specifications in an operational environment.

The process of development worked through from laboratory bread-board models to subsystem feasibility models (SSFM) and then on to system feasibility models (SFMs). At the end of the development phase, all of the hardware, firmware and software are tested together during formal system proving (FSP). Because the development work was being conducted at a number of different locations, it became essential to establish a mechanism to ensure that the equipment on each model was kept at the same development level. This mechanism has been carried through to the manufacturing and operational environment. To achieve control, the item code for each unit contains a suffix which gives details of its build. This suffix, known as the *ENU level*, indicates the level of essential (*E*) and non-essential (*N*) changes that have occurred, and the number (*U*) of *E* changes that have been made since the last clean build. Each time clean board artwork is generated the *U* level returns to zero.

The requirements of each build are established by BT, in conjunction with the prime contractor (PC), as a system performance specification. This is then translated into subsystem designs and defined in terms of hardware, firmware and software modules; the hardware being realised as slide-in-units, and the software and firmware modules as structured code.

The current situation is an evolving one where new designs are essentially enhancements of existing designs.

HARDWARE BUILD

When the development of a realisation has progressed to a stage where the development team is satisfied that the product can meet the realisation specification, a period of FSP is initiated. If this is satisfactory, the product is offered to BT, which then requires the developer to demonstrate that the development objectives have been achieved. This process is known as *product acceptance testing* (PAT). The build of the particular realisation consists of the unit codes and their ENUs at the completion of PAT. This is called the *target build*, and all works-order exchanges in a particular realisation (for example, TE1) are provided to the same build: in effect, the target build becomes part of the exchange specification. Once PAT has been achieved, BT is responsible for ensuring that all in-service exchanges are maintained at the latest target build; this is change control.

BT is committed to the modernisation of its network by using digital techniques. This commitment led to an urgency to get digital switching systems into the operational network as soon as possible which, in turn, necessitated the manufacture of equipment in advance of its development being completely finalised. To minimise the risk associated with this situation, a procedure, known as *authority to make*, was established whereby the suppliers submit requests for the manufacture of the various codes of equipment required on each exchange system before they are put into initial production. This gave BT the opportunity to minimise the risk of having to pay for equipment which did not meet its design specification. Authority-to-make was effectively given against a specific artwork layout of a printed-wiring board and for a given volume of production. As the system design stabilised, so the need for the authority-to-make procedure diminished.

Once equipment has been manufactured and delivered to a works-order site (on the early orders in advance of PAT), it becomes necessary to raise the *E* values to keep in step with the latest target build. Hence, there is a need to track the build level of each item of equipment delivered to site and for the supplier to make the necessary modifications to achieve the new target build.

The modifications can be carried out either on site or at the manufacturer's plant, depending on the quantity and complexity of the modification. The instructions, documentation and parts required to upgrade from one *E* level to another are all incorporated in a modification action pack (MAP).

This process is broadly illustrated in Fig. 2.

The commissioning and demonstration exercises, which proceed the acceptance of a works-order exchange, are part of the quality assurance process. One of the requirements placed on the supplier prior to entering the commissioning and demonstration phase is to show that the exchange is at the target-build level.

In order to maintain a record of the exchange once it comes into service, a computer system, known as *BLAISE* (build level at in-service exchanges), has been developed, and adopted for use at exchanges during the works-order period to record the target and actual *E* levels on site.

BLAISE

The BLAISE system is ideally suited to the size of the trunk exchange programme, which is controlled centrally. Although this centralised control does not exist for local exchanges, the BLAISE system is being used to control their build. However, the greater population of local exchanges and their different organisational arrangements will necessitate investigation into its long-term use.

The BLAISE system consists of a database, which details the equipment provided at an exchange, and a number of programs that allow the data to be manipulated according to the needs of the user.

The database is a listing of the equipment by common code, the target ENU and the actual ENU, together with the serial number of the individual units and their suite, rack and shelf positions within the exchange. This information is used on site to obtain details of the differences between the actual and target ENUs and to record the whereabouts of the individual units. A *remarks* column provides the facility to record, for example, whether the unit has been returned to the factory for modification or repair, the dates at which events happen etc. It is envisaged that the maintenance staff on site or at the operations and maintenance centre (OMC) will also use the BLAISE system as a means of keeping a maintenance history of the hardware; that is, as a replacement for fault record cards. On site, the BLAISE system operates on a small business computer (SBC) using floppy-disc storage.

BLAISE also plays an important role in the change-control function. This is operated centrally and requires details of all exchanges.

For change-control and monitoring purposes it is not essential to have details of serial numbers and the location within the exchange: it is necessary, however, to know how

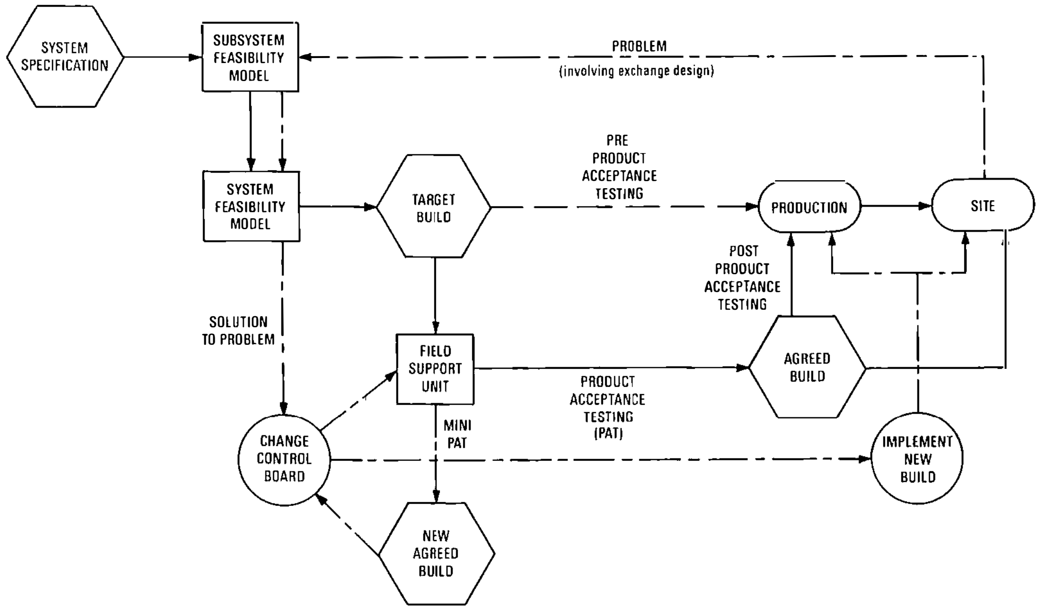

FIG. 2—Build control arrangements

many items exist at what build level and at which exchanges they are located. BLAISE, held centrally on a hard-disc system because of the amount of information involved, gives this information. One of the management programs associated with BLAISE, known as *BUDGET*, enables the costs of a modification to be evaluated.

The change-control monitoring facility has use both within the works-order and in-service periods of the exchange.

Another application of BLAISE that is currently being investigated is its use for monitoring the utilisation of equipment. Certain parts of the exchange (for example, digital line terminations (DLTs)) are provided for specific design periods. These are allocated for either internal or external applications, and their utilisation has to be monitored for a number of reasons; BLAISE fulfills this need.

SOFTWARE BUILD

The main stages in the establishment and control of software builds are:

(*a*) software development,
(*b*) formal proving of software at system level,
(*c*) production of standard software suites, and
(*d*) in-service software upgrades

Throughout all these stages, it is essential to have an accurate and complete definition of the software.

Software Build Strategy

System X software is provided as a standard suite at system level for each exchange; and the exchange-dependent data, prepared by local operations staff, is added to this to create the exchange load tapes. A standard software suite (SLD) consists of modular subsystems (SLSs), each of which is broken down into separate processes (SLCs) (see Fig. 1). The modular approach is essential to maintain control of such a large software suite and to allow its development at a number of different locations. In order to secure the software suites, both BT and the PC operate their own compatible software libraries on mainframe computers located at different sites. These libraries hold secure copies of software and firmware files, software tools and software parts lists. The PC library is known as the *development library* (PCDL) and the BT library as the *field library* (BTFL). The delivery mechanism of software and firmware files to the BTFL is currently via magnetic tape.

Software Development

The software subsystems are built by separate design teams as software modules and tested at module interface level by using software emulators. After satisfactory emulation testing, the subsystems are tested on an SSFM to the original design specification. During this phase of development, implementation errors are corrected by producing new source code, which is recompiled on the PCDL to give new object code. After the individual subsystems have been tested satisfactorily, they are progressively integrated into a complete system build on an SFM, and interface tests are made between the subsystems. This enables an initial system build to be created. If integration testing detects errors, they are corrected by software fixes, which are added in machine-code form to the subsystem object code. If excessive fixes are added in this manner, the source code becomes difficult to maintain: a subsystem with excessive fixes requires recompilation of its source code to give a clean build. All fixes at subsystem level are entered on the PCDL as fix files. When integration is complete and a satisfactory degree of stability has been reached, a target build is established and FSP commences.

Formal Proving of Software at System Level

Proving is conducted across various SFMs to specified testing schedules. It is usual for each SFM to specialise in a particular feature of the build. At the conclusion of proving, a formal release/acceptance procedure for the transfer of software from the PCDL to the BTFL is initiated.

The delivered product consists of:

(*a*) a complete build definition including the software and firmware,

(*b*) a software parts list giving details of the source code, object code, fix files, support tools such as compilers, linkers, listers and process combiners for building and testing both

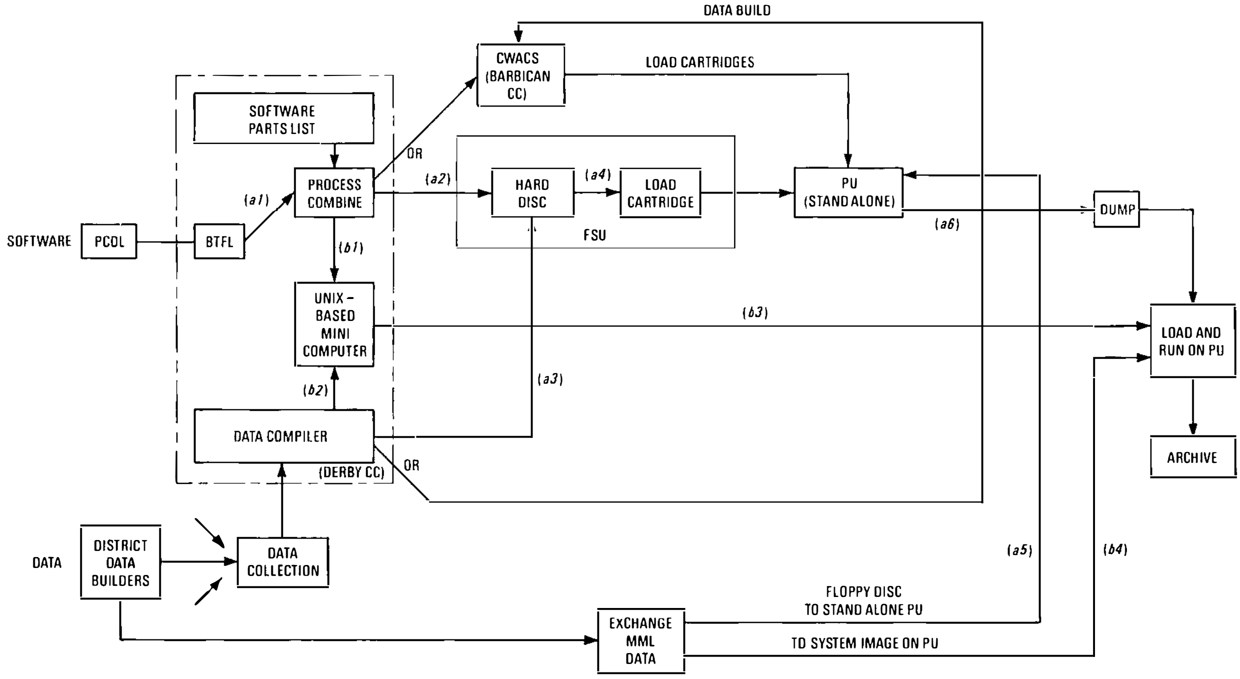

FIG. 3—Software and data combination

software and firmware,

(*c*) a copy of the built software product in source code, object code and system image,

(*d*) the necessary information to update BT's documentation database (DDB), and

(*e*) paper documentation to support the delivery (for example, design documents, operating instructions etc.).

The product then goes into the final stage of system proving, PAT, which is carried out on the field support unit (FSU) at Martlesham Heath. A total of approximately 2000 hours of testing at system level have been undertaken for the TE1 product alone.

Complete confidence in the accurate definition of the software to be used at exchanges is gained by BT performing the following:

(*a*) compiling all source code delivered to the BTFL and comparing the resultant object code with the object code delivered by the PC, and

(*b*) building a complete software image from the BTFL and comparing it with the software provided by the PC for PAT.

At the successful conclusion of PAT, the software required by the new product is deemed fit for in-service use and is submitted to the master area of the BTFL by the release/acceptance procedure.

Production of Exchange Load Tapes

The load tape consists of a standard software suite (SLD), plus the exchange-dependent traffic data for a specific installation. The load tapes are produced by using one of the two methods indicated in Fig. 3.

(*a1*) The first stage in the production of a system software image is to run the software parts list on the full system. This initial stage of process combining produces loadable object code files on the BTFL. This part of the procedure is common to both methods (*a*) and (*b*) below.

Method (a)*—System X Processor Production*

This method requires the use of an exchange processor to further combine the loadable object code files (*a1*) with the particular exchange-dependent data. The combining process is:

(*a2*) The process-combined load files included in the software parts list are downloaded from the BTFL[1] over a private-wire link to a microcomputer equipped with a hard disc.

(*a3*) The exchange-dependant data, produced by BT operational departments, is entered onto the BTFL and the data compiler is run. This combines all the various data input forms into two loadable object code files. These files are then downloaded to the hard disc as in (*a2*).

(*a4*) The load files held on the hard disc are then dumped individually to System X compatible cartridges, which are then loaded onto the processor in accordance with the system building information in the software parts list.

(*a5*) Certain data, such as password data, has to be loaded directly onto the exchange and not incorporated via the data compiler. This data is entered by man–machine language (MML) commands which, for speed of entry, are combined onto a floppy disc.

(*a6*) All of the data is combined on a stand-alone System X processor to produce a system image, which is then dumped to a single cartridge for transfer to a works-order exchange.

This method of load tape production has been used for the production of the single-cluster SEP exchange software, but is unsuitable for the multi-cluster systems planned to come into service during 1985.

An alternative approach to (*a2*), (*a3*) and (*a4*) above, which is used for local exchanges, is to feed the information to the cartridge write and control system (CWACS) at the Barbican Computer Centre, where a similar combining process takes place.

Method (b)*—Minicomputer Production*

The combining process used to produce exchange load tapes uses both a Unix-based minicomputer and a System X processor. It is an off-line system and involves:

(*b1*) Down loading the individual process load files from the BTFL onto a Unix-based minicomputer system.

(*b2*) Process combining and the addition of the exchange-dependent data. This takes place on the minicomputer, which is able to generate System X compatible cartridges.

(*b3*) The compatible cartridges are loaded onto a System X processor, usually the target exchange.

(*b4*) As (*a5*).

In the early stages of the works-order-exchange programme, load tapes are tested for software and data compatibility and loadability at the FSU before being applied to an exchange. In addition, the first few builds are tested at the FSU for basic call-handling capability.

IN-SERVICE EXCHANGES

Once an exchange goes into service, the responsibility for the exchange passes, within BT, from the works to the operational organisations. Changes to the build within a given realisation are most likely to arise because problems occur on in-service exchanges.

Problems which require amendements to the existing software have patches (machine-code amendments) produced. These patches are system tested on the FSU and, if satisfactory, raised to the status of a fix. The fixes go through the same formal accept/delivery mechanism as the full system software. The procedure of producing fixes rather than recompiling subsystems is used because of the time taken to recompile and retest the subsystems at system level. This approach is in common with many other large real-time computer projects. In the case of an urgent problem, an individual fix may be applied at sites, however, due to the thoroughness of FSP and PAT, such circumstances are now virtually unknown. The support given to in-service exchanges is detailed in another article in this issue of the *Journal*[3].

CONCLUSION

System X has a much more complex configuration than former systems. In order to preserve the compatibility of the software, the firmware and the numerous hardware items, it is necessary to control the timing of the implementation of modifications and to maintain a record of:

(*a*) the product status values and the quantities of equipment in each location, and

(*b*) the latest target system build levels which identify the approved compatibility. This enables checks to be carried out to ensure that an item is at the correct product status value, and to enable the implementation of changes to be controlled.

The use of the BLAISE system enables:

(*a*) these records to be maintained;

(*b*) estimates to be made of the total cost of proposed retrospective modifications and/or replacements;

(*c*) arrangements to be made for ordering and supplying the required stores; and

(*d*) co-ordination of production changes with the required in-service compatibility.

References

1 MILWAY, N. R. P., and ABRAM, J. System X: Build Control, Part 2—Data Management. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 265.

2 SHEEKEY, B., and COLLIER, Q. G. System X Design and Support, Part 4—Development Documentation Scheme. *Post Off. Electr. Engs. J.*, July 1980, **73**, p. 100.

3 HOOKER, D. S., and HOWE, J. In-Service Support for System X Exchanges. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 270.

Biographies

Malcom Fox is Head of the Digital Switching Support Group in British Telecom's Trunk Services Switching Works Division. He spent nine years in the Transmission Department at AEI Woolwich, where he was responsible for the development of terminal and line equipment. With the demise of AEI, he spent three years with Bell Northern Research, Canada, working on system engineering aspects of their 274 Mbit/s PCM system before joining BT in 1971. As a member of the private network planning group, he established the planning standards for many of the large private networks that exist in the UK. In 1980, he transferred to the System X Launch Department where he led teams responsible for economic and interworking standards. He moved to Trunk Services in 1983 and is now responsible for trunk exchange build control, commissioning and demonstration support.

Mike Storey is attached to the System Management Group in BT's Trunk Services Operations Department. He joined the Research Department at Dollis Hill in 1959, and, since 1968, has worked in Headquarters with computer-based systems, initially on hardware appraisal and acceptance testing and, later, on tender evaluation, implementation, and production of software applications programs for the Experimental Packet Switching Service. He moved to the System X Launch Department in 1979 and was a member of a team that specified a system management strategy for the Release 1 System X exchanges. His current work involves software building and control of in-service System X trunk exchanges. He is also public relations officer of the North West London branch of the British Computer Society.

System X: Build Control

Part 2—Data Management

N. R. P. MILWAY, B.TECH., A.M.I.E.E.†, and J. ABRAM*

UDC 621.395.34

System X exchanges require a large amount of data to be loaded with their software, for site configuration purposes. System X data requirements are substantially greater than for earlier systems, and correct operation depends on the accuracy of that data. This article, the second of two on build control[1]*, describes how data is built and loaded by using the available computer aids, and the records and systems that support data update in service. The impact of major rearrangements such as software enhancements, extensions and remote concentrator unit reparentings are also briefly discussed.*

INTRODUCTION

Stored-program controlled exchanges like System X bring with them the need for new approaches to many traditional installation and administration tasks. This is particularly true of the various aspects of managing installation data, which is information about the particular subscribers, network, equipment etc. held in the exchange processor memory and which relates to a specific site.

Previous types of exchange systems also had to be set up specifically for each site, but generally this was achieved by using hard-wired straps, jumpers and threadings. Where data has been required for exchange systems (for example, TXE4A exchanges), this has been on a limited scale.

The difference with System X and other such systems is that considerably more detail is flexible and held as data, and that data can easily be altered by using a terminal or by loading a magnetic cartridge. Because of this, the information has to be prepared in a particular way, and several previously separate tasks have come together under the heading of *data management*.

SOFTWARE AND DATA

It is important in an article on data management to be quite clear what is meant by *installation data* and how this differs from *software*.

Software represents the standard set of rules (application programs) held in the exchange processor memory for dealing with the task of handling calls. The software is identical for all exchanges of a given type; for example, all trunk exchanges have the same standard software suite. This software is closely managed by means of controlled releases from a central library. The detail of software management is beyond the scope of this article.

Installation data is the specific information required to fit with the software in order to tailor it to each particular exchange site. It is therefore exchange dependent. This data describes to the exchange software the equipment configuration and the interfaces to the network. It controls every aspect of exchange operation, as indicated in Table 1.

RESPONSIBILITIES

Responsibility for the supply of software and data for each new System X installation currently lies with British Telecom (BT) (although this is under review). In this way it is possible for BT to know precisely what is loaded in the processor when the exchange is accepted from the contractor. The data load is therefore prepared in advance and supplied to the contractor some four weeks before the final acceptance demonstration. This allows a period for the contractor to load the unit with the supplied operational data and to check that it operates satisfactorily and, in conjunction with BT staff, to prepare test schedules designed to demonstrate the exchange facilities. Clearly, if the exchange is to perform as required, it is vital that the data load is accurate and correctly prepared.

TABLE 1
Installation Data

Type of Data	Information Held
Subscriber data	Class of service Equipment number allocation Meter allocation Special facilities
Circuit information	Signalling information Time-slot allocations
Network data	Routes Code translations
Hardware configuration	Equipment interconnections Time-slot use Alarms Tones and recorded announcements
Call accounting data	Tariff rates Tariff programme
Administrative data	Management statistics Test routines/values

ORGANISATION

Two separate organisations within BT, namely, Local Communication Services (LCS) and National Networks (NN), manage local exchanges and trunk exchanges, respectively. The approach for each is, however, essentially similar, in that local staff are responsible for almost all aspects of data management. Expert support is provided by headquarter's teams, who also develop computer aids and produce detailed guidance. The involvement of local staff is essential because the planned high rate of System X installation is such that more data builds have to be produced than a central team

† Local Exchange Systems Department, British Telecom Local Communication Services
* Trunk Services Operations and Maintenance Department, British Telecom National Networks

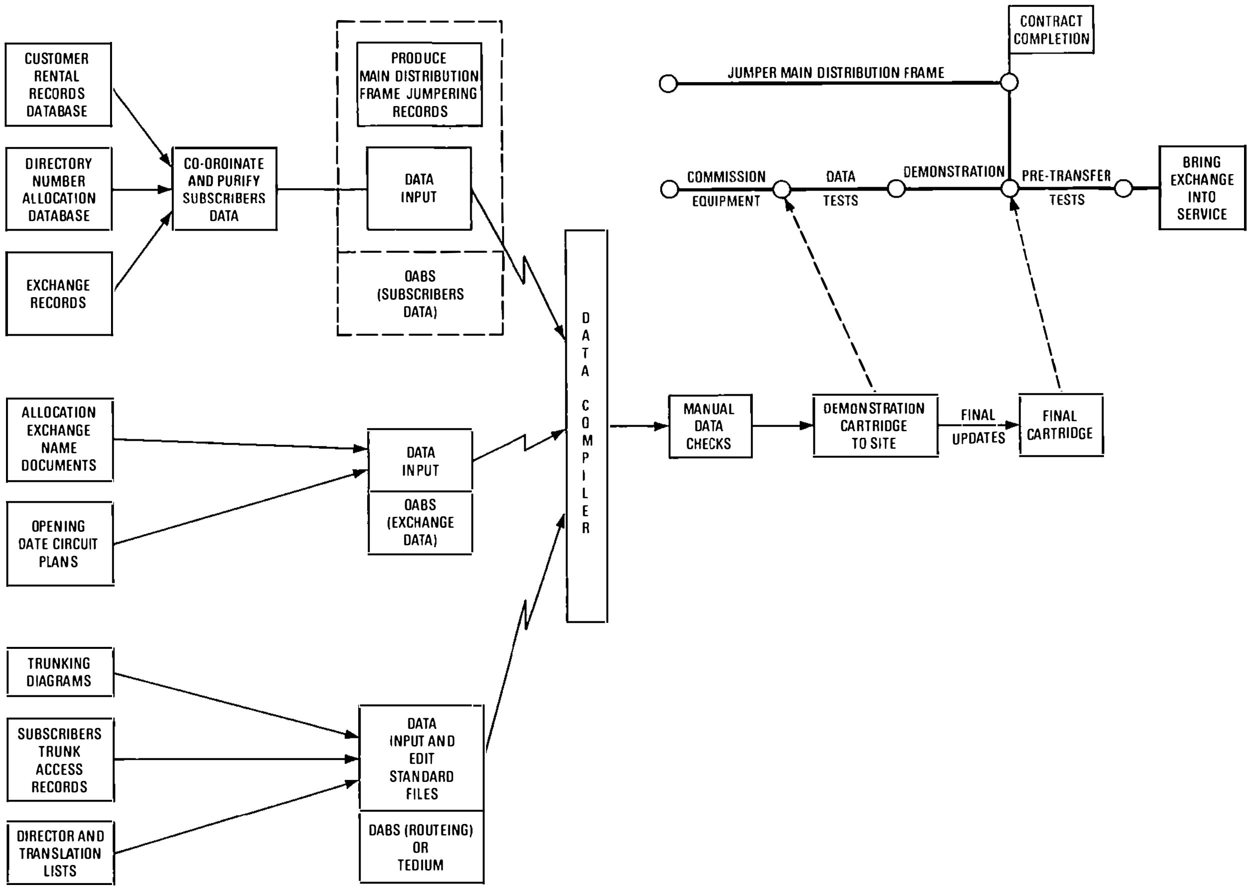

DABS : Data allocation and build system

TEDIUM : Trunk exchange data input and update medium

FIG. 1—Outline data preparation process

could support. Another benefit of this approach is that local experience is gained which is useful in the subsequent maintenance of the exchanges.

THE NEED FOR MANAGEMENT

Many factors make it essential that the production of data builds and their subsequent management are closely controlled. Among these are:

(*a*) the service offered by an exchange depends directly on its installation data and, hence, this must be correct in every detail;

(*b*) the size of the data is considerable (for example, approximately 90 man-days of effort are needed to prepare the one million words of data required for a 10 000-line exchange);

(*c*) certain aspects of the data, especially routeing data, are very complex, and compatibility with the network routeing strategy must be maintained throughout the network;

(*d*) the sources of information for the build are numerous and require considerable co-ordination to achieve the final build;

(*e*) a contractual obligation exists to produce a data-load cartridge by a specific date for formal exchange acceptance;

(*f*) unauthorised or spurious data changes after the exchange is in service must be detected and investigated; and

(*g*) data must be regularly archived and secured to prepare for catastrophic failure.

INFORMATION SOURCES

Comprehensive data-build handbooks are produced by the headquarters support groups in both LCS and NN. Training is provided by the BT Technical College at Stone in addition to Head Office seminars dealing with new developments and specific data-related topics. The required data is found in a wide range of manual records and computer systems. As these records can be inaccurate, the first task is to collect the data together and confirm its accuracy and consistency. This requires particular attention in the case of subscriber records, and for exchange replacements it is generally necessary to perform a 'jumper pulling' exercise to verify records. Certain information is provided at the contract stage, and must be incorporated into the build without change.

Fig. 1 illustrates the basic data-load collection and preparation process.

COMPUTER AIDS

Data Compiler

The principal computer tool used for building data loads is called the *data compiler*. This system was developed by Plessey Major Systems Limited (PMSL) under the System X development contracts, and provides powerful facilities for:

(*a*) accepting man-readable input;

(*b*) extensive vetting and checking;

(*c*) compiling data tables to match the precise requirements of the exchange software;

(*d*) producing load files that can be written to magnetic cartridge; and

(*e*) producing data listings.

A parallel facility, known as the *data decompiler*, provides the reverse function in converting an exchange archive tape

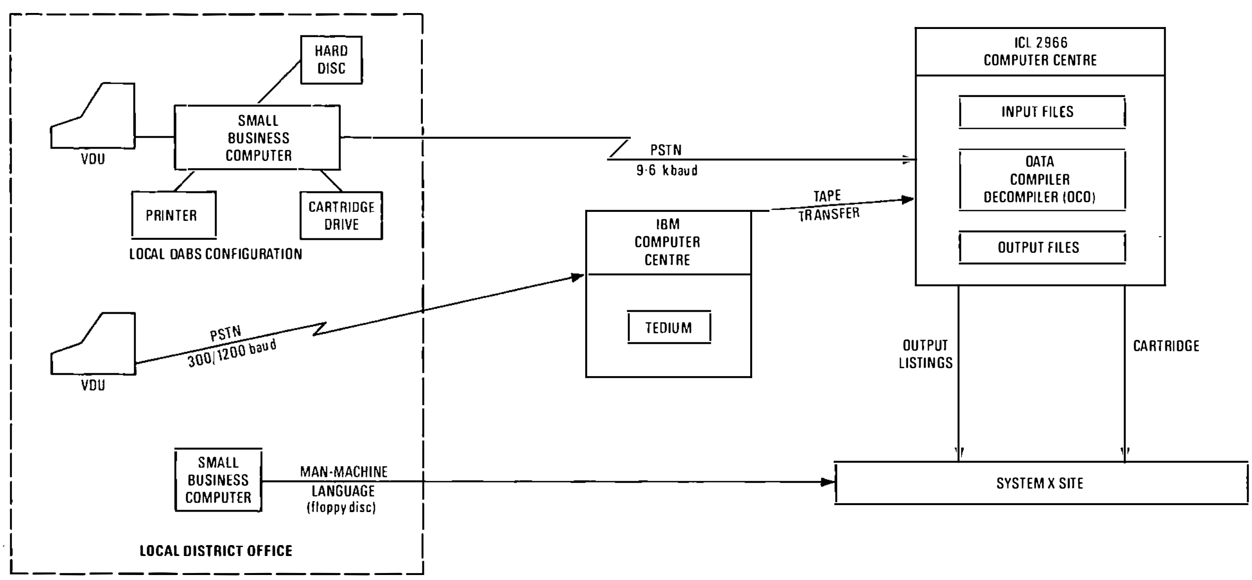

DABS : Data allocation and build system
TEDIUM : Trunk exchange data input and update medium
PSTN : Public switched telephone network
VDU : Visual display unit

FIG. 2—Data build tools

back to man-readable records, suitable for re-input to the data compiler if required.

The data compiler and decompiler (DCD) is a batch system which runs on an ICL 2966 DME mainframe computer. The system is characterised by long and expensive computer runs, together with very complex input to allow for all possibilities, some 80 different types of input form being used. However, both LCS and NN have developed front-end systems to simplify data input and vetting, save time and to reduce computer costs.

Data Allocation and Build System (DABS)

The data allocation and build system (DABS) is the system developed within LCS, and is shown in Fig. 2. DABS operates on locally-based 8 bit small business computers (SBCs) and provides

(*a*) formatted input with immediate vetting,
(*b*) shorthand input where possible,
(*c*) standard files,
(*d*) automatic allocation of subscribers to line cards (equipment numbers),
(*e*) file transfer to and from the data compiler mainframe computer,
(*f*) remote operation of the data compiler, and
(*g*) limited magnetic cartridge writing facilities.

DABS was initially intended for the preparation of local-exchange data, but with the increasing convergence of local and trunk data requirements, it will be enhanced to accommodate trunk exchanges.

Trunk Exchange Data Input and Update Mechanism (TEDIUM)

The trunk exchange data input and update mechanism (TEDIUM) is a group of programs available on an IBM mainframe computer which provide simplified data compiler input for early trunk exchanges with subsequent file transfer to the data compiler computer. DABS can be used for later versions of the trunk exchange system.

DATA PROVING

Accuracy and completeness are all important in preparing a data load; this can be checked in several ways, which, taken together, can give a high level of confidence:

(*a*) DABS and TEDIUM both perform extensive vets on data input;
(*b*) the data compiler additionally performs comprehensive cross checks between all related data items;
(*c*) the data compiler listings are manually checked by staff other than those who prepared the original data against the original records;
(*d*) the completed data load is tested for loadability and viability with the exchange software, either at the target site or at the field support unit;
(*e*) during testing prior to acceptance of an exchange from the contractor, the equipment is demonstrated by using the BT prepared data load; and
(*f*) before and during the exercise to bring a new exchange into service, detailed tests are carried out (pre and post transfer), to test compatibility with the network.

DATA LOADING

Data can be loaded into the exchange processor in two ways: either on an exchange terminal as man–machine language (MML) instructions, or on a magnetic cartridge as an image of the processor memory. The data compiler produces files of the latter type covering most of the data required, and these are input to the exchange using the 'loader' mechanisms.

Certain minor aspects of the data are not covered by the data compiler, and these are input directly as man–machine commands.

DATA ARCHIVE

The exchange system is able to archive data files to magnetic cartridge. These files are used for a number of purposes such as:

(*a*) recovery after catastrophic failure;
(*b*) decompilation for audit comparisons;

(*c*) decompilation to produce data listings for records; and

(*d*) decompilation to recycle data for extension or software enhancement purposes.

IN-SERVICE MANAGEMENT

Data management does not end when the exchange is brought into service. Network and customer changes will require alterations to the installation data for the life of the exchange.

Data changes are carried out by means of terminals connected via the operational and maintenance centre[2, 3] (OMC), or directly on site, by using MML instructions. MML instructions can also be loaded in bulk from floppy disc.

Exceptionally, where a very major change is required, the data can be dumped, amended and reloaded by using the data decompiler and the data compiler.

It is important to maintain records of the exchange configuration, and in the case of subscribers data (which changes most frequently) this is achieved automatically by means of a subscribers record system (SRS) database, which forms part of the OMC.

Records of other data aspects are much less subject to change, and are maintained manually. They are replaced as necessary by archiving the exchange data through the data decompiler to produce new listings.

Whenever changes have been made, an archive cartridge is dumped to provide an up-to-date back-up copy.

SOFTWARE ENHANCEMENTS

A more substantial need to manage change arises whenever a new software build has to be loaded. Generally, the new build is provided to add or improve facilities, and this will often impact on data structures.

Clearly, it is vital to maintain customer service during such an enhancement, and to provide for continuity of data, including that set by customers for star services, and current call-accounting meter values. This requires very precise and complex procedures, which are extensively proved on the captive model at the field support unit, before operational use.

In essence, the exchange data is archived to magnetic cartridge and then dispatched to the computer centre. Against a very tight timetable, the data is then decompiled, reformatted/augmented, and re-compiled to match the new software build. Once received back at the site, the enhancement procedure takes place, generally overnight, closely coordinated with any necessary operating system or hardware changes.

The procedure requires that the exchange processor backing store is partitioned and, whilst the exchange continues to operate on its existing software, the new data plus software is loaded into the other side of the backing store.

The exchange system is then restarted, with a short break during which new calls are not accepted. This break allows special software known as the *automatic transfer process* (ATP) to interrogate the 'old' memory to obtain current star-service and billing data (which will have changed since the original archive). This is then reformatted and written into the 'new' memory, thus continuity is preserved.

If the trial period is successful, the new software is written over the old and the enhancement is completed. At any point before that final step, the exercise can be aborted and the system restored to its original state in a controlled manner.

EXTENSIONS/REARRANGEMENTS

Any major rearrangement to an exchange once it is in service (for example, extending the exchange, or reparenting a remote concentrator unit to a different host processor) will require a procedure essentially similar to that used for software enhancements.

FUTURE DEVELOPMENTS

The major problem in establishing early mechanisms for data management is that, until development is complete, requirements are constantly changing. To cope with this problem, computer support tools tend to be designed with objectives of flexibility and simplicity, rather than ease of use or efficiency. Now that the design of System X is stabilising, requirements for computer support and procedural aspects are coming under review. In particular:

(*a*) the DABS and the data compiler/decompiler are likely to be re-developed to improve the support they offer, and to reduce operating costs;

(*b*) the feasibility of providing database systems to record aspects of data not covered by the subscriber record system is being examined;

(*c*) network management centres, with direct data update capabilities, are under development; and

(*d*) specific systems are being developed to support the preparation and maintenance of complex routeing data.

CONCLUSION

Data management is a complex and wide ranging discipline requiring high levels of skill and computer support. The performance of any stored-program control exchange depends directly on the quality of the loaded data.

A range of facilities now exists to support these activities for System X, and enhancements can be expected now that a level of product stability has been achieved, and field experience is becoming available.

References

1 Fox, M. J., and Storey, M. H. System X: Build Control, Part 1—Product, Hardware and Software Build. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 261.

2 Strickland, L. F., and Hewitt, M. A. New Operations and Maintenance Centres for Second Generation System X Exchanges. *ibid.*, Jan. 1985, **3**, p. 286.

3 Baty, R. M., and Sandum, K. N. System X: The Maintenance Control Subsystem. *ibid.*, Jan. 1985, **3**, p. 277.

Biographies

Nigel Milway joined the Post Office with an honours degree in Applied Physics in 1972. After a period in which he was involved with early System X specifications, he moved on to establish the configuration management mechanisms for System X development and production. Subsequently, he was concerned with the development of a range of computer support systems and, more recently, was responsible for defining data management requirements for System X and developing the DABS system. He is now head of the System X and TXE4 project management section in the Local Exchange Systems Department.

John Abram joined the Post Office as a Youth-in-Training in Shrewsbury Area in 1959. After a period with the External Planning Group, he joined the Network Planning Department in 1967 working on trunk network utilisation and routeing at Oswestry. In 1983, he joined the group in the Trunk Services Operations Division of NN responsible for the DMSU data build organisation. He is currently leader of the DMSU data build field support sub-group stationed at Oswestry.

In-Service Support for System X Exchanges

D. S. HOOKER, C.ENG., M.I.E.R.E.†, and J. HOWE*

UDC 621.395.34

This article discusses the philosophy behind the support organisation set up to assist field maintenance staff for System X exchanges and describes how this support is given for in-service exchanges.

INTRODUCTION

It has always been the objective of British Telecom (BT) to provide its customers with as high a quality of service as possible, consistent with economical provision of switching and transmission plant and combined with acceptable maintenance costs. With electromechanical switching systems, maintenance philosophy has centred on preventative maintenance, regular routining and special faults investigation. System X, with its high reliability, inherent diagnostics capability and control architecture, has required BT to look afresh at maintenance and support arrangements.

This article identifies the key issues in the support of System X exchanges and discusses the BT policy and plans which have been, and are being, put in place.

SUPPORT PHILOSOPHY FOR SYSTEM X

To provide for an effective support organisation for field staff, the following are necessary:

(*a*) the exchange system design must be tested fully and proven before introduction into the field (this activity is known as *system proving*), and

(*b*) the field maintenance staff must be trained fully to maintain the exchange system when it is in service, and have access to adequate documentation.

When these prerequisites have been met, then a support organisation can be built around the following philosophy:

(*c*) the field maintenance staff should have access to comprehensive diagnostic facilities; and

(*d*) the vast majority of exchange faults should be located and cleared by the field maintenance staff by using the diagnostic facilities and changing slide-in-units (SIUs) where necessary.

However, it is recognised that the complex nature of processor-controlled exchange systems means that the field maintenance staff may not be able to deal with every fault, especially during the early build up of exchanges in service and while expertise is still being gained. Also, a low fault rate militates against the aquisition of a high degree of expertise by all field maintenance staff; to achieve this, such expertise must be centralised. Thus, for those exchange faults that cannot be dealt with by the field maintenance staff, a support agency must be available to provide assistance.

Such a support agency must provide 24-hour all-year-round cover for urgent problems; that is, where the field maintenance staff have been unable to deal with an actual or potential major service failure of an exchange.

On those occasions when the BT support agency is unable to deal with a problem, the support agency must be able to call upon the exchange supplier, as the design authority, for additional support, which again must be available on a 24-hour all-year-round basis.

To implement this philosophy, the BT maintenance support agency (MSA) was set up in 1979.

EARLY SUPPORT ARRANGEMENTS

The basic functions of the MSA were:

(*a*) to provide a 24-hour 7-days-a-week contact point for field maintenance staff who had an exchange in trouble, so that such staff could be put in contact with an expert who could provide assistance;

(*b*) to provide assistance over the telephone and, where necessary, arrange a visit to the site by an expert;

(*c*) to solve problems in priority order by providing a single interface to the function responsible for design correction;

(*d*) to manage the problem solving activities, in particular to track the stage reached in the resolution of each problem;

(*e*) to monitor statistical and numerical information on the performance of in-service exchanges, so that any exchange degradation could be spotted before real trouble occurred; and

(*f*) to be present at the exchange during vulnerable periods such as the first few days of service, and during enhancements.

Field maintenance staff could contact the MSA at anytime, but there were specific instructions that the MSA must be contacted immediately when:

(*a*) the procedures for handling the fault given in the operations and maintenance manual or other documentation supplied by the support agency had been exhausted, or

(*b*) one hour had elapsed since the fault first appeared, or

(*c*) service was being severely degraded.

During the normal working day, contact with the MSA could be made directly by telephone or facsimile. Outside normal hours, staff would ring a dedicated answering machine to record their name, telephone number, exchange in difficulty and the nature of the problem. The duty officer would then interrogate the answering machine after being radiopaged and contact the original caller; it is an objective that the field maintenance staff should be in contact with an expert within 10 minutes.

If the problem was of a serious nature (that is, service affecting) and could not be dealt with over the telephone, the MSA duty officer would send his best-qualified available

† Trunk Services Operations and Maintenance Department, British Telecom National Networks

* Local Exchange Systems Department, British Telecom Local Communications Services

expert to the exchange to get it back to normal. Further back-up support from the exchange suppliers was available if necessary.

The majority of emergencies did not require a visit to site and, when the problem did not necessitate immediate action, the field maintenance staff would write a full description of the fault on an exchange incident report (EIR), which would be sent to the MSA. Here, the course of action to be taken would be decided and a responsible officer appointed whose function was to see the problem through to solution in an appropriate time-scale. The usual procedure was for the problem to be investigated by the field support unit (FSU) at Martlesham Heath, where the BT model could be configured to match the exchange with the problem. If the FSU could not solve the problem, a model incident report (MIR) describing the problem in detail would be sent to the function responsible for design correction (see below). Depending upon the priority of the fault, the problem would be solved either immediately, within 2 or 3 days, or the solution would be incorporated into the next system build (say 6 months). All solutions are tested on the BT model before release to the field through configuration control.

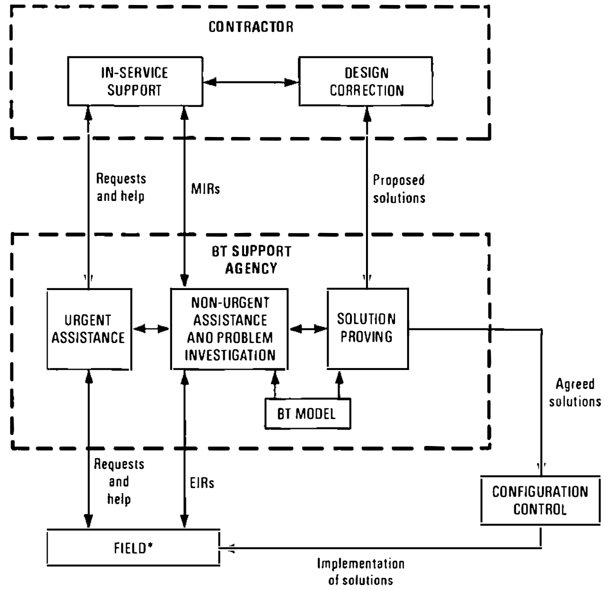

* Includes District organisation where appropriate

FIG. 1—Basic support organisation

LESSONS LEARNED FROM EARLY EXPERIENCE OF SUPPORT

The major lessons learned were:

(*a*) the architecture and design of the system provided the basic framework for a reliable and maintainable system;

(*b*) the system could be maintained by existing field staff if they had reasonable training and good maintenance documentation;

(*c*) the system proving activities cannot test the design fully and some design defects will emerge only when the exchange is put into service;

(*d*) second- and third-line maintenance support to in-service exchanges is required, as the field maintenance staff do not have the expertise to deal with the more obscure faults;

(*e*) a single MSA to cover the whole range of systems and number of exchanges coming into service is impractical;

(*f*) if a specialist is required on site at very short notice, a geographically centralised support unit does not fit the bill in all cases;

(*g*) the MSA needs access to specialist experts to deal with the more obscure problems;

(*f*) to avoid as far as possible calling out specialists to site, much more emphasis must be put upon solving problems remotely, and the development of tools to assist in this fuction is necessary;

(*h*) in the early days, the support function must, to a large extent, rely on the ongoing development teams; this can lead to conflicting priorities and competition for limited resources, hence it is necessary to separate the support and development functions;

(*i*) the support requirements from the contractors must be defined clearly; and

(*j*) improved telecommunications and computerised databases are required.

ONGOING SUPPORT ARRANGEMENTS

The reorganisation of BT into the two separate businesses of Local Communications Services (LCS) and National Networks (NN) provided management with the opportunity to establish support agencies which were more independent of the development departments. Staff with the necessary expertise were transferred to LCS Headquarters and NN Head Office along with ownership of the BT models. The latter is vital as even expert staff lose their expertise gradually unless they are able to work on a model (or exchanges) on a regular basis.

For LCS, the imminent establishment of District offices provides the opportunity for the establishment of the necessary geographically dispersed maintenance support teams, while for NN, the density of System X trunk exchanges means that a single Head Office second-line support function is adequate. The present support agencies are the field support agency (FSA) for LCS, and the trunk exchange support agency (TESA) for NN, the former MSA having been absorbed into the current FSA. Although the LCS and NN arms of BT now have their own respective support agencies, the support philosophy and principles and functions in the early arrangements outlined above have not changed. Fig. 1 shows the basic features of the support arrangements.

INVOLVEMENT OF CONTRACTORS

As the design authority, the contractor is responsible for producing new exchange builds and providing solutions for all design defects uncovered during operation of the exchanges. Such solutions may be in the form of software fixes, hardware/firmware modification action packs, or complete new builds, depending on the complexity of the solution. In this context, it is necessary that the contractor has awareness of exchange performance in service and, to this end, certain performance criteria are monitored. The results are fed back to the contractor so that any deterioration in exchange performance can be detected and dealt with before there is any serious degradation to customer service.

As well as monitoring performance criteria, at least during the early life of an exchange, the contractor provides BT with in-service operational support (ISOS) and product design support (PDS) services. The ISOS service provides for telephonic advice by the contractor to an exchange site in serious difficulty, that is, when there is an actual or potential major service failure, and an immediate visit to site by the appropriate expert where this proves necessary. In addition, the contractor is able to provide advice to the LCS Headquarters and NN Head Office support organisations for exchange problems of a less urgent nature. The

ISOS arrangements are called upon when either:

(*a*) BT has insufficient expertise to deal with the problem itself (for example, a complex processor problem requiring a processor expert); or

(*b*) it is more sensible, in the first instance, for the contractor to deal with an urgent site problem, for example, a site visit by an expert is necessary, and the contractor's expert is closer to the exchange site than the BT expert, hence the contractor can be on site more quickly.

The PDS service provides for correction of design defects found by in-service experience, which includes temporary software patches, and the formal release of new builds of software, firmware and/or hardware to BT for application at in-service exchanges as described below.

DESIGN CORRECTION

As with every newly launched exchange system, once the exchanges go into service and carry live traffic, minor design weaknesses that have remained undetected during system proving become apparent. These defects need to be corrected on exchanges which have yet to be installed, as well as those already in service.

All problems found at in-service exchanges which, it is felt, may be due to design defects, are investigated by the system experts within the support agency. However, as BT is not the design authority for System X—this resides with contractors—then BT's responsibility for getting design defects corrected lies in defining the problem, bringing it to the attention of the contractor, and testing any design correction submitted by the contractor to verify that the original defect has been cured. Once this is established, then the normal configuration-control procedures are followed to ensure the design correction is implemented in the field in a controlled manner.

Problems brought to the attention of system experts are investigated, and if the problem proves to be a design defect, then the problem is defined as fully as possible and passed to the contractor. Once a design defect has been notified to the contractor, it is his responsibility to produce a solution that is acceptable to BT. The normal procedure is for the contractor to test the solution to his own satisfaction using whatever resources are available, and then pass the solution to BT to test as necessary the proposed solution on its own model(s) before releasing the solution to the field via configuration control for implementation on in-service exchanges. In addition, the contractor applies the agreed solution to those exchanges still in production and to those installed but not yet handed over to BT.

In practice, all such solutions are in the form of changes to the build level of the exchange and, as a build is defined in terms of hardware, software, and firmware, solutions to problems are in terms of changes to any (or a combination) of these. In addition, documentation will need to be updated in line with any changes to hardware, software and/or firmware, although exceptionally, solutions to problems will comprise only changes to documentation; for example, corrections to incorrectly documented procedures.

USE OF BT MODELS

Because of the complex nature of processor-controlled systems, on-site maintenance action comprises the use of comprehensive diagnostic facilities and the consequent changing of SIUs and other hardware items of equipment. Detailed design-defect investigation must be carried out at a captive installation; that is, one that carries no live traffic and is also capable of reflecting the build of in-service exchanges. Such an installation is located at the BT Research laboratories at Martlesham Heath, and this installation contains a model for each type of System X exchange in the field. Also, there is a Release 1 operations and maintenance centre (OMC), two basic OMCs (LCS and NN variants), and a test bed of existing exchange types, namely Strowger, crossbar, electronic, auto-manual centres etc. including a comprehensive range of signalling systems and customer apparatus; this enables interworking tests to be performed on all the different types of exchange equipment to be found in the field to verify compatibility with respect to interworking. In addition, artificial traffic generators are used to simulate traffic flow in a network environment. All in-service problems arising on System X exchanges are investigated on these models, before being passed to the contractor for a solution (assuming a design defect is brought to light), and all solutions to design problems proposed by the contractor are tested and proved on the model before being released to the field for implementation.

Although the *raison d'etre* of the models is for supporting the in-service exchanges in investigating field problems, the models are used also for the acceptance of new products before they go into service, and such new products will vary from enhancements which provide new facilities with minimal or no hardware changes, to complete new systems requiring a new model.

In addition to the BT models, there are also models located at the contractors' works and, although these have been used as development tools, some will be retained by the contractors for their own use in investigating in-service problems brought to their attention by BT, as well as being used for ongoing development activities.

CURRENT DEVELOPMENTS

To help improve the maintenance and support of System X exchanges, various developments are either under way or being considered, these include:

(*a*) *Paperless EIRs* Instead of the exchange maintenance staff at the operations and maintenance unit (OMU) filling out a paper EIR and forwarding it to the support organisation(s), on-line access over a dial-up link to a central EIR database will be provided, and enable exchange maintenance staff to enter EIRs directly onto the database from a suitable terminal. Similar on-line access will be available also to the support organisation, including contractors. This database is being developed to include analysis of more than one EIR into a common problem, electronic messaging, management information relating to EIR progressing, and the transmission of software patches directly to configuration control in order to avoid transcription errors that tend to arise when the convertion from one media to another takes place; for example, paper to floppy disc. Eventually, this facility may be extended to allow the direct loading of software fixes onto the exchange from the database.

(*b*) *Remote Access Terminals* These will enable members of the BT support organisation, as well as the field maintenance staff in the OMU, to gain access to an exchange from a remote point by using the man–machine interface; such access will be allowed only under strict control of the exchange maintenance staff. Initially, the access will be restricted to the interrogation of a buffered rolling 2-hour record of the exchange transaction log; subsequently, access will be extended to on-line interrogation and also WRITE capabilities, although the latter is still under consideration. This interrogation feature will be used to enable an expert at a remote location to gain a picture of the events leading up to a particular exchange difficulty. The WRITE feature will enable remote insertion of diagnostic commands to investigate further exchange problems, and manual re-configuration of the exchange to preserve the best possible quality of service in times of difficulty. Appropriate security features and procedures will be employed to ensure correct use of the system.

(*c*) *Exchange Maintenance and Control System* This will provide for enhanced terminal facilities, and it can be added to a basic OMC so that all terminal users have access to the enhanced facilities, or be used as an enhanced terminal in its own right. Phase 1 of the development provides for numerous databases and will allow local command composing and editing, including MML command assistance, combined with simultaneous reception of information from the exchange; MML command recall for editing and retransmission; fault output analysis; and the operations-and-maintenance manual stored in a paperless form on one of the databases with easy access and retrieval. Phase 2 of the development will further enhance the system to provide a knowledge base to assist in fault localisation; that is, a step towards an intelligent knowledge-based system.

CONCLUSION

In-service support experience to date has demonstrated the need for a carefully defined fault escalation procedure involving field maintenance staff, support agency staff and contractors. Particular emphasis has to be placed on ensuring that vigorous procedures are followed when an exchange is in difficulty. In addition, the need for a high degree of expertise, combined with adequate back-up tools and documentation has been highlighted. To provide customers with the high quality of service they expect, it is essential that BT has in place the necessary support arrangements, which will be continually reviewed to ensure that they match the needs of the District Managers.

Biographies

Denis Hooker joined BT in 1960 as a Youth-in-Training in London South East Area. After training, he worked on exchange construction and exchange maintenance on ferrite-core register-translators. In 1968, he joined the Headquarters Development Department as an Assistant Executive Engineer and, ultimately as a Head of Group, worked on register-translator development during which time he designed a complete IC for use in the MOST director. In 1976, he moved to the Network Strategy Department and six years later joined the National Networks Trunk Services Department as Head of the trunk switching maintenance policy and procedures section.

Jim Howe joined BT in 1961 as a Youth-in-Training in Glasgow Telephone Area. After graduating from the University of Strathclyde, he joined Headquarters Development Department in 1969 as an Executive Engineer working on the specification of common-channel signalling systems. From 1976, as Head of Group, he led the System X reliability and maintenance team responsible for the specification of the maintenance control subsystem. From 1979 to 1984, he managed the System X proving team responsible for the Release 1 local exchange. He is currently Head of Section in the Digital Maintenance Support Division of the Local Exchange System Department.

Early In-Service Experience of System X Exchanges

G. E. MATHIAS, C.ENG., M.I.E.R.E., K. N. SANDUM, MA., C.ENG., M.I.E.E.,† and S. BARBER, M.SC., B.SC., C.ENG., M.I.E.E.*

UDC 621.395.34

This article describes the introduction of the early System X exchanges into British Telecom's network, discusses, in broad terms, their performance and operation, and reports on the experience of users of the system from both the administration's and the customers' viewpoint.

BACKGROUND

British Telecom (BT) announced its new range of digital exchanges at the Geneva International Telecommunications Exhibition in September 1979. The range comprised a remote concentrator unit (RCU), a medium and a large local exchange (MLE and LLE), a junction tandem exchange (DJSU), a combined local and trunk exchange (DPLE) and a trunk exchange (DMSU). The first exchange to go into public service was Baynard House digital junction switching unit (DJSU). This was an early version of System X and made use of a GEC Mk. IIBL processor; it opened in July 1980. It was configured as a tandem exchange serving director exchanges within Central London, and provided them with access to exchanges to the south-west of London (see Fig. 1). This exchange was replaced in July 1981, by which time it had provided a very useful experience in the operation of a digital exchange in the existing analogue network.

The first System X local exchange to open was at Woodbridge, in Suffolk, in July 1981 (see Fig. 2). Shortly afterwards, in December 1981, a local exchange at Arrington and a digital main switching unit (DMSU) at Cambridge were brought into service (see Fig. 3).

The largest of the System X local exchanges in service is at Hale, near Liverpool, and has 2400 connections. It is configured as a director exchange (see Fig. 4); Woodbridge and Arrington are both non-director exchanges.

System X is evolving in terms of hardware, software and facilities available by way of a system enhancement programme (SEP). The first trunk version of the new SEP exchange was brought into service at Coventry in December 1983. It is equipped with a new processor[1], which is planned for use in all future System X exchanges.

† Local Exchange Services, British Telecom Local Communication Services

* Trunk Services Operations and Maintenance, British Telecom National Networks

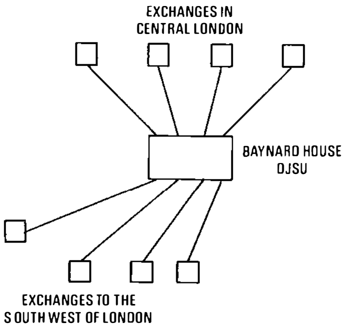

FIG. 1—Baynard House DJSU network configuration

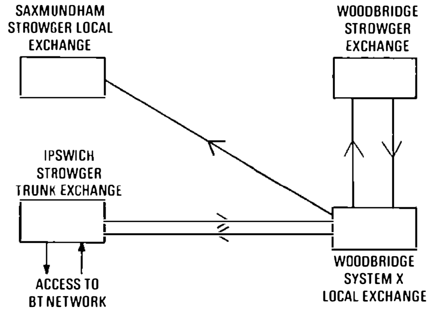

FIG. 2—Woodbridge local exchange network configuration

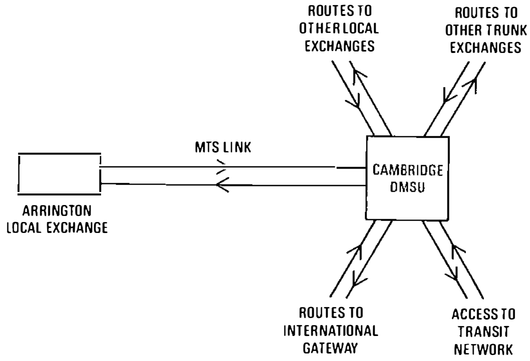

FIG. 3—Arrington local exchange and Cambridge DMSU network configuration

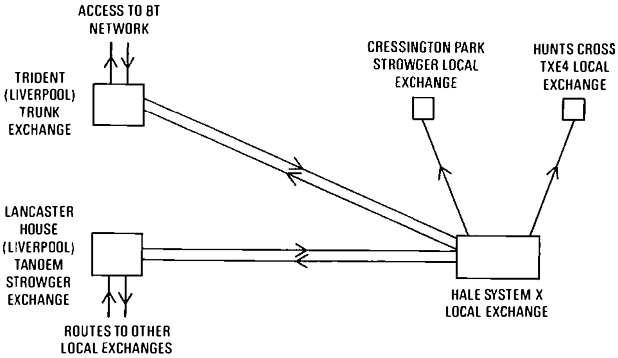

FIG. 4—Hale local exchange network configuration

Brief details of the early System X exchanges are given in an Appendix to this article.

EVOLUTION AND ENHANCEMENT

Early Local Exchanges

The early local exchanges have undergone some major planned enhancements during the first three years in service. The original build at Woodbridge, when it opened, was capable of providing all the standard telephony services, but lacked the new star-service facilities to be provided by System X. The star services to be provided by System X are:

Abbreviated dialling
Call diversion
Advice of duration and charge
Reminder call
Repeat call
Call baring

The first enhancement (E1), applied to Woodbridge in December 1981 and to Arrington during February 1982, enabled BT to introduce two of these star services at Woodbridge. The star services offered were abbreviated dialling and three types of call diversion (that is, permanent, no reply and busy).

The second enhancement (E2) was loaded at Arrington in February 1983 and at Woodbridge in July 1983. This enabled the use of bubble memories to secure billing data and customer-alterable data associated with the star services. The enhanced software also provided an improved maintenance control subsystem (MCS)[2], and permitted direct digital interconnection between Arrington and its trunk exchange, Cambridge DMSU, using common-channel signalling.

A third enhancement (E3) was used as the opening build for Hale exchange in June 1983, and was loaded on Arrington in February 1984 and Woodbridge in March 1984. This build was dimensioned to give the greater exchange-line capacity needed for Hale and to provide improved maintenance facilities; for example, digital-switch diagnostics.

The current local exchange software was loaded at all three local exchanges during June 1984 and allows BT to offer itemised billing.

Early Trunk and Digital Tandem Exchanges

Baynard House DJSU and Cambridge DMSU both opened in 1981 with a similar 200 erlang software build. The first enhancement provided improved MCS facilities and the ability for common-channel signalling interworking with Arrington. This enhancement was applied to Cambridge DMSU in December 1982 and to Baynard House DJSU in March 1983. In May 1983, a further enhancement, incorporating improved maintenance features, was applied to Cambridge DMSU and forms its current build.

A new software build was loaded at Baynard House DJSU in October 1983 to give an increased traffic capacity of 500 erlangs and an improved processor operating system.

A further two enhancements are planned for the early trunk/tandem exchanges to increase their nominal traffic capacity to 1800 erlangs, and to allow common-channel signalling interworking with the new SEP exchanges.

SEP Trunk Exchanges

The initial SEP trunk exchange was brought into service at Coventry Spires, in December 1983, and has sucessfully demonstrated the reliability and the power of the basic SEP processor architecture and design. The unit used a single-cluster four-CPU processor.

This design has been enhanced to support common-channel signalling and a number of maintenance-related features. In addition, this system has the capability of carrying a nominal load of 3600 erlangs and 135 000 BHCA†. This capacity will be extended further with the introduction of the first multi-cluster processor units, which will support a nominal load of 18 500 erlangs and 500 000 BHCA.

MAINTENANCE AND OPERATIONAL EXPERIENCE

These first System X exchanges have provided BT with valuable maintenance and operational experience of an on-line updateable common-control digital exchange in the national network.

During the early in-service period of each exchange, assistance was provided to the local maintenance staff by on-site experts from both BT headquarters and the manufac-

† BHCA—Busy hour call attempts

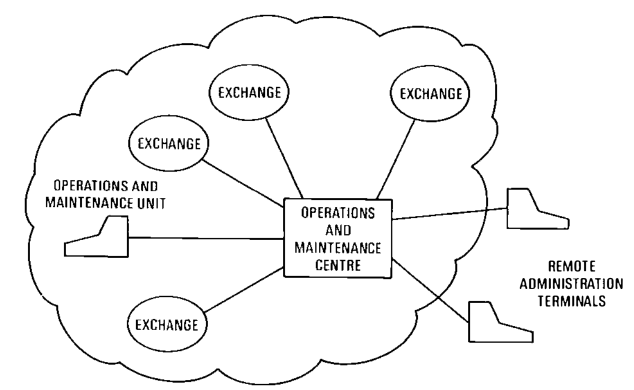

FIG. 5—Remote operations and maintenance making use of an OMC

turers. This was phased out as the experience of the local maintenance staff increased. The speed at which this occurred has been good and has increased with each new installation.

The ability to perform operations and maintenance activities remotely paved the way for the introduction of an operations and maintenance centre (OMC)[3] (see Fig. 5). This is, in essence, a concentration/interconnection point for both input to and output from a number of System X exchanges by users (both engineering and non-engineering), within a geographical area. All future local System X exchanges will, when brought into service, be connected to an OMC.

To gauge the reaction of potential non-engineering users, a trial using remote terminals was carried out in the Cambridge Telephone Area. BT staff have adapted well to the new procedures and techniques necessary for the maintenance and operation of System X exchanges. A further trial, this time making use of an OMC, is planned to take place in Central London.

System X exchanges also identify problems encountered in the network to which they are connected; investigation of these problems is assisting greatly in improving the performance of the network and the quality of service given to customers.

PERFORMANCE

Three key indicators are used in this article to illustrate the exchange performance of System X; these are

(*a*) locally-sited test-call senders—these give an accurate indication of call failures attributable to the exchange (see Fig. 6),

(*b*) the result of test-call sending performed by measurement and analysis centres (MACs)—this too gives an indication of own exchange call-failure rates and, since the same equipment is used to assess other BT switching systems, it enables comparison with other systems (see Fig. 7), and

(*c*) customer reports (see Fig. 8).

The local exchange development specification calls for an overall own-exchange call-failure rate of not more than 0·08%; this is achieved and bettered. Both the number of customer complaints attributable to exchange faults and the number of out-of-hours call outs are currently better than those experienced on TXE4 exchanges. The hardware failure rates are also better than the predicted values.

The introduction of any new exchange system inevitably brings with it changes in procedures, practices and techniques. BT staff have responded well to these changes and the challenge that System X has brought. Once over the critical familiarisation phase, maintenance staff have adapted well to the changes and are able to deal competently with any problems that have arisen on the system.

The service experienced by customers on System X exchanges is generally superior to that on existing exchanges.

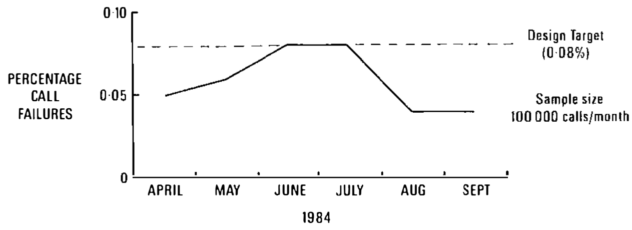

(*a*) Hale local exchange

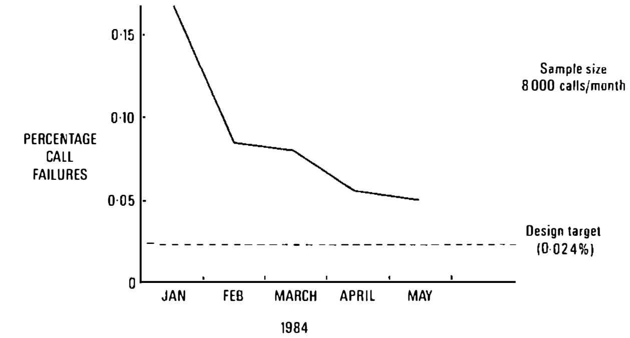

Note: Results from locally sited call senders

(*b*) Coventry Spires trunk exchange (first five months in service)

FIG. 6—Typical figures (cumulative) for own-exchange call failures

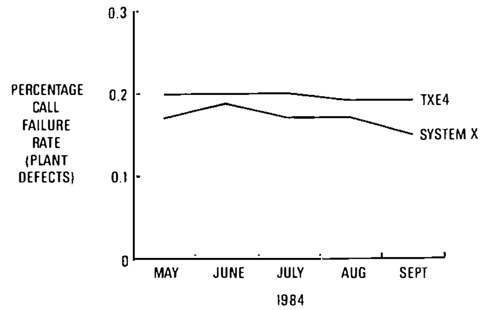

Note: The figures are cumulative from April. The System X results are from Hale local exchange

FIG. 7—Typical own-exchange failure rate obtained by measurement and analysis centre—comparison with TXE4 national results

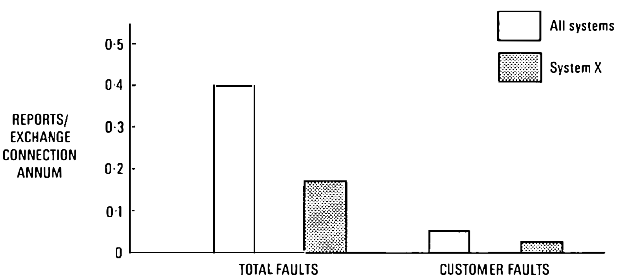

Average values taken over a sample period April to August 1984

FIG. 8—Reported faults

The automatic announcements, together with the star services, are key features of System X and have been well received (particularly their quality) by customers. However, the full potential of System X facilities, and the improved service possible, will not be realised until there is a significant penetration of System X exchanges. Only then will the full benefits of a fully integrated digital network be made available to the customer.

CONCLUSION

In introducing System X, BT has undertaken a major step forward to a completely new generation of processor-controlled digital switching systems. There have inevitably been problems, but overall, BT's experience of its early in-service System X exchanges has been good. They have exceeded the performance of other exchanges in service despite being subject to an almost continuous programme of phased enhancement, and are only now reaching their planned realisation.

These early exchanges have provided BT with valuable knowledge and experience in the operation, maintenance and support of a modern digital switching system, and staff have adapted well to the change and the challenge.

This early experience and the rapid penetration of the SEP exchanges will provide a sound basis on which to build the future telecommunications network of the country, and provide a good-quality service to customers.

References

1 TROUGHTON, D. J. *et al.* System X: The Processor Utility. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 226.

2 BATY. R. M., and SANDUM, K. N. System X: The Maintenance Control Subsystem. *ibid.*, Jan. 1985, **3**, p. 277.

3 STRICKLAND, L. F., and HEWITT, M. A. New Operations and Maintenance Centres for Second Generation System X Exchanges. *ibid.*, Jan. 1985, **3**, p. 286.

Biographies

Glyn Mathias joined the then Post Office in 1964 as a Youth-in-Training in the Portsmouth Telephone Area. He joined the Transmission Maintenance Division of the Engineer-in-Chief's Office, on promotion, in 1966, where he was involved with analogue transmission maintenance and, later, with the maintenance requirements for switched digital data networks and System X. Since the introduction of System X in 1980, he has been involved with in-service support and performance. He is currently Head of Group responsible for performance monitoring, assessment and improvement of public digital switching systems.

Keith Sandum joined the Post Office in 1963 as a Youth-in-Training. After a period on transmission maintenance, he joined the Computer Systems Branch where he was responsible for acceptance testing large computer systems and in establishing the London Airport Cargo Electronic System (LACES) and the Direct Keying System for entry of bills at the National Girobank. He was then involved with packet switching for five years, including working at the ARPA network measurement centre in California, the specification and management of the MMI and control system for the experimental packet switching system (EPSS) and, finally, as the BT member of the project team for EURONET. For the past four years, he has been responsible for various aspects of the operation and maintenance of System X, including software maintenance, testing, build control and performance measurement. He is now Head of Section responsible for the maintenance, repair, support and operations policy for System X, UXD5 and other digital switching systems.

Stephen Barber gained a first-class honours degree in Physics at Manchester University in 1972. He joined BT as an Executive Engineer in the teletraffic division and, in 1975, under BT sponsorship, gained an MSc. in Digital Electronics at Manchester University. After work on the design and development of special call detail monitoring equipment for use in dimensioning the System X network, he moved to the System X Development Department where he was involved in both hardware standards and system integration. He is currently in the Trunk Services Division of National Networks with responsibility for digital trunk network switching maintenance support and has been closely involved with the introduction into service of the initial SEP System X trunk units.

APPENDIX

Early System X Exchanges

Date In Service	Exchange	Type	Brief Details	Notes (October 1984)
July 1980	Baynard House DJSU	Tandem (non standard)	Non-standard multi-processor (GEC Mk. IIBL) Nominal capacity 200 E Average busy-hour traffic 100 E	Recovered July 1981
July 1981	Woodbridge	Local Non-Director	Duplicated main/stand-by processor 900 customers 34 analogue junctions 56 digital junctions	In service at final build
July 1981	Baynard House DJSU	Tandem	Multi-processor Nominal traffic capacity at opening 200 E Average busy-hour traffic at opening 100 E Enhanced nominal traffic 500 E	Being enhanced to 1800 E
December 1981	Arrington	Local Non-Director	Duplicated main/stand-by processor 700 customers 5 analogue junctions 63 digital junctions (MTS signalling)	In service at final build
December 1981	Cambridge DMSU	Trunk	Release 1 multi-processor Nominal traffic 200 E Average busy hour traffic 100 E	In service at final build
June 1983	Hale	Local Director	Duplicated main/stand-by processor 2400 customers 20 analogue junctions 129 digital junctions	In service at final build
December 1983	Coventry Spires DMSU	Trunk (SEP)	Single-cluster version of SEP processor Nominal traffic capacity 1000 E Average busy-hour traffic 100 E	Being enhanced to 3600 E

DJSU : Digital junction switching unit DMSU : Digital main switching unit SEP : System enhancement programme

System X: Maintenance Control Subsystem

R. M. BATY, M.TECH., B.SC(ENG)., C.ENG., M.I.E.E.†, and K. N. SANDUM, B.A., C.ENG., M.I.E.E.*

The maintenance control subsystem (MCS) provides a centralised control to co-ordinate the maintenance activities within the subsystems that make up System X. This article discusses BT's maintenance philosophy for its digital exchanges and outlines the facilities provided by the MCS.

INTRODUCTION

System X is made up of a series of modular subsystems, both hardware and software. Although the subsystems themselves have internal maintenance facilities, there is a need for a centralised control to co-ordinate subsystem maintenance activities. This is provided by the maintenance control subsystem (MCS), which consists of software and hardware, and gives the following main facilities:

(*a*) a man–machine interface (MMI) that enables maintenance staff to determine and change the state of the system hardware and data from a terminal;

(*b*) a collection point for all fault reports and alarm indications, together with an MMI to output this information to enable maintenance action to be taken; and

(*c*) an automatic control which prevents the inadvisable use of software and hardware under fault conditions, and which, if necessary, reconfigures the system to achieve this. When such fault conditions are removed, the hardware and software can be brought back into service. The automatic control also rejects maintenance actions that would jeopardise system performance or security.

MAINTENANCE PHILOSOPHY

British Telecom's (BT's) objectives for its digital exchange (TXD) maintenance organisation are:

(*a*) to minimise the degree to which faults become service affecting by the use of surveillance of exchange processor output etc.,

(*b*) to restore customer service at the earliest time when faults are service affecting and the system has not automatically re-configured itself to safeguard service, and

(*c*) to use manpower and other resources in a cost/benefit-effective manner.

The TXD maintenance organisation of BT Local Communications Services (LCS) is structured around two operational entities:

(*a*) *Operations and Maintenance Centre* (OMC) This is a location at which, or through which, various operational and maintenance functions are carried out. The users of the OMC cover several operational disciplines and are dispersed throughout the service area.

The OMC comprises a processor, associated port switching hardware, modems and peripheral equipment.

(*b*) *Operations and Maintenance Unit* (OMU) This is an engineering base within an OMC area; it comprises a number of work positions with visual display terminals/hard-copy printers, together with other engineering support facilities, and is the focal point for the maintenance of all the TXDs in a defined geographical area.

An OMC can have one or more OMUs connected to it. The first OMU, which is usually co-located with the OMC equipment, is called the *main OMU*, and the second and subsequent OMUs *satellite OMUs*. An OMU, together with all the TXDs with which it is associated, is called an *OMU area* (OMUA).

The responsibility for the service given by the TXDs rests with the engineering maintenance manager in charge of the OMUA whose responsibilities are to monitor the service given by all the TXDs in the OMUA and to ensure that the maintenance and allied engineering activities are executed in the most expeditious manner.

Exchange alarms and automatic fault report output are presented at both the exchange and the OMU at all times. Surveillance of these alarms and outputs is carried out at the TXD and the OMU; if the exchange maintenance officer is in attendance at the TXD he will take the initiative for appropriate action.

If the maintenance officer is not in attendance, the OMU initiates remedial activities and fault action. If, after initial attention, a site visit is deemed necessary, the OMU contacts the maintenance officer and advises him accordingly. On arrival at the TXD, the maintenance officer assumes responsibility for remedial activities.

The OMU is the fault record and documentation centre for the OMUA; it has facilities for carrying out in-depth analysis of problems/faults and, as such, provides technical support to maintenance officers when required.

The OMU is the communication point for all matters affecting the integrity of the network as a whole at all times, and for all other matters concerning individual TXDs when the maintenance officer is not in attendance.

INTERFACE WITH THE SYSTEM

The prime interface with the system, be it at the OMU or on site, is via a terminal—the man–machine interface (MMI). The input from a terminal is checked by the MMI and is routed to the appropriate subsystem for action; in most cases, this is the MCS.

When problems or irregularities are detected within the system, they are reported by the owning subsystem to the MCS, which logs the incident and arranges for an appropriate output to both the local terminal and the OMC.

THE OPERATIONS AND MAINTENANCE MANUAL

Changes to the system data (for example, provision of service) are carried out by using man–machine language (MML) commands. The possible changes and the MML commands necessary to perform them are fully detailed in the operations and maintenance manual (O&MM). The O&MM also provides information for the maintenance staff to assist them in the interpretation of fault output and the necessary MML commands to perform reconfiguration or restoration in the event of total or partial system failure.

EMACS

To assist the maintenance staff in interpreting fault output and carrying out maintenance procedures involving the use of MML, an intelligent terminal is being developed. This consists, in essence, of a microcomputer, on which the contents of the O&MM are loaded, associated with the terminal to give on-line fault interpretation. The concept has been assigned the name *EMACS* (exchange maintenance and control system).

† Trunk Services Operations and Maintenance Department, British Telecom National Networks

* Local Exchange Systems Department, British Telecom Local Communications Services

OVERVIEW OF THE MCS

Resources

The exchange configuration is built up from a set of resources; each resource being identified by a three-letter mnemonic. The resources specify either:

(*a*) a specific piece of exchange hardware (for example, digital line termination (DLT)),
(*b*) a hardware concept that itself consists of more than one item (for example, circuit (CCP)), or
(*c*) a software item not directly related to hardware (for example, digit decode (DDC)).

Entities in categories (*a*) and (*b*) are termed *hardware resources*, those in (*c*) *software resources*.

The data relating to a resource may be distributed across a number of subsystems; for example, a circuit is distributed across the digital switching subsystem (DSS), the call processing subsystem (CPS) and the message transmission subsystem (MTS). In order to provide security, some of the resources are duplicated or triplicated.

Resource States

A resource must always be in a particular state (see Table 1), the state being defined as:

(*a*) *Maintenance State* This defines the availability of the resource to call handling and maintenance actions. These states can be changed by maintenance action or by maintenance software under fault conditions.

(*b*) *Qualifiers* These provide additional information about the maintenance state of resources. Most qualifiers are associated with particular maintenance states. Some qualifiers are changed during normal call handling (for example, BUSY, FREE), others during maintenance procedures (for example, CAMP-ON-BUSY). Some qualifiers are mutually exclusive (for example, BUSY, FREE), others are not (for example, IN USE, FREE).

(*c*) *Parameters* These specify the data associated with a resource.

TABLE 1
Resource States

Maintenance State	Typical Qualifiers	Typical Parameters
Equipped (EQ) Not equipped (NE) In Service (IS) Out of service (OOS) Temporarily out of service (TOOS) Test traffic allowed (TTA)	Free (FREE) Busy (BUSY) Camp on busy (COB) In use (IUSE) Not in use (NUSE) Worker (WKR) Stand-by (SBY) Parked (PRKD) Blocked (BLKD) etc.	Time switch identity Ringer pair number Number of seconds per unit charge Route type etc.

RELATIONSHIP OF THE MCS WITH OTHER SUBSYSTEMS

System X exchanges are divided into a number of subsystems, each providing separate functions. Each subsystem is responsible for the maintenance, update and provision of fault-handling mechanisms for its own resources. The centralised control necessary to co-ordinate such actions is provided by the MCS, which handles actions on resources distributed across a number of subsystems and co-ordinates the necessary changes. Additionally, when overload conditions occur, the MCS sheds low-priority jobs.

SYSTEM RECOVERY

Software and data errors are corrected by rollback, which is part of the operating-system process. This is a procedure whereby data thought to be corrupted is replaced by an original (rollback) version held on a backing store. If the fault persists, the level of rollback is increased; eventually, this may escalate to a system restoration where the complete system software and data are reloaded, and the hardware brought back into service under the control of the MCS in conjunction with rollback.

RESOURCE MANAGEMENT

Resource management controls and co-ordinates the alteration of the maintenance states, qualifiers and parameters of the system resources. Request for such alterations can be from on-line updates (OLUs) input by maintenance staff via the MMI, or from the MCS fault-handling action requiring system reconfiguration.

Validation of On-Line Updates (OLUs)

Resource management requests are validated by the MMI against the password level of the person inputting the request. Access to combinations of resources and their state transitions is controlled by MMI passcard data.The requests for change are input to the MCS and passed to the relevant subsystem(s), where checks are made to ensure that all actions are valid for the current maintenance states of the exchange resources; that is, OOS on a DLT carrying traffic would not be valid. The maintenance states of related resources are considered so as not to endanger system security, and all resource updates dependent on the original request are automatically implemented.

Resource Hierarchy

The exchange resources are dependent on one another, and these dependency relationships must be maintained in a consistent manner. The relationships are defined by means of resource hierarchy diagrams, a typical one being shown in Fig. 1. These show the parent–dependant relationships between resources. In general, parents must be in a higher or equal state than their dependants to ensure correct parent–dependant relationships across subsystem boundries,

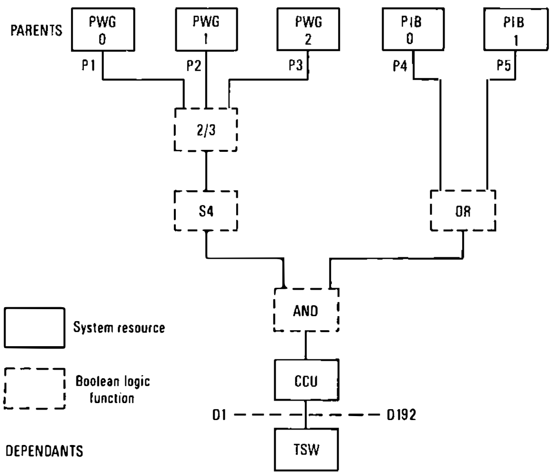

The example shown is a central control unit (CCU), which is a combination of hardware and firmware performing the function of setting up paths across DSS time switches (TSWs). Such paths set-ups are reported from the CCU to the central switch control software via the processor input/output buffers (PIBs). the paths can only be set up if the primary waveform generators (PWGs) are functioning correctly. The CCU has five parents (3 PWGs and 2 PIBs) and up to 192 dependants (TSWs). The logical relationship between parents and dependants can be complex; in addition to AND and OR, there are ten special logic functions such as the 2-out-of-3 logic function and the S4 logic function shown above.

FIG. 1—Typical resource hierarchy diagram

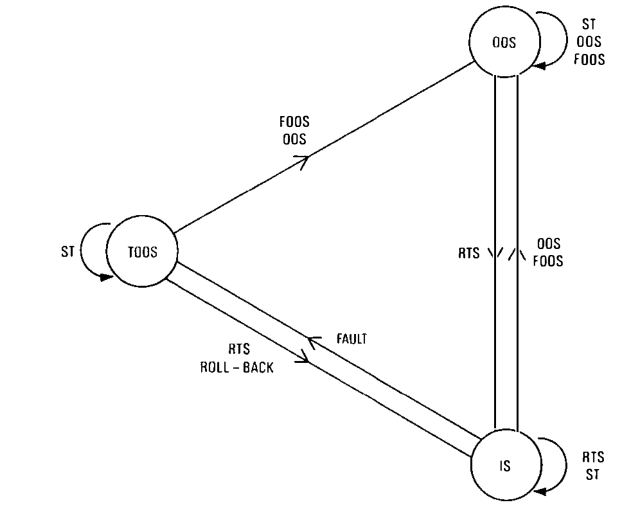

The diagram shows the valid maintenance states for a CCU, that is, TOOS, OOS and IS, together with the effect of valid active and passive actions on those states. Additionally, the CCU may go from IS to TOOS under fault conditions and from TOOS to IS following a rollback.

FIG. 2—Typical resource state transition diagram

although this depends on the exact logic relationship defining the hierarchy. The order of precedence of states is IS, OOS followed by NE. It is essential to maintain a common database within the MCS of the entire system resource hierarchy.

When a change request is processed, the MCS generates the necessary sequences of subsystem requests by examining this common database. Inserting or removing resources alters the data defining the resource hierarchy; OOS, RTS etc. do not alter the hierarchical relationships, they change only maintenance states. The MCS controls all active actions to ensure that all resources are kept in a hierarchically consistent state, if necessary by changing the states of dependent resources by propagating down the hierarchy and checking the states of other parent resources whilst doing so. Each resource has a state-transition diagram defining its valid states; a typical one is shown in Fig. 2. Transitions from one state to another are as a result of OLUs or fault-handling action. OLUs can be *passive* or *active*, see Table 2.

TABLE 2
On-Line Updates

Active Actions	Typical Passive Actions
Insert (IN)	Status (ST)
Remove (RM)	Read (RD)
Return to Service (RTS)	Add (AD)
Allow Test Traffic (ATT)	Delete (DE)
Out of Service (OOS)	Change (CH)
Force Out of Service (FOOS)	List (LI) etc.

MCS transactions conform with CCITT† syntax rules. The commands are in two classes:

(*a*) *Class 1* These are action-resource type used to manipulate, interrogate or amend resources of the form

<action-resource,> <resource number,> <parameters;>

(*b*) *Class 2* These are used to manipulate or obtain system-level data of the form

<action nmenonic:> <parameters>

A typical response to a resource management request is shown in Fig. 3.

† CCITT—International Telegraph and Telephone Consultative Committee

```
    1376001  08  0079  84-06-21  0910

ST-GLK:17;
    1376001  08  0079  84-06-21  0910
    STATUS REPORT

RESOURCE: MTMP GLK    17             ( 1)  STATE: OOS  MAN
    TRANSACTION COMPLETED

    1376001  08  0080  84-06-21

RS-GLK:17&&18;
    1376001  08  0080  84-06-21  0911
    TRANSACTION IN PROGRESS

----1376001  08  0080  84-06-21  0911
    STATUS REPORT

RESOURCE: MTMP GLK    17             ( 1)  STATE:  TOOS
    TRANSACTION COMPLETED

----1376001  08  0080  84-06-21  0911
    STATUS REPORT

RESOURCE: MTMP GLK    18             ( 1)  STATE:  OOS  AUTO
    TRANSACTION NOT PERFORMED, ACTIVE FAULT REPORT 0052
```

The dialogue shows a status request on GLK 17 (MTS link number 17), and a reply that it is OOS due to manual action. This is followed by a request to return to service GLK 17 and 18. GLK 17 is placed in the TOOS state, but GLK 18 remains in the AUTO OOS state due to an outstanding fault on it.

FIG. 3—Typical resource management input and output

FAULT HANDLING

Under fault conditions, the system is required to:

(*a*) detect the fault conditions;
(*b*) analyse the symptoms;
(*c*) take action to minimise the consequences;
(*d*) report the fault; and
(*e*) return the system to normal operation once the fault has been rectified.

Subsystems are responsible for identifying, analysing and reporting faults to the MCS. The MCS itself takes action to minimise the consequences of faults and outputs fault reports via the MMI, and returns affected resources back into service once the fault has been removed. For some resources, the MCS is the owning subsystem and takes responsibility for detection, analysis and fault management. In the case of severe faults, the MCS activates alarms at the local control point (LCP) and the OMU, as shown in Fig. 4. The MCS receives fault indications on shared hardware resources from common services units (CSUs) and forwards these to the owning subsystems. The MCS may then receive fault reports from these owning subsystems to be handled in the normal way.

Fault reports from the subsystems are of the following types:

(*a*) *Equipment Faults* caused by a hardware resource

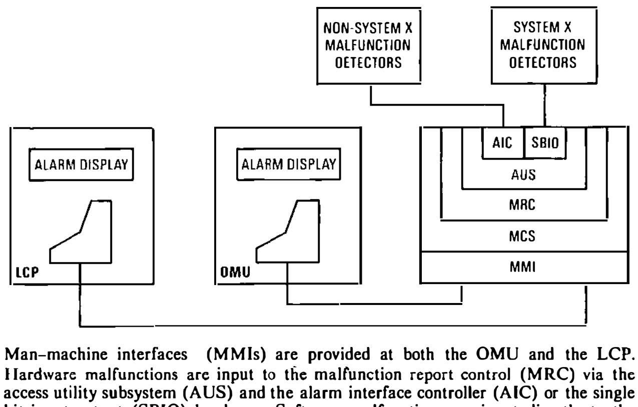

Man-machine interfaces (MMIs) are provided at both the OMU and the LCP. Hardware malfunctions are input to the malfunction report control (MRC) via the access utility subsystem (AUS) and the alarm interface controller (AIC) or the single bit input output (SBIO) hardware. Software malfunctions are input directly to the MRC.

FIG. 4—System X man-machine and alarm interfaces

giving persistent failures;

(*b*) *Circuit Faults* caused by incorrect message protocols during call handling, possibly caused by a hardware fault;

(*c*) *Monitor Resource Faults* caused by the number of in-service resources of a particular type falling below its minimum allowable level;

(*d*) *Non-System X Alarm Faults* caused by faults on non-System X hardware, reported via a CSU;

(*e*) *Access Violation Faults* caused by repeated invalid log-on attempts to an MMI terminal;

(*f*) *Transient Faults* caused by non-persistent faults recurring at a rate insufficient to unduly affect system operation;

(*g*) *Software Faults* caused by errors detected during limit and consistency checking or any other defensive programming checks;

(*h*) *Rollback Reports* output when a process has been rolled back to use its last secured data, caused by escalation of fault reports;

(*i*) *Remote Concentrator Unit Faults* similar to non-System X alarms, except that they originate from a remote concentrator unit; and

(*j*) *Emergency Call Reports* output when an emergency (999) call is made, giving the directory number of the call originator.

Each fault causes an output to be sent to the LCP and OMU terminals and administration back-up system (ADBUS) cartridges (unless deliberately suppressed by the administration). If necessary, fault reports contain supplementary information from software tasks to enable detailed analysis. Typical fault reports are shown in Fig. 5.

Some faults are sufficiently serious that the MCS immediately takes the affected resources out of service. When such a fault is cleared, the MCS automatically returns the resources to service. Faults of this type are:

(*a*) equipment faults,
(*b*) circuit faults,
(*c*) non-System X alarms,
(*d*) access violations, and
(*e*) monitor reports.

Such faults are allocated a serial number and recorded in a secure fault report list as shown in Fig. 6.

```
AAAA1376001       1599  84-06-20  1045
    CIRCUIT FAULT REPORT            84-06-20  1044-58

NUMBER:  0160  PRIORITY: 4  CLASS: 06  SYMPTOM: H'02
RESOURCE: TCMP CCP     3-  196         ( 1)  STATE: OOS  DEP
AUTO-RTS: YES  INHIBITS SET: NONE
    FM CODE: 000   RM CODE:  1    0 24
TIME OF CLEAR: 00-00-00  0000-00


++++1376001       1426  84-06-21  0853
    SOFTWARE FAULT REPORT           84-06-21  0853-28

PROCESS: PUM  (H'011)  MODULE:  80 FAULT:    5
INVALID EXTERNAL MESSAGE    RESTART SENDER:          H'060
H'0060 H'086E H'0006 H'0021
H'0000 H'0003 H'0000 H'0000 H'0000 H'0000


AAAA1376001       1427  84-06-21  0900
    EQUIPMENT FAULT REPORT          84-06-21  0900-09

NUMBER:  0101  PRIORITY: 1  CLASS: 06  SYMPTOM: H'03
RESOURCE: POSM MMD     0-    9         ( 1)  STATE: OOS  AUTO
AUTO-RTS: NO   INHIBITS SET:              DIAG
    FM CODE: 000   RM CODE:  1    0 14
TIME OF CLEAR: 00-00-00  0000-00


++++1376001       1674  84-06-21  1032
    TRANSIENT FAULT REPORT          84-06-21  1031-58

RESOURCE: MTMP GLK    18               ( 1)  STATE: TOOS NUSE
SYMPTOM: H'15
H'E001 H'000C H'8006 H'0040 H'0002 H'8100
```

This shows typical circuit, software, equipment and transient fault reports. the reports indicate the associated resource, together with the associated fault data. The fault data can be interpreted in the O&MM fault directory, or directly by EMACS in a user-friendly manner.

FIG. 5—Typical fault reports

1376001 OS 0207 84-06-21 1547
LIST OF FAULT REPORTS REPORT PAGE: 1

FLT	RESOURCE		DATE	TIME	PTY	CLASS	SYMPTOM
0043	POSM BTA	0-H'FFFF	84-06-20	1638-20	1	06	H'03 ( )
0044	DSOL TSW	2- 0	84-06-20	1641-52	2	06	H'70 ( )
0059	POSM PSA	0-H'0022	84-06-20	1856-47	1	06	H'05 ( )
0060	POSM BTA	0-H'FFFF	84-06-20	1856-48	1	06	H'03 ( )
0061	POSM BTA	0-H'FFFF	84-06-20	1856-48	1	06	H'03 ( )
0062	SISC PCM	38	84-06-20	1948-19	1	06	H'02 ( )
0063	TCPM CCP	2- 259	84-06-20	1951-56	4	06	H'03 (CCT)
0066	TCMP CCP	2- 257	84-06-20	2303-31	4	06	H'03 (CCT)
0067	TCMP CCP	2- 82	84-06-20	2303-56	4	06	H'03 (CCT)
0068	TCMP CCP	2- 258	84-06-20	2303-59	4	06	H'03 (CCT)
0069	TCMP CCP	2- 260	84-06-20	2304-12	4	06	H'03 (CCT)
0108	POSM MDA		84-06-21	0901-03	1	06	H'01 ( )
0109	POSM BTA	0-H'FFFF	84-06-21	0901-22	1	06	H'01 ( )
0110	POSM MMD	0- 13	84-06-21	0905-25	1	06	H'40 ( )
0113	MTMP GLK	19	84-06-21	0918-09	1	06	H'01 ( )
0114	MTMP GLK	17	84-06-21	0918-30	1	06	H'01 ( )
0115	MTMP GLK	18	84-06-21	0919-00	1	06	H'01 ( )
0116	MTMP GLK	20	84-06-21	0919-43	1	06	H'01 ( )

This shows the current active fault reports; the gaps in the fault serial numbers are where faults have been cleared. Each fault report indicated the date and time of its occurence, together with its priority and associated fault data.

FIG. 6—Typical fault report list

Fault Reporting Operation

On receipt of an equipment or circuit fault report, the MCS allocates a record serial number and returns this to the reporting subsystem. If the fault is alarmable, appropriate alarms are raised.

If automatic out of service (AUTO, OOS) is set against a resource, the faulty resource and its dependants are taken out of service by propagating down the resource hierarchy; the resources are then marked AUTO, OOS. Faults which can be cleared automatically are marked by an AUTO, RTS indicator. On completion of the resource management actions, the equipment or circuit fault report is output. In a short-term overload condition the above action may be abandoned. If AUTO, OOS is not set (OOS INHIBIT), then all the MCS does is to record the fault report and report it via the MMI.

Fault Clearance Operation

There are two mechanisms for clearing faults:

(*a*) *Manually* This applies to exchange hardware. It is initiated by manually clearing a fault report subsequent to the fault condition being physically removed. Resource management then starts the OLU action to commence checking that the resource and its parents have acceptable state-qualifiers.

(*b*) *Automatically* This applies to circuit and non-System X faults. These are initiated by the MCS detecting the clearance of fault conditions either via the CPS or by an alarm key being reset.

On completion of the clearance action, the AUTO, OOS qualifier is removed and the fault record cleared.

Interrogation Facilities on Fault Records

Fault reports are stored in the fault record list and enable the MML interrogation facilities to:

(*a*) output a single fault report by serial number,
(*b*) remove particular fault reports from the list,
(*c*) read alarm status, and
(*d*) list fault reports against:
  Status (ACTIVE or CLEARED)
  Resource Type
  Resource Identity
  Circuits
  Non-circuits
  Before or after a specified time

```
1376001  08  0051  84-06-22  0705
FAULT SUMMARY REPORT                              PAGE:  1

FROM 84-06-21  2352-16 TO 84-06-22  0705-14

TRANSIENT FAULTS
COUNT  RESOURCE  SYMPTOM
  26    SISC PCM  H'01
   2    SISC TSA  H'4B
   2    SISC TSA  H'4D

EQUIPMENT FAULTS
COUNT  RESOURCE  SYMPTOM
   8    SISC CCT  H'0F
   1    SISC PCM  H'0B
   1    MCS2 TCC  H'05
   1    SISC TSA  H'4B
   1    SISC TSA  H'4D

SOFTWARE FAULTS
COUNT  PROCESS MODULE  FAULT
   1    H'11     80       5
   3    H'16     0       21
   3    H'16     2        1
   8    H'16     2        2
   1    H'16     2        5
   1    H'16     2        7
   2    H'16     3        1
   2    H'2E     3        1
   1    H'31     2     >254
  39    H'33     3       50
 101    H'33     7        4
   2    H'33     7       20
```

This shows counts of transient and equipment faults with their associated resource and symptoms (coded in hexadecimal) and software fault counts with their associated process, module and fault number.

FIG. 7—Typical fault summary

Additionally, a storage area is allocated in the MCS to give fault summary information on the following:

(*a*) equipment fault reports,
(*b*) transient fault reports,
(*c*) rollback reports,
(*d*) software fault reports, and
(*e*) redundant fault reports.

This allows the fault summary, which gives numbers of occurrences of each type of fault report, to be output. Facilities exist to clear this summary and record only subsequent reports. A typical fault summary is shown in Fig. 7.

Malfunction Reporting and Alarm Control

The malfunction reporting and alarm control activates alarm lamps and bells in response to fault report information. It also links the telephony subsystems with the malfunction detectors; that is, fuse and power supply alarms.

The access utility subsystem (AUS) controls the alarm interface controller (AIC) hardware, which gives alarm outputs (lamps, bells, etc) and receives inputs from non-System X alarm detectors and various keys. System X malfunctions are detected by single bit input output (SBIO) hardware, which is controlled by the AUS. The SBIOs also provide outputs to control equipment; for example, alarm control.

Audio-visual alarms are provided for the following:

(*a*) exchange alarm bell,
(*b*) System X fault lantern,
(*c*) exchange failure lantern,
(*d*) alarm classification lamps,
(*e*) alarm priority lamps, and
(*f*) fault condition indication lamp.

All fault reports are given alarm classifications (1–15) and, for System X resources, priority numbers (1–7). Priorities 1–4 ring the exchange alarm bell, which is extended out as in non-System X exchanges. Non-System X maintenance staff indicate their presence or absence by means of a series of STAFF PRESENT keys. System X maintenance staff indicate their presence by man–machine commands at the terminal. The presence or absence of staff determines whether or not alarms are extended out of the exchange. The MCS hardware is capable of monitoring hardware associated with another exchange and raising alarms when System X staff are present.

After system initialisation, all SBIO outputs are set to the OFF condition; alarm indicators are set to the state of the fault archive; SBIO inputs are read rapidly and the states are updated in the MCS records. After a process restart, all alarm indicators are set in accordance with the state of the fault archive.

The following activities occur periodically:

(*a*) the system watchdog timer is set every two seconds,
(*b*) the audio-visual alarms are set in accordance with the fault archive every four minutes to ensure that the two remain in step, and
(*c*) the SBIO inputs are read every second, and subsequently processed by the malfunction report control (MRC).

Fault Correlation

The MCS can examine fault reports that are received close in time and select from them those likely to represent the root causes of faults and suppress secondary fault reports. This is based on their proximity in time and their relationship within the resource hierarchy. Not all fault reports are correlatable. For equipment fault reports, if the resource is in the TOOS state, the MCS attempts to return it to service, treats the faults as root causes and so limits the output of spurious fault reports.

Resource Monitoring

Certain resourcc types arc marked in the MCS data as being monitored. This enables action to be taken when, for a particular set of resources, the percentage in the IS state falls below a specified level; for example, when the number of circuits in service on a route falls below say 75%, an alarm is given.

Fault Management Options

The following options can be switched on and off by man–machine commands:

(*a*) return-to-service attempts,
(*b*) software fault reports,
(*c*) transient fault reports,
(*d*) redundant fault reports, and
(*e*) suppressed fault reports.

and for particular resource types only:

(*a*) return-to-service attempts,
(*b*) fault correlation, and
(*c*) overwriting of action fault reports.

MAN–MACHINE COMMUNICATION

The fundamental MMI is provided by the operating system. In addition, the MCS man–machine communication function provides for the validation of input commands for syntax parameter values and logic as well as compiling input commands from the MMI to send them on to subsystems as tasks. It also formats the messages from subsystems prior to their being output by the MMI and controls the queueing and destination of messages to be output by the MMI.

Outputs from the MMI are of three basic types.

(*a*) *Solicited* These are in response to transactions from the maintenance staff.
(*b*) *Unsolicited* These are initiated from subsystems as fault reports.
(*c*) *Scheduled* These are in response to requests entered via schedules by the maintenance staff.

All input transactions receive outputs within five seconds; in some cases, this will be a confidence message such as *transaction in progress* prior to the final output.

Macro Interpretation

A facility that enables a single command to generate a list of MCS MMI actions is provided. This enables actions such as the addition of a customer, which requires eight transactions in three subsystems, to be initiated by a single command. The facility includes the ability to read data tables and use the data subsequently to produce input commands. It is particularly useful for frequently used command sequences where errors could occur.

Obey Tapes

A facility exists to input resource-management commands from cartridges under the control of the MCS; the mechanism is a secure one which includes hand shaking to verify correct input of commands.

DIAGNOSTICS

Facilities exist under the control of the MCS for carrying out diagnostic tests on hardware to gain additional information on reported faults or to test newly replaced equipment. This enables faults to be located with a resolution finer than that provided by fault reports. The diagnostics are initiated by manual action.

Two levels of diagnostics exist:

(*a*) *First Level* These are preset test sequences which are run against resources and which locate faults down to slide-in-unit level.

(*b*) *Second Level* These are manually driven with a set of test commands to give a resolution finer than is achievable with first-level diagnostics. Use of second-level diagnostics is normally as a result of faults found while first-level diagnostics are being run.

Diagnostics can only be run on hardware in the MANUAL OOS state. The tests are provided by the subsystems that own the resources on which the tests are run. The control for the tests is provided by the MCS, which determines whether or not the diagnostic tests, subsystem and hardware are free to run the tests. Once diagnostics have commenced, the MCS monitors their progress to ensure their completion within a specified time. If satisfactory progress is not maintained, the tests are aborted. The MCS formats the diagnostic reports from the subsystem for output via the MMI.

ROUTINING

Facilities exist under the control of the MCS to run test schedules on in-service hardware and so detect dormant faults. There are two main applications of this: general routining, and trunk and junction routining (TJR).

Routining is allowed only on hardware in the IS or TTA state, and is run in the following circumstances:

(*a*) after restarts to check out subsystem hardware;
(*b*) periodically, following built-in schedules; and
(*c*) after a manual request from the MMI.

Condition (*a*) is under subsystem control, (*b*) and (*c*) are controlled by the MCS routine control.

Periodic routines are set up in a schedule set by the MMI. The MCS validates manual requests to run routines to ensure that

(*a*) the routine is not already running,
(*b*) the routine has not been inhibited,
(*c*) the current workload on the exchange will allow routines to run, and
(*d*) the resource is in the correct maintenance state.

The MCS monitors the progress of routining to ensure its satisfactory completion or controlled abortion. All manual routines and failed scheduled routines produce outputs that are formatted by the MCS. A number of routines are allowed to run in parallel depending on the exchange workload.

Local Line Automatic Routining (LLAR)

Local line automatic routining (LLAR) enables the automatic testing of a single customer's line by using the test network subsystem (TNS) facilities under the control of the MCS in conjunction with the digital subscribers switching system (DSSS). It is possible to schedule the testing of all local lines on a concentrator.

Trunk and Junction Routining (TJR)

TJR gives the same facilities as those on existing analogue systems; that is, the routiner-sender calls distant answering equipment, detects the *called subscriber answer* (CSA) signal and exchanges tones to verify correct transmission levels.

As well as being able to routine all circuits, it is possible to routine a single circuit, a network band of circuits and a network route of circuits. The MCS controls these tests by passing messages to the system interworking subsystem (SIS). Before running a test, the MCS, which holds circuit data to control the routines, checks the circuit to verify that it is in the IS FREE state. If a circuit is in the OOS unknown state, the MCS attempts to return it to service before proceeding; if the RTS attempt is unsuccessful, it is reported as a failure, otherwise the MCS proceeds as normal. Busy circuits are dealt with in one of three preset ways:

(*a*) *Step Over Busy* (SOB) Circuits are stepped over to be tried later. Each circuit has a busy count record; if a retry fails, then it is reported as a failure.

(*b*) *Camp On Busy* (COB) Circuits are taken through the following sequence OOS, TTA, routine, RTS. The OOS action is taken only after waiting for the circuit to become free; a time-out is set on this wait period. Repeat attempts are allowed before the circuit is reported as a failure.

(*c*) *Speech Detect* A speech-detect mechanism is applied to the circuit; if speech is not detected, it is reported as a failure.

The SOB and COB require the MCS to return repeatedly to circuits found in the BUSY state; to do this the MCS keeps data tables of the states in which circuits were found and the number of repeat attempts.

SYSTEM RECOVERY CO-ORDINATION

The system has the ability to invoke recovery action under various fault conditions. There are a number of levels of recovery action, namely

(*a*) system restoration,
(*b*) system initialisation, and
(*c*) subsystem rollbacks.

System restoration is carried out under the control of the MCS. This involves the MCS sending messages to all the subsystems, which then initiate recovery action. The exact restoration sequence depends upon the exchange architecture; a typical sequence is given below.

(*a*) The MCS instructs the subsystems to initialise their hardware.

(*b*) On completion of (*a*), the MCS has a first attempt at returning resources to service and clearing fault reports.

(*c*) Semi-permanent paths are set up.

(*d*) A second attempt is made to return resources to service.

(*e*) Outgoing circuits are cleared.

(*f*) Subsystems start normal running and confirm this back to the MCS. Routining on the DSS commences.

(*g*) Subsystems are informed that restoration is complete. A successful restoration message is output via the MMI.

During the restoration process, confidence messages are output. The restoration sequence is controlled by a series of time-outs; should restoration not succeed within the overall time-out, then a failure message is output and manual restoration must be invoked. Fault clearance during recovery is carried out in a controlled manner so that hierarchical integrity and data consistency across the subsystems are maintained.

HARDWARE-ONLY SUBSYSTEMS

Some subsystems consist of hardware without a dedicated software subsystem; for example:

(*a*) analogue line terminating subsystem,
(*b*) digital multiplexer,
(*c*) tones and recorded announcements,
(*d*) three-party connection subsystem, and
(*e*) network synchronisation subsystem.

In these cases, the MCS is required to hold the subsystem resource states and validate any changes; invoke the detailed action required in the hardware associated with state changes and maintain the resource hierarchy; and generate fault reports against the subsystem hardware. In addition, the normal subsystem control functions (that is, resource management, fault handling, man–machine communication, diagnostics, and routining) are carried out.

HOLD AND TRACE

System facilities are provided under the control of the MCS to hold or trace a call path through an exchange and give malicious-call identification. The early System X exchanges have exchange-only hold-and-trace; these will subsequently be enhanced to provide network hold-and-trace. The hold or trace request is input to the MCS and, if necessary, queued before passing on to the CPS. The HOLD conditions and any necessary inter-exchange communication are handled by the CPS. The results of trace and hold are sent back to the MCS for output via the MMI. A malicious-call identification is available where a preset called party can initiate the printout of the directory number of the calling party at the OMU.

PRIVATE CIRCUIT CONTROL

A facility is provided within the MCS to enable permanent or part-time call paths to be set up across an exchange. These paths are set up manually, and remain until cleared manually unless fault conditions prevent their continuity. Such circuits can utilise existing public circuits or be set up as new circuits in system data. The control of part-time private circuits is via the MCS scheduler. The commands to set up and clear down private circuits are built into the MCS macros, which send OLU requests to the DSSS, DSS, CPS, SIS, MTS and MCS as necessary.

MULTITHREADING

In some cases, the MCS allows multiple invocations of the same transaction; and these can be in different states between initiation and completion. In such cases, transactions can be completed in a sequence different to that in which they were requested. In other cases where multithreading is not allowed, transactions are queued and processed serially.

In order to control multithreading, the MCS must control the queuing of transactions at modules allowing their processing when modules are free and, if necessary, time-out such actions under busy conditions. The MCS must also control the allocation of free records in the common data areas.

The extent to which multithreading is allowed is restricted by the necessity to prevent undue complexity.

OVERLOAD CONTROL

To cope with overload conditions, the MCS takes the following actions:

(*a*) it suppresses fault reporting and recording when a backlog of fault reports accumulates,
(*b*) it rejects requests received from the MMI when buffering or multithreading limits are exceeded,
(*c*) it prevents the selection of MTS circuits if an overload condition exists at the next exchange, and
(*d*) it inhibits the initiation of routining actions (general routining and TJR) when an overload message is received from operating system overload control.

AUDITING

The MCS has the overall control of the maintenance states of resources. However, because of the complexity of the system, it is possible for inconsistencies to arise in the relationships between resources. This may be caused if, say, a resource goes OOS and its associated fault report is lost due to overload conditions. To overcome such inconsistencies, resource-state auditing functions are carried out within the MCS by means of status checks, and attempts are made to return to service certain resources.

SCHEDULER

A facility is provided to enter commands that are to be actioned at a later time or date, or periodically. The commands are entered via the MMI and stored within the MCS. The stored commands are executed at the specified times exactly as if they had been actioned manually at that time. Commands can be stored in the scheduler for the following MCS functions:

Resource Management
Diagnostic Control
Routining Control
Macros
Trunk and junction routines

There are two types of schedule commands:

(*a*) *Deferred Requests* These are actioned only at a specified time and date, which are deleted from the scheduler's store once they are actioned.
(*b*) *Periodic Requests* These are actioned at a specified time every day or set days of the week.

It is possible to interrogate the stored commands in the scheduler by listing all stored commands or those against a specified resource or a resource type.

The listed commands include both the resource identity and a scheduler reference number. Scheduled commands can be deleted only by reference to both of these identities. The scheduler operates in the single-threaded mode handling only one operation at a time.

REPLICATION STEERING

Some subsystems contain replicated processes, each of which controls some of the subsystem resources enabling a reduction in the data required per process. Replication steering sends resource messages to the correct replicated process within the subsystem. With multi-cluster systems, a subsystems resource can be spread over several clusters, necessitating multi-cluster steering to direct resource messages to the correct cluster. These steering functions require destination information on the cluster and replication process to be built in to some resources. Not all resources are replicated in either processes or clusters.

TEST NETWORK SUBSYSTEM (TNS)

The MCS handles the man–machine input and output for the TNS. It also provides resource management information from other subsystems to enable MCS actions via read or status actions.

The control of TNS threads is part of the MCS MMI function; multithreading of TNS commands is allowed up to a limit dependent on the number of other current actions being handled by the MCS.

PRACTICAL EXPERIENCE WITH THE MCS

Operations and Maintenance

The initial impression of the MML and the MCS is that they are cumbersome and difficult to use. However, most users soon become familiar with a sub-set of the commands and quickly build on these.

The MCS is very powerful, particularly its capability to reconfigure large parts of the system, as this action can be performed remotely.

Within the initial software, there was no provision for automatic manipulation of the resource hierarchy. This meant that a detailed knowledge of the hierarchy was required to perform resource manipulation without adverse effect on the system. The introduction of automation enhanced the potential security of the system by eliminating some of the out-of-sequence commands that could possibly have been attempted.

To perform simply and quickly major (and even minor) changes to the system via an on-line terminal is a significant improvement over existing hard-wired logic common-control systems. It makes possible the realisation of remote operations and maintenance of exchanges.

Maintenance Statistics

In addition to outputting system fault reports both locally (and in some cases, remotely), all such information is recorded on the ADBUS cartridges. These cartridges are sent to a computer centre where they are analysed. The analysis has enabled comprehensive statistics to be compiled, and has highlighted potential problem areas warranting further investigation.

Biographies

Bob Baty joined BICC Power Cables in 1964 as a Student Apprentice and gained an honours degree in Electrical Engineering at University College, London. After graduating, he worked on the design of cable-making machines, computer peripherals, and computer systems before joining BT in 1972. In 1973, he completed a M.Tech.course in Electronic Engineering at Brunel University. Since 1972, he has worked in the Development Department on electronic component standards, and in the System X Development Department on the local exchange proving team. Currently, he is Head of Group in the TS Operations Division concerned with maintenance support for digital trunk exchanges.

Keith Sandum joined BT in 1963 as a Youth-in-Training. After a period on transmission maintenance, he joined the Computer Systems Branch with responsibility for acceptance testing large computer systems and in establishing LACES at London aiport. Later, he was involved in packet switching systems, working at the ARPA network centre in California; in the specification of the MMI and control system for the EPSS; and as the BT member of the project team for EURONET. Latterly, he has been responsible for various aspects of the operation and maintenance of System X, including software maintenance, testing, build control and performance measurement. He is now Head of Section responsible for the maintenance, support and operations policy for System X and other digital switching systems.

The British Telecom Operator Services System

Q. G. COLLIER, B.SC.†

The System X operator services system is a highly-flexible system that is able to support a number of network configurations. In addition, the basic system is fully capable of enhancement to meet the needs of any telephone administration. This article briefly outlines the additional facilities being provided for British Telecom's inland service.

INTRODUCTION

An earlier article in this *Journal** described the salient features of a basic operator services system (OSS) using modern technology which was compatible with System X and capable of enhancement to meet the specific requirements of any telecommunications administration. This article briefly outlines the enhancements needed to meet the requirements of British Telecom (BT).

BACKGROUND

In 1983, it was decided that, without prejudice to ongoing discussions on the modernisation of its operator service, BT would proceed with a development-and-supply contract to perform the necessary adaptive engineering to meet BT's facility requirements for an OSS, and to supply three pilot installations. The work is being undertaken by Plessey Major Systems Ltd. (PMSL), and the pilot installations are expected to be brought into service in the first half of 1986.

FACILITIES

The additional facilities needed to meet BT's requirements can be briefly summarised as follows:

(*a*) the translation of exchange names to all-digit numbers;

(*b*) the validation of credit-card numbers;

(*c*) the translation of Freefone numbers to telephone numbers;

(*d*) the provision of local dialling lists;

(*e*) the derivation of caller information from path-of-entry data for traffic from non-System X exchanges;

(*f*) the provision of a completed-call file to permit billing records to be examined locally; and

† Local Exchange Systems Department, British Telecom Local Communications Services

* PASHLEY, M. System X: The Operator Services System. *Br. Telecommun. Eng.*, Oct. 1983, **2**, p. 216.

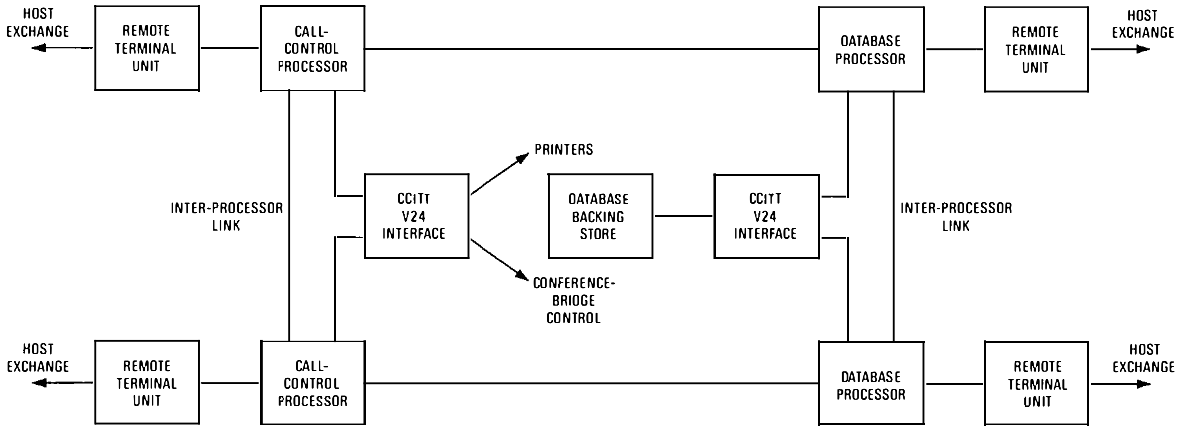

FIG. 1—Revised operator central control arrangements

(*g*) the provision of a suspended-call file for delay, booked fixed-time and reminder calls.

NETWORK ASPECTS

Initially, three networking options were identified for the OSS:

(*a*) a stand-alone unit, with its own dedicated switching, control and interworking equipment;
(*b*) a unit hosted on a System X local exchange, using that exchange for switching and interworking and, if necessary, sharing common routes and interworking equipment with that exchange; or
(*c*) a unit hosted on a System X trunk exchange, digital main switching unit, or international exchange, using that exchange as briefly outlined in (*b*) above.

However, it has been decided that the OSS equipment will be hosted on System X local exchanges for the three pilot installations, because this reduces the implementation cost when compared with the stand-alone unit. On the other hand, it does mean that particular care will have to be taken to ensure continuity of access to operator services when failure occurs in line plant or switching equipment in the network. This will involve the extensive use of automatic alternative routeing (AAR) and similar advanced network facilities made available by modern switching systems.

SYSTEM ARCHITECTURE

The basic system architecture described in an earlier article has not been changed by the additional BT facility requirements; the salient features set out below remain unchanged:

(*a*) the use of a host processor and digital switch;
(*b*) the use of a dedicated processor, communicating with the host processor by CCITT‡ No. 7 signalling, to provide queuing, operator-specific facilities etc.;
(*c*) the use of a three-party conference bridge that remains in circuit for the duration of the call for all calls involving the operator; and
(*d*) the use of 2 Mbit/s digital links to interface the operator call-handling centres to the system for speech and control purposes; this allows operators to be located remote from the host exchange, and potentially eases administration problems such as recruitment.

However, the additional facilities required by BT have resulted in some changes to the system realisation.

The operator central control (OCC) has been functionally split to provide separate call-control and database processors.

‡ CCITT—International Telegraph and Telephone Consultative Committee

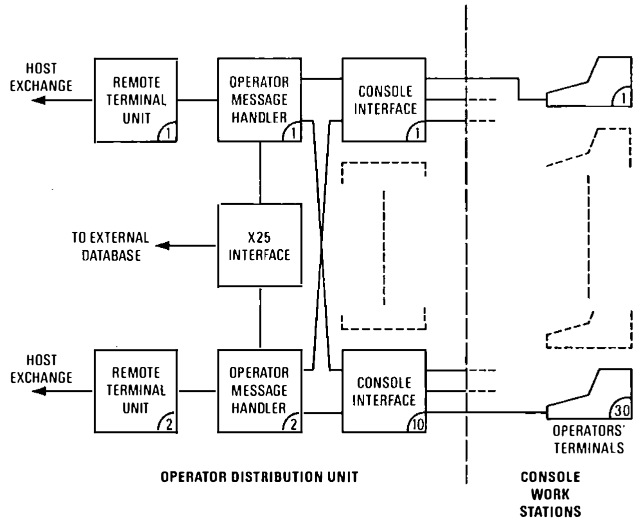

FIG. 2—Revised operator interface structure

This is because many of the BT specific-facility requirements are concerned with the manipulation of data stored internally within the OSS. The revised structure of the OSS is shown in Fig. 1.

The operator interface (OI) function has been enhanced to provide an interface to external databases, such as the directory assistance system (DAS); it uses the CCITT Recommendation X25 protocol. Adoption of the X25 protocol minimises the difficulties of enhancing the system to interwork with other database systems. The revised operator interface structure is shown in Fig. 2.

CONCLUSION

The adaptive engineering of the OSS to meet BT's requirements will give BT a system capable of providing customers with an operator service in line with the more sophisticated facilities being provided elsewhere in the network.

Biography

Quin Collier is a Head of Group in the Digital Systems Evaluation and Planning Division of Local Exchange Systems Department. He joined BT under a University Studentship Scheme in 1969 and graduated from Southampton University with an honours degree in Electronic Engineering. He worked on the development of TXE4 and subsequently System X from 1973 until 1980, when he moved to the System X Launch Department, where his duties involved providing support to Regions on the planning aspects of System X introduction. He has been responsible for the supervision of the BT OSS contract since early 1984.

New Operations and Maintenance Centres for Second Generation System X Exchanges

L. F. STRICKLAND, and M. A. HEWITT†

UDC 621.395.34 : 621.395.743

An earlier article[1] *described the local administration centre, subsequently renamed the* operations and maintenance centre *(OMC). This article explains how OMC philosophy has evolved through experience gained in the field of use, identifying the need for a new version of the OMC. The introduction of an administration network is described, together with the background to developments leading to a basic OMC in advance of the OMC2, and offering the prospects for adminstration users other than exchange maintenance staff to be given access to System X exchange databases.*

INTRODUCTION

Since the earliest days of System X, it has been the intention to have an operations and maintenance centre (OMC); that is, a computer support system for exchange maintenance engineers, plus other users in the administration. In an earlier article[1] this was called a *local administration centre* (LAC). The realisation described in that article was of a computer centre that used an exchange switching processor, the pre-processor utility (PPU), and was engineered in System X telephone equipment practice. That particular OMC has now been designated as the *OMC1* to distinguish it from the other OMCs that follow it.

Thinking on OMCs has changed radically since the original ideas and this has been backed by experience in certain key areas, namely a trial of remote terminal users interworking with the Arrington System X local exchange and known as the *Cambridge User Trial*. In addition, experience has been gained on the OMC1, one of which is currently in service at Cambridge. The division of the British Telecom (BT) business into Local Communication Services (LCS) and National Networks (NN) has had a further impact on the OMCs, since separate maintenance organisations have led to the identification of different requirements to fulfil these separate needs.

It has also been recognised that the OMC1 does not provide a satisfactory OMC solution for LCS use, and this has led to the development of two further types of OMC, known as the *basic OMC* and *OMC2*.

The purpose of this article is to put the OMC into its wider context, to update the reader, particularly on the development of the basic OMC, and to mention the longer term OMC2.

SYSTEM X EXCHANGES AND THE ADMINISTRATION NETWORK

At this stage, it is helpful to introduce the idea of the administration network, since the OMC forms part of it. For this purpose the administration network can be defined as a network that enables interconnection between administration terminal users or support computers and telephone exchanges, to enable transactions to take place as required by the administration. The administration terminal user can be an exchange maintenance engineer or, say, a member of the sales staff.

Fig. 1 shows the administration network as a black box, with exchanges and administration terminals connected through separate interfaces. A System X exchange is shown, since this affords the most comprehensive interworking with the administration network, but other exchanges could be connected. The System X exchanges themselves are interconnected to form the public switched telephone network (PSTN), but for the purposes of this article, it is not important to show this.

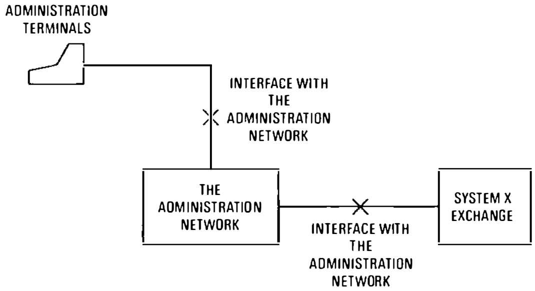

FIG. 1—The administration network

For the System X exchange interface to the administration network, there are two main facets to be considered: the first is the output that the exchange can deliver to the administration network; the second is the interaction between the administration terminal user and the exchange via the administration network.

The exchange output to the administration network can be:

(*a*) *Man-Readable* (ASCII†) *Output* This is capable of being transmitted directly to some device such as a visual display terminal, and understood by a trained person.

(*b*) *Non-Man-Readable Output* This has a binary format, which requires processing on a computer before it is capable of being output to a device, where it can be understood.

(*c*) *Alarm Output* This is output via the maintenance control subsystem (MCS) of the exchange, indicating, for example, an equipment failure, and can be used to light a lamp, to ring a bell, or be formatted into a message for transmission to a distant reception point via the administration network.

The interactive capability between administration terminals and the System X exchange must enable users to:

(*a*) interrogate data held at the exchange,
(*b*) change data held at the exchange, or
(*c*) carry out a maintenance action.

From this very simplified description, it can be seen that the administration network can be realised in a number of different ways. Indeed, there are already many support

† Local Network Strategy Department, British Telecom Local Communications Services

† ASCII—American Standard Code for Information Interchange

computer systems for different applications (for example, the administration of repair service control by computer (ARSCC)) and further computer systems, of which the customer service system (CSS) is a major example, are likewise undergoing development. It is not the purpose of this article to delve too deeply into the administration network, since many items could be discussed such as inter-support computer communication, but rather to concentrate on one part of it; that is, the OMC.

GENERAL REQUIREMENTS OF AN OMC

In determining the requirements of an OMC (that is, a computer centre within the administration network), consideration had to be given to many factors, including:

(*a*) the functions to be performed by the OMC,
(*b*) the level of interface with the exchange,
(*c*) the level of interface between the OMC and the terminal user of the OMC,
(*d*) the data to be held by the OMC,
(*e*) the size of the OMC in terms of the number of exchanges and terminals to be connected to it,
(*f*) the types of exchange with which the OMC should interwork,
(*g*) the administration duties that should have terminals connected to the OMC, and
(*h*) the way the OMC should relate to other centres in the administration network.

It is clear from consideration of all these points that an OMC can be realised in many different ways. Cost is obviously a big factor in deciding the requirements that will actually be implemented.

However, it has become apparent that an OMC is required when the first System X exchange is introduced in an area or district to provide maintenance surveillance from an operations and maintenance unit (OMU), which may be remote from the OMC and include terminals from the OMC. Also required is a user-friendly man–machine interface between terminal users and System X exchanges.

It is also clear that the OMC1 either does not provide the requirements, or imposes severe restrictions on essential facilites required of an OMC for LCS use (although this is not the case for NN).

This had led to the specification of a new OMC, known as the *OMC2*, which is seen as the longer-term OMC solution for LCS. In order to provide an OMC in advance of the availability of the OMC2, it was decided to develop the basic OMC.

There are two versions of the basic OMC:

(*a*) The Mark 1 basic OMC, which interworks with System X Release 1 exchanges, is the first operational basic OMC. The pilot installation has been installed at Lancaster House, Liverpool, to serve Hale local exchange and will provide valuable operational experience.

(*b*) The Mark 2 basic OMC, which will interwork with System X system enhancement programme (SEP) and Release 1 exchanges.

Unlike OMC2s the basic OMC has been designed to cater for a small catchment area of up to five System X exchanges serving 50 000 connections (as well as supporting the required terminal users). It can therefore be used, in the future, in situations that do not justify the provision of the OMC2s.

TERMINAL USERS OF THE OMC

The OMC is primarily a computer support system for the exchange maintenance engineer; but it also does other things. Furthermore, no distinction has so far been made between the various duties of terminal users of the OMC, which makes a great deal of difference as to what the OMC must do; the main difference arises from whether or not the duties involve exchange maintenance.

Terminal Users of the LCS OMC

For exchange maintenance, an OMU is provided within a defined boundary, and acts as an exchange maintenance co-ordination and alarm concentration point for the local network exchanges. There will also be terminal access to the OMC from a terminal located at an exchange local control point (LCP), even if that exchange is a remote concentrator unit (RCU).

Other administration terminals may be required for many duties which have been identified as terminal users of the LCS OMC. However, as other support computer systems in the administration network are further developed, it is likely that terminal users connected to the OMC in the early days will become users of other systems in the longer term. The Cambridge User Trial gives some idea as to which users could be involved.

Cambridge User Trial

This trial, which took place in the Cambridge Telephone Area between October 1983 and April 1984, did not involve an OMC, but was a trial of remote terminal usage. Many users were involved, including:

(*a*) auto manual centre (AMC),
(*b*) repair service control (RSC),
(*c*) customer services (CS),
(*d*) territorial accounts group (TAG),
(*e*) sales,
(*f*) installation control (IC),
(*g*) trunking and grading (T&G), and
(*h*) OMU—In this case, using a remote terminal, connected directly to Arrington exchange.

With the exception of the OMU, the above users had terminal access through small business computers (SBCs) directly to the Arrington Release 1 System X local exchange. These SBCs provided a user-friendly interface to the user, and this feature was later incorporated as part of the administration support function (ASF) in the basic OMC.

Terminal Users of an NN OMC

Within a defined boundary, there will be an NN Trunk Services (TS) OMU, which includes terminals for the OMC.

The TS OMU is the focal point for operations and maintenance activities of all trunk services, and is the point from which the servicing functions of the OMU catchment area are co-ordinated. OMUs are all located within walking distance of the digital main switching units (DMSUs) for which they are responsible. The strategy, in general, is for an OMU to be associated with each DMSU, so that, in many cases, the TS OMU caters for only one DMSU. However, in larger conurbations, where the DMSUs are close together, more than one DMSU could be associated with one TS OMU. There will also be terminal access from a terminal at the DMSU LCP to the OMC.

BACKGROUND TO THE BASIC OMC AND ITS INITIAL DEVELOPMENT

A basic OMC development was initiated within the Local Network Strategy Department (LNSD):

(*a*) to satisfy the essential LCS requirements,
(*b*) to be ready for operational use in time to satisfy the early introduction of the second generation SEP System X exchanges (in advance of the OMC2), and
(*c*) to provide an OMC solution for use in areas (or districts) with only a small number of System X exchanges.

Development of the basic OMC began in earnest in April 1983. The hardware chosen had to be readily available, and use had to be made of software packages that were already written for other purposes and which could be incorporated with minimal changes. This still left a large amount of application software to be written and the engineering of the physical realisation.

The phase 1 development of the basic OMC led to the production of a Mark 1 basic OMC. The aim of this first phase was to design and build a basic OMC, and to connect it to a working local System X exchange. This was to gain early operational experience of the viability of the design and to test some of the essential features and the hardware. The Hale System X local exchange was chosen for this purpose, with the basic OMC hardware being located at Lancaster House, both in the Liverpool Telephone Area.

The BT Factories Department 410 Multi-User Multi-Processor (or 410 MUMP) was chosen as the processor for the system. Factories Department undertook the modifications to the operating system and production of further software which, together with the LNSD application programs, made the 410 MUMP suitable for use as a basic OMC.

The result was the Mark 1 basic OMC which, after exhaustive testing to System X models, went into service at Liverpool in April 1984. Experience gained with the Mark 1 basic OMC since that time, together with studies and experience gained on the Cambridge User Trial, has led to a greater awareness of what is required of the basic OMC to interwork, particularly with the SEP System X exchanges.

Phase 2 of the basic OMC development is geared to the development and production of the Mark 2 basic OMC, with BT Factories now becoming the contractors for the main system hardware.

FUNCTIONAL DESCRIPTION OF THE MARK 2 BASIC OMC

Before the functions of the Mark 2 basic OMC are described, it is necessary to explain the exchange interface.

Exchange Interface to the Basic OMC

As shown in Fig. 2, the SEP System X exchange has, as an interface to the administration network, a maximum of 6 ports to the CCITT V24 standard specification. These are all connected through modems and private circuits to the basic OMC. Any man-readable (ASCII) output required at a remote point, or any man–machine communication between a remote administration terminal and the exchange under these circumstances, is via the basic OMC. There is also a physical alarm interface to enable alarms on the exchange to be passed to the basic OMC.

The remote V24 ports of the exchange are labelled (*a*) to (*f*) and the exchange is programmed to deliver its ASCII output to ports (*a*) and (*b*). In practice, all fault report output is normally sent to port (*a*), and other ASCII output (for example, management statistics) to port (*b*). Under certain exchange fault conditions, however, all of this output may appear on either ports (*a*) or (*b*). Ports (*c*) to (*f*) allow intercommunication between the exchange and any terminal the basic OMC permits to be connected to the exchange.

The functions of the basic OMC are also shown in the block diagram in Fig. 2.

The Terminal Switching Function (TSF)

The terminal switching function (TSF) is the front-end of the basic OMC as far as the terminal user is concerned. It provides for access between a large number of terminal users and a target processor, which may be a processor of the basic OMC, the System X exchange, or even a foreign database; that is, not part of the basic OMC, or the System X exchange.

The TSF is programmable, so that users have access only to target processors for which they have authority; for example, non-exchange maintenance users must not be permitted to access the exchange maintenance function (EMF). The TSF is also required to present an initial display showing the identity of the basic OMC and, via a HELP facility, to show the options open to the user.

The Administration Support Function (ASF)

The administration support function (ASF) provides an intermediate stage between a user connected to the TSF and the exchange. The main purpose of the ASF is to provide a user-friendly interface to the terminal user, the need for which has been described in a previous issue of this *Journal*[2] and demonstrated in the Cambridge User Trial. To do this, the ASF must provide for translation between the System X exchange and the format required by each terminal user duty, of which there may be many.

The ASF also provides a database of customer record data; for example, the type of line, the supplementary services operative etc. This is the database interrogated by an administration user, rather than by accessing the exchange directly, and this reduces terminal activity on the exchange. With the potential existence of so many administration terminals, the ASF database also allows the administration to keep check of what changes have been made on customer's data.

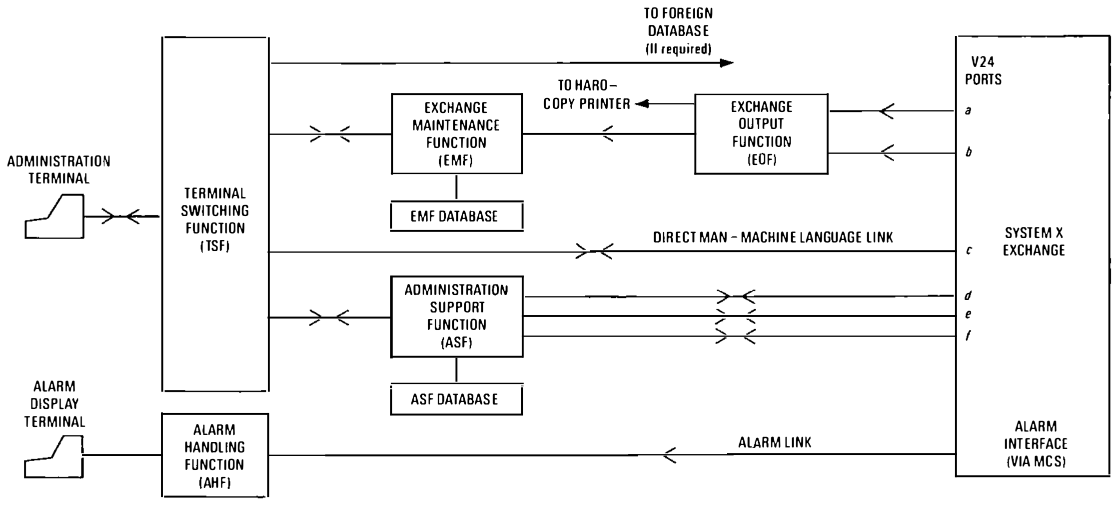

FIG. 2—Functions of the Mark 2 basic OMC

The Exchange Output Function (EOF)

Man-readable output from the exchange is sent to the exchange output function (EOF) of the basic OMC. The EOF recognises the message format of the received output by means of a syntax analyser, and decides whether:

(*a*) it is proper to the EMF (see below), in which case it passes the message to the EMF (a target processor), or

(*b*) if it is proper to a dedicated printer (sited at some remote point), in which case it directs it there; for example, management statistics for the LCS basic OMC are directed to the trunking and grading duty.

Note that a fault message, such as an alarmable event, can be directed both to the EMF and to a dedicated printer in the OMU.

The Exchange Maintenance Function (EMF)

As described above, the EMF receives exchange fault output directed to it through the EOF. The EMF has many files, which are tools for the exchange maintenance user; for example, a current fault file. Processing of fault information into easily understood files is thus automatically implemented by the EMF. The exchange maintenance user (for example, in the OMU) is permitted access to the EMF through the TSF and can manipulate the fault, or other information held in the EMF.

The Alarm Handling Function (AHF)

The alarm handling function (AHF), the last of the functions of the basic OMC, provides a central overlay alarm reception point independent of the basic OMC processor, mainly for the System X exchanges, and also for the basic OMC itself.

CONCLUSION

This article has discussed the background to OMCs and explained how a basic OMC is being developed. A subsequent article will be published in the *Journal* and will concentrate on the Mark 2 basic OMC, including detail on its design and the hardware used. Other articles on OMCs (for example, OMC2) are also anticipated.

ACKNOWLEDGEMENT

The authors acknowledge the assistance given by colleagues in LCS/LNS3 and NN/TSO2 in the preparation of this article.

References

[1] WALIJI, A. A. The Local Administration Centre. *Post Off. Electr. Engs. J.*, Apr. 1980, **73**, p. 36.

[2] JOHNSON, D., and NATHAN, T. D. A Communication Program for On-Line Systems and its use in Man–Machine Language Translation Work. *Br. Telecommun. Eng.*, Oct. 1983, **2**, p. 198.

Biographies

Len Strickland is Head of the OMC implementation group within the Local Network Strategy Department. He joined the then Post Office in 1955 as a Youth-in-Training and worked in the London Telecommunications Region on exchange maintenance. He transferred to the Engineer-in-Chief's Office in 1965 and worked on transit network studies and the definition of national signalling systems. Since 1978, he has been involved on OMC work and the organisation of the Cambridge User Trial. He is currently concerned with OMC planning, testing and implementation and remote terminal user matters.

Mike Hewitt is an Executive Engineer in the Local Network Strategy Department. He joined Telecommunications Headquarters in 1962 from exchange maintenance duties in London North West Area, and was associated with relay design, exchange maintenance and national signalling system facilities. Since 1980, he has been concerned with OMCs, initially determining the facilities for the OMC1 and co-operating with the development teams for pre-service testing. Current responsibilities include agreeing terminal user requirements and managing the installation and commissioning of basic OMC systems.

System X Exchanges—The Ordering Programme

D. C. HOLMES, C.ENG., M.I.E.E.†, and J. D. STOREY, B.SC., A.M.I.E.E.*

UDC 621.395.34

This article outlines British Telecom's plans for building up digital switching capacity in the UK network under its modernisation programme, and discusses some of the provisioning and ordering aspects of these plans.

INTRODUCTION

Since the opening of the first System X local exchange in July 1981, there has been a period of consolidation whereby initial design features within the system have been updated in line with later technology; and production, installation and commissioning facilities capable of handling the volume of equipment required to support British Telecom's (BT's) modernisation programme have been established. Of particular importance has been the need to stabilise system design as quickly as possible in order to allow equipment manufacture to build up to the levels necessary to support this programme.

Under the modernisation programme, BT's objectives are to establish the digital trunk network over the next four years, so that it can handle all trunk traffic no later than 1990, and to provide, by the same date, some 80% of the total local exchange connection capacity from either System X or TXE4/4A exchanges.

The first System X exchanges reflecting the new design aspects were completed and accepted from the contractors in 1984.

PHASING OF DEVELOPMENT AND PRODUCTION

The assessment of when a development design is sufficiently stable to enable a production line to be committed to full-scale manufacture is a critical decision. Clearly, there is a great incentive to bring a new system into service as quickly as possible, and thereby reap the benefit of better service and facilities. However, if a design is manufactured too early, there is a risk of high modification levels being incurred, as a result of further design proving, with possible adverse effects on the service life of the product.

For System X, the problem has been approached by adopting an evolutionary principle and setting discrete design stages. By this means, the development has been phased and designs have been released to the production line against the achievement of each phase. The discrete design stages produce a stable system build, defining both hardware and software levels, against which exchanges can be commissioned and accepted for operational service.

Control of these phases from a production point of view has been exercised through procedures which call for an evaluation of the state of the design before any production commitment is made.

Aiming for the shortest possible time-scales for realisation of designs for field use always carries the risk of abortive work, but the approach adopted has helped to minimise such problems.

THE PROGRAMME

To gain in-service experience, a small number of exchanges based on the earlier design work have been commissioned and accepted for public service use. The range includes trunk, local tandem and local exchanges.

In addition, two trunk exchanges at Coventry and Leeds have now been completed and these are the front runners of a series of units to be installed, using the first main system build level for trunk exchange application. This system build, known as *TE1*, is based on a single-cluster processor and gives a switching capacity of 3500 erlangs. A multi-processor build, known as *TE2*, will become available during 1985 and will give a switching capacity of up to 18 500 erlangs[1]. This later system build will be required at all trunk exchanges in due course under the 60-centre strategy[2]. Plans therefore cater for an in-service enhancement of the TE1 build exchanges to TE2. To meet future facility requirements within the trunk network, the processor store capacity will also be upgraded to 3·5 Mwords (from the TE1 realisation of 1+1 Mwords) as a part of the TE2 build.

The size of the trunk programme in terms of capacity and number of units up to December 1985 is illustrated in Figs. 1 and 2.

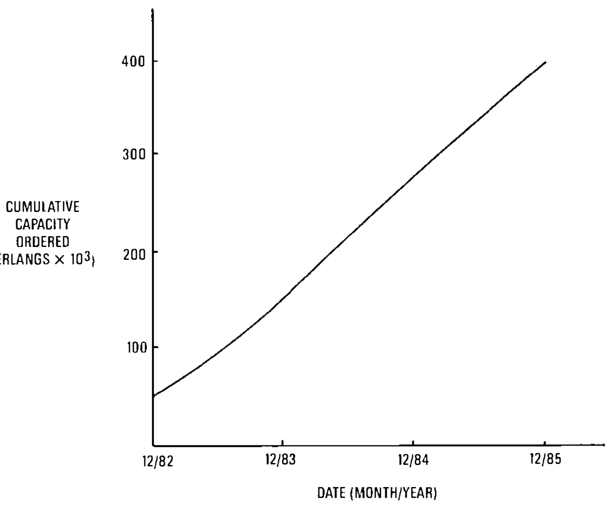

FIG. 1—Trunk capacity ordering position to 31 December 1985

† Trunk Services Planning and Works, British Telecom National Networks

* Local Exchange Systems Department, British Telecom Local Communications Services

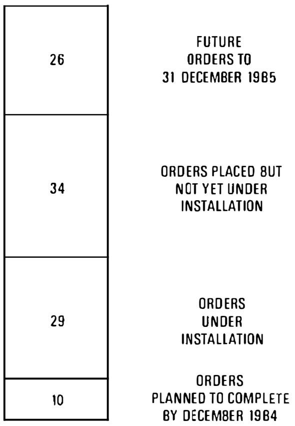

FIG. 2—Progress on trunk exchange orders

The local-exchange programme makes similar use of single- and multi-cluster system builds; that is, a single-cluster processor build having a capacity of 16 000 customer connections, followed by a multi-cluster build taking the connection capacity to 60 000. Associated with this evolution are the provision of a tandem/local exchange capability and significant improvements in the customer switching subsystem area in terms of both design and cost[3].

A significant feature of the provision of local exchanges is the extensive use of remote concentrator units (RCUs) parented on distant processor-controlled local exchanges[4].

Fig. 3 shows the current and future local-exchange programme in terms of exchange connection capacity ordered up to June 1985.

LEAD TIMES FOR EXCHANGE ORDERING

So that the placing of exchange equipment orders can be planned, the required lead times must be established; allowances can then be made to ensure that all ordering and provisioning activities take place in time to achieve the required completion dates. Lead times have traditionally been broken down into 'A' periods—the time allowed to engineer the work and to manufacture sufficient equipment to start and, thereafter, to sustain installation; and 'B' periods—the time allowed to install, commission and demonstrate the exchange, including integration, if appropriate.

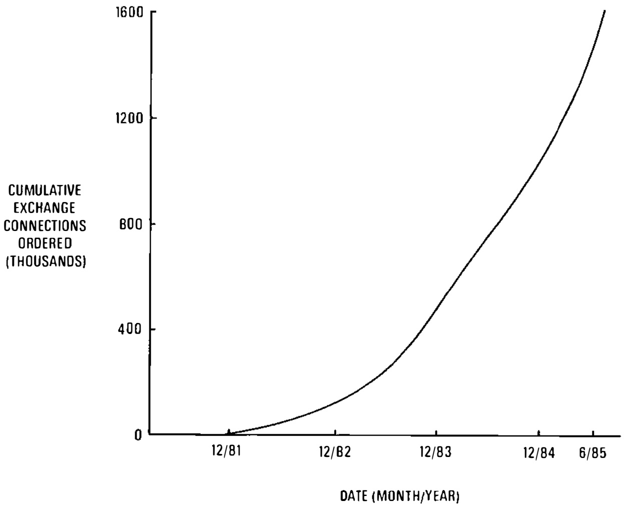

FIG. 3—System X local exchange connections on order to June 1985

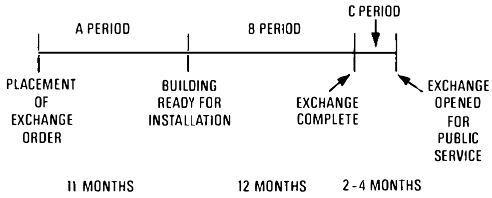

FIG. 4—Average ordering lead times appropriate to the early System X programme

The end of the A period coincides with the availability of the accommodation; by this time, all building and services work required to support the start of the exchange installation should have been completed.

A typical example, equating to a trunk or large local exchange and applicable to the early System X orders placed under non-competitive contracts, is shown in Fig. 4.

The 'C' period covers the time required to test the external network via the digital exchange after acceptance of the unit from the contractor.

There is a need to reduce lead times substantially in order to respond more quickly and accurately to customer and network requirements. The potential for equipment modularity and the installation and commissioning techniques adopted for System X have already yielded significant reductions in time-scales compared with those shown in Fig. 4. Consideration has been paid to speeding up the supply and interchange of documentation required during the engineering phase of a contract, particularly information related to building work, such as the provision of cable holes. This activity requires accurate planning, is of importance to the exchange engineering work and must take account of the practical building constraints.

Orders for trunk units and large local exchanges placed during 1984 have been based on lead times geared to achieve an overall A + B period of 18 months, and it is expected that further reductions to around 15–16 months will be aimed for over the next 2 years.

This improvement in lead time for System X becomes particularly significant when it is compared with lead times for equivalent analogue exchanges, which have required an overall A + B period of some 3 years. As can be seen from Table 1, overall lead times for all types of System X exchanges are considerably shorter than for an equivalent size capacity provided by a modern analogue switching system.

TABLE 1
Comparison of Typical Analogue and Digital Exchange Order Lead Times

Exchange Type	A Period (months)	B Period (months)	C Period (months)
TXE2	13	9	2
System X RCU	9	4	2
TXE4A	15	22	3
System X large local exchange	9	10	3
TXK1 main network switching centre	12	24	3
System X digital main switching unit	9	9	3

DESIGN MODULARITY

Advantage has been taken, where possible, of the opportunity to standardise exchange design and provisioning rates. For example, the provision of trunk exchanges is based on a series of standard packages, covering core equipment such as the processor and switchblock. Choice of package is related to the required traffic capacity of the exchange in question. Some flexibility is allowed in peripheral equipment, such as signalling units, to enable site requirements to be more closely and economically met. The provisioning concept is still, however, to aim for a modular approach. Each package size can be accommodated within a limited choice of suite and rack layouts to allow for local conditions.

To take the principle a stage further, a study was initiated into the potential for introducing a degree of modular design into local exchanges. Because the local exchange is directly connected to the customer, wide variations must be faced in such critical parameters as average calling rate, terminal apparatus type mix and traffic flow patterns, over which the administration has only a limited degree of influence. Studies were initiated into the customers concentrator area of local exchanges, since this represents a high proportion of the total local exchange equipment. It quickly became apparent that the additional equipment costs incurred in applying a limited range of standard packages on a national basis to the wide variations in exchange characteristics experienced in the local network more than cancelled out the manpower savings resulting from simplified planning, engineering and installation procedures. On this basis, standard packages for local exchanges have not been introduced on a national basis. However, the results of this study have been applied by some Districts to produce RCU packages to meet their specific local requirements.

NON-COMPETITIVE ORDERING ARRANGEMENTS

System X trunk exchanges up to March 1983 were ordered on a non-competitive basis with the business shared between the contractors. An exchange specification, listing the equipment to be provided on the order, was produced and maintained by BT, and issued to the contractors as part of the contractual documentation. Similar arrangements were also applied to local exchanges, except that the non-competitive procedures were used for a further year until March 1984.

During the early stages of ordering, the system design was still evolving, and it was necessary to ensure that the exchange specification reflected changes in equipment dimensioning and provisioning rules. Specification amendments resulting from these changes were periodically forwarded to the contractors, to ensure that the specification and system realisation were kept in step.

Each exchange order required contracts to be placed with a lead contractor, called the *main system contractor*, having overall responsibility for the installation of the exchange, plus supporting subsystem contracts with the appropriate contractor. This procedure recognised the split of the development between the contractors, and the shared responsibility for the development and supply of equipment to exchange sites. For example, GEC Telecommunications Ltd. developed the processor as a subsystem and was responsible for processor provision at all the early exchange sites, regardless as to who was the main system contractor.

Payment terms reflected the spirit of the early collaborative development stage, and were initially phased on a monthly basis against detailed contract resource plans, with the later non-competitive business moving to formal phased payments against rate schedule pricing.

These early arrangements helped to achieve the earliest possible introduction of equipment in the field and the build up of adequate production and installation resources to meet the main programme.

COMPETITIVE TENDERING

Major changes were promulgated in the Autumn of 1982 affecting the future development and supply arrangements for System X.

In summary, Standard Telephones and Cables plc (STC) withdrew from System X, and the ongoing programme was shared out between the two remaining manufacturers, GEC and Plessey Major Systems Ltd. (PMSL). PMSL took over design-authority responsibility for System X and became the prime development contractor. Coupled with these changes, BT made clear its intention to introduce competitive procurement procedures as soon as possible, commencing with trunk exchanges as from April 1983.

So that an invitation to tender could be issued, it was necessary to produce a functional specification against which a contractor could tender, and to establish and document the contractual and commercial terms required under competitive arrangements. The emphasis of the specification was placed on defining the functional requirements and left the onus on the contractor to specify the detailed provisioning level against BT-supplied traffic data.

The payment terms agreed with the contractors, although still phased, are much more related to progress as seen at sites, and include a bonus payment for timely completion, or liquidated damages in the event of delay. Overall, the procurement arrangements reflect BT's ongoing commercial policy towards the System X provisioning programme.

Fig. 5 illustrates a typical tender cycle showing the main activities and time-scale leading up to order placement. For trunk exchanges, contractors have been invited to tender against a tranche of orders, each tranche covering a forward 6 month period.

Competitive tendering for local exchanges has been in force since April 1984 and uses similar procedures to those outlined above. However, because of the size of the local programme and the need to maintain flexibility on ordering options, each tender cycle, or tranche, covers 3 months worth of orders in comparison with the trunk-exchange cycle of 6 months worth of orders.

PROJECT CONTROL

A computerised project control scheme has been introduced to progress the work on each exchange contract.

The database contains a set of key activities and interchanges identified against the appropriate time-scales for completion. Planned dates are entered once a firm order requirement has been identified, and achievement is recorded as progress occurs. The scheme incorporates an early-warning procedure whereby activities occurring within a sct time-scale are flagged up for priority attention, giving time for

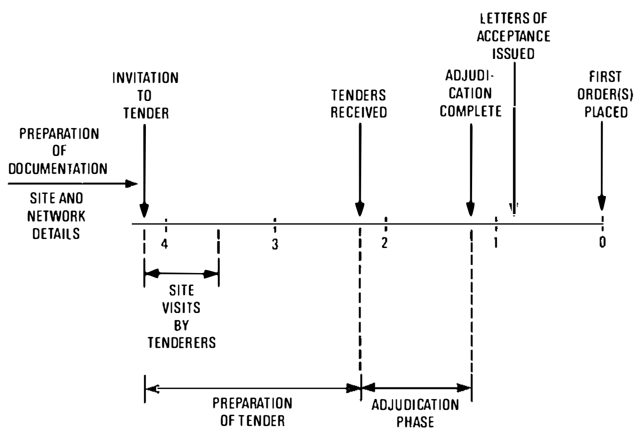

FIG. 5—Typical competitive tender cycle

corrective action to be taken, if appropriate, to avoid potential delay.

The project control procedures operate over the whole period beginning before order placement, once planning has identified the need for switching capacity at a centre, through to contract completion and initial introduction into service.

The control scheme briefly mentioned above is particularly relevant to the trunk-exchange programme, where it has been introduced on a national basis. Similar control arrangements have been introduced within Regions to cover the local-exchange programme.

CONCLUSIONS

This article has referred to the build up of BT's modernisation programme and has outlined the approach taken. It has touched on the factors which have influenced the early planning and ordering phases.

There will be a need in the future to adapt procedures to cater for private-venture variants within the overall design and perhaps to recognise a move towards development and supply contracts for the longer-term future of System X. Both will need careful management to ensure that the integrity of system design and change control procedures are maintained.

References

[1] BRANDER, R. W., and BURVILLE, P. J. Evolution of System X—A Review. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 223.

[2] GARBUTT, B. N. Digital Restructuring of the British Telecom Network. *ibid.*, Jan. 1985, **3**, p. 300.

[3] WARD, R. C. System X: Digital Subscriber Switching Subsystem. *ibid.*, Jan. 1985, **3**, p. 241.

[4] RABINDRAKUMAR, K. The Function and Use of Remote Concentrator Units. *ibid.*, Jan. 1985, **3**, p. 245.

Biographies

David Holmes joined the Post Office Engineering Department in 1954 as a Youth-in-Training in the Brighton Telephone Area. He gained experience on both switching and transmission maintenance before moving to the Engineer-in-Chief's Office, London, in 1964, where his initial involvement concerned the commissioning of video links. This was followed by a period on exchange planning and economics before he joined the System X Launch Department, and subsequently National Networks, where he assumed responsibilities for works progressing on System X exchanges.

John Storey graduated from Southampton University in 1966 with an honours degree in Electronics. After working for the Ministry of Technology, he joined the Post Office in 1970. Since that time, he has worked in Telecommunications Headquarters on the formulation of planning standards for crossbar main network switching centres, and the introduction of System X into both the main and local networks.

Quality Assurance for System X

P. E. GILLAM, B.SC.(ENG.), F.I.Q.A.†

UDC 621.395.34 : 658.5

Quality needs to be 'designed in' at the earliest stages of any project. These and all subsequent activities should be part of an all-embracing quality-management system. This article discusses the quality practices adopted on System X.

INTRODUCTION

Quality, as a term, can mean many things to many people. For example, it is often said that a Rolls Royce is a better 'quality' car than a Mini when what is really meant is that it is a more 'luxurious' car.

It may be useful, therefore, to try to define what is meant by 'quality' before considering it in relation to System X.

WHAT IS QUALITY?

The national standard for terminology used in the quality field[1] describes quality as: 'the totality of features and characteristics of a product or service that bear on its ability to satisfy a given need.'

This has been summarised in a Government publication[2] as follows: 'Quality is the sum of

knowing the customer's needs,
designing to meet them,
faultless construction,
reliable bought-in components and sub-assemblies,
certified performance and safety,
clear instruction manuals,
suitable packaging,
punctual delivery,
efficient back-up service, and
feedback of field experience.'

The outline and basic requirements of a system of management that will enable this to be achieved have been published as a national standard[3]. Additional background material is provided in reference 4.

This article indicates how the quality-management philosophy has been applied to the design, manufacture and installation of System X.

BACKGROUND

Some years ago, much of the work undertaken by British Telecom-Quality Assurance (BT-QA) was concerned with the inspection of products, generally at the manufacturer's works but also in BT support laboratories.

Over the years, this situation has evolved so that now manufacturers are required to operate an effective quality-management system that ensures and demonstrates that their products conform to contracts.

With the introduction of the electronic exchange system TXE4 in the early-1970s, this concept was extended into the installation activities that the contractor undertook. It was further extended into design activities with the introduction of System X.

The objective is that the quality-management concept should embrace all activities, from planning and design through to maintenance and eventual recovery, to ensure that each activity integrates into a coherent whole, thereby minimising the risk of failure/breakdown and optimising cost effectiveness.

DESIGN

The first link in the quality chain of any project is its conceptual phase, since it is pointless designing and producing something that the customer either does not want or will not use because it is unreliable, cumbersome, too expensive etc.

Recognising this, a joint Advisory Group on Systems Definitions (AGSD) established the basic criteria for a new switching system which could provide customers with greater facilities by using the existing telecommunications network. Micro-electronic technology, modular design, stored-program control and integrated digital switching and transmission were just some of the criteria identified as being essential to provide an economic reliable service capable of rapid response to changing customer needs.

An exploratory phase followed, the result of which was a report defining a facilities schedule for the new system.

A development programme involving BT and the major switching-equipment suppliers was instituted. For the first time on a major project there was a formal involvement and commitment by both BT and the suppliers to introduce quality-management disciplines into the design programme. The QA team (suppliers and BT) worked very closely with, and in support of, the development teams, with a formal role in design reviews as part of the design approval process. QA operates somewhat as a catalyst to ensure that all the requirements of the design which are explicitly or implicitly stated are being met. A properly documented design process follows from this. A development quality manual and specific development quality plans detail the quality procedures and indicate the particular roles of development QA and production functions during the development work. This ensures that the evolving designs and documentation are suitable for bulk production using the latest production technology.

Standards for components, equipment practice etc. were established and all documentation was issued in a common controlled numbering series. The whole activity was subject to review by the Technology Working Party. Individual subsystems were the responsibility of nominated Post Office (now BT) Liaison Officers (POLOs).

The result of the design process was a series of specifications for a switching system that would meet the customers' (changing) needs, would be reliable and easily maintainable, would be capable of immediate manufacture by the partici-

† Major Systems Procurement, British Telecom Development and Procurement

pating manufacturers, and would be subject to rigorous documentation control.

Thus was System X created.

SOFTWARE

Software forms a major part of System X and its design phase is critical. Even more so than hardware, any deficiencies detected at a later stage (for example, replication) can be extremely difficult to rectify if the software is structurally deficient. Consequently, the modular approach to software was adopted so that small elements of software could be prepared and checked as an entity. These elements could then be integrated into larger and larger units until the final software package for any situation was realised.

To ensure that quality was built into the software, the purpose of each module, together with its input and output requirements, was carefully defined and allocated to a group comprising highly skilled software designers. Reviews and checks were instituted throughout the process to verify not only that the coding was correct and in accordance with the appropriate rules, but also that the module specification itself was correct. Close co-operation was maintained with groups preparing associated modules.

On completion of a software element, it was submitted to the POLO for assessment and authority to release. It was then put into the a database and subjected to compatability tests by the groups preparing associated elements. Only on satisfactory conclusion of these tests was it put into the software master library, from where it could be called to make up subsystem and system software, which was also tested.

During the course of development, a software quality package evolved[5], and this is now available from BT-QA for use on other software development projects.

HARDWARE

The manufacturers had, of course, been operating quality-management systems to control manufacturing operations long before the introduction of System X. However, quality-management arrangements were required by BT-QA to be enhanced in line with the national standard[3]. In general, altogether new production facilities were involved using highly-automated methods of production and testing to match the new technology of System X and to derive maximum cost saving and quality benefits. Quality-management arrangements were revised by the supplier to match the new requirements. Consequently, the manufacture of System X was undertaken in a disciplined manner, with quality plans and associated documentation being prepared and implemented.

Components form the major part of the hardware cost, and replacement costs are significant if they fail in service. The quality of the components used is important, therefore, if a reliable economic system is to be realised. It is not sufficient merely to rely on component testing during manufacture to obtain the necessary assurance: a total programme is essential. The major elements of this programme include:

(*a*) rigorous component and purchasing specifications,
(*b*) thorough testing of samples of a component supplier's product,
(*c*) assessment of the component supplier's technological capability to sustain continued production,
(*d*) assessment of the component supplier's quality-management system,
(*e*) goods-inwards testing by the System X manufacturer,
(*f*) subsequent testing to a determined test strategy,
(*g*) reliability monitoring and stress testing, and
(*h*) feedback and investigation of failed devices.

As an illustration, it may be useful to consider the typical quality procedures associated with components at one of the manufacturers (the others being similar).

The purchasing group has access to the register of approved sources for components, and so the suppliers for a particular device are known and orders can be placed accordingly.

When the components are received, they are diverted to the goods-inwards department and verified against the associated documentation. It is also confirmed that they have come from a registered supplier. Procedures for handling components are followed; for example, special precautions for static-sensitive devices. All integrated circuits (small-scale integration, large-scale integration, programmable read-only memories, random-access memories, etc.) are tested 100% at an elevated temperature on automatic test equipment (ATE), which is itself subject to verification/calibration procedures.

The results are logged, and conforming components are passed to the stores. (The log takes the form of a history record, one of which is maintained for every supplier and item purchased.) Details regarding the origin of the component are maintained so that, if a problem occurs, the component can be traced to source and components from the same batch identified. Components are handled and stored in accordance with defined procedures to ensure that they are not degraded.

The subsequent assembly processes have been designed to minimise incorrect placement of components, the basic philosophy being that once the design has been established then, if the right components are put in the right place on the printed-wiring boards, the resultant unit will function correctly. The factory personnel are given clear, concise instructions, together with the tools that are required to do the job correctly. Regular audits are carried out to ensure that these instructions are adequate and that they are being followed.

After assembly and soldering, all slide-in-units (SIUs) are electrically tested.

The wired shelf groups follow a similar process, and are then integrated with the appropriate SIUs to form equipped shelf groups. These are 100% factory commissioned and then have the contractual strapping and labelling applied, which is again 100% inspected. Prior to packing, a sample is checked as an audit, and a further audit sample is taken after packing.

INSTALLATION

System X exchanges are installed on a modular and functional-area basis, and quality has been enhanced by reducing the number of variants to a minimum. For each type of System X exchange, from the remote concentrator unit, through small and medium local exchanges, to large local exchanges (including their trunk versions), standard layouts have been agreed. Each exchange is designed and dimensioned for each exchange function; for example, interworking (analogue function equipment, remote concentrator units). This has reduced the number of exchange configurations and cabling types and has standardised cable lengths and connections.

The design of the interconnecting cables, and their length and identification labelling are now performed by computer. The use of computer-aided running-out lists on site has removed many of the on-site cabling and interconnection problems.

A typical sequence of operations at an installation is:

(*a*) receive and handle goods,
(*b*) mark out floor,
(*c*) construction,
(*d*) run and test cables, and
(*e*) equip racks and frames.

The necessary standards, construction methods and quality controls are identified in an installation-methods manual and in an installation-QA manual, produced by the contractor. New methods and procedures have been produced to safeguard product quality; for example, a multi-pack for transporting (and including electrostatic protection for) the 40 SIUs which make up a wired shelf group. Much work was done to ensure that these packs, plus others for racks, were able to withstand the stresses imposed by transportation.

The testing strategy for System X is one of a continuous process from factory to site to an agreed test schedule with the minimum of repeat testing on site. The on-site testing phase can be divided into two elements:

(*a*) commissioning (that is, those tests which the contractor decides to do), and
(*b*) demonstrating the exchange (to BT by the contractor).

This includes the quality-of-service test and demonstrates that the exchange facilities work to the requirements of the works order.

The contractor ensures the quality of the completed exchange by adherence to the agreed procedures and test.

The BT Clerk of Works' activity is to oversee the implementation of the contractor's quality operations and to formally witness the exchange demonstration on behalf of BT.

THE FUTURE

The preceding paragraphs have outlined how formal QA procedures have been introduced to cover all of the suppliers' activities associated with System X; that is, design, component procurement, manufacture and installation.

For the future, particularly in view of the increasingly competitive nature of the telecommunications business, there is a need to consider extending the concept of these QA procedures over all those activities undertaken by BT during the planning, installation, bringing into service and maintenance of System X exchanges.

An activity which is deemed to be of equal importance to the supplier's installation and commissioning procedures is the testing work undertaken by BT to ensure that a System X exchange can be correctly integrated into the public network; that is, pre- and post-transfer testing. Consideration is now being given to the enhancement of BT's QA system to encompass these activities. The associated quality plan defines:

(*a*) objectives,
(*b*) responsibilities,
(*c*) procedures,
(*d*) preservation of product quality (for example, handling) procedures,
(*e*) calibration (for example, test equipment),
(*f*) audit procedures and responsibilities,
(*g*) corrective action, and
(*h*) defect reporting.

The introduction of the above will be a further step towards the implementation of the concept of total QA for all activities, covering those performed by both BT and supplier, related to System X.

ACKNOWLEDGEMENTS

The author wishes to thank Mr. P. W. Blake (BT National Networks) and Mr. R. L. Williams (BT-QA) for their help in the preparation of this article.

References

[1] British Standards Institution. BS 4778: 1979 Glossary of terms used in quality assurance.
[2] The Case for Quality. Produced by the Department of Trade and Industry as part of the National Quality Campaign.
[3] British Standards Institution. BS 5750: 1979 Quality Systems.
[4] GILLAM, P. E. A Guide to Quality Management Systems. British Telecom-Quality Assurance.
[5] STACE, W. E. Project management package for software. British Telecom-Quality Assurance.

Biography

Phil Gillam is Controller of QA Services in BT-QA. In 1967, he joined the then I-Branch Cable Test Section at Arnos Grove after graduating from London University with an honours degree in Electrical Engineering. Subsequently, he worked in the Research Department at Dollis Hill on the 60 MHz cable project and headed a team developing test equipment for QA laboratories. He has worked on the TXE4 Project Review Committee/Task Group and has headed the BT Calibration Laboratories. He is currently involved in the implementation of quality-mangement systems throughout BT, and is involved with both the Department of Trade and Industry and the British Standards Institution on national and international quality standards and practices.

Local Network Strategy—Today's Plans for Tomorrow's Network

K. R. CROOKS, C.ENG., F.I.E.R.E., F.B.I.M.†

UDC 621.395.743

This article outlines the strategy being adopted by British Telecom to modernise its local network with digital plant. It briefly describes the principles and objective that have been adopted in the creation of a network master plan.

INTRODUCTION

The modernisation of British Telecom's (BT's) local network with digital plant is a task of enormous proportions. It is taking place in an environment of accelerating change that includes rapid technological evolution, increasing variety of customer apparatus, competitive commercial pressures relating both to equipment supply and service revenue, as well as a fundamental reorganisation of the Business and its methods of control.

The local network at present comprises over 6000 exchanges interconnected by some 1·5 million junction circuits and serves 20 million exchange lines with nearly 30 million telephones. The modernisation of this mostly analogue network with digital plant is clearly a huge undertaking by any standards, one that requires not only massive investment, but also careful planning to ensure its most effective evolution.

This article outlines the principles and objectives of the modernisation strategy being adopted, describes the mechanism that has been established to enable the building of a network master plan (NMP), and indicates the direction of the evolving plans.

MODERNISATION OBJECTIVES

There are a number of business objectives that provide the foundation for the local exchange modernisation strategy. Briefly, these are:

(*a*) to protect existing markets against competition,
(*b*) to improve the quality and provision of present services,
(*c*) to minimise capital and operating costs,
(*d*) to develop and introduce new services, and
(*e*) to create an advanced and flexible network to meet the needs of an information society.

It is imperative that an effective commercial approach is applied to the deployment of modern plant. This is so that BT should both gain an adequate return on its investment, and also ensure that its services will be preferred to those provided by competitors. Thus, the provision of new or enhanced services should be made available to those customers who will place an appropriate value on them. Consequently, it is necessary for the market sectors to be identified with respect to the classes of new service potentially available. To this end, a network marketing plan is in preparation which, with its support studies, will provide the essential base on which to build the switching and transmission plans.

The limitations imposed upon BT's ability to provide a sophisticated telecommunications network stem largely from the high proportion of analogue plant in service. Fig. 1 shows the present distribution of exchanges by size, and Fig. 2 shows the proportion of exchange capacity that is being planned in each switching system. Strategic target dates have also been set for key plant sectors, to achieve the fastest possible rate of modernisation commensurate with the above objectives. In the 1984 NMP, these have covered various service-related categories of local network plant. For example:

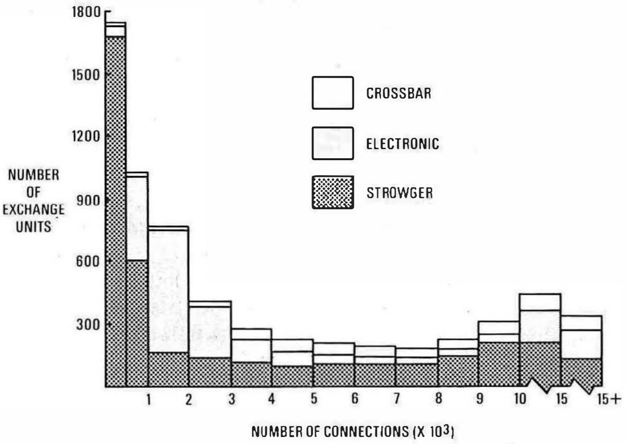

FIG. 1—Distribution of local telephone exchanges by connection capacity at 6/84

(*a*) providing the digital principal local exchanges (DPLEs) at major business centres,
(*b*) completing the replacement of large (over 7000 connections at March 1990) Strowger exchanges,
(*c*) closing analogue group switching centres (GSCs) and director tandem exchanges,
(*d*) providing a digital exchange presence at all sites with major marketing needs,
(*e*) completing the digitalisation of the junction transmission network, and
(*f*) taking account of the latest dates for replacement of non-Strowger (TXK1, TXE2, TXE4/4A) exchanges.

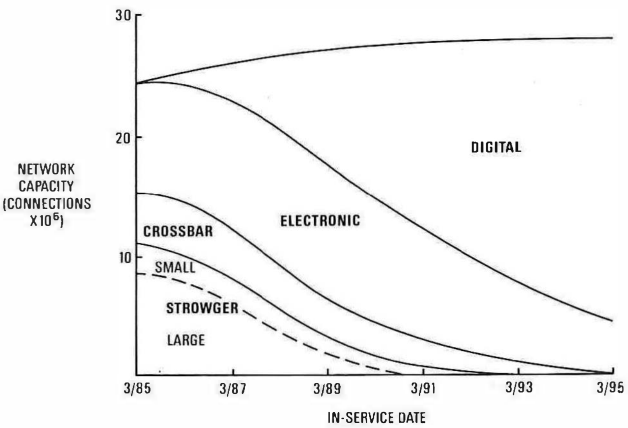

FIG. 2—Planned exchange capacity

† Local Network Strategy Department, British Telecom Local Communications Services

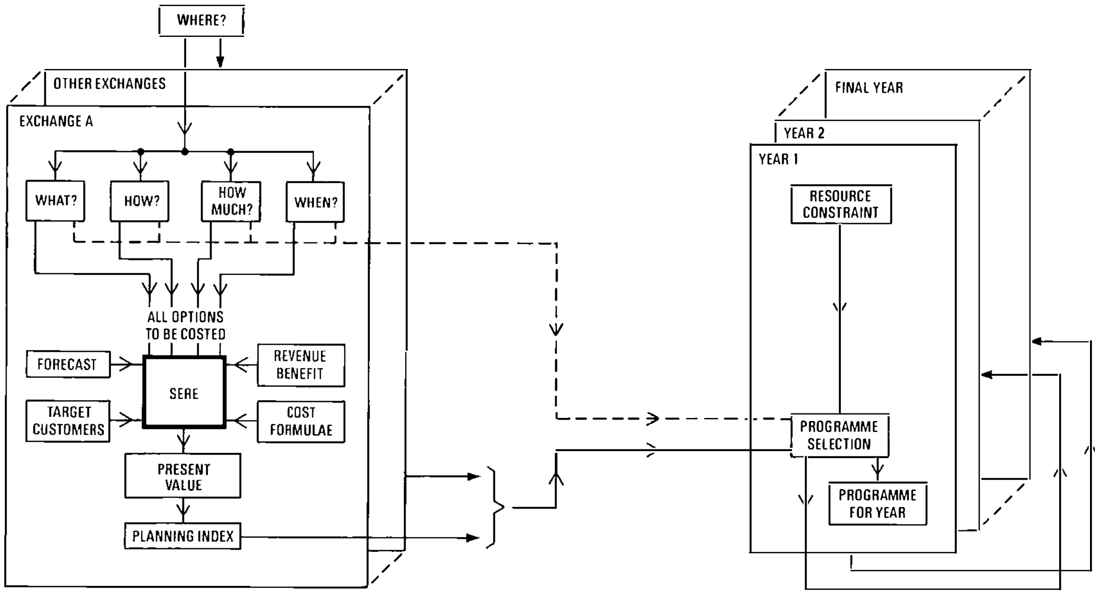

FIG. 3—Role of strategy evaluation and ranking of exchanges (SERE) model in exchange order selection process

The strategy is regularly reviewed, and it is likely that the needs of the customers and of the Business may be better met by a more sharply focussed pattern of resource deployment.

PLANNING OPTIONS

The Network Planning and Strategy Division of BT's Local Communications Services (LCS) Headquarters provides each Region/District with the key modernisation criteria, and such further tools and guidance as are required to enable detailed plans to be built for the main plant items.

The digital-exchange deployment strategy is driven by the needs of the commerical market. Hence, a commercial input, provided by the marketing/commercial divisions, or other local expertise, is required with regard to the need for a digital presence at a given site. It is vitally important that investment is properly controlled at this stage. Therefore, informed judgements are required as to the categories of customer most likely to avail themselves of the new services, the numbers of them involved, the types of service needed and, most importantly, the revenue expected from the investment made in the plant.

It is anticipated that services such as integrated digital access (IDA) will be favoured mostly by the business community, with multi-line IDA employed by large-user customers. Nevertheless, it is expected that exchange-based supplementary services will also be valued by certain sectors of the residential market. The accommodation of these enhanced services requires the rapid development of a digital exchange and transmission infrastructure to provide the bearer network for these and other evolutionary facilities.

It is thus of crucial importance that the need for a digital exchange presence at a given site is based upon the best commercial judgement possible. After the identification of such a need, a number of planning decisions are required. These fall into the following categories:

(*a*) *Member of the digital exchange family to be used* This requires a decision as to whether the processor exchange unit should be of the large, medium or small type, or whether a remote concentrator unit or multiplexer should be used. In addition, it is necessary to determine the switching network structure in which the unit will operate. Namely, whether a processor exchange can be justified at all, or, consersely, whether the facilities offered by a DPLE are required.

(*b*) *Method of introducing exchanges into the network* It is possible to provide a digital unit that will operate alongside an existing analogue exchange and provide full digital capability for a proportion of the customers. This is referred to as *overlay*, but it is more appropriate to describe it as the first stage in the replacement of the analogue exchange. A second option is to replace partially or fully the existing exchange with a digital unit. It is important to consider fully the ramifications of telephone number changes with these options, with particular emphasis on the availability of adequate changed-number information equipment.

(*c*) *Size of unit to be provided* The number of connections to be installed at an exchange is subject to economic analysis. This analysis will, nevertheless, include a consideration of the commercial and service factors. Small units serving the immediate market may be preferable to wholesale exchange replacement. However, where potential revenue is expected to be significant, a larger unit may well be justified on these or on other grounds such as accommodation or growth needs.

(*d*) *Timing of introduction into service* A number of factors impinge upon this crucial aspect. It is necessary to ensure that there is adequate capacity to meet the needs of normal growth, but equally, consideration must be given to timing so that potential custom is not lost to competitive service providers. There are also significant economic factors that have an influence on the timing of an installation. Initial units must be viewed in the light of the overall plan for a given site, so that overall business objectives are met.

The problems involved when considering the above factors in combination are complex, but a recently-enhanced planning tool to assist in their evaluation is available. This planning tool, known as the *strategy evaluation and ranking of exchanges* (SERE) computer model, enables a wide variety of optional data to be considered, and provides a quantified input to the planning decision[1]. The role of the SERE model in the planning process is shown in Fig. 3.

NETWORK MASTER PLAN PROCEDURE

The importance of preparing a master plan for an exchange, charge group and Region/District cannot be overestimated. In essence, this is because of the lead-time periods involved in the provision of accommodation and equipment which, despite optimistic expectations for improvement, remain of significant importance. Thus, there is a need for forward

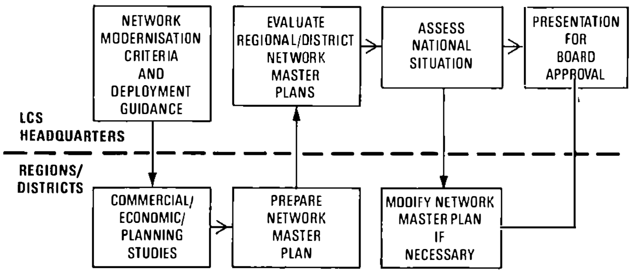

FIG. 4—Network master plan process

planning to be documented so that a structured and practical plan for deploying plant resources can be prepared. The NMP provides the fundamental framework for the notation of resource deployment proposals covering a number of plant and service sectors. As a consequence, it is possible to conduct global analyses of the plan over a wide range of deployment aspects to assist in resource management. The NMP process, as applied to the local network, is shown in Fig. 4.

Fig. 4 outlines the key stages of the procedure which, following the availability of strategic guidance from LCS Headquarters, is conducted primarily by computer interaction. The NMP is an annual exercise rolling forward from the previous plan and embracing the strategic or operational changes required by the developing environment.

Customers are the key to successful operations, hence the requirement to give planning priority to those exchanges with customers who are likely to place sufficient value on the availability of digital services. To this end, the prime objective is to provide digital exchanges where there are customers with large installations, say, with over 10 lines, and who are already aware of the value of high-speed voice and data communication. Other priorities follow, such as the need to examine the customer demand from the remaining classes of customer, and also to ensure that other market competitors are adequately matched. Decisions at this level, along with the many other economic and planning options, are most effectively taken by those nearest to the customer interface; hence, the NMP process requires these issues to be resolved at Regional/District level.

In all major and complex operations of this nature, there is inevitably a measure of risk, the minimisation of which is a normal planning requirement. The deployment of digital capacity needs, therefore, to embrace a measure of flexibility that is not normal to contract installation; namely, a rapid response to changing circumstances. To this end, it is essential that the Business develops the capability to install equipment rapidly wherever it is required. A programme has therefore been prepared to enable BT installation staff to provide digital capacity by using package procurement of the relevant hardware, as and where the need arises, but with short lead times.

At the conclusion of the NMP process, a presentation is made to the Business Directors with a view to the subsequent endorsement of the plan by the Management Boards. The NMP then becomes the approval in principle to enable operational programmes of plant procurement to be developed.

The NMP has evolved over recent years and is now tuned to meet regional and national requirements. It will readily serve the needs of Districts as they become operational. The database which forms the foundation of the NMP can be interrogated to provide a very wide range of plant deployment analyses. Summaries that can be generated include:

(*a*) exchange equipment orders,
(*b*) exchange equipment brought into service,
(*c*) exchange systems by type,
(*d*) exchange capacity by unit and connections,
(*e*) analyses per annum or cumulative,
(*f*) junction circuit channels, digital line systems,
(*g*) junction equipment ordering requirements, and
(*h*) time-division multiplex penetration.

The flexibility of the database is such that the above and other analyses can be provided singly, or in combination, according to need.

EMERGING PROCUREMENT AND DEPLOYMENT STRATEGIES

The evolution of the Business into a public limited company is just one of the many reasons why a review has been necessary of the way in which BT procures its plant. In the exchange-switching field, manufacturers of digital exchange equipment have been invited to supply and install plant by the competitive-tender process, and this has led to a rearrangement of a number of planning and procurement procedures.

There is, too, a growing need for effective visibility of other digital switching systems that have characteristics amenable to BT's emerging competitive position. Aspects such as price, delivery, performance, facilities, maintainability, track record and embodiment within BT's network are key factors in examining the suitability of digital systems from the international market. The deployment of a digital exchange of alternative design and, possibly, facilities etc. could require a fundamental reappraisal of a network's strategy. In order that the feasibility of introducing such an alternative system may be properly considered, studies are being conducted in the Network Planning and Strategy Division in co-operation with Planning Support and Procurement Divisions. Preliminary results suggest that the combined circumstances of technological advance, customer and business needs, competitive position and evolution of the network into one of the most advanced in the world will require BT to embark on a new telephone exchange switching policy that will lead it into the next century[2].

CONCLUSION

The effective planning of a local exchange and transmission network of the magnitude discussed requires strategic policies to be based on informed data regarding the deployment of plant. The NMP provides, through its various sections, the information necessary to give a picture of the present and planned disposition of resources. This is the essential basis upon which major decisions are built regarding the future direction of the Business's largest proportion of capital investment.

References

[1] CROOKS, K. R. Programming the Digital Modernisation of a National Local Exchange Network. Second International Network Planning Symposium—NETWORKS 83, Brighton, Sussex, UK.

[2] BROWN, D. W., and LIDBETTER, E. J. The Future Network. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 318.

Biography

Ken Crooks has held a variety of posts since joining the BPO in 1952 as a Youth-in-Training. On promotion to Assistant Executive Engineer, he was appointed to the BTL training school, where he lectured on a number of engineering subjects and then developed programmed-learning techniques based on behavioural psychology. As an Executive Engineer, he was engaged on designing the London sector switching and also international exchanges. Within Headquarters, he has had Head of Group responsibility for TXK3 system planning and the TXE4 local exchange modernisation strategy. He has been responsible for System X programming policy and latterly local exchange ordering programmes. He is currently Head of the Network Modernisation Planning Section in the Network Strategy Department, Local Communications Services, responsible for the deployment plans for digital exchanges and the operation of the NMP process.

Digital Restructuring of the British Telecom Network

B. N. GARBUTT, B.SC., C.ENG., M.I.E.E., M.B.I.M.†

UDC 621.395.34 : 681.32

This article, which is based directly on a paper presented to the Technical Symposium at TELECOM 83 in Geneva, reports on recent studies that have led to plans for restructuring the British Telecom network concurrently with its rapid conversion from an analogue-switched to an all-digital network. In particular, it reviews the reasons for changing from a 360- to a 60-trunk-node structure, describes the method of evolution from one to the other, and outlines further structural developments, now under active consideration, made possible by the adoption of the restructured network.*

INTRODUCTION

In 1980, the British Telecom (BT) network was basically a 2-wire analogue-switched network employing progressively increasing quantities of digital transmission plant, but with very little stored-program control (SPC) equipment in the trunk network. Extensive studies had shown that an accelerated modernisation programme using digital switching equipment was not only desirable but economically viable. Accordingly, BT adopted a policy of modernising the entire trunk network by 1992, and the local network by the end of the century. This was a challenging and daunting undertaking since it necessitated the replacement of one of the 360 existing analogue trunk units with a digital unit, on average, every two weeks.

With this commitment to evolve from an analogue-switched to a digital network, it was important to devise methods that:

(*a*) minimised the need to provide large quantities of analogue/digital (A/D) interworking equipment, which would have only a short useful life;

(*b*) minimised the need to rearrange repeatedly individual circuits as parts of the network were progressively converted to digital working;

(*c*) minimised the need to update 'intelligence' at remote parts of the network whenever individual nodes and links were changed from analogue to digital plant since, in the absence of SPC equipment in the existing analogue network, this 'intelligence' was hard-wired and, hence, both difficult and costly to alter;

(*d*) maximised the use of the large economic modules of switching and transmission equipment; and

(*e*) reduced the highly complex interactions between achieved modernisation rates in different parts of the network, thereby facilitating control of the modernisation process.

Earlier studies[1] by the United Kingdom Trunk Task Force nearly a decade before had also shown that an all-digital network having a significantly decreased number of trunk switching nodes was potentially more economical than the current network of 360 nodes. These studies had lead to the provision of increasing quantities of digital transmission plant in the UK during the 1970s, but plans to reduce the number of trunk switching nodes had been thwarted by the complexities of implementation.

RECENT STUDIES

Against this background and the changing technological and political environments, both of which demanded increased network flexibility, studies were undertaken to establish the optimum structure for the network as it evolved from the analogue-switched network to the all-digital network.

Trunk Network Nodes

The number of trunk nodes had remained largely unchanged since trunk switching units were first established; their locations had been largely determined by the prevailing economics of transmission and switching systems at that time. Since then, as a result of technological changes such as increased module size and the use of multiplexing techniques, transmission and switching costs have changed significantly, particularly relative to each other. Transmission costs have steadily declined in real terms, whereas switching costs have steadily increased until, with the introduction of digital systems, a step function downwards occurred.

Modelling studies in early-1981 confirmed the work of ten years earlier: namely, that overall network costs would be minimised by having a very much fewer number of very much larger trunk switching nodes. The relationship that was derived between overall network costs and the number of trunk nodes is shown simply in Fig. 1.

A specific minimum network cost occurred at approximately 55 nodes, but the following more general conclusions were drawn from the work:

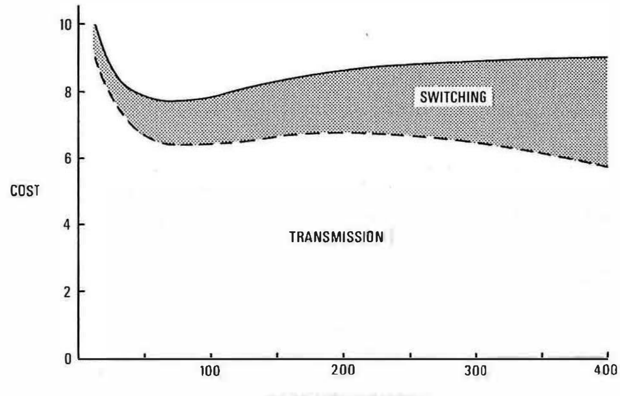

FIG. 1—Network costs

† National Networks Strategy Unit, British Telecom National Networks

* B. N. GARBUTT. The Digital Restructuring of the British Telecommunications Network. Fourth World Telecommunications Forum Part II, Geneva, 1983.

(*a*) cost penalties for departing from this minimum number of nodes rose more rapidly for decreasing the number than they did for increasing it;

(*b*) in the area of minimum network cost, transmission costs accounted for approximately five-sixths of the total network costs;

(*c*) the results were affected more by the efficiency of utilisation of the transmission systems rather than by the length of links or by the quantity of switching equipment required; and

(*d*) over a wide range of sensitivity analysis, the minimum network cost always corresponded to a number of nodes in the range 50–80 and was predominantly in the 50–60 range.

In view of the above, it was decided that the objective digital network should contain nominally 60 trunk switching units.

Interconnection of Trunk Nodes

With 60 trunk nodes, the maximum number of interconnecting bothway traffic routes would be 1770 when all were fully interconnected. Analysis of likely traffic flows between these 60 large centres (average size 9000 erlangs at 1992) indicated that all but a very few of these routes could be justified if conventional justification principles were used. It was therefore decided that the marginal cost of providing all these 1770 routes, including the few small ones, would be more then offset by the major benefits in the simplification of network planning and plant provision that would result from a fully-interconnected network. Accordingly, the network is now being planned on the basis of full interconnection. Such an arrangement, as well as giving the benefits of simplicity in planning and implementation, establishes a framework of routes which will provide extensive scope for the progressive introduction of dynamic routeing techniques such as automatic alternative routeing (AAR), and for the use of novel techniques for providing network resilience (see the later section on future developments).

NETWORK EVOLUTION

As mentioned previously, earlier suggestions for reducing the number of trunk nodes had always foundered on the problems of evolving from the large number of analogue nodes to the smaller number of digital ones in the absence of existing SPC equipment. During mid-1981, a suitable method of evolving from the tiered analogue structure with 360 nodes to the fully-interconnected digital structure with 60 nodes was devised. In addition to meeting the objective of achieving an economically-optimised network, the method meets the key requirements outlined in the second paragraph of the introduction.

Outline of Evolution

The basic principle of the method is one of overlay but with severe constraints upon the points at which transfer between the new evolving digital network and the existing analogue network is permitted to occur. The two networks are to be kept separate with digital local exchanges (DLEs) (or analogue exchanges served by digital transmission) connected to the digital trunk units and with analogue local exchanges (ALEs), which are still served by analogue transmission, remaining connected to analogue trunk units. Transfer to the other network is only allowed for the final junction link in a call once the destination trunk unit in the originating network identifies that the called exchange is within the other network.

This has the effect of:

(*a*) concentrating (and minimising) the A/D interworking equipment at existing analogue trunk unit sites;

(*b*) eliminating the need for remote parts of the network to know to which network any particular exchange is connected at any particular time;

(*c*) allowing the use of the existing analogue network to run down, basically unchanged, as traffic is transferred onto the growing digital network; and

(*d*) ensuring that end-to-end digital connections are achieved between any two digital exchanges.

Basic Trunk Network

The structure of the trunk network and its evolution from the analogue network are best described by a series of diagrams with commentary as follows.

Fig. 2 shows three representative ALEs currently parented on an analogue trunk unit, which in turn is connected hierarchically to other such units. Lines and arrows between exchanges represent traffic flows. For simplicity of illustration non-hierarchical routes are ignored.

For the new structure, a number of analogue-trunk-unit catchment areas (six on average) are allocated to each of the proposed digital trunk units. With the conversion of a local exchange from analogue to digital working, digital transmission is provided from the DLE to a flexibility point, known as a ***digital distribution frame*** (DDF), at the site of the analogue trunk unit and from there to the (usually, remote) digital trunk unit, shown as route **1** in Fig. 3. Digital trunk units are of course interconnected digitally (route **2** in Fig. 3).

All traffic originated by the DLE remains in the digital network until it reaches its destination trunk unit. Terminating routes (shown as route **3** in Fig. 3) from there to each of its allocated analogue trunk units enable calls to analogue destinations to be completed. A/D interworking equipment is located at the site of the analogue trunk unit between this unit and the DDF.

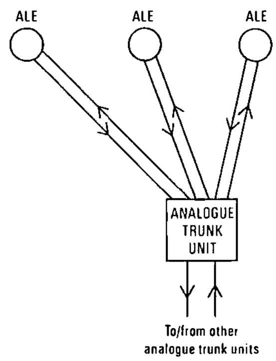

FIG. 2—Analogue network

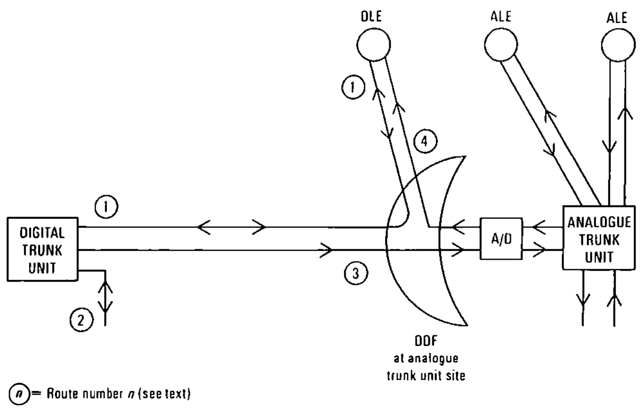

FIG. 3—Overlay digital network

Traffic originated by an analogue exchange remains in the analogue network until it reaches the destination analogue trunk unit. Terminating routes (shown as route **4** in Fig. 3) from there to all DLEs in the catchment area of the analogue trunk unit enable traffic to be routed to digital destinations. Traffic to analogue destinations remains throughout in the analogue network.

Fig. 3 therefore shows the basic overlay digital network which is necessary to allow any call to be routed through the network(s). In summary, for interworking between the analogue and digital networks, the network requires

(*a*) a route from a digital trunk unit to each of the analogue trunk units allocated to it (route **3**), and

(*b*) a route from an analogue trunk unit to each of the DLEs in its catchment area (route **4**).

Local Tandem Network

In principle, the network shown in Fig. 3 can handle both trunk and local traffic, but this can involve a considerable amount of tromboning of local traffic from the DDF to the digital trunk unit and back. In most cases, however, a DLE is being provided at, or very near to, the analogue trunk unit to serve the customers in that locality. This exchange could be connected into the network in the same way as any other DLE; that is, with a bothway route to the digital trunk unit (route **1**) and a terminating route (route **4**) from its analogue trunk unit.

However, it is usually preferable to connect this particular DLE into the network in such a way that it can carry out the two extra functions of tandem switching local traffic between digital exchanges and concentrating traffic that requires to pass to and from the A/D interworking equipment. This local exchange, known as a *digital principal local exchange* (DPLE), has connections as shown in Fig. 4.

In summary, it can be seen that

(*a*) the route **1** from other DLEs is split at the DDF with some 2 Mbit/s modules connected through to the digital trunk unit as before, but with the remainder diverted to the DPLE;

(*b*) instead of routes from the analogue trunk unit to each and every DLE in its area (routes **4** in Fig. 3), a single route (route **5** in Fig. 4) is provided via the interworking equipment to the DPLE; and

(*c*) instead of routes from the digital trunk unit to each analogue trunk unit in its area, a route from the digital trunk unit to the DPLE (route **6** in Fig. 4), together with a short route (usually within the building (route **7** in Fig. 4)) from the DPLE to the analogue trunk unit, is provided.

In this way, not only is the interworking equipment confined to the site of the analogue trunk unit, but its use is confined to short routes between nearby units. Access to, and egress from, the digital trunk network is wholly digital with no A/D interworking routes.

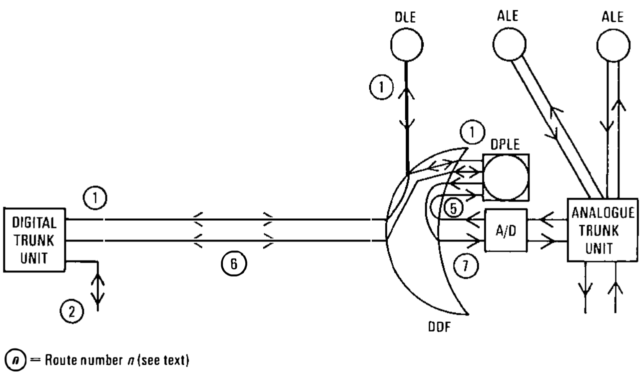

FIG. 4—Initial use of digital principal local exchange

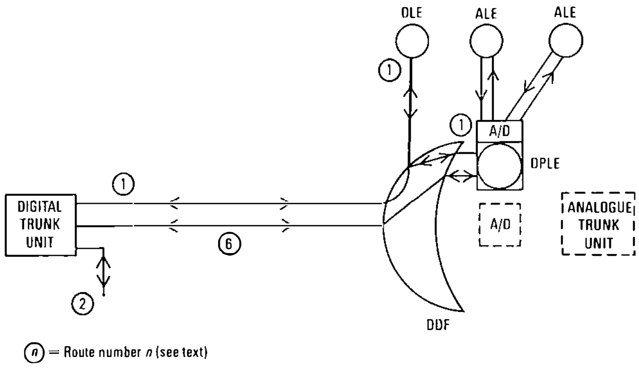

FIG. 5—Later use of digital principal local exchange

Ultimately, the DPLE undertakes the further functions of terminating analogue lines from ALEs, providing them with charge-band determination and metering-over-junctions, concentrating their trunk traffic and distributing terminating traffic. This allows the analogue trunk unit to be closed before all the analogue transmission plant has been removed from the network. The end result is shown in Fig. 5. Further information about the role of the DPLE is given in an accompanying article in this issue of the *Journal*[2].

The concept of the evolutionary structure has been reduced to its barest essentials for the purposes of the preceding illustrations and commentary, and it is important to elaborate on certain key aspects which have been omitted in the interests of simplicity; in particular, detail of the DDF and the need for network security and resilience.

Digital Distribution Frames

Reference has been made to the DDF only at the site of the analogue trunk units. The structure envisages them also at a whole range of locations including DLEs, digital trunk units and branch points in the transmission network. Furthermore, the discussion has implied operations at the 2 Mbit/s level. Ideally, the principle will also be applied at each hierarchical level, with multiplexing provided between the various levels as necessary. The DDF plays a vital role within this network structure because it allows the network to be reconfigured at and above the 2 Mbit/s level, either to avoid switching-unit exhaustions, or to load up new, or spare, routes or switches. It ensures that, in the process of modernising the network from analogue to digital working, circuits have to be rearranged only once at the circuit level, with all subsequent network rearrangements taking place at the 2 Mbit/s level or above.

Network Security/Resilience

No reference has been made in the foregoing description of the evolutionary structure to the need to maintain an adequate degree of network security and resilience. A fundamental review of these aspects is currently in progress and is considering two separate, but related, topics: firstly, the level of network security/resilience that should be provided and, secondly, the optimum method or methods of providing this. The former aspect is fundamentally a commercial question for BT to answer, concerned as it is with determining what the customer is prepared to pay for a given level of service—in this particular instance at times of traffic overload or plant failure.

In determining the optimum method(s) of providing for security/resilience, the mutual interactions of the large range of individual techniques need to be, and are being, considered. For example, the use of dynamic routeing techniques to route round a problem can make the connecting in of stand-by plant in the problem area superfluous.

While these studies are under way, various techniques are being applied to enhance the security/resilience of the network on an interim basis; in most cases, they constitute a continuation of well-established practices from the analogue network. These can be summarised as follows:

(*a*) Final-choice routes are generally dimensioned to cater for a 10% or 20% traffic overload.

(*b*) A service protection network of stand-by 140 Mbit/s digital line systems is provided. This will be of use for most of the routes between trunk units and from DPLEs to their trunk unit.

(*c*) The fully-provided inter-trunk-unit routes employ diverse routeing; this means that half (or, at worst, one third) of each route is provided over a different physical path to its destination.

(*d*) The access link from the DDF at the analogue trunk unit site is reduced in size, and access is provided direct from the DDF to an alternative digital trunk unit. This alternative security link is capable of carrying a minimum of one third of the traffic to/from the catchment area and offers protection against failure of either the trunk unit or the normal access link.

(*e*) Between local exchanges and this DDF, diverse routeing is applied to all traffic. Additionally, as indicated in Fig. 4, the route is split at the DDF with some 2 Mbit/s modules being connected through to the digital trunk unit with the remainder to the DPLE. All traffic, though normally passing to the trunk or local unit as appropriate, has access to either switchblock in the event of traffic surge or link/node failure effects.

FUTURE DEVELOPMENTS

Active consideration is now being given to the possible automation of the DDF aimed at identifying the viability and the hierarchical levels at which it should apply. In effect, such a development would constitute a wideband switch but operating in a structural mode rather than under individual call control. Paths through the network could be set up or broken down at short notice to ensure that segments of transmission plant were not dedicated to particular subsets of traffic at a time when there was little of that particular traffic available. For example, the situation could be envisaged where, when traffic from a given local exchange to a remote part of the country had built up to a sufficient level (or was forecast to do so), that exchange processor could arrange with the automated DDF control for the traffic to bypass its normal trunk unit and, instead, be connected through the DDF at the appropriate transmission rate (2 Mbit/s, 8 Mbit/s etc.). Such an arrangement would ensure that plant was dedicated only to specific traffic when it could be effectively and efficiently used and, at other times, reverted to the normal route to the trunk unit where it would be accessible to all traffic.

A number of significant benefits would accrue from the adoption of DDF automation:

(*a*) All traffic would have access to all plant at short notice. Consequently, all longer-term plant planning could be based on bulk forecasts of total demand for telecommunications services—an inherently less-volatile demand than for each individual service.

(*b*) Novel forms of network security would become possible. At the time of, say, node failure, the lack of automated DDFs in the network means that line systems terminated on the failed unit are not usable and traffic has to be carried on alternative plant; this requires the provision of redundant plant. With automated DDFs, the line systems connected to the failed unit could be switched over to an alternative unit, or some could be used to extend others to remote switching capacity which happens to be idle at that particular time. This would offer the major advantage of being able to reduce the level of provision of normally redundant plant without having to reduce the level of service.

(*c*) Some of the control arrangements that would be necessary for the above could potentially be made accessible to the customer. This would effectively provide for a switched wideband service fully integrated with narrow-band services.

PROGRESS TO DATE

Having embarked on the program to produce the more structured but, at the same time, more flexible network outlined above, BT has further accelerated its proposals and now plans:

(*a*) by 1986/87, to have digital switching capability at the 60 trunk unit sites with full interconnection between units by digital transmission systems;

(*b*) by 1988/89, to have adequate capacity to permit the off-loading of all trunk traffic from the analogue network;

(*c*) by 1988/89, to have provided most of the DPLEs to undertake tandem switching of local traffic and to concentrate trunk traffic from residual analogue exchanges still served by analogue transmission systems at 1990; and

(*d*) by 1990, to have closed down the analogue trunk network.

CONCLUSION

This article has indicated the main reasoning behind BT's plans to restructure its network concurrent with its change from a 2-wire switched non-SPC network to an all-digital network, and has outlined the method of evolution from one to the other. Major benefits that could accrue in the future from the adoption of this restructured network have been suggested.

References

[1] BREARY, D. A Long-Term Study of the United Kingdom Trunk Network. Parts 1 and 2. *Post Off. Electr. Engs. J.*, Jan. 1974, **66**, p. 210, and Apr. 1974, **67**, p. 37.

[2] ROLFE, M. A. Role of the Digital Principal Local Exchange. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 304.

ACKNOWLEDGEMENT

The author is indebted to his many colleagues throughout BT who have, both directly and indirectly, contributed to this work and, in particular, to the Head of the National Networks Strategy Unit and the International Telecommunication Union (in whom the copyright of the original TELECOM 83 paper is vested) for permission to publish this paper.

Biography

Nigel Garbutt joined BT (part of the British Post Office at that time) after graduating in 1970. He worked on audio cable and transmission equipment installation practices before being sponsored on a Diploma of Management Studies course in 1973/4. He returned to work on junction planning practices for pulse-code modulation (PCM) transmission systems and, in 1978, was instrumental in establishing the policy for the advance provision of PCM in the junction network in readiness for the arrival of digital exchanges. He developed the strategic model of the junction network, which assisted in the formulation of the 1980 modernisation plans for the whole of the BT network. Since that time, he has been concerned with the network structure studies and resultant policy decisions reported in this article, and now leads the Network Structures, Modelling and Economics Section in the National Networks Strategy Unit.

Role of the Digital Principal Local Exchange

M. A. ROLFE, B.SC.†

UDC 621.395.722 : 621.374

This article describes the role of the digital principal local exchange in the evolution of British Telecom's digital switched network.

INTRODUCTION

The digital principal local exchange (DPLE) is a medium or large local exchange (MLE/LLE) designated to provide additional network functions within a defined catchment area. The functions essential to the status 'principal exchange' are:

(*a*) switching tandem traffic between exchanges within its own catchment area;
(*b*) terminating analogue line plant, and performing call charging for analogue exchanges; and
(*c*) concentrating and forwarding main network traffic from within its catchment area to a digital main switching unit (DMSU).

The capability to perform these functions is an intrinsic part of the MLE/LLE design; the decision to limit their use primarily to the DPLE and to use the DPLE site as the main location for interworking between analogue and digital networks is dependent upon network economics.

LOCATION

The 1984 network master plan (NMP) indicates that DPLEs will be established at the majority of existing GSC sites, and will, in many cases, assume the catchment area of the GSC. In suitable circumstances, the opportunity is being taken to combine GSC areas to form one larger DPLE area. The installation of a DPLE at an existing GSC site places the unit at or close to a focal point of the existing transmission network and thus gains benefits in network costs and flexibility.

TIME-SCALE OF PROVISION

The installation plan for DPLEs is linked to the digitalisation of the main transmission network, which is being planned for completion by 1988/89; consequently, DPLEs are planned to be in service at all the major business centres by this time. It is expected that, by 1990, all other DPLEs will be in service and the transfer of traffic from the analogue switched main network (ASMN) will be complete.

ROLE DURING INITIAL AND EVOLVING STAGES OF NETWORK MODERNISATION

While the digital main transmission network and DMSUs are being brought into service, the DPLE has a fundamental role as an interface between the analogue and digital networks. During this period, traffic will be off-loaded from the GSCs on to the digital network. Some analogue local exchanges (ALEs) within the GSC catchment area may be replaced by digital units; other ALEs, as they are provided with time-division multiplex (TDM) line plant, will gradually be reparented on to the DPLE and/or the DMSU. Eventually, those ALEs with analogue line plant will also be transferred to the DPLE.

The DPLE will therefore be required to (see Fig. 1):

(*a*) switch tandem traffic between the digital local exchanges (DLEs) and ALEs served by TDM transmission in its catchment area, and provide charging information for ALEs as necessary;
(*b*) concentrate traffic from the exchanges in (*a*), destined for ALEs still parented on the GSC, via the digital distribution frame (DDF) and interworking equipment;
(*c*) distribute traffic sent to the DPLE by the GSC, on interworking links, either from exchanges still parented on the GSC, or incoming from the ASMN;
(*d*) concentrate traffic from ALEs served with TDM line plant parented only on the DPLE, for onward routeing over the main network as appropriate, and distribute traffic to these exchanges incoming via the DMSU;
(*e*) distribute digitally originated traffic incoming from the digital switched main network (DSMN), destined for ALEs parented on the GSC, via the interworking link to the GSC, in cases where no DMSU-to-GSC link is available;
(*f*) forward any traffic overflowing from direct ALE/DLE-to-DMSU routes, as appropriate; and
(*g*) have the ability to tandem switch local traffic for more than one linked numbering scheme or charge group.

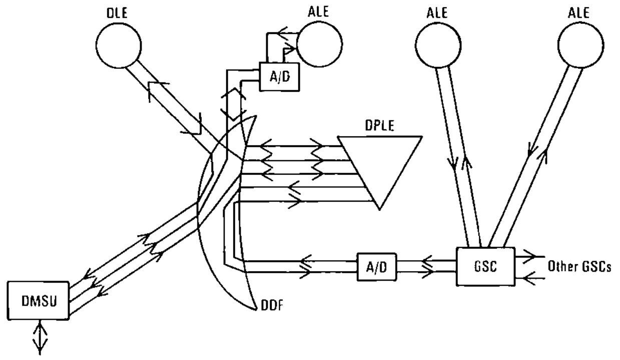

A/D: Analogue/digital interworking equipment

FIG. 1—DPLE during evolving digitalisation

ROLE AFTER MAIN NETWORK MODERNISATION

After the installation of the DSMN, but while junction network modernisation is still in progress, the DPLE will assume the parenting responsibility for those analogue local exchanges remaining when the trunk function of the GSC is ceased. The DPLE will continue to act as the main point of interconnection between analogue exchanges and the digital network, and will switch tandem traffic between all ALEs and DLEs within its catchment area. Furthermore, it will have to (see Fig. 2):

(*a*) concentrate traffic from all ALEs without direct DMSU links, for onward routeing through the digital network;
(*b*) distribute incoming traffic from the DSMN to ALEs;
(*c*) terminate routes with analogue transmission and provide charging information; and
(*d*) perform discrimination for combined level routes from unit automatic exchanges (UAXs).

In addition, the DPLE will make an important contribution to the security and resilience of the DSMN during

† Local Network Strategy Department, British Telecom Local Communications Services

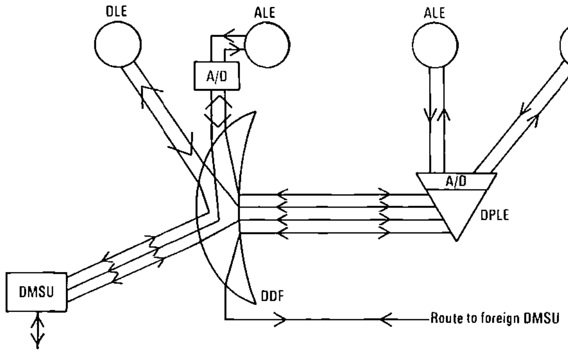

A/D: Analogue/digital interworking equipment

FIG. 2—DPLE after closure of the GSC

traffic surges or when line plant or the DMSU has failed. This is achieved in two ways:

(*e*) Traffic normally routed directly between the DLE/ALE and DMSU will be allowed to overflow via the DPLE, and thereby will be provided with an alternative means of access to and from the DSMN.

(*f*) The DPLE will have at least two traffic routes to the main network: to its parent DMSU and to at least one 'foreign' DMSU. Each route will provide a degree of security for the other. All traffic, whether from a DLE or ALE, which arrives at the DPLE, will share the DPLE's access and will thus be provided with security/resilience.

FURTHER EVOLUTION

During the first stage of modernising the network to digital working, the DPLE will have an essential role as the interface between the analogue and digital networks and as a concentration and distribution point for the traffic in its catchment area. As the digital network evolves, the DPLE will remain a node of particular significance, but the emphasis will move from the functions of the exchange at the DPLE site to its disposition as a major confluence point of the transmission network.

Biography

Margaret Rolfe works in the Network Planning Projects Section of BT's Local Network Strategy Department, where she is currently engaged on network structure studies. After graduating in Mathematics from the University of Bristol in 1971, she joined the Telecommunications Management Services Division of the Post Office, working on the development of a management accounting and information system. This was followed by a period in Operational Programming Department, on the design and introduction of the 'new' ASCE system, before she moved to her present post.

British Telecom Press Notice

SIGNALLING STANDARD FOR DIGITAL PABXs ON THE ISDN

Last October, British Telecom (BT) announced that it is to adopt the new signalling standard Digital Access Signalling System No. 2 (DASS No. 2) for use between digital call-connect systems (PBXs) and System X telephone exchanges. The new standard is intended for use on the next generation of PBXs linked to BT's integrated services digital network (ISDN), and its adoption will enable extensions on integrated services PBXs (ISPBXs) to be used as digital information-technology workstations.

The ISDN, which is based on the System X family of digital electronic telephone exchanges, will set up digital communications paths between customers connected to it. These paths will be able to carry information-technology services—data, text, facsimile and graphics—integrated with speech.

Customers will connect into the ISDN through integrated digital access (IDA) links in the local network. These will provide high-speed information paths between local System X exchanges and customers' premises.

IDA is available in two types to meet the wide range of customers' requirements: one, called *multi-line IDA*, serves PBXs, and the other, called *single-line IDA*, is for smaller installations. Single-line IDA will provide two parallel digital paths, one capable of being used for speech or data and operating at 64 kbit/s, and the other for data at 8 kbit/s. This would allow an investor, for example, to talk to a stockbroker over the 64 kbit/s path and simultaneously to refer to a computer database of share prices on the 8 kbit/s path. Single-line IDA will provide this dual speech and data capability on one telephone line.

ISPBXs will use multi-line IDA at 2 Mbit/s, which will provide a block of 30 communications paths, each of 64 kbit/s, between the PBX and the exchange. Each path in multi-line IDA can be used to support speech, data or text.

BT is introducing DASS No. 2 in the absence of any internationally agreed standard, because it feels that there is a need for customers and manufacturers to gain early experience and to benefit from the advanced features of an ISDN. These advanced features include multi-service access, high-speed call set up and network facilities such as calling-line identification.

DASS No. 2 is almost identical to the inter-PBX signalling system used on private digital networks—digital private network signalling system (DPNSS)—except that it incorporates features at the higher level needed for interworking with a public exchange. BT intends to implement DASS No. 2 on local exchanges for PBX groups. This announcement will enable equipment suppliers to make an early start on the software development needed to implement DASS No. 2 on ISPBXs.

Moreover, BT will continue to support DASS No. 2 after any international signalling standard has been agreed and subsequently implemented on System X exchanges. In this way, BT expects to get the best of both worlds: encouraging suppliers to make a start on implementing digital services now while keeping the door open for later developments based on a possible different international standard.

Loading the Digital Network

A. E. WALPOLE, and H. S. J. DAVIES†

UDC 621.395.34 : 681.32

The conversion of BT's network from analogue to digital working must be carefully managed to ensure that the aims of modernisation are achieved and the needs of the customer are met. This article outlines the procedures that are being adopted for implementing and loading the digital network.

INTRODUCTION

Converting a complete telephone network from analogue to digital working is not something that can be done quickly because of limitations on manufacturing capacity, manpower availability and capital expenditure, among others. These dictate that the network must be converted over a period of time. Accompanying articles in this issue of the *Journal*[1, 2, 3] summarise how British Telecom (BT) plans to achieve this conversion.

The changeover of traffic from the analogue to the digital network must be carefully controlled to ensure that plant is utilised effectively while, at the same time, the quality of service given to the customer is maintained. An added complication is that BT has been divided into separate businesses: Local Communications Services (LCS), National Networks (NN) and British Telecom International (BTI). Effective liaison and co-ordination between these three bodies is essential if the conversion plans are to be implemented successfully.

THE PLAN

Once the network strategy has been formulated, it must be converted into an established plan for each switching unit. For each digital main switching unit (DMSU), a catchment area plan (CAP) is produced that contains details of all the switching units in the area, event dates and loading/offloading traffic profiles. The traffic profiles reflect the agreements made between LCS and NN in the traffic agreement procedure (TAP), where traffic levels crossing the business boundaries are agreed. The profiles should also align with the annual schedule of circuit estimates (ASCE), which is a definitive statement of circuit requirements for the network. The CAP and the ASCE provide the backbone of the planning information for the network and its switching units, and are used as the planning input to the operational implementation process.

IMPLEMENTATION OF THE DIGITAL MAIN NETWORK

A responsibility of the Trunk Network Operations Division in NN is to translate the network plan into a practical implementation plan that makes best use of the available resources. Forward planning and provisioning is not always as successful as it might be, and it is essential that the provision and utilisation of plant is closely monitored to ensure that resources are put to best possible use.

The plant allocation and circuit provision activities are started 13 months prior to the planned completion date (PCD) of the DMSU. The plan is represented in a document called *digital opening requirements* (DORs). One set of DORs is produced for each DMSU and shows the circuit requirements for integrating that DMSU into the network. The Trunk Network Operations Division allocates plant and issues circuit advices and, in effect, produces an implementation plan based on available resources. This plan, called the *opening date circuit provision schedule* (ODCP), highlights those routes where shortages have been identified and expedient action is necessary to overcome the problem, as well as the circuits that will be available for the opening of the DMSU. Such expedient action may be to advance the provision of the plant in shortfall or to route the traffic in a different way. The ODCP is produced 11 months prior to the PCD to allow sufficient time for the necessary circuit provision work to be completed and for the data load tape, which tailors the DMSU to its position in the network, to be prepared.

THE NEED FOR CONTROLLED LOADING

It could be considered that the most convenient way of introducing digital technology into the network would be to do it overnight! But it has been recognised that this is not possible because of the enormity of the task and, in particular, the limitations on capital expenditure and manpower availability; the strategy for modernisation aims at providing customers with digital facilities at the earliest possible time, within the framework of these limitations, and taking account of competing market forces.

LCS is proposing to modernise all the local exchanges by 1995, but to load all originating trunk traffic (both analogue and digitally originated) onto the digital main network by 1990. NN is proposing to provide and fully interconnect the DMSUsby 1986/87. Between now and 1990, theDMSUs will have to be extended several times to provide the necessary capacity and, during that time, the routeing and flow of traffic within and between the analogue and digital networks will be complex. Careful monitoring and management of the fitted capacity will be required to ensure that the aims of modernisation are achieved and that facilities are provided when the customers want them.

Although BT has been split into separate businesses, the network is still fully integrated, and activity in one part of the network can have a direct effect on the service provided in another part. Again, this emphasises the need for close liaison between NN, LCS and BTI to ensure the successful introduction of digital facilities.

HOW THE NETWORK IS TO BE LOADED

When the DORs are prepared, they contain requirements for connecting local exchanges to the DMSUs. The quantity of traffic and the year in which it should be loaded onto the digital network will have been agreed between the businesses as part of the TAP, which forms the basis of the traffic forecasts. However, between the date of the forecast and the proposed implementation date, many factors are liable to change. Planned completion dates of switching units

† Trunk Services Operations and Maintenance, British Telecom National Networks

or the provision of line plant may not be on target or, indeed, market forces may result in the local exchange being loaded at an earlier date than was previously planned. Whatever the situation, NN must have a digital main network operational to meet the demands and must be able to monitor the sufficiency of the fitted capacity.

The loading of traffic onto the digital main network will be controlled and monitored by using the network traffic loading system (NTLS). The software for this system is being developed and will be available soon. In outline, the system will be able:

(*a*) to hold data on forecast demand and fitted working capacity and compare the two,

(*b*) to remodel the demand in the light of changing events and make a new comparison, and

(*c*) to have rapid access to measured traffic quantities and make comparisons to identify critical network components.

The fitted capacity will be held on a per-route basis and will be transferred into the system from the main network utilisation system (MANUS), which is the mechanised system currently being introduced. The NTLS will hold information on line plant and exchange terminal equipment fitted, planned and in use; it will prompt the issue of advices to circuit provision controls to provide, cease or rearrange circuits.

The forecast demand again will be held on a per-route basis and will be derived from the latest available ASCE information.

The ability to model the demand is provided from the digital unit circuit estimates (DUCE) system, which has been used for the last 2–3 years to generate route estimates for the digital main network. The system operates on a matrix of traffic between group switching centres (GSCs) derived from the national traffic matrix. Originating traffic fed into the system is distributed over the network by using the matrix proportions for each unit in turn, and the results are summated on a unit-to-unit basis to create an overall picture of the network.

Measured traffic will also be collected via the network traffic management system and fed straight into NTLS.

The NTLS will enable those routes which are at, or approaching, critical capacity to be identified so that action can be taken to restore or maintain the quality of service given to the customer.

CONCLUSION

The theme of close inter-business liaison is carried into the forecasting and planning activities by means of the TAP, which ensures that planning and investment processes within BT are all aligned. The theme must be carried through into the operational area if BT is to be successful in achieving its aims. BT cannot afford to load traffic onto the digital main network if there is insufficient capacity. This would result in customers being dissatisfied with the service and could ultimately lead them to transfer their business to a competitor. NN will need to have constant and close contact with both LCS and BTI to ensure that changes in plans and events are reflected in the main network. Procedures are being developed to achieve this and at this stage it is envisaged that NN District staff will have the responsibility for maintaining a close association with LCS District and Sub-District staff to ensure that the information is gathered. At the centre of these activities will be the NTLS, which will be run by the Network Loading Group at Oswestry. Information fed into this group from NN District staff will be crucial for the successful introduction of the digital main network and its facilities to BT's customers.

References

[1] GARBUTT, B. N. The Digital Restructuring of the British Telecom Network. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 300.

[2] ROLFE, M. A. Role of the Digital Principal Local Exchange. *ibid.*, Jan. 1985, **3**, p. 304.

[3] CROOKS, K. R. Local Network Strategy—Today's Plans for Tommorrow's Network. *ibid.*, Jan. 1985, **3**, p. 297.

Biographies

Arthur Walpole was until recently Head of the Network Traffic Loading Group in the Trunk Services Operations Division of NN. He joined BT, in the Coventry Telephone Area, as an open-competition level-1 entrant in 1965. Since 1971, he has worked on the planning and implementation of the main network, including the design and implementation of the transit network. Over the past three years, his work has included devising and designing the DUCE system for producing route estimates for the digital main network and, latterly, co-ordinating the operations aspects of introducing DMSUs into the network.

John Davies joined BT in 1971 as an open-competition level-1 entrant in the London North West Telephone Area. Since transferring to Headquarters in 1976, he has been involved in the design of private networks and, more recently, in the management and use of main-network resources. He is now Head of the Digital Network Loading Sub-Group in the Trunk Services Operations Division of NN, with particular responsibility for ensuring that the transfer of traffic from the analogue to the digital network is co-ordinated and efficient.

Network Management in the Digital Network

S. HEAP, B.SC.†, and J. D. ARTHUR, B.A.*

UDC 621.395.74 : 621.374 : 654.01

The technical features of the digital network planned by British Telecom will enable traffic on the network to be controlled remotely and in real time. This article discusses why such control is essential for the network to be managed effectively and describes the controls that will be possible. It goes on to indicate how traffic data available on-line from the digital exchanges could be processed to enable the operation of the network to be monitored easily.

INTRODUCTION

The British Telecom digital network is being built up according to all the rules for grades of service, transmission, routeing and so on. It has been planned to carry certain forecast traffic and designed to give a good fault-free performance. If all the activities that have an impact on the provision and use of plant and equipment are also well co-ordinated and effectively achieved, the management of the network is likely to benefit both the customer and the administration. However, situations in which the network is not able to provide a satisfactory performance will occur; for example, because of failure, late provision of plant or changes in types or patterns of traffic. These situations fall into the realm of the network manager.

The network manager's tasks are threefold. The first is to monitor the network. This is done by measuring traffic flows, lost calls, failures, quality of service and so on. The second task is to analyse the data gathered, and estimate its importance and the likely situation in the network over the ensuing period, which may be hours and days, or weeks, months and years. This may be done by determining likely trends of traffic on routes and comparing them with available capacity, or by looking for a pattern of faults that could indicate a potential major fault. The third task is to take action either to prevent problems from occurring or to minimise the seriousness of their consequences.

The remedial action taken generally falls into two categories: initially, in the short term, to mitigate the immediate effects of a situation and, subsequently, in the longer term if necessary, to alter the network or the network plan to a new configuration. The first, together with its monitoring and analysis phases, is termed *network management*; the second is an aspect of network development and planning.

Within the network management activities, two main functions are evident. The first is to take action to provide alternative network paths in the event of failure; the other is to alter the way the network handles traffic either when failure occurs or simply when traffic flows are such that the available capacity is insufficient. The second of these, known as *traffic control*, is the subject of this article, as this function in particular is brought to prominence by the technical facilities of the digital network.

THE NEW NETWORK

The reasons for moving towards a digital network are well known, but those features that are of most relevance are worth highlighting: stored-program control (SPC), digital switching and transmission, and common-channel signalling.

SPC provides great flexibility in the way the network can handle calls. Routeings can be changed, circuits taken out of service, even call set-ups blocked—all from a remote terminal in real time. Integrated digital switching and transmission minimises the equipment required at the interface between switching and transmission, and provides a transmission performance that is virtually independent of the distance and the number of exchanges through which a call is routed. Moreover, common-channel signalling can convey more complex messages around the network more quickly than any current signalling system, and this opens up the possibility of sending administrative messages relating to the state of the network between processors.

In addition, facilities such as automatic alternative routeing (AAR) and automatic rerouteing (ARR) will be available in the new digital network, although the rules for their use in the top tier have not yet been fully defined. Data on the performance of the network, both in terms of its serviceability and the traffic and call flows, will be available on-line for either immediate analysis or later main-frame processing. A digital system such as this has many benefits for the administration and the customer that have been well covered elsewhere; but possible problems with the network will also arise, and this makes some kind of real-time control essential. There are three main areas of concern:

(*a*) *Vulnerability* In an exchange with distributed processing (for example, a Strowger exchange) the probability of a complete exchange failure is very low, unless something like the power area is destroyed. An SPC system, however, relies on a central unit to control calls through the switch, and failure of this central area could cause the unit to fail. This failure could be because of a corruption in the program from which the exchange is unable to recover, or because of a failure of some hardware that seriously affects the call-handling capacity of the unit. A third problem, which again arises because of the common-control nature of the device, is the reaction of the unit to overload. Here again, a Strowger exchange would continue to switch calls on some routes while experiencing heavy overload on a different route, whereas a common-control unit has to handle all calls through its central area. A large volume of call attempts or clear downs, caused perhaps by the failure of a route or simply by very heavy demand, could overload the call-handling processes and cause the unit to shed call attempts.

(*b*) *System Failure* The digital network is built up from a small number of large switching units interconnected by large capacity routes. A failure of either of these would take a large amount of capacity out of the network. This would

† Trunk Services Operations and Maintenance, British Telecom National Networks

* Local Network Strategy Department, British Telecom Local Communications Services

not be the end of the problem, as a failure could cause all calls currently in progress at that time to be released, and this would generate a large number of *clear* signals to adjacent units. If the failure had occurred during a particularly busy period, there is a good chance that the adjacent units would not cope with the additional demands for clearing calls and the units would be overloaded. The problem would then spread and affect call attempts well away from the original failure. It is not suggested that these failures would happen frequently, but their possibility must be taken into account in the management of the network.

(*c*) *Traffic Conditions* The third reason for real-time control comes from the changing nature of the traffic on the network. As the number of customers connected to a network increases, the potential demand on the network is increased, although it may not ordinarily be turned into actual demand. In other words, although most of the new customers will be residential and will generate little traffic during the working day, a potential exists for a very large increase in traffic if some event happens to stimulate these customers to make calls; for example, a television advertisement. This would not be too bad if the only customers affected are those trying to call the quoted number; unfortunately, this is not the case. The large volume of traffic blocks access to other customers in the exchange and, if the problem is sufficiently severe, will cause serious problems on incoming routes.

These problems occur now in the existing network, but are likely to be more pronounced in the new digital network. To some extent, the analogue network is self limiting, because, once a route to the required destination is congested, no more traffic can be routed there. In addition, traffic being focussed on some out-of-the-way place is limited by the capacity of any intermediate tandem exchanges. In the digital network, this will not be the case. The routes themselves will be much larger and will collect traffic from much larger catchment areas. Facilities such as AAR have been provided to help traffic to its final destination even if the first-choice route is busy; thus, there is a danger that extra traffic having little chance of being successful will be routed into the problem area. These facilities would improve the service when conditions are normal, but they only worsen the situation when congestion already exists.

More stimulation of traffic, particularly residential traffic by television advertising, will also be seen, and this will make the provision of an acceptable service more difficult. In addition, the integrated digital network will allow the introduction of many new types of service that will quickly be introduced and may have very different characteristics; for example, shorter holding times, which will increase the call demand on the processors for a given traffic level, or different busy hours to the normal conversational traffic. These events will happen in time-scales which are much too short for customary planning procedures to take place.

These three aspects demonstrate the case for being able to alter the way the network behaves in real time. In many instances, the effect of a new service cannot be fully anticipated and, to keep the network in control, the ability to react quickly is necessary. The new network provides remote facilities for changing its operation in such a way.

CONTROLS

There are two main types of control in the network management field: expansive and protective. Expansive controls are the preferred option as these attempt to reroute those calls that are likely to succeed around the problem area; protective controls will be used when this is not sufficient and some pressure on the network must be removed. Expansive controls will include changing routeing schedules, or implementing dormant ones, to increase the number of available route choices above the normal level. Often an overload will be too severe to be dealt with in this way, as it will have the effect of spreading the congestion further into the network. If this is the case, then the introduction of protective controls would be considered. A simple protective control would be to change the recorded announcement given when a call fails, to explain the reason for the failure, in the hope that this will deter repeat attempts. For a general overload, raising the trunk reservation factor on traffic to a route can improve service to other parcels of traffic, although, if all else fails, the blocking of calls entering the network would have to be considered. This could be broadly based if there were very severe problems affecting the network but, more usually, it would be used to bar perhaps 90% of calls to a particular problem telephone number. In this way, the number of ineffective calls in the network would be drastically reduced, while some calls would still be allowed to get through and possibly succeed.

TRAFFIC ASSESSMENT PARAMETERS

If it is assumed that these types of control can be applied, the problem of knowing when and which control to apply remains. Three main parameters register the behaviour of the network, as follows:

(*a*) Seizures per circuit per hour gives an indication of the effectiveness of the service being offered. If there are a large number of ineffective calls on a route, the holding time drops and the number of call attempts and seizures increases.

(*b*) The answer/seizure ratio gives a measure of the success rate of calls at their final destination. It indicates calls that are being routed to their destination, but are not then succeeding.

(*c*) The all-circuits-engaged measurement indicates that all the circuits on a route are engaged. New call attempts either overflow on to an alternative route or are lost if this is the final route. As such, it is a simple measure of the calls failing to get out of the exchange itself.

All these parameters are calculated in real time, in addition to other less essential ones, but they obviously cannot be examined on all routes simultaneously. A threshold system, where thresholds are set for each parameter, is therefore necessary; an indication, called an *exception report*, is generated when one or more of the thresholds is exceeded.

TRAFFIC MEASUREMENTS NEEDED

A wealth of information will be available on-line from the digital exchanges. It is currently thought that an output frequency of five minutes gives the optimum balance between sample size and delay in being informed of a problem, although this could well be changed with experience.

The first essential is for call counts to destinations ranging from national number groups down to an individual telephone number for both offered calls and successful calls. Measurements of traffic intensity to indicate congested routes and the availability of spare capacity to take rerouted traffic are also required. Finally, the performance of the exchange itself must be monitored.

A large volume of data must therefore be gathered. How can it be handled from a human point of view? A method of sifting information and displaying only those items required is needed so that the network manager can see at a glance what is happening in the network. Equally important, the network manager needs to be able to relate the parameters in each individual unit, to produce a picture of the real problem, which could, of course, involve many units in the network. The obvious answer is to present the

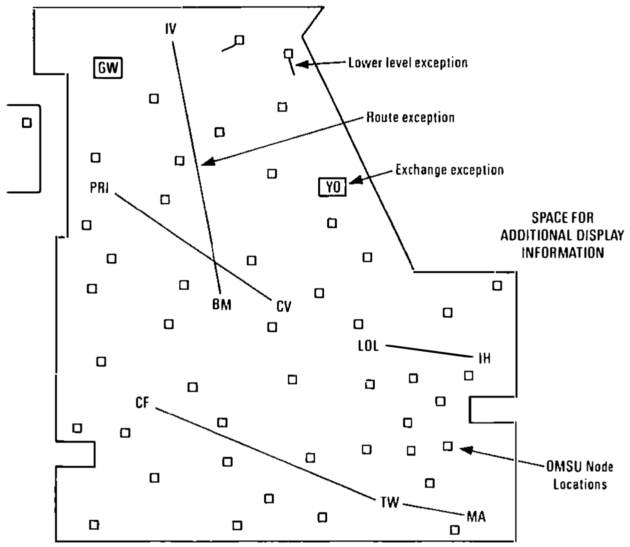

FIG. 1—Global-map display

data graphically, and Fig. 1 shows an example of the type of display envisaged.

This display would be updated every five minutes and would give a picture of the situation in the network. By using the same basic format, the display could be used to indicate those routes with less than say 60% occupancy, to discover whether there was sufficient spare capacity to re-route traffic away from a problem area. It would also be able to display the state of routes and exchanges at a lower level in the hierarchy. This would be shown on a grid based on a nominated digital main switching unit (DMSU), and any exception reports from the lower-level exchanges would be shown.

Once this overview of the network had been obtained, however, access to the actual data that generated the exception report would still be needed; normal displays would give in detail all the relevant parameters and enable backtracking through previous measurement periods, to see how the situation had developed.

BENEFITS

The benefits fall into three main categories: service, financial and network resilience.

The service benefits that arise from adopting this approach are self-evident. The performance of the network can be monitored in almost real time, and awareness of the reaction to fault or overload conditions will be much quicker than at present.

On the financial side, the aim behind real-time control is to maximise the number of call completions in the network. Repeat attempts could well use a competitor's network in the future; thus network management can only serve to increase the revenue from those calls that otherwise would have failed. The second way to obtain financial benefits is by tightening the dimensioning rules to take account of the fact that the real-time control is present and should be able to handle any sizeable overloads of traffic. The current provisioning tables include an element of over-provision to ensure that a reasonable service is given under various degrees of overload, and this over-provisioning could be reduced with capital savings of plant.

Finally, the introduction of the type of control outlined would have a marked effect on the overall resilience of the network. An automatically-switched digital service protection network is already being introduced to switch in spare plant when a transmission system fails, and the introduction of the real-time traffic control element will complement this as well as extend this protection to both exchanges and other areas of the network where spare plant is not available.

Biographies

Steve Heap is Head of the Network Management Section of National Networks's Trunk Services Operations Division in Oswestry, Shropshire. He joined Telecommunications Headquarters in London in 1972 after graduating from London University with an honours degree in Physics. Since that time, he has been involved in the introduction of advanced traffic-recording equipment and, more recently, in the production of exchange-dependent data for DMSUs. The section he leads is responsible for the day-to-day management of the trunk network.

John Arthur is Head of the Network Management Group in the Local Network Strategy Department of BT Local Communications Services. He joined Birmingham Telephone Area in 1965 and moved to London five years later to work in the Advisory Group on Systems Definition on early requirements for the digital network. He later worked on the requirements for modern PABXs, including those for Monarch. Since 1979, he has been involved with the specification and introduction of network management techniques for the digital network. He has an honours degree in Biology and Geology.

Integrated Services Digital Network

C. D. E. PRICE, B.SC., C.ENG., M.I.E.E.†, and R. A. BOULTER, B.SC., C.ENG., M.I.E.E.*

UDC 621.395.34

This article describes the initial implementation of an integrated services digital network (ISDN) in the British Telecom network. It describes the technical developments that form the basis of a pilot service, based on four System X local exchanges and due to be brought into service early this year, and discusses some of the plans being made for a more extensive coverage.

INTRODUCTION

British Telecom (BT), in common with all major telecommunication administrations throughout the world, is implementing a multi-purpose integrated digital network (IDN) based on processor-controlled digital exchanges, interconnected with digital transmission and supported by fast inter-processor common-channel signalling[1]. With the introduction of digital local exchanges the IDN will provide a fully digital connection from the input side of the local exchange on which a call originates to the output side of the local exchange on which that call terminates.

The integrated services digital network (ISDN) is created by extending the IDN to the customer, by means of integrated digital access (IDA) (see Fig. 1), in such a way as to permit customers to utilise the end-to-end digital network to satisfy their full diverse range of data and voice telecommunication requirements, both switched and non switched, at up to, and possibly in excess of, 64 kbit/s.

The IDN necessarily will take some years to implement. For a number of years at least, it will be necessary for the ISDN to interwork with other established BT networks—for example, KiloStream and Packet SwitchStream—in order to provide customers with some of the facilities they require. Indeed, such separation of networks may remain a feature of the BT telecommunications service. However, to the customer on the ISDN, such separation of networks is not visible; the ISDN customer has, in IDA, a single means of accessing all network facilities, however they are provided.

† Local Exchange Systems Department, British Telecom Local Communication Services

* System Evolution and Standards Department, British Telecom Development and Procurement

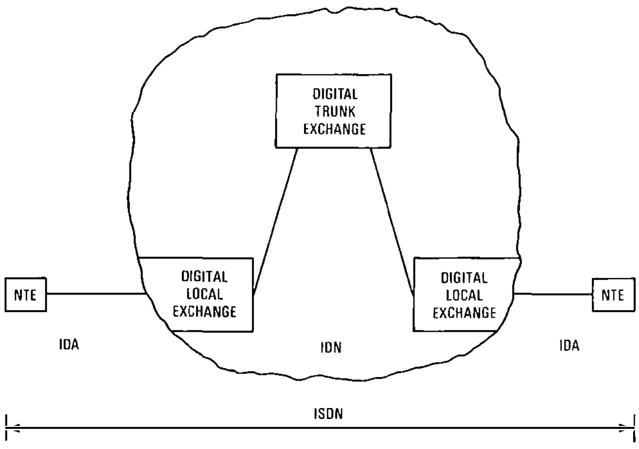

FIG. 1—Relationship of ISDN, IDA and IDN

A previous article[2] reviewed the basic philosophy of the ISDN and discussed the services and facilities, in particular data services and facilities, offered by the ISDN in the context of those available from existing telecommunications networks. The same article discussed in some detail the provision being made for interworking between the ISDN and those existing networks. This article describes the technical basis on which the initial realisation of the ISDN in the BT network is built, the developments that have been undertaken and the plans that BT is formulating for the implementation of an ISDN. The initial phase of implementation takes the form of a pilot service and, in the main, the description given below relates to the provision being made for that pilot service. A second phase, based on the same equipment design, except for the use of the new network terminating equipment (NTE) described below, is expected to begin during 1986. A third phase, planned for 1987/88, is currently being specified and, as far as possible, will provide for the implementation of internationally agreed standards.

GENERAL PRINCIPLES

Examination of the features of voice and non-voice services shows a high degree of commonality. The objective with an ISDN is to make the economies of scale created by the bulk telephone traffic available to other services. This can be done if the network capabilities are largely independent of particular services or implementation arrangements.

The general capabilities required to create an ISDN are:

(*a*) a fully digital network with service-independent transmission and switching media, and fast message-based inter-exchange signalling;

(*b*) customer digital access through the local network; and

(*c*) message-based customer-to-network signalling to provide the speed and repertoire appropriate to new non-voice services and enhanced voice services.

Of these, the first is inherent in the creation of the IDN, and the second and third are features of the IDA. In addition to the above capabilities, a practical realisation of an ISDN necessitates further capabilities from the network as a whole; for instance, digital route identification for calls requiring wholly digital, 64 kbit/s, paths through a network which, for an interim period, will include analogue paths. Such further capabilities imply an overhead which the ISDN places on the IDN but, in practice, this overhead is not great, being limited to revised call-processing routines and network-signalling protocols.

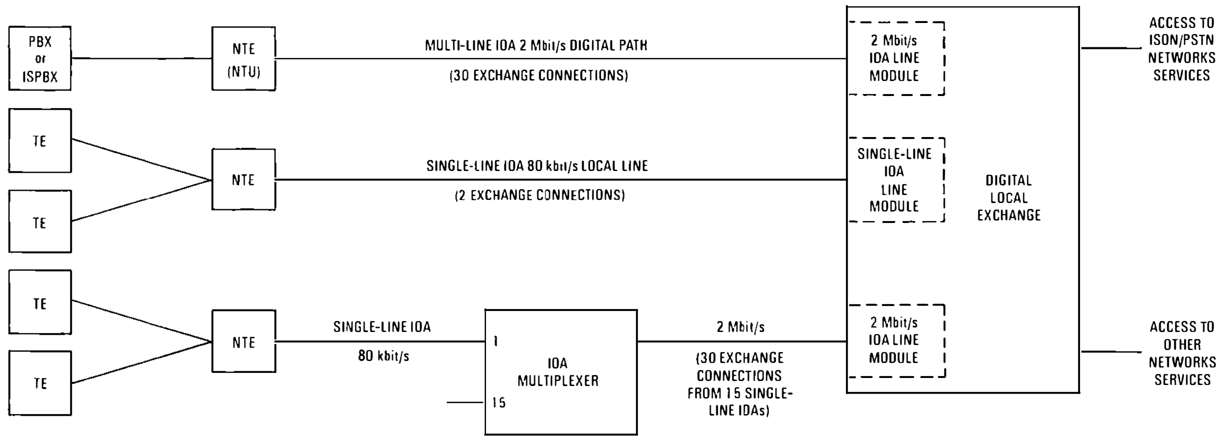

TE: Terminal equipment

FIG. 2—Integrated digital access

ACCESS OPTIONS

Two IDA options are to be made available to customers: single-line IDA and multi-line IDA.

Single-Line IDA

Single-line IDA utilises a single pair of the existing local cable network to carry the two directions of transmission for two traffic channels, each representing one exchange connection, together with the signalling information relating to both traffic channels. The initial design of single-line IDA comprises one traffic channel of 64 kbit/s, suitable for both voice or data at rates up to 64 kbit/s; a second traffic channel of 8 kbit/s, suitable for data only; and a signalling channel of 8 kbit/s. Thus, the total information rate required is 80 kbit/s in both directions.

In the customer's premises, a single-line IDA NTE provides X- and V-series standard CCITT† interfaces to terminal equipment. In the local exchange, an ISDN-compatible design of customer's termination extracts the 8 kbit/s signalling channel and provides for the separate connection of each traffic channel to the concentrating switching stages of the digital subscriber switching subsystem (DSSS) of System X, the 8 kbit/s traffic channel being rate adapted, by re-iteration, to 64 kbit/s for transmission throughout the IDN.

Multi-Line IDA

Multi-line IDA provides up to thirty 64 kbit/s exchange connections in time-slots (TSs) 1–15 and 17–31 of a 2 Mbit/s digital path. The 64 kbit/s channels provided by TSs 0 and 16 are allocated for alarms/synchronisation and for signalling in the normal way. Any requirement for rate adaption would be a function of the equipment connected to multi-line IDA, and this equipment would have to interface directly to the signalling in TS 16.

The multi-line IDA connection is terminated at the customer's premises on a network terminating unit (NTU), which constitutes the NTE for multi-line IDA and presents a standard CCITT 2 Mbit/s interface to the customer. At the exchange, a further design of ISDN customer's termination extracts the signalling and provides for the connection of the traffic channels to the DSSS concentrating switching stage. Multi-line IDA could be used for any purpose involving the connection of a number of ISDN traffic channels to the BT network. However, the principal use initially is expected to be the connection of digital PBXs to the ISDN, thus permitting the provision of the ISDN capability to the extensions on that PBX. A digital PBX designed to take advantage of this capability is known as an *integrated services PBX* (ISPBX).

† CCITT—International Telegraph and Telephone Consultative Committee

Multiplexer Access

Single-line IDA, exceptionally, can be provided to customers in exchange areas still served by analogue exchange equipment by means of an IDA multiplexer. This multiplexer extends single-line IDA connections for up to 15 customers to a digital local exchange over a 2 Mbit/s digital path. Use is made of the same access to the digital exchange as the multi-line IDA connection.

The three methods of IDA connection are illustrated in Fig. 2, and the following paragraphs provide an outline of the individual elements of these three access methods.

NETWORK TERMINATING EQUIPMENT (NTE)

The principal functions provided by any single-line IDA NTE are:

(*a*) standard X- and V-series terminal interfaces,

(*b*) protocol conversion (interworking the signalling over the terminal interface with that used in the 8 kbit/s signalling channel to the exchange),

(*c*) rate adaption as required (to bring the customer's terminal information rate up to the 64 kbit/s or 8 kbit/s data rate of the traffic channel to be used),

(*d*) the multiplexing of the two traffic channels and the 8 kbit/s signalling channel (total 80 kbit/s), and

(*e*) the interface to the full-duplex digital transmission system.

In order to accommodate the wide range of terminals that might be connected, the NTE also provides for the input and update of information relating to the particular terminals actually connected to each of the data ports; for example, the information data rate of the terminal.

The rate-adaption function of the NTE involves the mapping of the CCITT standard synchronous data rates (2400, 4800, 9600 and 48 000 bit/s) and the recognised asynchronous rates presented on the data port into the 8 and 64 kbit/s traffic channels. At the time of development, no internationally recognised standard technique for adapting these rates to both the 8 and 64 kbit/s rate was available. BT therefore implemented its own techniques (based upon certain standards that did exist) which enabled all recognised synchronous rates up to 64 kbit/s to be mapped onto the 64 kbit/s bearer, and rates up to 4·8 kbit/s to be mapped onto the 8 kbit/s bearer channel. Similarly, asynchronous rates of up to 9·6 kbit/s and 1·2 kbit/s could be mapped onto the 64 and 8 kbit/s bearers, respectively.

Two NTEs have been developed initially for the pilot service. The NTE1 has been designed as a desk-top instrument and includes a digital telephone, keypad, display and a single data port. This data port can be configured with a CCITT X21*bis* interface (or V24 via an external adapter), or the leased circuit version of X21, to operate over the 8 or 64 kbit/s traffic channels to the exchange; the 64 kbit/s channel is not available for data when a telephony call is in progress.

Call set-up is by means of the in-built keypad, and an alphanumeric display, which assists the user with call set-up and facility programming; that is, the input of terminal dependent data. On-hook dialling is provided, together with a number of special function and service request keys (some programmable by the customer), to simplify the use of supplementary services; for example, closed user groups and short-code dialling.

The NTE1 was designed during the early stages of the ISDN development programme and prior to the regulatory decisions embodied in the recently published BT licence. As the NTE1 incorporates features of a terminal as well as those of an NTE, it does not meet the conditions for the liberalisation of customer equipment laid down in the licence. Consequently, it will not form part of the generally available ISDN equipment; its use is to be restricted to the pilot ISDN implementation programme only.

The NTE3 is designed for applications where there are a number of different terminal equipments requiring access to the BT network and where all are capable of controlling call set-up procedures. For this reason, the NTE has no built-in telephone, keypad or display, but it has six terminal ports, each of which may be configured to support terminals with one of the following: an X21 interface (leased or circuit-switched variant); an X21*bis* interface (or V24 via an external adapter); a 2-wire analogue interface or an X24/V28 interface. The latter provides access via a modec (a *MO*dem and co*DEC*) to public switched telephone network (PSTN) basic modems at rates of up to 300 bit/s. Because no more than two of the six terminal ports may be connected to the exchange at one time, the NTE resolves contention between them, acting effectively as a traffic concentrator.

New standard rate-adaption techniques have now been agreed by the CCITT for synchronous rates (CCITT Recommendation X30). In addition, the European Computer Manufacturers' Association (ECMA) is about to agree a technique for mapping the asynchronous rates up to 19·2 and 4·8 kbit/s onto the 64 and 8 kbit/s bearers, respectively. BT is now proposing to adopt these techniques in a new cost-reduced NTE, NTE4, which is being developed for wider use when the ISDN is extended beyond the pilot during 1986. At the same time as the NTE4 is introduced, the NTE1s and NTE3s will be either modified or recovered, because the same rate-adaption standard must be employed in all NTEs throughout the network.

The basic NTE4 will be a data-only NTE offering two X21 ports supporting both circuit-switched and leased-circuit variants for synchronous data rates up to 64 kbit/s. There will also be a V24 port for facility programming. In order to facilitate interworking to V-series terminals and to provide a telephony capability, a range of terminal adapters is also being developed to connect onto the X21 port of the NTE4. The first adapters to be developed will be those providing an X21*bis* interface (or V24 via a passive adapter) and a telephony adapter, either an analogue variant offering the facility to connect an analogue telephone and/or modem, and a digital variant offering the possibility of a digital telephone.

LOCAL NETWORK ACCESS

The existing local cable network has already been established to provide telephony to customers and represents a large proportion of BT's capital assets. Although this cable network has been installed to carry analogue signals with a bandwidth of less than 4 kHz, it can be exploited to carry the wider-bandwidth digital signals of a single-line IDA. The bandwidths for the initial single-line IDA, with an information rate of 80 kbit/s, will be of the order of 100–200 kHz, depending on the modulation technique employed. As a consequence, the signals will be subjected to much greater attenuation (up to 60 dB) than the telephony signal (now a maximum of 15 dB on lines served by digital exchanges). Notwithstanding these high attenuations, it is still possible to provide full-duplex transmission over a single pair of wires by the adoption of suitable modulation techniques. Initially, BT will be using two techniques for providing 80 kbit/s transmission over the local network. Most customers will use a simple burst-mode technique[3], but a more complex echo-cancelling technique[4] will be used on a smaller experimental scale.

Burst-Mode Technique

In the burst-mode technique, the 80 kbit/s input data is loaded into a buffer and then clocked out onto the cable pair in bursts of 22 bits (20 information bits and 2 marker bits) at an increased rate of 256 kbit/s. The interval between the bursts of data, 250 μs, provides ample time to receive bursts of data from the remote transmitter. Full-duplex transmission is therefore achieved without a low-level receive signal having to be identified in the presence of a high-level transmit signal. When the transmit bursts of a number of systems in the same cable are synchronised, the problems of near-end crosstalk are also avoided. (However, it should be pointed out that, if systems on pairs with a long reach are mixed in cables with a short reach, this advantage is, to a certain extent, nullified.) A WAL2[5] line code is used, which ensures that sufficient timing information is present and that no equalisation of the characteristic distortion of the line is required at the receiver. The system is therefore relatively simple and consequently cheap to implement. The disadvantage of the system is that it suffers high attenuation, because of the higher frequency at which it operates.

Echo-Cancelling Technique

With the use of an echo-cancelling system, it is possible to transmit and receive simultaneously. The system also uses a WAL2 line code, but the basic line rate is only 88 kbit/s (comprising the 80 kbit/s information rate plus an 8 kbit/s framing signal), with the frequency spectrum centred on 88 kHz. Fig. 3 illustrates the principles of the more complex

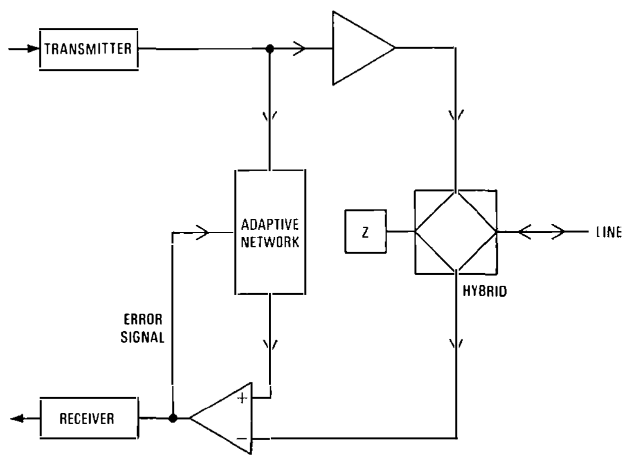

FIG. 3—Principle of echo cancelling

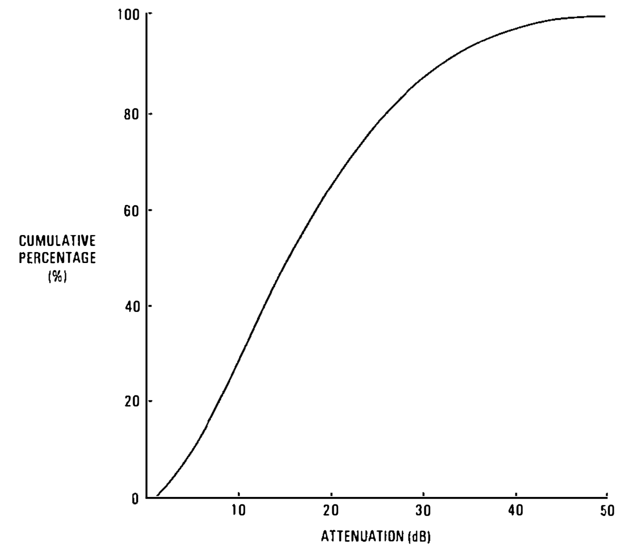

FIG. 4—Distribution of local line loss at 100 kHz

echo-cancelling system. A hybrid provides some initial reduction of the transmit signal passing into the receiver, but inevitably the transmit signal will still be much larger than the receive signal on the longer lines. An adaptive network therefore models the echo-path response between the transmitter and receiver, which results from the imperfect hybrid and the reflections on the line. The signal generated is then subtracted from the receiver input to leave the true receive line signal. The adaptive network derives its signal by minimising the residual echo components on the wanted receive signal; hence, the system's name.

Transmission Limits

The planning limit to which service to the majority of telephony customers has been provided is 10 dB loss at 1600 Hz, and this equates to a maximum loss for digital systems operating around 88 kHz and 256 kHz, such as those to be used initially for single-line IDA, of 40 dB and 60 dB, respectively. The current burst-mode system with a maximum permissible loss of 34 dB at 256 kHz will give a range of approximately 2·5 km over 0·4 mm copper pairs, and will therefore enable 78% of customers to be connected directly to their local exchange. The experimental echo-cancelling system, with a permissible loss of 30 dB at 88 kHz, will provide a slightly longer reach of approximately 3 km and would enable 89% of customers to be connected directly. Fig. 4 shows a graph of the percentage of customers having a loss less than, or equal to, the value shown at 100 kHz. More technologically advanced transmission systems under development, designed to provide the CCITT recommended 144 kbit/s information rate referred to below, have considerably better performance and will be capable of reaching more than 98% of customers.

SIGNALLING

The introduction of the ISDN is dependent on the provision of a new signalling system between the customer and the exchange. This new signalling system is necessary to provide the speed and repertoire appropriate to the full range of services and facilities which will be provided by the ISDN. At the time when BT was preparing its specification for the pilot service, no CCITT Recommendations on customer access signalling systems were available and so a totally new digital signalling system, digital access signalling system (DASS), was defined[6]. The structure of DASS is based upon the levels of the International Standards Organisation (ISO) model for Open Systems Interconnection (OSI)[7].

The introduction of this new message-based signalling system into the local network enables a much wider range of facilities to be provided, and allows more information to be provided by the network on the progress of calls. One item of additional information is the service indicator code (SIC) used to reject incompatible call attempts; for example, incompatible because of different data rates or different services. However, DASS No. 1, as defined for the pilot ISDN, will initially support only the same supplementary services as defined for telephony customers and normally accessed by those customers using multi-frequency (MF4) signalling (such as call diversion, abbreviated address, three-party call etc.), along with two new services for data customers (closed user groups and originating- and terminating-line identities[2]). This is in order to minimise the exchange software development required for the pilot ISDN. Users must also indicate at the beginning the type of call; for example, telephony or data.

The definition of the DASS—which can be used on both single-line IDA and multi-line IDA to PABXs, the development of digital PABXs and the use of digital leased circuits to form digital private networks led to a need for a digital inter-PABX signalling system. BT and a number of UK PABX manufacturers have collaborated on the definition of a signalling system based upon the DASS, but enhanced to meet the inter-PABX signalling requirement. This further signalling system is called *digital private network signalling system* (DPNSS)[8]. During the definition of the DPNSS, it became apparent that it was desirable to align more closely certain of the messages of the DASS with those of the DPNSS. At the same time, proposals were being made to enhance the repertoire of signalling messages in order to provide more facilities to the user. An enhanced version of DASS No. 1 was therefore defined, called *DASS No. 2*, which would enable PABXs to have a common type of signalling system for both inter-PABX links on private circuits and PABX-to-network signalling[9].

In the pilot ISDN, DASS No. 1 will be used on single-line IDA by NTEs, and DASS No. 2 on multi-line IDA by ISPBXs. However, since the full features of DASS No. 2 were not defined in time for the pilot development, ISPBXs will be able to use only a subset of the features of DASS No. 2 during the pilot. Further developments are in hand fully to support DASS No. 2 for ISPBXs and, when this is introduced, the common features of the DPNSS and DASS No. 2 will enable them to be interleaved on a common signalling channel such as TS 16 in the 2 Mbit/s multiplex structure. This will allow some of the 30 traffic channels from the ISPBX to be used for private circuits providing direct PBX-to-PBX connections, while others are used for switched public network connections.

LOCAL EXCHANGE

Fig. 5 represents the major functional blocks of the concentrator of a System X digital local exchange, the digital subscriber switching subsystem (DSSS)[10].

Single-line IDAs are connected to digital line units (DLUs), a number of which are terminated on a line controller to form a single-line IDA line module. Similarly, multi-line IDAs, and 2 Mbit/s digital paths for IDA multiplexers, are connected to digital line terminations (DLTs) which, with the signal converter, form a 2 Mbit/s IDA line module.

The interfaces into the module controller and the concentrator switch from these line modules are identical to those of the analogue line modules on the exchange. Thus the provisioning of IDAs necessitates specialist ISDN equipment

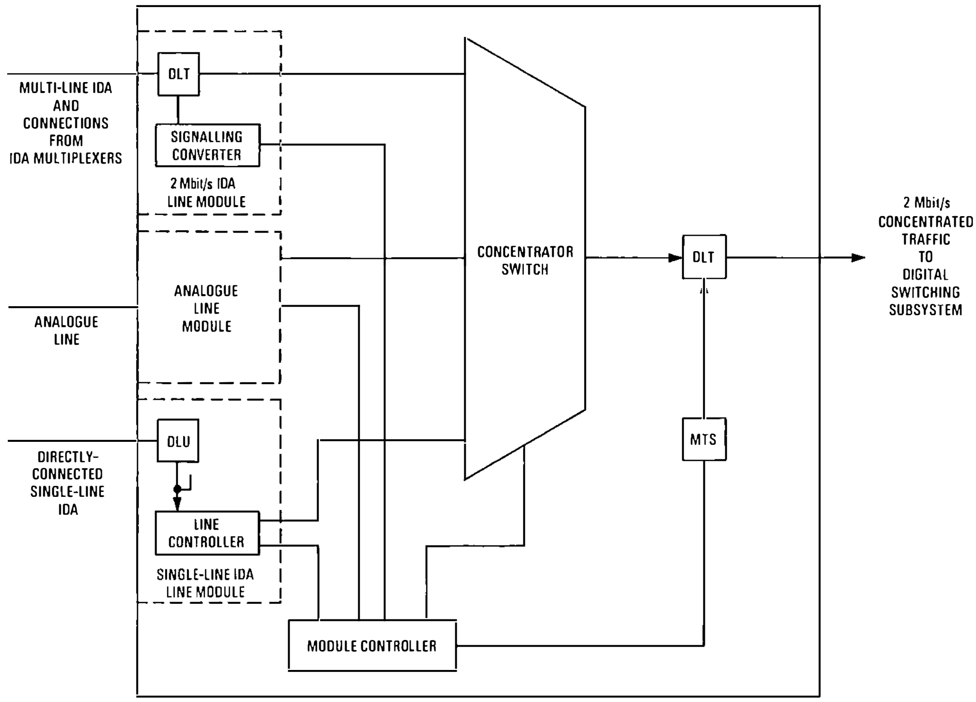

FIG. 5—Generalised view of the digital subscriber switching subsystem for IDA connections

ment only in the form of shelves of IDA line modules. A single DSSS can be equipped with a number of shelves of analogue line modules as well as shelves of single-line IDA and 2 Mbit/s IDA line modules. The implementation of the ISDN in the exchange necessitates no action other than specifying those alternative IDA line shelves and dimensioning the DSSS according to the traffic loads expected on the mixture of analogue and IDA lines.

Out-of-Area Access

In the early years of the ISDN, not all exchange areas will be served by digital local exchanges, and the IDA multiplexer has been developed to provide service to customers outside the catchment area of a digital local exchange, where this can be economically justified.

The IDA multiplexer houses DLUs, identical to those used by customers directly connected to a System X exchange, supporting single-line IDA. The multiplexer supports 15 single-line IDAs and multiplexes the 64 kbit/s traffic channels into TSs 1–15 of a standard 2 Mbit/s structure (conforming to CCITT Recommendation G732) and the 8 kbit/s traffic channels are reiterated up to 64 kbit/s before occupying TSs 17–31. The signalling messages contained within the 8 kbit/s signalling channels are statistically interleaved in TS 16. An IDA multiplexer is connected by a 2 Mbit/s digital path to a parent System X local exchange, where it terminates on the 2 Mbit/s IDA line module.

A further means of access from exchange areas not served by digital local exchanges is via the KiloStream network. By using two 64 kbit/s bearer circuits, and appropriate interworking equipment, a customer's 80 kbit/s IDA circuit can be carried to the parent ISDN exchange. This facility is seen as an exceptional expedient, and its use is expected to be very limited.

CALL TYPES

There are three main types of ISDN calls: the ISDN telephony call, the data call and the voice/data call. Each of these types is identified to the network by means of the SICs referred to earlier.

ISDN Telephony Call

The ISDN customer with a telephone can make telephony calls to other ISDN customers who have telephones, or to any other customer connected to the world-wide telephone network. The calling customer experiences the usual tones and announcements that an ordinary System X telephony customer would hear.

ISDN Data Call

When a customer requires a data call, the NTE sends the appropriate SIC to the network. This tells the network that an all-digital transmission path is absolutely necessary and that common-channel signalling must be employed throughout the routeing.

For a data call, the routeing digits from the NTE can be sent only in a single block. The network routes the call, checking as it goes for an all-digital routeing. The network passes the same SIC to the called NTE that the first NTE sent to the network. This will tell the called NTE the type of call being offered and, hence, the NTE can deduce which of its terminals would be able to accept the incoming call. If that terminal is free, then the terminal is alerted and the call can be answered.

At this stage, the network has provided a low-error-rate uninterruptable path, which is protected from intrusion by telephone operators. In contrast to the ISDN telephony call, only supervisory messages are sent to the customer and these are confined to the signalling channel.

In the rare situation where an analogue route is encountered on a data call set-up, the network clears down the call and sends back a message giving the reason why the call is being rejected.

ISDN Voice/Data Call

The calling customer indicates that a voice/data call is required. The network sets up the call as a voice connection but ensures, in a similar way to the straight data call, that an all-digital path is provided. If, during the call, the customers decide that they need to send a facsimile—a diagram, for example—then they will be able to switch over to the data and the appropriate terminals are connected to line. The customers are then free to change between the voice and data modes.

IMPLEMENTATION

ISDN Pilot

At an early point in the definition of the ISDN, it was decided that BT should obtain practical experience of an ISDN by providing a pilot service at the earliest possible date. The principal objectives of such a pilot were set as:

(*a*) the acquisition of experience on the technological, development and operational aspects of providing a national ISDN; and

(*b*) the stimulation of interest in the ISDN among customers, as well as the suppliers of terminal equipment capable of taking maximum advantage of the ISDN.

The pilot is to be based on four System X digital local exchanges: two in London and one each in Birmingham and Manchester.

The first unit to open for ISDN service will be Baynard 489 (in the City of London) early this year. The remaining pilot units, due to open over the following seven months, are Midland 631 (in Birmingham), Maida Vale 328 (in London) and Blackfriars 831/3 (in Manchester).

Each of the four exchanges, in addition to providing service to several thousand analogue customers, will be capable of providing digital service to about 250 single-line IDAs and 10 multi-line IDAs. Thus the pilot will comprise a total capacity of some 1000 single-line IDAs and 40 multi-line IDAs. Of these, only a few will be in the exchange areas of the four ISDN exchanges. The majority of customers with single-line IDA will be in some 50–60 exchange areas served by IDA multiplexers. In this way the pilot service, whilst based on only four ISDN exchanges, will involve customers spread geographically over much of London, Birmingham and Manchester as well as a number of other major centres of population and business activity in England. The geographical coverage planned for the pilot is illustrated in Fig 6. However, it will be possible to add further multiplexers to meet demand for single-line IDA identified as a result of IDA marketing. Multi-line IDA will be available, as required, by the provision of 2 Mbit/s digital paths between the customer and the most appropriate ISDN local exchange.

ISDN Demonstration

In order to provide potential customers, including terminal manufacturers, with an early indication of the impact that the ISDN will have on the future of telecommunications, as a prelude to their participation in the pilot service, a small demonstration ISDN has been created in London. This is based on a single DSSS effectively running as a remote concentrator unit (RCU) and serving only ISDN traffic. Normally, some 20 single-line IDAs are in use, serving various modern terminals typical of those capable of taking advantage of the capabilities offered by the ISDN. Fig. 7 shows the demonstration room and some of these terminals in use. The majority of the IDAs are within the same building as the exchange equipment. However, a few make use of pairs in the ordinary local cable network to connect facilities in other parts of London approximately 1·5 km away, and others make use of KiloStream circuits to extend single-line IDA to other facilities in Martlesham Heath some 100 km distant.

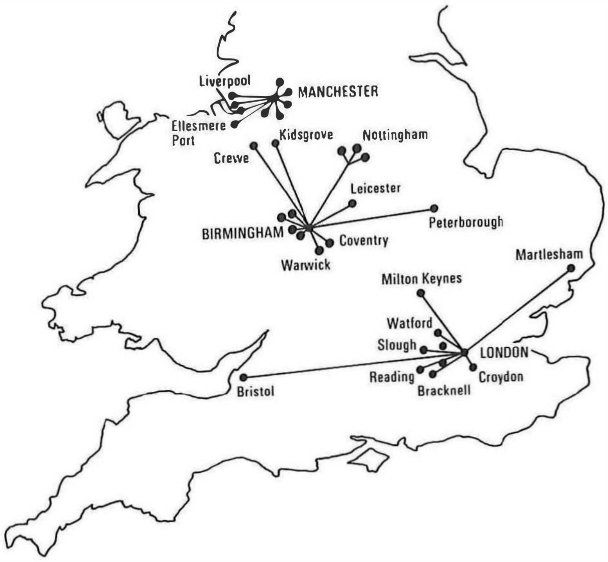

FIG. 6—Pilot ISDN—geographical spread of single-line IDA

In addition to its use as a marketing demonstration, the same equipment has from time to time been augmented with further single-line IDAs—some directly connected, some connected via multiplexers—to provide working exhibition displays in London's Guildhall, Birmingham, Wembley, Brighton and as part of the Martlesham '84 exhibition. Advantage has also been taken of the availability of a working ISDN to carry out tests, in particular to verify the transmission planning standards for the provision of service over the local cable network. During the first nine months of operation, some 1000 representatives of major customers have visited the model, as have a number of senior members of foreign administrations from all over the world.

National ISDN

The pilot ISDN is not a field trial. It is the nucleus on which a national ISDN will grow. All the digital local exchanges being purchased by BT have the inherent capability to

FIG. 7—ISDN demonstration

terminate both single-line IDA and multi-line IDA. All the digital local and trunk exchanges have the inherent capability of providing those few additional network features that are necessary to realise the full benefits of an ISDN. A decision to include IDA line modules in the normal procedures for planning, dimensioning and ordering local exchanges will come only after marketing and operational experience has been obtained from the pilot. Nevertheless, IDA line module equipment, IDA multiplexers and NTEs will be purchased centrally by BT Headquarters, and procedures are being established by which each Area/District will be able to provide IDA services on the basis of their market forecast, making use of these central stocks of equipment as necessary.

STANDARDS

Considerable progress has been achieved on the standards for the ISDN during the 1980–1984 Plenary of the CCITT, and a new series of Recommendations, the I-series, was ratified at the Plenary Assembly meeting in October 1984. The most important of these Recommendations relate to the definition of the access structure and a new customer interface to the network.

In particular, agreement has been reached on CCITT Recommendation I412, which describes a structure for single-line IDA. This has a total information rate of 144 kbit/s comprising two 64 kbit/s traffic channels, either of which may be used for data or voice calls, and a 16 kbit/s signalling channel which may be used also for packet access.

Other CCITT Recommendations that were agreed in October 1984 relate to the customer-to-network interface (CCITT Recommendation I420). However, these recommendations are still insufficient to allow independent implementation of compatible systems. Nevertheless, there are major international initiatives being taken in an attempt to achieve rapidly a specification in sufficient detail for this purpose, and there are good prospects of major advances during the coming months. BT is playing a full part in this activity and is looking towards the implementation of the resultant international agreements at the earliest practical date. Their introduction offers the opportunity of making economies in terminal and NTE design, at the same time creating an opportunity for common designs of terminal equipment for connection to the ISDNs of all those telecommunication administrations implementing the standards.

Despite the air of optimism and the undoubted advantage of achieving common international standards, it is not BT's intention to delay offering the advanced facilities of the ISDN to its customers. If agreements are not reached in the timescales required to meet BT's plans, then those plans will rely on the interim standards described above, international standards being implemented at the earliest opportunity. Whatever happens, BT will continue to support the interim standards for as long as equipment designed to those standards remains connected to the network.

CONCLUSIONS

The ISDN provides the switched network capable of up to and, with the use of more than one traffic channel and appropriate terminal design, in excess of 64 kbit/s. The use made of that capability to provide data and/or voice services is determined by the customer in choosing the terminal equipment to be connected to the ISDN.

It is easy to limit the appreciation of the ISDN to those terminals currently available and being demonstrated on the demonstration ISDN. As the ISDN becomes a reality, so the terminal manufacturers are recognising the importance of providing an interface to their equipment which is aimed at its use on the ISDN. In addition, ideas will emerge, indeed have emerged already, on future possibilities for further terminals and services that depend on the facilities offered by the ISDN.

What has been described above relates principally to the first stage of ISDN implementation in the BT network—the pilot ISDN—the first ISDN implementation of any size by any telecommunications administration, anywhere in the world.

Planning and development work is already proceeding with the objective of providing a more extensive coverage of the ISDN and, at the same time, implementing internationally agreed standards, as and when they become feasible.

There is no doubt that, in the ISDN, the beginning of a major revolution in public telecommunications networking capabilities can be witnessed, with BT in the vanguard of the technology that makes it possible.

References

[1] SHURROCK, C. J. R., and DAVIS, E. Inland Network Strategy. *Br. Telecommun. Eng.*, Oct. 1982, **1**, p. 143.

[2] LISLE, P. H., and WEDLAKE, J. O. Data Services and the ISDN. *ibid*, July 1984, **3**, p. 79.

[3] BIRD, J. R., BYLANSKI, P., and TRITTON, J. A. Aspects of Transmission and Distribution in the Local Network, Telecommunications Transmission—into the Digital Era. IEE Conference Publication No. 193.

[4] ADAMS, P. F., GLENN, P. J., and WOOLHOUSE, S. P. Echo Cancellation Applied to WAL2 Digital Transmission in the Local Network, Telecommunications Transmission—into the Digital Era. IEE Conference Publication No. 193.

[5] BOULTER, R. A. A 60 kbit/s Data Modem for Use Over Physical Pairs. International Zurich Seminar on Digital Connections, March 1974.

[6] BRITISH TELECOMMUNICATIONS NETWORK REQUIREMENT Digital Access Signalling System, DASS No. 1. BTNR186, 1983.

[7] ISO Data Processing—Open Systems Interconnection—Basic Reference Model. ISO TC97/SC16/537.

[8] BRITISH TELECOMMUNICATIONS NETWORK REQUIREMENT Digital Private Network Signalling System, DPNSS. BTNR188, 1984.

[9] BRITISH TELECOMMUNICATIONS NETWORK REQUIREMENT Digital Access Signalling System, DASS No. 2. BTNR190, 1984.

[10] WARD, R. C. System X: Digital Subscriber Switching Subsystem. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 241.

Biographies

Chris Price is a Head of Section in the Works Support Division of the Local Exchange Systems Department of BT. He joined the then Post Office as a probationary Assistant Engineer in 1958 and, after a period of initial training, attended a sandwich degree course on a Post Office scholarship. For most of the period from 1963–1979, he worked on the development of electronic exchange systems, initially the forerunner of TXE4, latterly System X. From 1979, he has been responsible for testing and direct labour works standards for System X and, more recently, for TXE4. He assumed responsibility for the implementation of the ISDN in March 1983.

Richard Boulter joined the then Post Office in 1963 as a Student Apprentice. After graduating from Birmingham University in 1967 with an Honours Degree in Electronics and Electrical Engineering, he joined a data transmission research group at the Post Office Research Station, London, and investigated new transmission techniques for use in the local cable network. In 1974, he was promoted to Head of Group with the task of looking at the problems of network synchronisation and was subsequently responsible for specifying the synchronisation equipment for the UK digital network. After further work on line-card developments for digital exchanges and a six-month consultancy period with SHAPE Technical Centre, he was promoted in 1981 to Head of Section concerned with systems evolution. His main responsibilities at present are to make proposals for the future development of BT's ISDN.

The Future Network

D. W. BROWN, B.SC., M.I.E.E.†, and E. J. LIDBETTER, M.B.C.S.*

UDC 621.39 "313"

This article discusses the new political environment for communications operations in the UK, the rapid progress in technological development now taking place, and the ways in which these factors may influence the provision of communications services in the future.

INTRODUCTION

British Telecom (BT) is currently passing through a period of rapid change. Not only does the modernisation of BT's networks permit the introduction of new services and improve the quality of service given to the public, but the loss of BT's monopoly to provide telecommunications services sharpens the need to offer what the customer wants both at work and in the home. Present trends indicate that advances in computer-based technology will continue at the same pace achieved over the last decade, and that their application to both work and leisure activities will be one of the most important factors shaping the networks of the future. For BT to maintain the lead it presently has over its rivals calls for an awareness of the changed environment in which BT now operates. Innovation within this environment will be the backbone of organisations' long-term growth. Scientists and engineers have the ability to influence this perhaps more than any other section of the workforce; how successful they are will greatly influence how future services and networks will develop.

ENVIRONMENTAL FACTORS

BT has always had well-established demand models to extrapolate historical demand for telephony in regard to general economic circumstances and the practical realities of housing and industrial development. While new service demands have always proved difficult to assess and as engineers we would simply ask 'what service, how big a demand and where', the environment now formed in the UK for the development of information technology creates new opportunities and uncertainties. The UK public switched telephone network (PSTN) currently operated wholly by BT is the major earner of BT's network revenues. Two key factors that affect the evolution of this network are:

(*a*) the requirement for BT's networks to co-exist with other licensed network operators in a competitive environment for communications networks and network services, and

(*b*) the liberalised supply of attachments to the network.

Each of these factors is strengthened by the technological progress in communication system development. It is this gamut of technical possibilities that the UK wishes to exploit by the creation of a more flexible structure for the provision of innovative services.

The regime created by the new licensing arrangements already comprises two national network operators—BT and Mercury Communications Ltd. (MCL), and two cellular radio operators, together with other mobile radio and value-added network operators; and the numbers will continue to grow. Thus, progressively, communications services will be provided by means of the interconnection of separately operated and often very different networks.

† Local Network Strategy Department, British Telecom Local Communications Services

* Trunk Services Planning and Works, British Telecom National Networks

TECHNOLOGICAL CHANGE

The rate of change of technology and the market place makes it conceivable that the future network structures will not build on existing network structures. For example, the innovative application of technology to mobile telephony services has resulted in the recent new cellular radio systems, currently being installed in the UK. Their existence, unthought of a few years ago, should make one consider the future of traditional solutions for the local network, which today represents the bulk of telephone network investment. Who can say that a low-cost radio access system will not become the future local network, given suitable cell sizes and available bandwidth?

Any latent communications opportunities to meet the ever-increasing demands of both commerce and society will surface from innovation in the design of communication systems, given the continued advancement in micro-electronics and information processing.

The vastly increased quantity of complex software necessary to encompass the new applications made possible with basic technological change will call for the software generation process to be simplified and increasingly mechanised if the timescales for introducing new products are not to be prejudiced. The cost of major developments has far-reaching implications for those involved in the manufacture of telecommunication systems, who increasingly will need access to world markets to recoup these product development costs, and for the network operators, who will wish to minimise their share.

NEW SERVICES

Because of the risk element with any new service, market judgement, albeit supported by research, will be an essential prerequisite before venture capital is committed. BT's local network is a key factor, as most new services are constrained by the economics of providing access to customers' premises. (Cable television is a good example of a service demanding a major re-engineering of the local network.) However, the historical example of the growth of telephony networks stimulated by the subsidy of local access from call revenues is unlikely to be repeated, and there are considerable difficulties in establishing major new investment in support of a new service capability outside the business sector.

At higher levels in the network, the aggregation of demands has always justified the investment in technology to obtain efficient use of plant. At the periphery of the network, the low utility of the customer-dedicated portion of the network has demanded a low-cost and essentially low-technology solution for access. Without radical changes in technology or service demands, the greatest scope is to increase the utility of existing investment at marginal cost.

NETWORK SERVICE PLANNING

The people who plan BT's networks should always be on the look-out to exploit a gap for an untapped market demand.

A range of special voice services that require universally available common access codes, together with charging arrangements to enable the charge for calls to be debited to the customer for which the call is destined, is one typical example. These types of services already enjoy considerable commercial success in the USA and Europe, but are presently provided only by operator intervention (that is, Freefone) in the UK. Without some short-term expedient, it would not have been possible to introduce these services until the early-1990s. The solution is a derived services network (DSN), which will form an overlay within the existing PSTN; it is expected to be in service this year. The DSN will carry mainly trunk traffic, and will not compete with similar services being provided within local fee areas.

For the future, major growth is foreseen in data, text, visual and mobile services. The particular nature of the service requirement, the density and location of traffic sources and destinations, activity factors associated with the information transfer and, moreover, association with other service demands will determine the economic solution. The test of the solution will be whether it is commercially viable.

The range of basic bearer services spans private circuits, packet-switched data, circuit-switched voice, circuit-switched data either as voice or data, and local area networks (LANs). As basic information transfer elements, each will be used to satisfy particular applications, or groups of applications, as part of both a network and service hierarchy. Effective use of these various capabilities is dependent on the ability to interconnect and enhance them to form the solution to a customer's application needs; this requires joint consideration of both terminal and network roles.

STANDARDS

No network can be accessed without a suitable terminal or customer apparatus, whether it be a simple telephone or a more complex digital PBX/terminal. The separation of terminal provision from the network places increasing importance on the use of agreed standards for such apparatus. With interconnected networks, defined interfaces and network protocols are required to support the development of networks with multi-service potential. The network operator(s) will increasingly need to educate the public on how to obtain the most suitable equipment, and on how to use the available facilities in a way which does not degrade services. In particular, recognition must be given to the potentially detrimental effect of unsuitable apparatus and network interconnections. Incidentally, this terminal apparatus is the source of further competition which affects the network engineer; that is, the competition which exists between providing services at a common point in the network or in the customer's terminal equipment.

The standards applied must aim to offer a wide range of options for integrating different applications and technologies to give easy and rapid communications. The most commercially acceptable approach to inter-network standards is to reach appropriate agreement between the network parties involved, lest some less than commercial standard be imposed. The future growth of communications, however, will depend on the emergence of commercially valid standards accepted across the world, and international fora will continue to play an important part.

PLANNING OBJECTIVES

Whatever technology and service innovation within the market place can together offer customers, two major factors must predominate in the minds of major operators: the need to achieve satisfactory levels of quality of service, and the need to achieve economy of operation. These are the only effective tools of competition.

NETWORK CONTROL

Effective control of network resources will become of increasing importance to deliver high-availability communications, advanced service features, and economy. This control of resources requires powerful signalling and routeing capabilities.

Signalling

Stored-program control (SPC) of switching and networks enables the signalling logic for a large number of information circuits (for example, speech) to be concentrated to reduce costs. Also, a much more efficient way of transferring information between SPC exchanges is to provide a bi-directional high-speed data link between the two processors, over which signals are sent in a digital form by means of coded-bit fields.

Thus, many messages controlling many hundreds of circuits can share a common-channel signalling (CCS) link in time-shared mode. This form of signalling will play a major role in future communications networks at both national and international level.

The CCITT No. 7 common-channel signalling system, which is about to be introduced, has been designed to carry signalling information for both voice and non-voice services, and is a fundamental step in the evolution towards a future integrated services digital network (ISDN). The main features of the signalling system are as follows:

(*a*) it is optimised for the digital telecommunications network in conjunction with digital SPC exchanges;
(*b*) it is designed to meet present and future information-transfer requirements for call control, remotc-control network management and maintenance;
(*c*) it provides a reliable means for the transfer of information in the correct sequence without loss or duplication;
(*d*) it is suitable for operations over analogue channels below 64 kbit/s (for example, 4800 bit/s); and
(*e*) it is suitable for use on point-to-point terrestrial and satellite links.

As the signalling is divorced from the switching and transmission paths, signalling information can be transferred between nodes in a number of different ways, which will call for separate signalling networks to be designed to give the necessary security. The combination of SPC and CCS will enable networks to be controlled, exploited and managed to a much greater extent than previously attainable.

Many administrations are already actively engaged in studying the extent and desired features of centralised services; typical of these are:

(*a*) network management of voice services to initiate temporary changes of traffic routeing patterns to deal with congestion, catastrophic failure etc;
(*b*) signalling network management;
(*c*) remote maintenance of networks; and
(*d*) centralised call accounting.

Other CCS systems which provide powerful error detection and correction functions will be introduced. One example, known as the *digital access signalling system* (DASS), is being designed for communication between the customer and the exchange; it should be suitable for both single-line and multi-line integrated digital access (IDA), and the many services which they may support. DASS is a message-based CCS system using variable length frames, which can be transferred over the 8 kbit/s signalling channel provided in each single-line IDA, as well as the 64 kbit/s signalling channel (time-slot 16) provided in each multi-line IDA.

New services, both voice and non-voice, will become economically viable as technical innovation and the level of intelligence in BT's network grow. Signalling systems will need to keep abreast of these developments to ensure that the full potential of multi-purpose networks is realised.

Network Routeing

During the introduction of the digital trunk network, automatic alternative routeing (AAR) will be used only on a limited basis. It will be used where dual digital main switching units (DMSUs) are required and in the routeing of traffic from the Channel Islands and the Irish Republic. BT is not looking to establish a trunk network using both AAR and automatic re-routeing (ARR) in the short term. Although the facilities are now becoming available with network modernisation, BT does not plan to use them until there is a convincing case that the network implications can be handled, and the measures necessary for managing and augmenting the facilities are available and tried.

In the longer term, BT needs to establish a trunk network which is flexible when diverse and unforeseen requirements are placed upon it. This could be achieved by planned over-provision on all major highways; the resulting increase in costs would have to be offset by gains in network quality and resilience.

The other way to achieve network flexibility is reconfiguration of the network to meet the new demand from the use of temporary spare capacity. This can be achieved by exercising control over the routeing of traffic. The intelligence necessary to effect the change can be spread throughout the network on a distributed basis (that is, switchable digital distribution frames) or centralised within the trunk switching units (that is, dynamic routeing). In the latter case, the trunk switching units would have to be supplied with information relating to the total state of the network and then would have to route the traffic accordingly. Both methods would need network management information which is not presently available.

A final policy is not expected to emerge until BT has field experience of the performance of digital switches and their interconnecting digital transmission systems.

Network Traffic Management

Computer-based information systems now hold out the prospect of obtaining a real-time picture of traffic flows across the network. Networks are often subject to overload arising from disasters, through to phone-in programmes, and unusual calling patterns on festival days such as Christmas. Congestion can also occur from network plant failure. With an immediate picture of the network, action could, in principle, be taken to reduce demand selectively or to re-route traffic on spare capacity, the intention being to maximise call completion on the overall network. While only a manual system is being experimentally developed for application within the integrated digital network (IDN), computer analysis will assist the identification of possible derived courses of action, which will be largely empirical, to be taken in particular circumstances. One day, a knowledge-based, 'expert', system hosted on a computer may be developed to initiate automatic action which could not be built into nodal routeing features such as AAR.

NEW NETWORK ELEMENTS

Integrated Services Digital Network

Customers will increasingly wish to view their service needs as being satisfied by a common premise entry point. Local network elements of the future include the extension of digital circuit switched channels to customers' premises which, with appropriate network terminating equipment and terminal apparatus, will form an ISDN available to all who are connected to a digital exchange. The ISDN pilot service provides access at 80 kbit/s (64 + 8 + 8) and 2 Mbit/s (32 × 64 kbit/s). The 80 kbit/s is split into two service channels and one signalling channel. These channels are compatible with those which could be provided within a higher-bit-rate transmission system (for example, at 144 kbit/s), which may emerge as one world standard. The prospect of optical-fibre capacity and optical components being available at low cost, as well as signal processing based increasingly on very-large-scale integration, (VLSI) makes it particularly difficult to predict the future local network architecture and the level of integration of broad and narrow bandwidth services on common plant. The local network will certainly comprise interfaces and gateways to private or semi-private networks, be they PABXs, LANs or metropolitan area networks (MANs). It is likely, however, that the use of common high-capacity transmission bearers will be of increasing importance in obtaining an economic solution. One approach is to provide an optical-fibre equivalent of the wire pair network with flexibility points that optically multiplex channels onto main optical-fibre cables. The switching of channels and the provision of advanced services could be at a point associated with a principal local exchange or the head-end of a cable television network. The processing of traffic is likely to be further concentrated in fewer nodal points than there are currently local exchanges, with the prospect of an enlarging and infinitely more powerful local network providing local communications, information services and access to major network bearer services.

Local Area Networks

Offering a general purpose communications medium through which a number of different devices can talk to each other, LANs are simple in concept, effective and flexible. All LANs are based on essentially the same principle; that is, they use addressed messages broadcast over a common path using time sharing or frequency multiplexing, rather than dedicated switched paths between devices. The technique has particular value in coping with bursty-type communication between computers. LANs are fundamentally different from any large switched network, where the LAN principle generally becomes uneconomic compared with a switched system, either of the circuit-switched or packet-switched type. How large is large depends on a number of factors, and private so-called *wide area* and *metropolitan area* networks are now emerging.

Cable Television

The current development of new cable television networks in eleven franchise areas is the result of government initiative to provide for the development of cable television across the UK. The least-cost tree networks have limited evolutionary potential, but may predominate until such a time as switched-star wideband networks become lower in cost or can be demonstrated to satisfy a wider user need for interactive services. Cable television networks are predominately local access networks and, unless the networks attract sufficient share of the revenue from the services offered, growth will be slow, as was once the case for telephony networks, particularly with competition from video cassette recorders and direct broadcast by satellite. An evolutionary path might, however, anticipate the emergence of high-capacity links to business premises and homes.

Wideband

For the transmission and reception of television-quality pictures, video telephone and high-speed data, a broadband network is required. By using the narrow-band network control concepts as the basis for building a supplementary wideband ISDN, additional services could be offered. Optical fibres and communications satellites, together with variable-bit-rate switching nodes, could form the basis of this network. In the case of business customers, these circuits could be connected to a new breed of PABX so that the

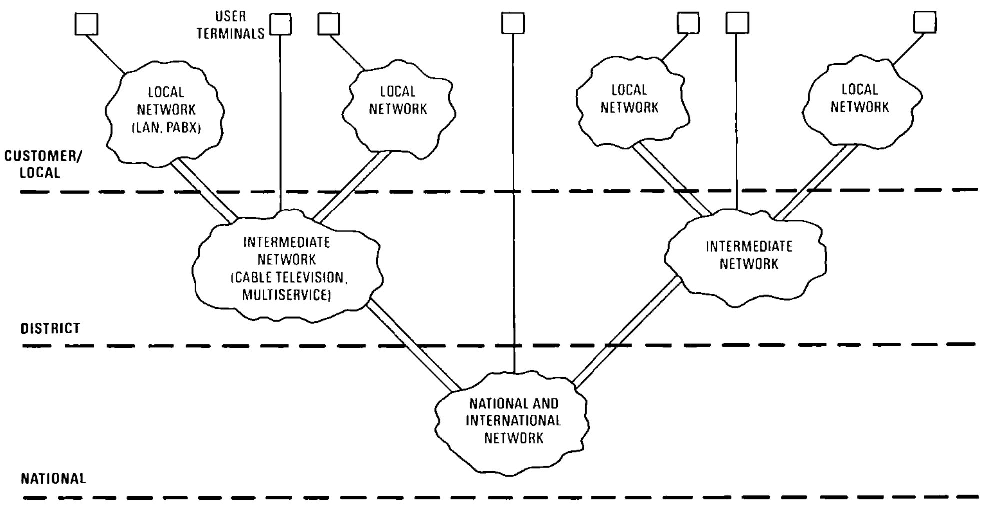

FIG. 1—Future network structures

bandwidth can be used for speech, data and video. Introduction of wideband services could lead to a revolution in home entertainment, leading to interactive communications using the television screen. It could enable an economic version of a viewphone service to be introduced at some point.

Cellular Radio

World-wide the demand for mobile radio telephony exceeds supply. The reason for the shortfall in supply is mainly the limitations in available radio spectrum. There are several solutions to the problem, of which cellular radio is the most promising. The advent of new technology—that is, large-scale circuit integration in electronics—has also made possible the concept of cellular radio.

Commercial operations of cellular radio are beginning in the USA. In Western Europe, the world's first international service has started in the four countries in Scandinavia, and other European countries are due to follow with large-scale nationwide commercial operations beginning in this year.

Cellular radio within the UK will be provided initially by two companies, Racal and TSCR, using the marketing product names of *Vodafone* and *Cellnet*, respectively. These companies are operating under the Government's liberalisation policy of the communications market, which permits private competition in the supply of equipment and services. The service within the UK, which was successfully trialled at the World Economic Conference in London last summer, is expected to start early this year. By the end of this year, cellular radio will be available in at least six major cities. Total cell coverage will exceed 25 000 km^2 and is expected to take in nearly 50% of the population, serving in the region of 10 000 customers. By the end of 1987, it is projected that the number of customers in the London area will be 30 000 and the number of customers nationwide could reach 80 000.

At the end of 1989, cellular radio will offer service to 90% of the population, and the total number of customers could be in excess of a quarter of a million. The figures for 1989 are considered by some sources as conservative. Racal has projected 250 000–300 000 customers on its network by 1989. Other information from those countries who are already operating cellular radio, Scandinavia, the USA and Hong Kong, already indicates that demand far out-paces original market projections.

As the cellular radio network grows, so will the interconnect links with BT's PSTN, affecting both traffic and routeing with the PSTN. It is therefore beneficial that cellular-radio operators and BT co-operate actively in system planning and network building to provide a profitable network.

THE FUTURE NETWORK

As telecommunications networks become more advanced, the use of new developments in technology will have a significant impact on the requirements and possibilities for mobile and fixed network facilities. Cellular radio will not be confined only to voice: in the future, digital information paths will also be established via mobile cellular radio. Greater integration of line, switch and radio communications networks will be required in order to allow development of advanced mobile communications. The line-based networks could well be developed to a higher level of integration than is currently envisaged, with full sharing of the line resources by dynamic allocation of capacity between a mix of services; a true multiservice network.

The generic hierarchical structure will be of the form shown in Fig. 1, which brings together all the major network elements now conceived. Given the pace of technological development and uncertainty mentioned earlier, only time will test which elements will predominate. It is certain, however, that innovative engineering will play a key part in the continued development of communications.

Biographies

David Brown is Head of the Network Planning and Strategy Division in BT's Local Network Strategy Department. After post-graduate experience in the manufacturing of semiconductor devices with Texas Instruments and GEC Semiconductors, he joined the Post Office Research Department at Dollis Hill in 1971 to work on the fabrication of highly reliable 10A- and 40-type transistors for submarine systems. Ultimately, he was responsible for service quality assurance and reliability assessment. In 1979, he joined the team evaluating modernisation strategy, with particular interest in financial evaluation. He is now responsible for the network master plan, network planning projects including traffic management, and the development and evaluation of local network strategy.

John Lidbetter is Head of the Traffic Forecasting and Digital Switching Planning Section in National Networks Trunk Services. After a period in a Telephone Area, he joined the Engineer-in-Chief's Office in 1957 to work initially on the development of customers' PMBXs and, later, PABXs. He transferred to the exchange computer design project in 1969, which pioneered the use of computers in BT for exchange planning and design. In 1976, he became Head of Section in Network Planning Department, and was responsible for the installation and works aspects of trunk exchanges, including the implementation of switching equipment for the maintenance and analysis centre (MAC) project and the UK radiopaging nctwork. He became responsible for his present section when the trunk planning switching and transmission functions were combined to form a single division, on the etablishment of National Networks in 1982.

British Telecommunications Engineering

Those wishing to subscribe to *British Telecommunications Engineering* can do so by completing the relevant section of the order form below. British Telecom (BT) and British Post Office (BPO) staff should complete the upper section and send it to their local *BTE Journal* agent or, in case of doubts, to the address shown; other subscribers should complete the lower section. A photocopy of this form is acceptable.

SUBSCRIPTION ORDER (BT and BPO STAFF)

To: The Local Agent *(BTE Journal)*, or
British Telecommunications Engineering Journal (Sales), Post Room,
2–12 Gresham Street, London EC2V 7AG.

Please arrange for me to receive *British Telecommunications Engineering.* I have completed the form opposite authorising deductions from pay to cover the special annual subscription of £2·04.

Date .... Signature ....

Name and initials (in block capitals) .... Rank ....

Official address ....

....

....

....

British Telecommunications Engineering: Authority for Deduction from Pay

I, ...

(name in full, surname first in block capitals), authorise the deduction from my salary or wages, until further notice, of the sum of 17p per month/4p per week for payment to the Managing Editor, *British Telecommunications Engineering,* on my behalf.

Date .... Signature ....

Official address ....

....

Rank .... Pay No. ....

Area/Dept. ....

Pay group reference ....

YEARLY SUBSCRIPTION ORDER (NON-BT/BPO STAFF)

To: *British Telecommunications Engineering Journal* (Sales), Post Room,
2–12 Gresham Street, London EC2V 7AG.

Please supply 4 quarterly issues of *British Telecommunications Engineering.* I enclose a cheque/postal order for the sum of £6·00 (Canada and the USA: $10) to cover the yearly subscription.

(Cheques and postal orders, payable to **'BTE Journal'**, should be crossed '& Co.' and enclosed with the order. Cash should not be sent through the post.)

Name ....

Address ....

....

....

Please state with which issue you wish your subscription to commence (April, July, October or January) ....

INDEX TO ADVERTISERS

Advertisements

Communications should be addressed to Mr A. J. Pritchard, The Advertisement Manager, *BTE Journal*, Room 608, 2-12 Gresham Street, London, EC2V 7AG.

No responsibility is accepted by the Journal for any of the private or trade advertisements included in this publication.

British Telecommunications Engineering

Copyright

Subscriptions and Back Numbers

The *Journal* is published quarterly in April, July, October and January, at £1·00 per copy (£1·50 per copy including postage and packaging); annual subscription: £6·00; Canada and the USA: $10·00.

The price to British Telecom and British Post Office staff is 51p per copy.

Back numbers will be supplied if available, price £1·00 (£1·50 including postage and packaging). At present, copies are available of all issues from April 1974 to date with the exception of the April and October 1975; January, April and October 1976; and January and April 1982 issues.

Orders, by post only, should be addressed to *British Telecommunications Engineering Journal* (Sales), Post Room, 2–12 Gresham Street, London EC2V 7AG.

Employees of British Telecom and the British Post Office can obtain the *Journal* through local agents.

Binding

Readers can have their copies bound at a cost of £8·75, including return postage, by sending the complete set of parts, with a remittance, to Press Binders Ltd., 4 Iliffe Yard, London SE17 3QA.

Remittances

Remittances for all items (except binding) should be made payable to *'BTE Journal'* and should be crossed '& Co.'

European customers can pay by means of a direct bank mail transfer. Customers using this method of payment should instruct their bankers to remit payments to the *Journal's* bankers—Midland Bank plc, 2 Gresham Street, London, EC2V 7JD, England—specifying that the beneficiary is to be *British Telecommunications Engineering Journal*.

Distribution and Sales

Correspondence relating to the distribution and sales of the *Journal* should be addressed to *British Telecommunications Engineering Journal* (Sales), Post Room, 2–12 Gresham Street, London EC2V 7AG.

Communications

With the exceptions indicated, all communications should be addressed to the Editorial Office, *British Telecommunications Engineering*, NN/CMS 2.2, Room 704, Lutyens House, Finsbury Circus, London EC2M 7LY. (Telephone: 01-357 4313).

Advertisements

All enquiries relating to advertisement space reservations should be addressed to Mr. A. J. Pritchard, The Advertisement Manager, *British Telecommunications Engineering Journal*, Room 608, 2–12 Gresham Street, London EC2V 7AG.

Published by *British Telecommunications Engineering Journal*, 2–12 Gresham Street, London EC2V 7AG, and printed in Great Britain by Unwin Brothers Limited, The Gresham Press, Old Woking, Surrey

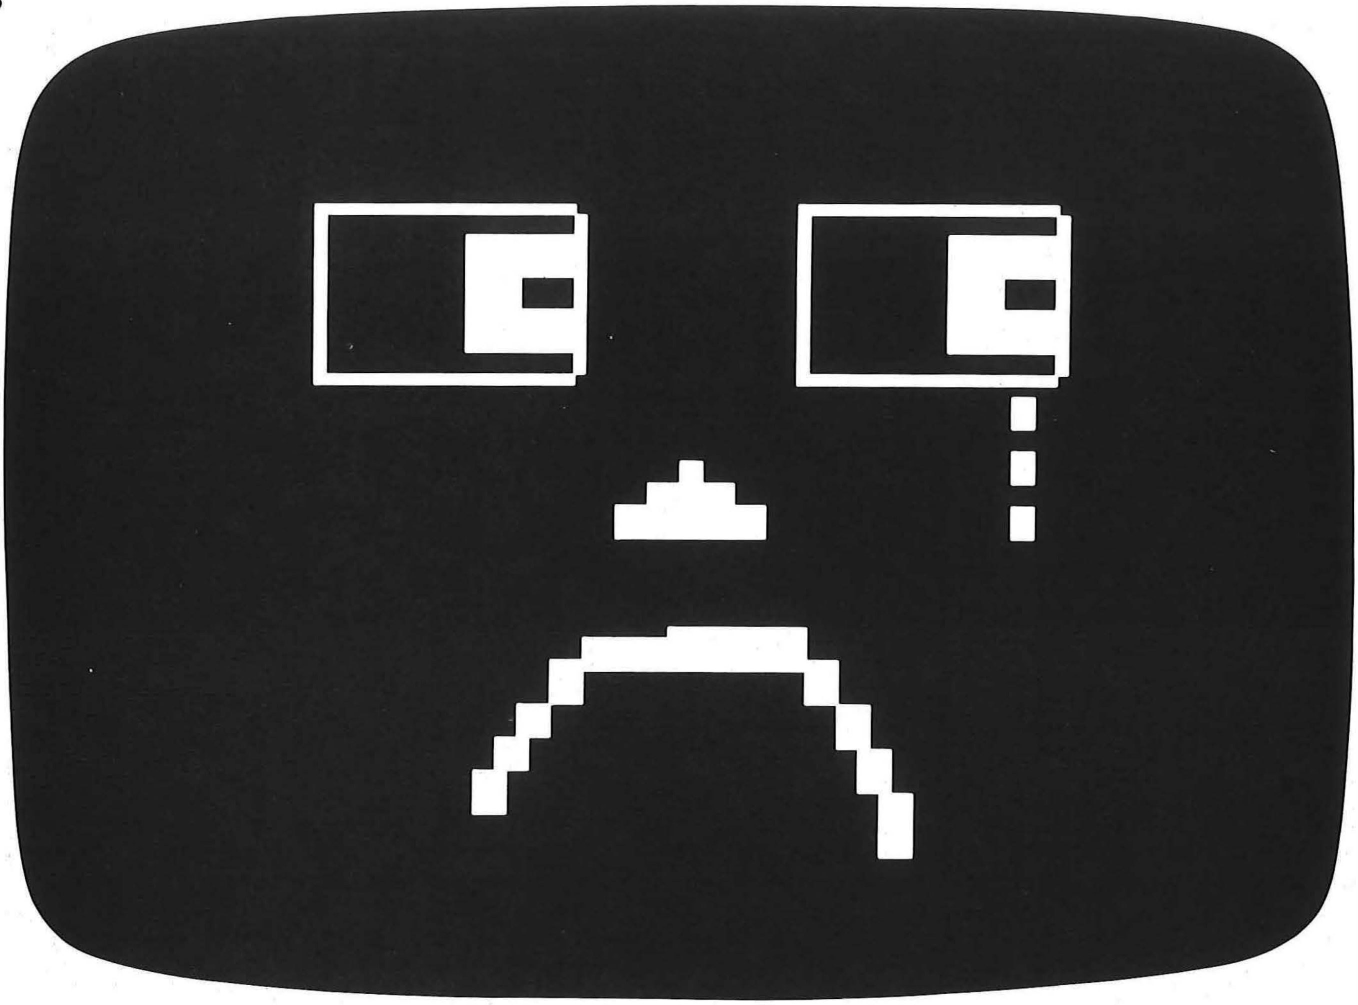

OTHER LINE-TESTERS ONLY GIVE YOU THE PROBLEMS.

The problem with some automatic line-test systems is that they're designed to do nothing but test lines.

Some will give you a rough idea of fault conditions, but some only display test results. Not the LRS-100. It goes quite a few steps further.

As a complete test system it is, of course, able to perform all of the standard functions you'd expect it to.

Such as demand-testing, when faults are reported. Routine-testing groups of lines overnight, to locate potential problems.

And it'll even carry out follow-up tests ('Robot Testing') on problem lines at regular intervals, to fmd intermittent faults.

But the difference is in what the LRS does with all the information after it's been collected.

For example, it cross-references reports, to build up patterns and recognise common faults.

Also, the LRS compares the condition of the line it's testing with other available information, and produces a System Recommended Action.

And in its full configuration LRS-100 will even keep an exact record of the total workforce available and its current workload, and assign each repair (according to priority) to the appropriate faultsman.

It can carry out the whole operation, from line-testing to assigning the repair, so quickly that you're able to make firm appointments with customers as and when they report faults.

This enables you to speed up clearing times and reduce fault report rates.

So your Repair Service Centre runs at

PAPST

SuperQuiet fans from Papst

Users of today's electronic equipment demand ever-quieter VDU's, Computers main-frames etc. To meet this need PAPST have introduced their new **SuperQuiet** range of fans.
By **SuperQuiet** we mean a noise level of less than 28 dB(A). Super quiet by any standards — for a fan, really excellent. This range is growing rapidly and the first six devices are:

4½" AC — The 4890N (220v) and 4840N (115v). both deliver 80m³h (48cfm) on 50Hz supplies. Noise level in free air 25 dB(A).
4½" DC — The 4112GXL (12v DC) and 4124GXL (24v DC) to deliver 85m³h (50cfm). Noise level in free air 28 dB(A).
3" AC — The 8850N (220v) and 8800N (115v). Both deliver 37m³h (22cfm) on 50Hz supplies. Noise level in free air 26 dB(A).

The AC fans are rated for a 20,000 hour life at up to 70°C; the DC fans at 35,000 hours at up to 65°C. A long **SuperQuiet** life.

SuperQuiet–the sound of silence

Distributors: Radio Resistor Ltd. Tel: (0234) 47188 Tlx: 826251
HB Electronics Ltd. Tel: (0204) 386361 Tlx: 63478
BA Electronics Ltd. Tel: (0454) 315824 Tlx: 449150

PAPST MOTORS LIMITED,
East Portway, Andover, Hants, SP10 3RT
Telephone: Andover (0264) 53655 Telex: 477512 Telefax: (0264) 59448 Group 3 Auto

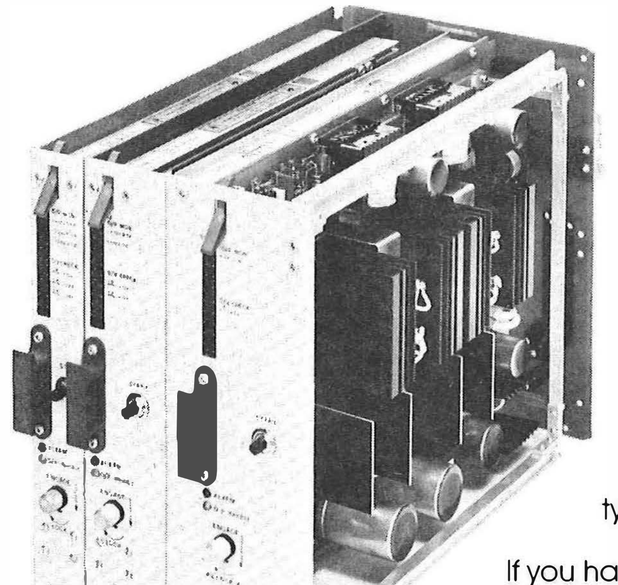

HIGH RELIABILITY POWER CONVERTERS

A.P.T. Electronics specialise in the design and manufacture of power converters for main, local and rural telephone exchanges, and for trunk, local and microwave transmission systems. These are switched mode power supplies in proven MOSFET technology with state-of-the-art performance, particularly in respect of conversion efficiency. Predicted mean time between failures (MTBF) typically exceeds 250,000 hours.

If you have an application where the reliability of your power source is a critical requirement, discuss your needs with the most experienced specialist company in the U.K.

A.P.T. ELECTRONICS LTD.
Darwin Close, Reading, Berkshire, RG2 0TB Telephone: Reading (0734) 862155 Telex: 848989

Manufacturers of equipment enclosures, metal fabrications and electrical assemblies: Fully approved by British Telecommunications. Ministry of Defence to DEF 05-24 British Standard B.S. 5750 and B.S. 9000 (Pending)

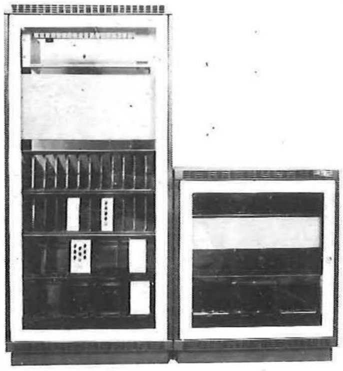

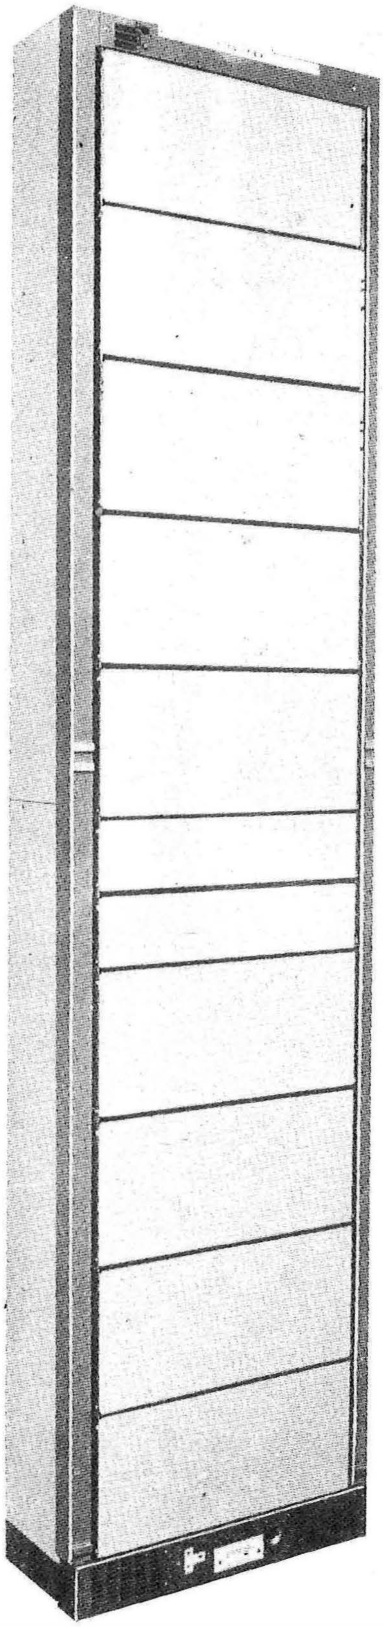

- ★ **3 and 7 Shelf Sizes**
- ★ **62 Type, TEP-1E, 19 inch Shelf Interchangeability**
- ★ **Earthing Bars**
- ★ **Fully Vented**

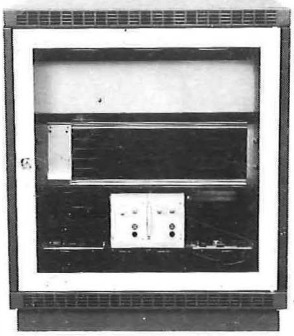

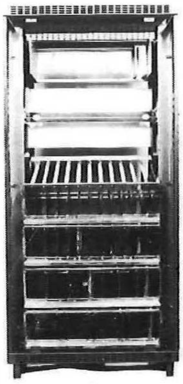

- ★ **Transparent or Solid Doors**
- ★ **Optional Rear Door**
- ★ **Choice of Colour**

HOWELLS RADIO LTD.

DENMARK ROAD,
MANCHESTER, M14 4GT.
TELE NO:- 061-226-3411

Approved for TEP-IE;* including racks, shelves and alarm units.

* TEP-IE is a registered trademark for British Telecommunications.