

BRITISH TELECOMMUNICATIONS ENGINEERING



Included in this Issue

BT Operational Support Systems

Architecture Framework

National Vocational Qualifications

for the Telecommunications

Industry



**The Journal of The Institution of
British Telecommunications Engineers**



BRITISH TELECOMMUNICATIONS ENGINEERING

Contents

VOL 15 ■ PART 2 ■ JULY 1996

Guest Editorial: Keeping Operations Ahead	113
Sinclair Stockman	
<hr/>	
<i>Operational Support Systems</i>	
The BT Operational Support Systems Architecture Framework	114
Nick Furley	
<hr/>	
<i>Making it Happen</i>	
Enhancement of the Digital Core Switched Network	122
Tony Grainger and Tony Poultney	
Cashless Services Network Operations Unit	128
Steve Ellett	
CallMinder™—The Development of BT's New Telephone Answering Service	135
Keith Beacham and Simon Barrington	
MeterLink™—BT's New Telemetry Service for Meter Reading	142
Kevin Welsby, Jeremy Adams and Jeff Deslandes	
<hr/>	
<i>Computer Telephony Integration</i>	
Computer Telephony Integration—The Meridian 1 PBX	150
Philip Johnson, Andrew Catchpole and Laurie Booton	
<hr/>	
<i>Process Management</i>	
Driving on Customer Satisfaction	157
A Re-engineering Case Study in the Retail Line of Business	
Allan Drew and Valerie Jones	
<hr/>	
<i>Telecommunications in the 21st Century</i>	
Shared Spaces	162
Laurence Bradley, Graham Walker and Andrew McGrath	
Smart Rooms, Desks and Clothes—Towards Seamlessly Networked Living	168
Alex Pentland, Rosalind Picard and Pattie Maes	
<hr/>	
The Face of Talking Machines in a Multimedia World	173
Andrew Breen, Julian Page and Bill Welsh	
A Model for the Project Management NVQ	179
Michael Tuggey	
National Vocational Qualifications for the Telecommunications Industry	186
Alan Martin	
<hr/>	
<i>Regular Features</i>	
pcochrane@btlabs	194
Book Reviews	195
BT News	197
Industry News	199
IBTE Contacts	Inside Back Cover

Cover picture:
BT's new-look public
payphone.
Picture by Eye Catchers Press
for BT.

Board of Editors

Chairman Bill Mills, M.A., C.ENG., M.I.E.E.

Managing Editor Paul Nichols, B.SC. **Secretary** Les Brand, M.ENG., M.B.A., D.I.C. **Treasurer** Denis Taylor, F.C.M.A.,

Assistant Editor Richard Kemp, B.SC., C.ENG., M.I.E.E.; Alan Bealby, M.SC., C.ENG., F.I.E.E.; Martin Evans, B.SC., PH.D.; Warwick Flury;

David Greenop, B.SC., M.SC.; Sam Quinn, B.A., DIP.M.; Murray Trantor, B.SC., C.ENG., F.I.E.E., D.M.S., M.I.MGT.; Keith Ward, C.ENG., F.I.E.E., M.I.MGT.

Theme Editors

Operational Support Systems Sinclair Stockman **Process Management** Lowry Stange **Making it Happen** Alan Bealby

Computer Telephony Integration Doug Chesterman **Telecommunications in the 21st Century** Peter Cochrane and Bill Whyte

Sinclair Stockman

Keeping Operations Ahead



BT's operational support systems (OSSs) are the set of computer systems which support the operation of BT's network, the provision of service to our customers and the overall operation of the BT business. The scope of areas supported range from marketing and sales activities, through to the configuration and delivery of services to our customers, and the maintenance and performance monitoring of these services and the underlying network. They are also used to support the planning and build activities of BT's networks both in the UK, and globally. The support systems also provide core business management support, including financial systems, personnel systems and supply management systems.

The scale of the operation supported by the OSSs is vast. The UK telecommunications network consists of in excess of 29 million connections. Each day 70 000 orders are taken, 32 000 repairs are performed, and 400 000 bills are despatched. On average, there is a major software delivery each day. The system set currently consists of 1300 applications contained in 970 systems. There are a large number of different combinations of hardware and systems software, and in excess of 3000 system interfaces. Physically, the systems are deployed at 20 major nodes across the country, with access

and local computing capability being available at numerous office locations.

The OSSs represent a significant investment for the company, and there is hardly any activity which does not depend upon them to some extent. They enable the company to realise improvements in efficiency across a wide range of activities, and help our people to provide excellent service to our customers.

In the rapidly changing environment in which BT operates, a number of drivers for change are impacting upon the company's operational support systems.

Customer demand for a wider range of services, coupled with increasing competition, both in the UK and globally, means we have to target at reducing the time to market for new services, while continuing to be able to provide integrated service support to all of our customers. A key challenge for the OSSs is to allow the company to achieve a totally customer centric view, while continuing to be able to exploit efficiently our underlying network. We also have to improve our ability to support our marketing and sales activities, by ensuring they have available the information necessary to make them fully effective.

There are a number of regulatory pressures which require changes to be made to our support systems, not least of which are number portability, and increasing levels of support for other licensed operators (OLOs).

As the number of systems has grown, so has the level of interconnection between these systems, often driven by the need to increase automation across the business processes. This, coupled with the need to achieve higher levels of flexibility, means that it is necessary to re-engineer a number of our major systems, to make data accessible to the set of applications which need to use it. Failure to do this, leads to

data duplication and inconsistency, and a reduction in the quality of service available to our customers.

In addition, the underlying network technologies provided by the company continue to grow, and the OSSs must be enhanced to support these technologies. Most recent examples are the deployment of asynchronous transfer mode (ATM), synchronous digital hierarchy (SDH) and the move from narrowband to broadband services.

The OSSs are a key source of competitive advantage for the BT group. The challenge is to ensure that they continue to be key enablers to the company's success. Our goal is to ensure that OSSs are never on the critical path for the introduction of new services. To achieve this, we have to ensure that the system set is highly flexible so that they can continue to meet the business requirements of the BT group.

This edition of the *Journal* contains the first of a series of articles (p. 114) which will place the OSSs in a clear business context, showing their relevance to the business processes and relating the technical challenges faced, and overcome. The articles will describe some of the success stories, as well as giving details of some of the problems and opportunities which lie ahead. The first article describes the operational support system architectural framework, which has been designed to underpin further developments.

I hope you find the articles both informing and interesting. Having been fortunate to have worked in this area for several years, I hope that future articles will convey some of the excitement felt by those working both on the development and the operation of the systems.

Sinclair Stockman

Sector Manager, Computing Systems
BT Networks and Systems

The BT Operational Support Systems Architecture Framework

This article, launching a new series in the Journal, outlines the process and systems framework that BT has developed to assist in the design, development and procurement of operational support systems.

Introduction

BT currently invests hundreds of millions of pounds each year in the development, procurement and maintenance of its operational support systems (OSSs). In order to ensure good value for that investment, the design community needs an OSS design framework or architecture which will help in the production of good designs. The TMN (telecommunications management network) is the architecture currently accepted by the industry for this purpose. However, TMN has focused primarily on interworking between operators and between systems and networks. It does not currently provide adequate help in the implementation of system solutions for network and service management internal to an operator.

In order to overcome this, BT has developed an improved process and systems framework to assist in the design, development and procurement of OSS solutions for the management of all of its services and networks. This framework is comprehensive in considering all processes and all functions and data within network and service management and can be considered to be complementary to TMN.

The Time For Change

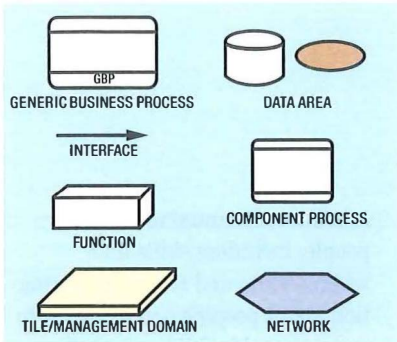
The TMN architecture, derived primarily from BT's cooperative networking architecture, was designed in the 1980s when telecommunications was very much simpler than it is today. The telecommunication environment is

becoming increasingly complex and demanding on network and service management with:

- rapid expansion of services;
- faster introduction of new services;
- increasing network complexity – overlay networks, and – multiple technologies;
- regulation;
- competition; and
- globalisation.

The problem is that almost every facet of this complexity impacts on the operators' OSSs. In order to deal effectively with this dynamic and complex environment, the design and development environment and the OSSs themselves need to be well structured and highly flexible, providing solutions in shorter time-scales at low cost. Often, a small-scale solution is required to suit the introduction of new niche services which may subsequently grow to match the market demand for the service and become integrated into the main operational processes of the operator.

TMN has concentrated on looking outwards covering, primarily, interworking between operators and interworking between OSSs and the network. Little has been done to address the design of OSS solutions required for the core of network and service management within an operator or service provider's company. It provides little detail of the



Key to symbols used in this article

functions and data at each of the layers and almost no model of the relationships between the functions and data. It also gave no view of the processes, although this area is now being addressed by the SMART (Service Management Automation and Re-engineering Team) initiative within the Network Management Forum.

The Legacy

This lack of a complete coherent design framework has resulted in poor solution designs which, over the years, have created a legacy in BT of:

- over 1000 different systems,
- islands of automation with little integration or flow through of processes,
- stove-pipe solutions for particular services or networks with high levels of duplication,
- poor data quality due to duplication and limited access to the data,
- multiple computing platforms increasing the problems of integration, and

Figure 2—Core components of Maintain and Restore Service generic business process

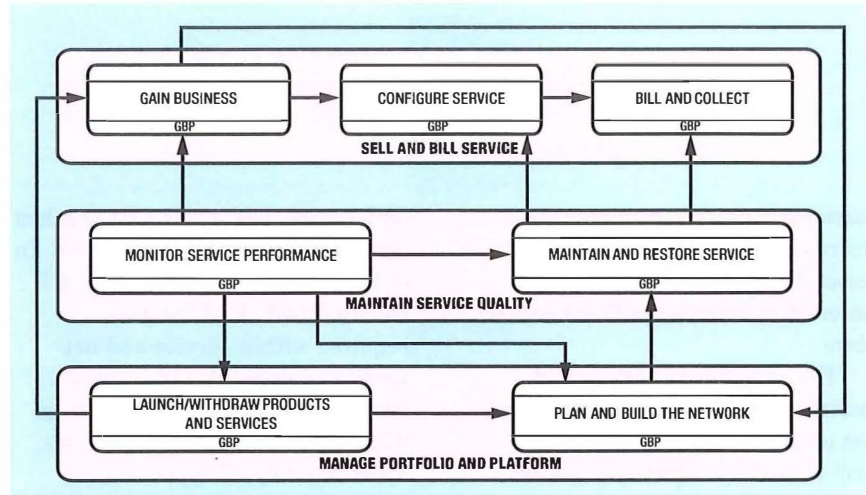
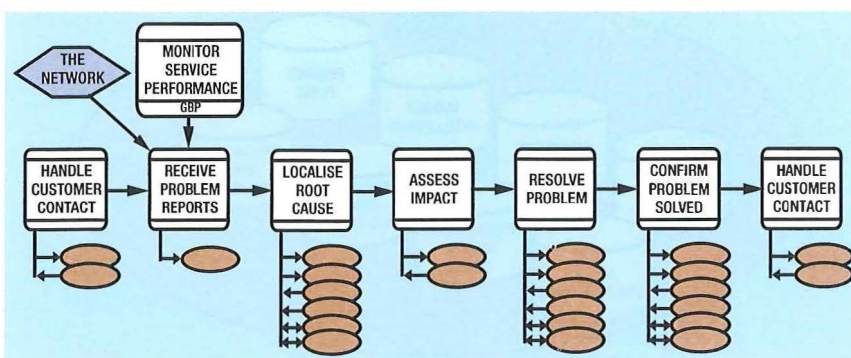


Figure 1—The process framework

- high cost of ownership in delivery and ongoing running costs.

Process Driven

In order to ensure that systems deliver the functionality that is required by the company, BT is committed to an approach where solutions should be driven from a process perspective with operational requirements being expressed in process terms. This framework therefore includes a set of fully defined business processes and all significant interactions between them. The objective is to use this process framework for the design and evolution of specific operational processes such that, with time, those operational processes become common and standardised across the business. In addition to facilitating common OSS solutions, it should also simplify the problems of operational organisation design and the integration of the management of new services and technologies into existing organisation structures. The

other aspects of the framework have been derived from this process model.

The process framework (Figure 1) comprises three very high-level processes which have been called *key business processes*:

- *Sell and Bill Service*—responsible for proactively selling services to customers and collecting the revenue for the use of those services;
- *Maintain Service Quality*—responsible for ensuring that the quality of each instance of service sold is within the service level agreement with the customer; and
- *Manage Portfolio and Platform*—responsible for managing the portfolio of products and services which the company sells and the network and OSS platforms on which those services are delivered.

Within each of these key processes are a number of lower level business processes, seven in total, each of which is called a *generic business process* (GBP), and within each of these are a set of interlinked *component processes*. These component processes themselves have been decomposed one further layer into what have been termed *actions*. Most importantly, these generic business processes also contain all of those items of data on which the component processes operate. This is considered to be a most vital aspect of the process framework.

As an example, Figure 2 shows a simplified view of one of these generic

business processes, including the more significant component processes, their interfaces and their interaction with the various items of data.


The component processes and actions were further analysed and a set of *functions* were defined which will deliver the requirements of the processes. The objective was to identify a set of reusable components of a scale and significance which could be considered to be the future building blocks of systems. The functions will be delivered by applications software in the operational support systems. Common functions will deliver parts of several different business processes.

Each component process within the business process operates on specific items of data and many of the component processes operate on the same data items. More importantly, component processes in different business processes operate on the same data items. This complete view of these relationships is vital in understanding the issues of dependencies between the systems which deliver these capabilities and the importance of making the data accessible. Historically, systems have been designed which implement specific processes or parts of processes together with the required data without making that data accessible to other systems. This has led to duplication of the data leading to inconsistency.

Data is Key

Data is now believed to be the single most important aspect of

Figure 3 – The six data groups



OSS design, but in the past this has been given insufficient attention. In order to address this problem, BT has analysed all of the data required within service and network management and created a structure within which to manage it. From analysis of the processes, approximately 35 types of data (called *data areas*) were identified, defined and grouped into six *data groups* (Figure 3):

- The *customers* group contains data about the customer, including name and address, role in a company, socio-economic group, the services provided, etc., together with a record of all of BT's interactions with the customer including current orders, problems and billing queries.
- The *bills* group includes customer accounts, invoices raised for use of services, information about billing problems, billing profiles and billing schedules which determine when invoices should be raised.
- The *jobs* group contains information about key activities which are managed by the job control functions including orders, trouble tickets and build orders. Some of the information in jobs will be manual tasks.
- The *other resources* group contains information about work, the

schedule of manual tasks, the people, including skills and whereabouts and security information about people's access rights to systems and buildings, including BT people and customers.

- The *products and services* group of data (Figure 4) describes the products and services, both in terms of the portfolio of types of service sold and the instances of services which have been sold to specific customers. This is broken down into more detail below.

The products and services portfolio contains information about each product and service in the portfolio including its service features and charging structure. Associated with this are marketing plans for promotion of specific products, market intelligence collected from customers about products and requirements and product plans for the launch of new products.

Service configurations contains high-level information about how the features of each instance of service are delivered by the network together with the service level agreements (SLAs). Associated with these are behaviour data about types and instances of service including usage, performance and problems.

- The *network* group (Figure 5) contains data which describes

Figure 4 – The products and services data areas

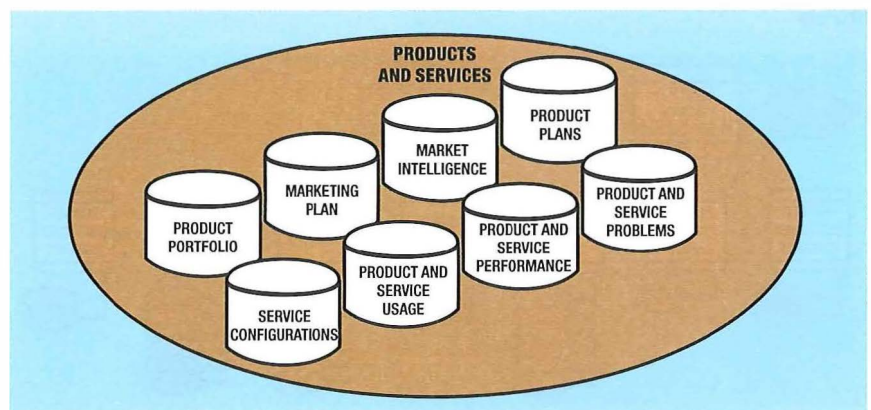
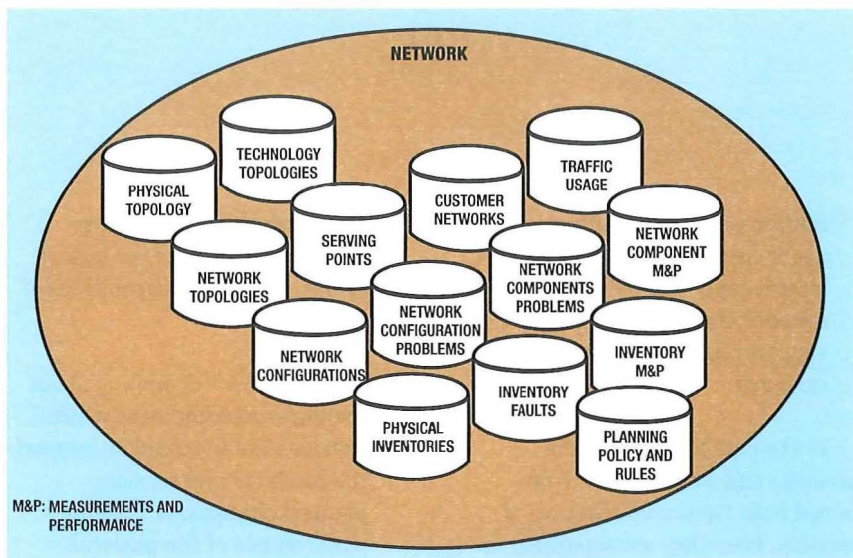


Figure 5 – The network data areas

the structure and behaviour of the network. The structure of the network is described in terms of topologies, configuration and inventories. The behaviour data includes information about usage, problems and performance and are associated with the relevant structure data. Finally, there are the planning policy and rules.



This structuring and grouping of the data will help in the design of databases and data servers due to the understanding of the relationships and affinities between the data areas with the processes or functions which operate on them.

Physical Systems Approach

The above description demonstrates the approach taken in designing the logical components of the architecture, the process and data frameworks. This logical framework has been designed to ensure that all future OSS solutions deliver known and required functionality and have appropriate access to a unique set of properly managed data. This was used as the starting point for the development of a physical systems framework.

This physical framework has been designed on the following basis:

- **Object-based (data-centred) approach** Whereas previous architectures have been primarily functionally (or process) based, this architecture puts a much higher emphasis on data and avoiding the data problems in the legacy systems.
- **Identify the key things to be managed** The best set of high-level business objects which the business needs to manage have been identified, considering where changes are likely to occur to avoid too many different types of change impacting on the same systems.

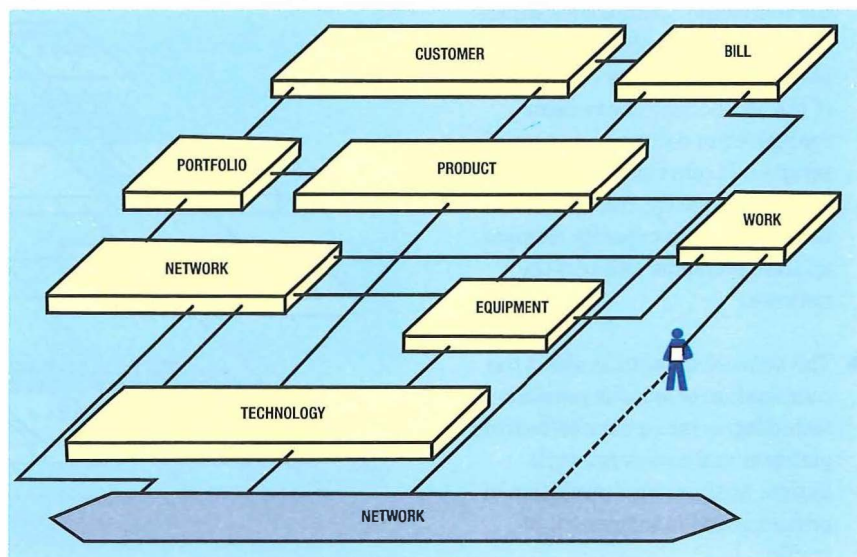
- **Identify management domains** Knowing the objects, management domains have been scoped and characterised within which the objects will be managed.
- **Identify the data which describes the objects** From the data analysis work, the data has been identified which describes those objects.
- **Identify the functions which operate on the data** From the process analysis work, the component processes and the functions which operate on that data were identified.
- **Encapsulate the functions and data within domains** The data and its functions were grouped together or encapsulated within the domain.
- **Minimise sharing of data across domains** With the functions and the data they need held together within the domains, there is little need for sharing across domains.
- **Maximise reuse of concepts, designs, code and interfaces** Within the domains, there will be opportunities to reuse concepts in terms of structures and relationships, designs, applications, database schemas and interfaces.

OSS management domains

Using the above approach, a set of eight management domains was developed, based around the previously defined data groupings. These management domains, illustrated in Figure 6, are defined below:

- The customer domain deals with all aspects of BT's customers and

Figure 6 – The OSS management domains (tiles)



contains all of the information about customers and their interactions with BT. It also contains the functions which support communication with the customer.

In the data framework, the products and services group contained both types and instances of services. Here they are separated as the issues in managing them are quite different:

- The *product domain* is about the instances of all products which have been delivered to customers, the SLA for each of them, how they were realised and how they have performed. It also contains the functions for delivering the services, managing problems and analysing performance and usage.
- The *portfolio domain* is about the range of products and services which BT sells, their performance, market intelligence and plans for new products. It also contains the functions for designing, performing trials and launching new products.

The network data in the data framework has been divided into three separate domains (technology, network and equipment):

- The *technology domain* is about the individual technologies which make up the complete network platform and how the capabilities of the technologies have been configured to deliver instances of services. It contains the functions to assign capacity, configure the technology, plan capacity changes, analyse problems and test the network.
- The *network domain* is about the combination of various network technologies into a service-bearing platform and covers network design, analysis and resolution of problems and management of traffic.

The network and technology domains deal only with the logical functions, features or capabilities of the network.

- The *equipment domain* is about the physical equipment which provides the functions to support the products and services, planned changes, faults and performance of the physical equipment.
- The *work domain* is about the manual work which needs to be carried out on the network and the workforce which performs that work. It contains functions to schedule and allocate work to the workforce.
- The *bills domain* is about the customer accounts, invoices for usage of products and billing problems. It contains the functions to raise charges for usage, generate invoices and handle payments, problems and debts.

These eight domains contain all of the functions and data to support the requirements of the seven generic business processes. The functions and data will be realised by a set of applications and data servers within each domain. These

applications and data servers will be supported by a set of closely coupled computers forming a computing infrastructure.

The domains are now known as *tiles* due to the icon that has been used to depict them.

OSS tiles—segmentation

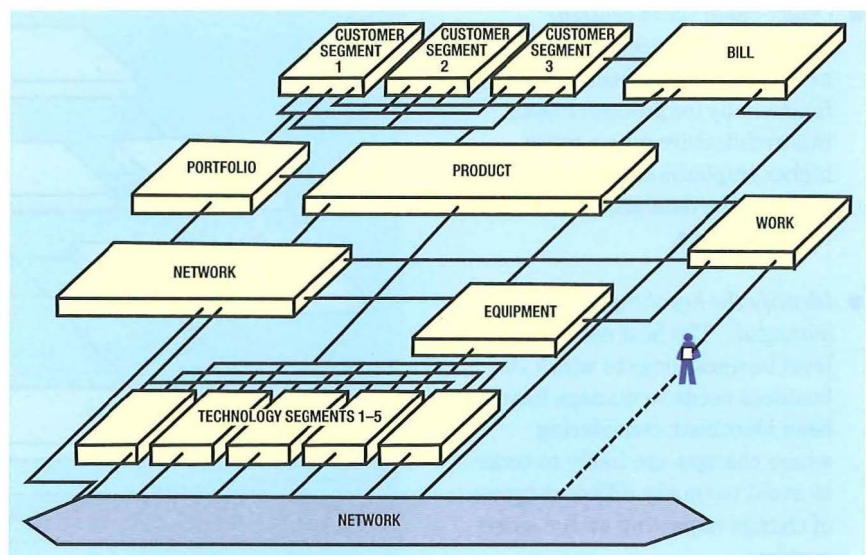
Some of the tiles have been segmented (Figure 7), supporting different physical solutions where there is a valid business requirement. The logic of the tile will remain the same in each case. For example, a complete network will consist of many different managed technologies which will require different management solutions. Parts of the relevant technology tile segment will be supplied by the technology supplier.

The customer tile is also segmented as it is believed that the physical solution will be different for different customer segments.

Tile detail

Each of the tiles contains a set of tightly coupled functions and data. In the case of the product tile (Figure 8), there are 10 data stores: products and services portfolio, product and service type performance, type usage and type problems, product and service instance measures and performance, instance usage and

Figure 7—Tiles segmentation



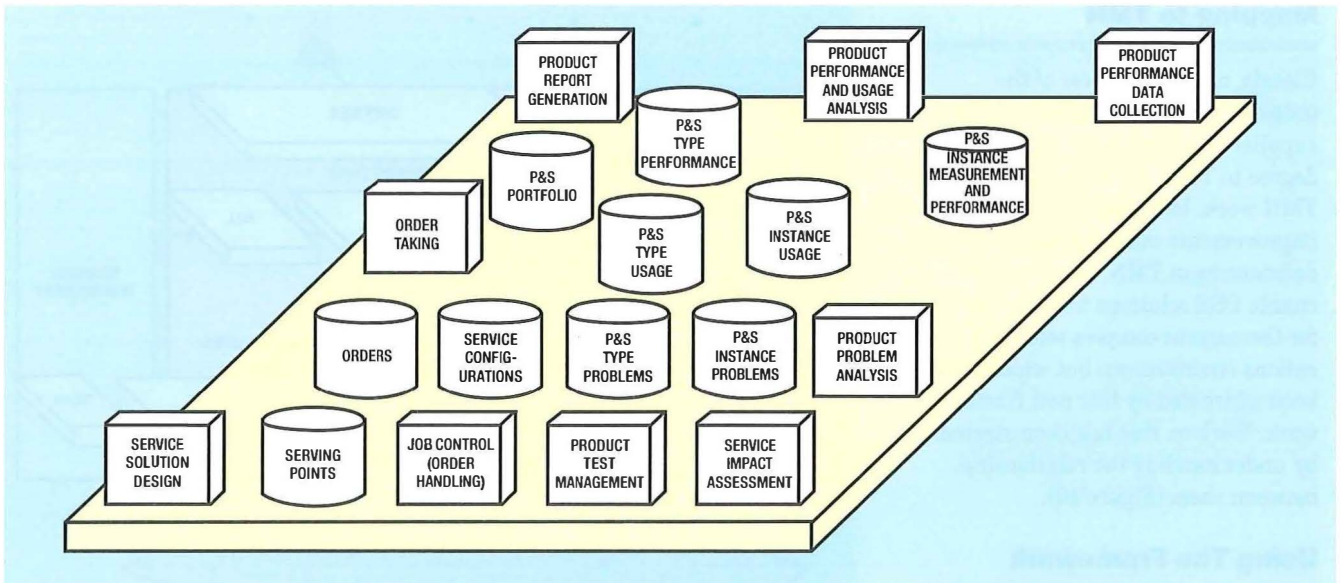


Figure 8—The product tile

instance problems, service configurations, orders and serving points.

These data stores are surrounded by and support nine functions which handle the collection, analysis and reporting of product type performance and usage. Order taking, service solution design, job control and product test management are involved in the Gain Business and Configure

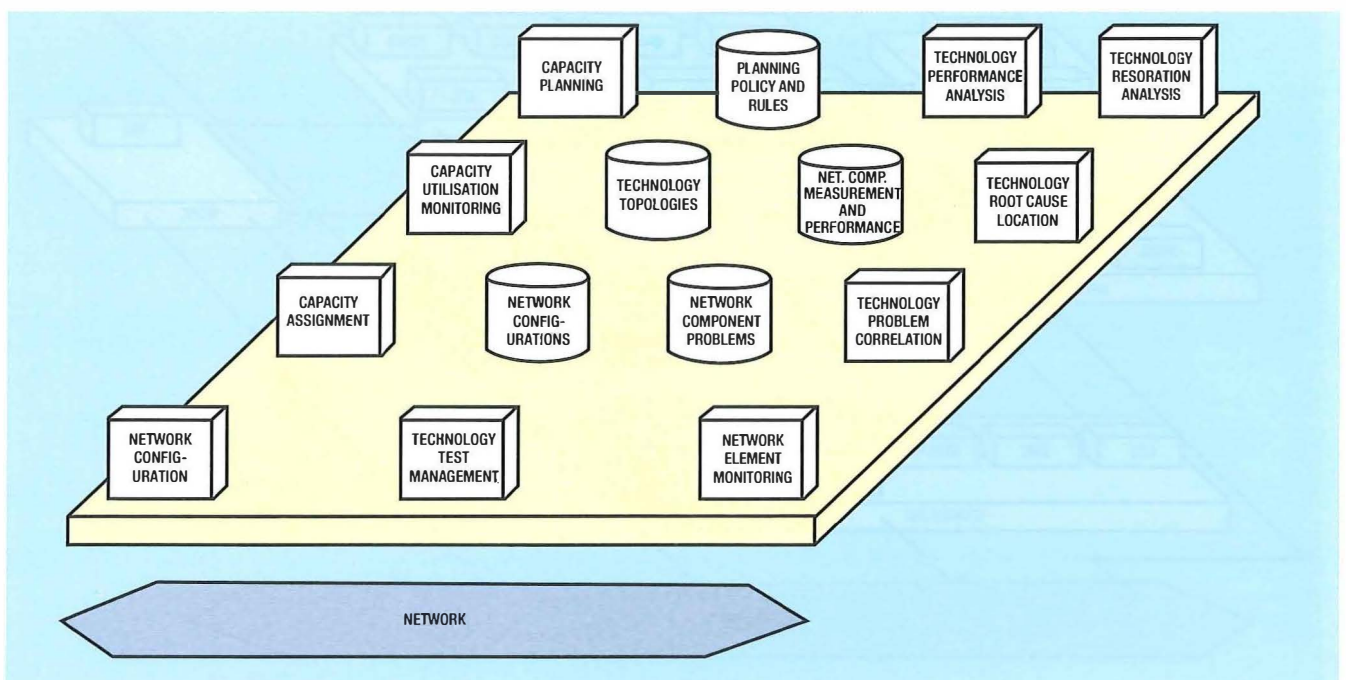
Service processes. Product test management is also involved in the Maintain and Restore Service process, together with service impact assessment and product problem analysis.

There is a high level of inter-connectivity between these functions and data stores reflecting the strong process and data affinity and supporting the existence of the tile.

There are also a number of linkages to functions on other tiles.

The other tiles have a similar set of functions, data stores and interfaces (see Figure 9 for the technology tile). The functions and data on the tile and across tiles cooperate together to deliver the end-to-end requirements of the business processes.

Figure 9—The technology tile



Mapping to TMN

Clearly, much of the rest of the industry in terms of operators and suppliers is committed to some degree to TMN. BT still supports the TMN work, but would like to see improvements made to address the deficiencies in TMN, which do not enable OSS solutions to be designed for the current complex telecommunications environment, but which have been addressed by this new framework. Work on this has been started by understanding the relationship between them (Figure 10).

Using The Framework

There are three main ways in which BT will use the framework.

For each of its current or legacy systems, BT can characterise that system in terms of the functions it performs and the data which it holds or uses. This will enable that system

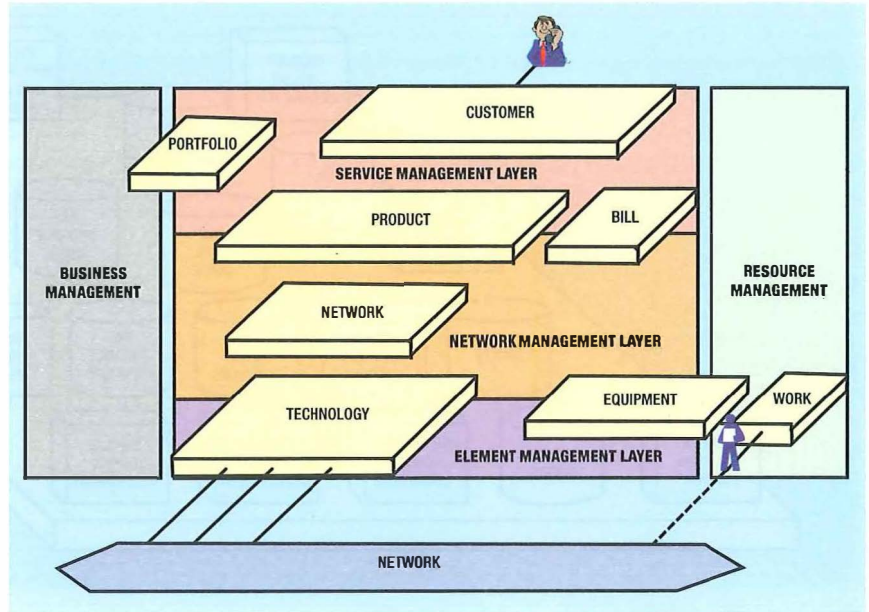
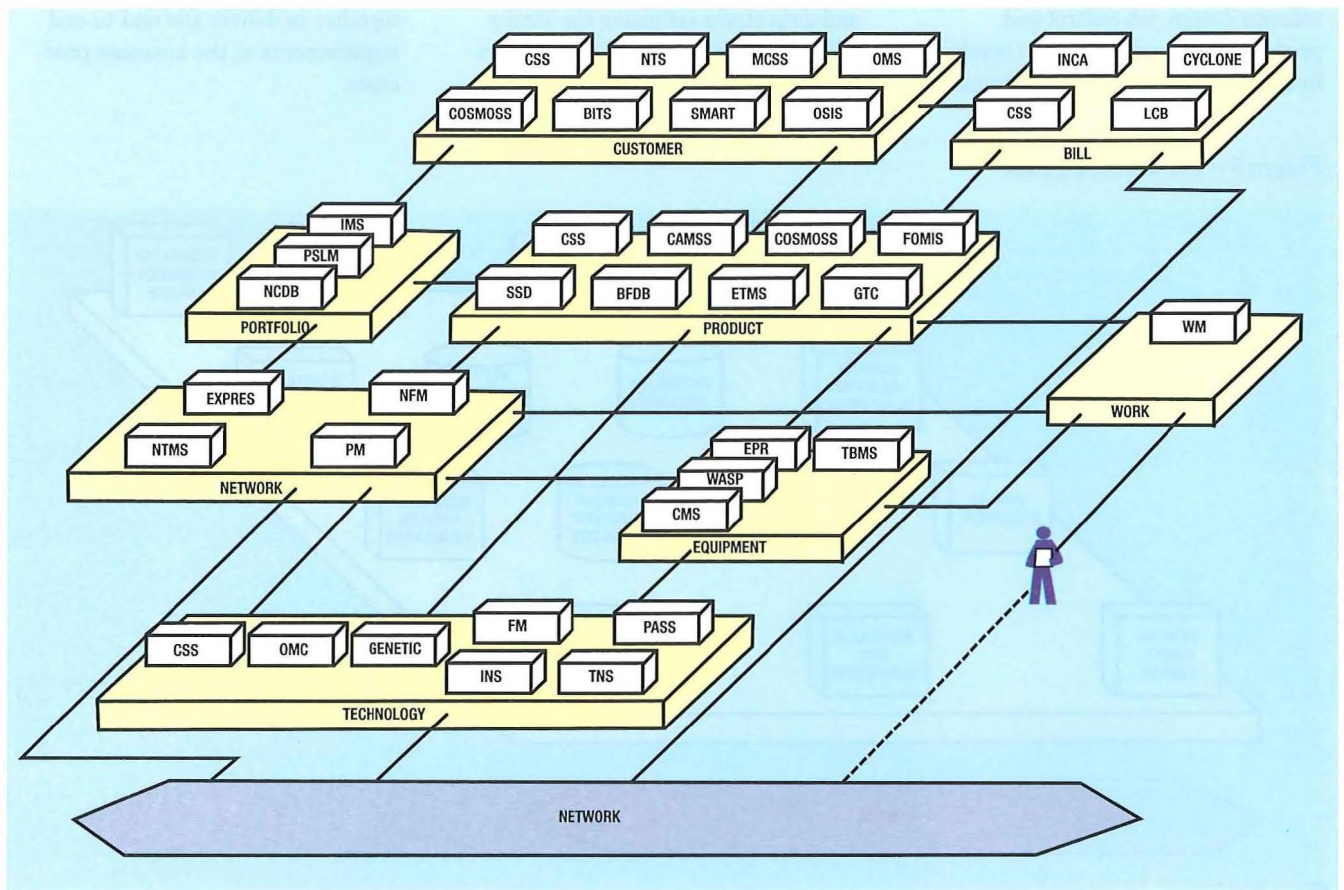


Figure 10—The relationship with TMN

to be placed on the appropriate tile (or tiles) in the framework. It cannot necessarily be expected that legacy systems will align with the frame-

work and often a legacy system will have an appearance on more than one tile. In future, with new systems, the aim will be to avoid this.

Figure 11—BT's key systems by tiles



When all of BT's current systems have been placed, this will help decisions to be made about their evolution and, in many cases, their rationalisation. An evolution plan has been produced for each of the tiles showing how the current systems will evolve and how gaps in functionality and data will be resolved.

The third and probably most important use will be in the design of solutions and systems for new business objectives and initiatives. The framework will help to ensure that each solution is consistent with others and that there is maximum reuse of solutions. It is a holistic framework, covering the whole of the service and network management scope, enabling designs to be produced for specific problems with a knowledge of the broader context. It also provides an essential checklist of all of the issues which need to be considered in a design.

Figure 11 shows BT's key systems mapped to the relevant tiles having characterised them by function and data types supported.

This framework will help BT to move its whole systems base forward in a controlled way by designing coherent solutions from a consistent understanding of the complete OSS picture used by all of its designers.

Further details on the value of the framework to BT and how it will be used will appear in a subsequent article.

Biography



Nick Furley
BT Networks and
Systems

Nick Furley graduated with a B.Sc. in Electrical and Electronic Engineering from the University of Manchester in 1974 within the Post Office Student scheme. He worked on the design and delivery of a number of BT's operational support systems including the digital exchange support system (DESS) and the district data collector (DDC), now known as the network mediation processor (NMP). In 1988, he took responsibility for the architecture and evolution planning of the network management systems and set up the Network Control Architecture Board. In 1994, he was seconded to the Breakout project, where he was a member of the team which developed the OSS architecture framework described here. In 1996, he took on the broader scope of architecture and evolution planning for all OSS.

Tony Grainger and Tony Poultney

Enhancement of the Digital Core Switched Network

The 1980s was a decade during which BT's network experienced massive changes in switching technology, bringing improved quality of service and many new services. The consequence of being at the forefront of change was frequently to be out of step with international standards. The 1990s have brought the need to align with the emerging standards and they will be seen to be delivering a step change in the quantity and complexity of services offered. Customer expectations continue to rise and the technology changes to the core network infrastructure must be achieved more quickly and seamlessly than ever before.

Background

The design of System X began in the 1970s and was a collaborative venture between the then British Post Office Engineering Department and its principal switch suppliers. A local exchange was demonstrated successfully at the Geneva Telecom 79 exhibition and, following this, a small number of local and trunk exchanges, known as *Release 1* units, were brought into public service in 1980, 1981 and 1982. Towards the end of 1979, a development programme was started to build on the initial design. The first task of this programme, known as the *system enhancement programme* (SEP), was to develop the exchange processor. Lessons learned from the experience of the Release 1 units were fed into the SEP development. In 1983, the design authority for System X had passed to GEC Ltd. and Plessey Telecommunications Ltd. (a joint venture with Plessey as the prime contractor) and later in that same year the first SEP exchange, Coventry Spires trunk exchange, was brought into public service. BT set out to modernise its network with the SEP System X as the primary building block. The suppliers also delivered regular upgrades to improve performance and deliver capabilities such as the integrated services digital network (ISDN). At the time design decisions had to be made in advance of the international standards; this was to be an important issue for later developments which will be discussed later in this article.

An international tendering exercise in the early 1980s resulted in a development and supply con-

tract being placed with Thorn Ericsson in March 1985. The contract covered the adaptive engineering required to modify the standard AXE10 digital exchange system to operate in the UK market and the supply of the first 100 000 lines of AXE10 into the local network.

The system was first introduced at Sevenoaks exchange in 1986, where it provided basic public switched telephone network (PSTN) services. Regular software enhancements followed to deliver features such as ISDN, together with hardware upgrades to increase the AXE10's processor throughput and datastore capacity.

The main emphasis of developments up to 1992 was to achieve improvements in quality of service and to reduce cost of ownership. The services offered by System X and AXE10 in the BT core network differed little 1984–1992. Table 1 indicates the way in which the number of services has grown, and continues to grow.

1992 Onwards—A New Approach

By 1992, many standards had achieved international agreement and more were emerging, particularly in the areas of ISDN and the intelligent network (IN), and the C7 signalling needed to deliver their services. In addition, new areas such as customer local area signalling services (CLASS) had been defined. The debate to bring these services and standards to the network was opened with GPT and Ericsson during the early part of the year. It rapidly became clear that both System X and AXE10 local ex-

Table 1: Services Offered

1984	Basic Telephony Network Services (Call Waiting, Call Diversion, etc.)
1985	Basic Telephony Network Services ISDN (Digital Access Signalling System No. 1 (DASS1))
1987	Basic Telephony Network Services ISDN (Digital Access Signalling System No. 2 (DASS2))
1994	Basic Telephony Network Services ISDN (DASS2) CLASS (Caller Display, Call Return) Other Licensed Operator (OLO) Interconnect
1995	Basic Telephony Network Services ISDN (DASS2) CLASS OLO Interconnect FeatureLine Number Portability
1996 Onwards	Basic Telephony Network Services ISDN (DASS2) ISDN (ETSI) CLASS FeatureLine Networking Number Portability Telemetry OLO Interconnect Intelligent Network Services V5.1 Fibre (TPON) Interface
Leading to:	Broadband ATM (Asynchronous transfer mode) TMN (Telecommunications managed network) Internet Access Optical Switching Fixed/Mobile Convergence

Table 2: Development Life Cycle

Pre-Order	Order	System Definition	Development	Verification	Acceptance	Roll Out
-31 months	-23 months	-21 months	-15 months	-10 months	0 months	+3 months
Establishing clients, business cases and budgets. Agreeing requirements, setting business priorities and evaluating tenders and SOCs.	Agreeing contractual terms and conditions, and payment and delivery.	Establishing full systems definition to meet client needs and to enable development to begin.	Supplier completes system design, coding and low-level testing.	Supplier verifies system functionality and performance. BT gives certification to signalling and charging.	BT accepts product as fit for roll out in the network, based on supplier test results and independent BT assessment in key areas.	Jointly managed by supplier and BT to deliver to network as quickly as possible and with minimal network disruption.

changes would need significant enhancement to their basic operating capability if they were to be able to support the growth of new and more complex services.

Before new services, facilities and infrastructure can be delivered, there must be a period during which requirements are captured and the new products are developed and tested. A typical life cycle for these activities is shown in Table 2.

With some urgency, a set of key requirements was prepared and presented to GPT and Ericsson aimed at establishing a common level of compliance coupled with a logical and closely aligned development programme. At the same time, the groundwork began on the preparation of the business case necessary to support such an investment. This was the beginning of Phase 1 of the digital local exchange (DLE) infrastructure programme.

For the first time, this was not the traditional task of adding more system functionality with the direct aim of introducing specific customer services. Instead the task was to enhance the basic infrastructure of the digital network as an essential enabler to service-oriented developments, as yet undefined.

Processor Upgrades

During 1990 and 1991, there had been great focus on the memory

there was a clear need to move to the next technology step in terms of processor capacity and system memory.

utilisation of System X and AXE10, and many exercises aimed at optimising the situation were carried out. In System X, these led to such things as the reduction of the replications of the signalling interworking subsystem (SIS) made possible by the shrinking number of analogue exchanges as network modernisation progressed. In the case of AXE10 there had to be constant monitoring of the use of system memory, especially as more complex services became available.

Central guidance and direction had been issued to Zones as early as 1987 regarding the monitoring of throughput and capacity. Although close monitoring was left to Zones, the broad picture was available at the centre from the measurement of the effects of each new software build. Judgements were different on System X and AXE10: on System X, throughput was not a problem and memory capacity was the key parameter to be addressed; on AXE10, there was a clear need to move to the next technology step in terms of processor capacity and system memory.

Signalling

BT's network modernisation programme had begun in advance of international standards. The signalling system in use, known as *CCITT No. 7 (BT)*, though C7 in format, was not capable of interworking with other administrations for the ETSI† standard ISDN services, nor would it support the emerging standards for the intelligent network (IN).

C7(BT) originally employed messages up to 62 octets in length whereas standards such as the ISDN user part (ISUP) called for a message length of up to 272 octets. This is an example of the penalty for being at the forefront of change. The signalling link hardware on

System X would not be capable of supporting 272 octet messages and therefore it would be necessary, at a later date, to make hardware changes and modifications at both ends of many thousands of signalling links.

Other new signalling requirements such as the signalling connection control part (SCCP) and the transaction capability application part (TCAP) would also be essential on both System X and AXE10, if the network was to evolve to handle the future deployment of IN.

The debate on all these issues, plus the impact of introducing ETSI ISDN, IN and CLASS, was opened with GPT and Ericsson in April 1992. There followed many months of intense activity—capturing requirements, building plans for software build deliveries and hardware upgrade programmes and negotiating contracts. The first orders were placed in December 1992 with others following in the months up to May 1993.

Infrastructure Enhancement: Phase 2

Late in 1993, and in the early months of 1994, the trialling of calling number delivery began. The service was launched in November 1994. As the first software builds resulting from the Phase 1 orders were deployed, their impact made it necessary to consider further memory enhancements. There was more enhancement of the CLASS services to come and the forthcoming developments to ISDN, which would bring ETSI standard services, would have a major impact upon memory usage. In addition, during the period of deliberation on these issues, a major upgrade to Centrex (FeatureLine) capabilities was commissioned. From all this, it was soon apparent that another two years would bring the need for yet more processor memory.

First calculations indicated that the System X need would be for more backing store but that the main store

could be held at 28 megawords (MW). The ability to hold the main store at 28 MW was dependent upon the reintroduction of memory management techniques (paging). However, these early builds also indicated that the processing power of System X was now being taken up at a significant rate and the potential impact of paging on processor throughput caused some alarm. After further reviews of projected memory usage, BT and GPT jointly decided to raise the main and backing store to 40 MW. This enhancement formed the major element of the digital local exchange (DLE) infrastructure programme Phase 2.

The Phase 2 infrastructure programme was also used to deliver the radical processor and system-memory enhancement already deployed to most AXE10 sites in Phase 1 to the remainder of the AXE10 switches in the network. This programme also set a framework for the future software deliveries and placed an initial order for the development of C7 ISUP signalling directly from the local switch. Once again, the key driver for both the System X and AXE10 programmes was the aim of enhancing the digital core network to support and enable the deployment of new revenue earning customer services.

All the foregoing seems to suggest that enhancements of System X and AXE10 occur only as a result of major activity at approximately two year intervals. This is far from being the case and during the intervening period many discussions had taken place with both suppliers: some leading to development orders being placed and others which were still unresolved when the Phase 2 project began. An example of this on System X would be the replacement of the automatic announcement facility. Orders had been placed for this during 1993 but by the time the Phase 2 project began a further enhancement of this facility had been identified. Equally on AXE10,

†European Technical Standards Institute

The infrastructure programmes have also had to consider the implications of the EC directive on electromagnetic compatibility

many of the requirements are very long term and the programme is subject to continuous demands for change or additions.

In the case of Phase 2 of the DLE infrastructure programme, all the outstanding issues were drawn into the project and a concerted debate opened with GPT and Ericsson in January 1994. Given the importance of correct resolution of the memory issues it is not surprising that it took until March 1995 to settle all the issues, obtain authority and place the orders.

Problems And Complexities

For the purposes of this article, the BT core network switching platform is considered to be comprised of digital switches. These elements are shown unshaded in Figure 1, which also gives an indication of the size of the network. These switches must appear as one platform to BT's customers but are manufactured by two suppliers, both of whom face demands from other telecommunications markets and operators. There is, therefore, a need to ensure that both suppliers will deliver features and facilities that will not only interwork but will do so in a manner which is transparent, and will present a consistent appearance to BT's customers throughout the UK.

It is also necessary to ensure that in deploying a feature or facility that it is done in a manner that takes account of the regulatory environment. Furthermore, while BT, and the British Post Office before it, has always gone to great lengths to ensure that charging conforms to the published costs and timings, the current regulatory situation demands greater evidence and recording of the verification of this conformance.

These constraints demand great diligence in both the capture of requirements and in the verification and integration of new software deliveries. Each software build is

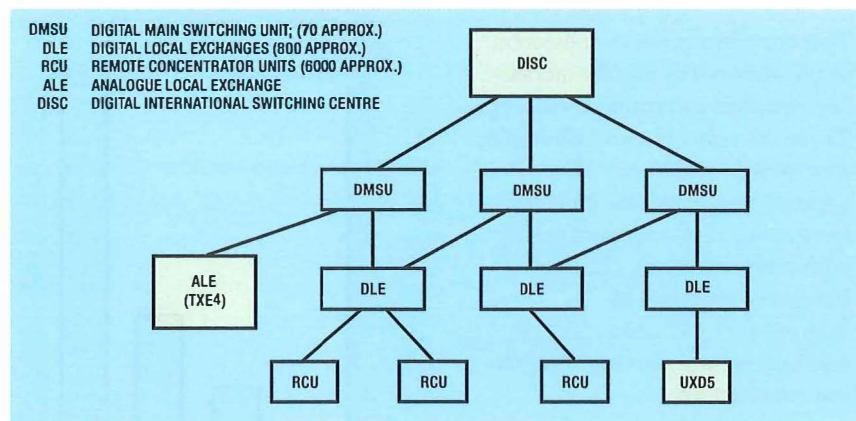


Figure 1 – Network schematic

subjected to an assessment of its impact on the charging process and tests are conducted to verify the truth of that impact. In addition, a period of testing is carried out to ensure that the new software integrates with the network and does not cause a network, as opposed to nodal, failure.

Probably, since the first common control exchange was introduced to the UK network, the objective of reducing downtime has been uppermost in many peoples' minds. Certainly, this has always been true of System X and AXE10. As the intrinsic performance of these systems has improved over the years, so the emphasis for this reduction has shifted to the downtime incurred as a result of upgrade or enhancement. The objective of zero downtime enhancements still eludes us and will almost certainly continue so to do until the next generation of processor technology is deployed. However, improvements have been made; for example, System X concentrators can now be upgraded without loss of service and AXE10 has undergone restructuring to allow new functionality to be added with minimal downtime. In addition, the host processor downtime incurred on loading a new software build has been progressively reduced on both systems.

The infrastructure programmes have also had to consider the implications of the EC directive on

electro-magnetic compatibility (EMC). For the supply of new exchanges the situation is clear: the responsibility is with the exchange supplier to deliver equipment which conforms to the directive. For the installed base, however, the position is far from straightforward. The directive suggests that any change to existing equipment can bring the requirement to make that equipment compliant to the new standards for emissions and susceptibility. Telephone exchanges have a planned life of more than 10 years and change of hardware is a major undertaking both in cost and management of customer service. It is therefore to be expected that hardware upgrades and extensions will have to be applied during the life of an exchange.

Programmes such as the DLE infrastructure programmes described above will bring hardware upgrades to exchanges which predate the availability of equipment practices which give compliance to the EC Directive, but which are no older than 12 years and some even as recent as 1995 installations. Clearly a means of deploying the required upgrades which both conforms with the directive and which is commercially viable must be found. This issue has caused much debate. There have been consultations with both suppliers, and processes which discharge BT's obligations have been established.


BT's network must evolve in a manner which ensures that its interfaces meet recognised standards.

This bar chart gives an indication of the scale of the rapidly increasing demands on switch processors. There are many factors influencing this growth in demand. The greatest impact to date has been the growth of interconnect traffic with other operators; this traffic brings the overhead of internetwork call accounting. In addition, we are now beginning to see the effect of:

- complex services with correspondingly large processor instruction counts;
- significant traffic growth from within BT and from OLOs; and
- a continuously changing 'call mix' as the proportion of calls using existing network services grows steadily.

Although these are the primary causes of growth in processor demand there are numerous secondary contributors such as the need for checksums in call-charge records required to meet stringent billing audit trail requirements.

It is vital also to take into account the fact that special



events such as phone-in programmes (for example, Children in Need) on national television and radio can produce high network loadings for relatively short times. There are also unplanned overloads in response to accidents and disasters. Throughout all of these, the network must remain stable and effective and both deliver network management information to the traffic management centre and respond properly to commands input to the switch.

Figure 2—Effect on processor demand on call mix and complex functionality

The Future

The impact of all the new functionality on both throughput and memory is now a major concern and it is apparent that the juggling of exchange loadings cannot be the sole answer to throughput problems for much longer. The impact of new software and the take-up of services are continually having to be reassessed. Figure 2 indicates how the processor demand has grown as more complex services are deployed.

There is also an additional factor in the take-up of processing power—the significant growth of traffic to and from other licensed operators

(OLOs). From a processing point of view, this traffic can be more onerous than that which is contained within BT's own network because of the additional call records that must be produced for commercial purposes. In the case of many of the customers who migrate to cable TV operators, although BT loses the customer, the network still carries much of that customer's traffic and at an increased cost in terms of the exchange processor's workload.

In addition to the capacity issues, the global nature of telecommunications and the multiplicity of operators in the UK mean that a C7(BT) situation cannot be allowed to arise

in the future. BT's network must evolve in a manner which ensures that its interfaces meet recognised standards. Equally, if BT is to meet its customers' needs, those standard interfaces will need to be made available at the right time. The adoption of more interfaces for optical-fibre transmission media and the migration of concentrator technology out of the exchange environment into the access network can be expected.

Both GPT and Ericsson have radical plans for the evolution of their products and discussions have begun on when and how their new products could, and should, be deployed. This is the majority substance of Phase 3 of the infrastructure enhancement programme.

Conclusion

In addition to the enabling technology changes, but no less important, is the need to reduce the time taken to develop new features and bring them to the market. Many, perhaps most, people will point to the IN as the way in which this will be achieved. However, a successful IN demands the right level of functionality and capability in the core network and the advantages of IN will be squandered if BT is unable to interact efficiently with its suppliers. Therefore, BT is not only addressing straight forward technology issues in regard to the infrastructure programmes but is also expending a great deal of effort on improving working practices with the aims of:

- reducing the time to market of new customer services;
- providing a consistent presentation to the customer across the total network; and
- using BT's multi-supplier platform to best advantage.

Undoubtedly, the task of upgrading the digital core switched network

over the period up to the end of the century will demand the highest level of cooperation between BT and its suppliers in addition to the most advanced technological development. The infrastructure programmes are key to the continued network evolution and will be the foundation upon which further advanced revenue earning services will be based.

Glossary

ATM Asynchronous transfer mode
AXE10 Digital switching system designed and manufactured by Ericsson Ltd.
CLASS Customer local area signalling services
C7 CCITT signalling system No 7
DASS1 Digital access signalling system No.1
DASS2 Digital access signalling system No.2
DLE Digital local exchange
DMSU Digital main switching unit
EC European Community
EMC Electromagnetic compatibility
ETSI European Telecommunications Standards Institute
IN Intelligent network
ISDN Integrated services digital network
ISUP ISDN user part
OLO Other licenced operator
SCP Service creation point
SEP System enhancement programme
SCCP Signalling connection control part
System X Digital switching system designed and manufactured by GPT Ltd.
TCAP Transaction capability application part
TMN Telecommunications managed network
TPON Telephony over passive optical networks

Biographies



Tony Grainger
BT Networks and Systems

Tony Grainger joined the Post Office Engineering Department as a technician in 1965. He was employed on exchange maintenance duties in the Manchester South Telephone Area until being promoted to AEE in 1978, when he moved to Headquarters in London. Initially working on project management for System X, he has since been involved with many aspects of digital exchange operations and development. For the past five years he has managed the System X Technology team.



Tony Poultney
BT Networks and Systems

Tony Poultney joined the British Post Office in 1963 as a Youth-In-Training in the Bedford Telephone Area. In 1971, he joined the TXK3 Crossbar team in Headquarters after gaining a B.Sc. in Electrical and Electronic Engineering at Hatfield Polytechnic. Subsequently, he worked with TXE2 and TXE4 reed relay systems before joining the AXE10 launch team in 1985. He is currently the programme manager for AXE10 technology responsible for system development and performance. He is a member of the IEE and a Chartered Engineer.

Cashless Services Network Operations Unit

The introduction of the cashless services replacement project, which includes the Chargecard service, called for a major change in the management of operations. This article reviews the progress of the project from inception to cutover, and examines the organisation of operational and technical support from the perspective of the cashless service network operations unit.

Introduction

BT's network operations unit (NOU) at Worthing first became involved with the cashless services (CS) platform in 1994 when it was asked to take on a lead NOU role for the network components of the 'old platform'.

The 'old platform' (Figure 1) consisted of 30 cashless services processor units (CSPUs) co-sited with, and accessed via, digital main switching units (DMSUs). The CSPUs were connected to a single cashless services database (CSDB) via KiloStream links. Syntegra managed the CSDB as well as providing a 24-hour service desk for problem reporting and progression across the platform. Worthing cashless services network operations unit (CSNOU) had responsibility for the maintenance and administration of the CSPUs and the KiloStream links that connected them to the CSDB.

The introduction of the cashless services replacement system (CSRS) changed the role and responsibility of the CSNOU significantly. The new platform was purchased from MCI and represented a significant increase in computer technology and data communications into core network call processing.

This article provides an operations perspective of the CSRS project, an overview of the CSRS platform and an outline of the roles and responsibilities of the CSNOU based in Worthing NOU (Figure 2).

† HOBBS, BILL. Cashless Services Replacement System Project. A Network Perspective. *Br. Telecommun. Eng.*, April 1996, 15, p. 20.

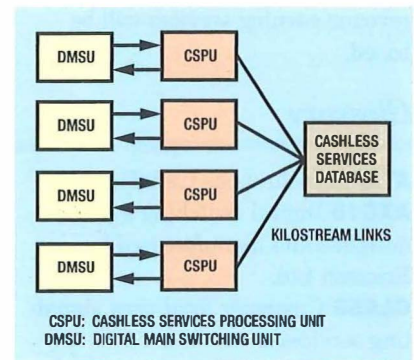


Figure 1—Old cashless services platform

A brief description of the cashless services portfolio has been published in a recent edition of this *Journal*†.

Cashless Services Replacement System Project—The CSNOU Role

As a result of the first meeting to determine the CSNOU role in the CSRS project, it was decided to focus on the following:

- the development of a firm understanding of the topography and technology to be deployed,
- the development of business processes that would establish the BT network as 'best in class', and
- the development of people policies that would deliver a well-trained and committed team.

Having completed the preparatory work in the above areas, a visit to MCI was undertaken. This had the objectives of seeing the MCI operation in the USA first hand and of testing the initial UK-based assumptions. From this visit and from the



Figure 2 – Worthing CSNOU

previous experience of managing the BT core network, emerged the work packages to support the people, business processes and systems used by CSNOU to manage the UK platform.

During the last quarter of 1994, the synthetics developed from the visit to the USA had determined the shape and size of the CSNOU, following which the recruitment and training of the team began in earnest. A total of four managers and 25 technical officers were recruited in the following months to provide 24-hour, 365-day-per-year management of the platform.

Because of the extensive range of skills required to maintain the platform the new members of the team had to be assigned specialisms within the platform; for example, switch and intelligent services network (ISN) expertise. Those people selected for training on the ISN completed on-the-job training with MCI in the USA (approximately 5 weeks), as well as a range of computing and data communications training in the UK. Switch people, too, underwent weeks of training with a full range of Ericsson AXE10 and Northern Telecom DMS maintenance courses to attend. Between training courses, CSNOU people supplemented their formal training by working with colleagues in BT Laboratories, Martlesham and MCI in the system test phases of the project, while continuing to develop business processes.

A key element of the new process was the need for the CSNOU to

develop the supporting infrastructure that facilitated direct escalation to third-party vendors. The plethora of vendors involved with the platform necessitated the identification of clear ownership, fast lines of communication and underpinned response times. All of these objectives were met with the support and involvement of colleagues across BT.

In mid-1995, CSNOU people became increasingly involved in the project to install and implement the new network, in particular the establishment of the customer service centre in Liverpool and the operator services centres in Dover and Leicester. The CSNOU accommodation became available in July, and there followed a period of intense activity with the team bringing to life the numerous support and surveillance systems required to manage the platform. Coincident with this, system testing was completed and, during a two-week period, the

platform was migrated from the test to the production environment.

High-volume call senders were deployed to generate sufficient calls to flex the system and measure the quality of service, followed by a period of operational readiness and customer acceptance testing. The beta trial with 9000 customers commenced in September 1995 when the CSNOU personnel were confident of the stability of the platform and their ability to manage it. Full cutover of some 6 000 000 Chargecard customers took place in February 1996.

Overview of the CSRS Platform

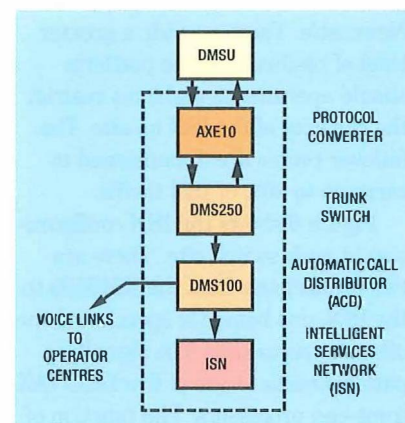
The CSRS hardware configuration is the same as that used by MCI in the USA, but with a different build of application software specific to the BT environment. The hardware is installed at five sites in the UK (see Figure 3), with two switch sites (Crawley and Newcastle), two operator centres (Dover and Leicester) and a customer service centre (Liverpool).

Figure 4 gives an overview of the equipment installed at the switch sites. Both sites are fully connected to DMSUs as either the first or second choice routing. The Ericsson AXE10 switch (AS36 Cyclone build) provides the interface between the BT PSTN network and the MCI CSRS platform. Its main function is

Figure 3 – CSRS sites



Figure 4 – CSRS switch centre overview



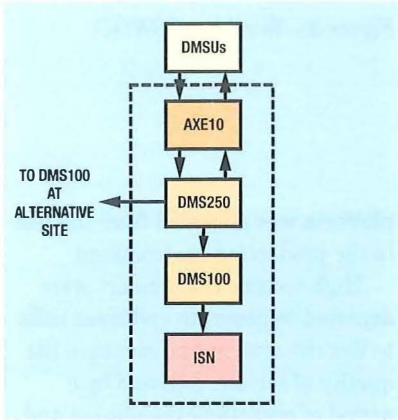


Figure 5—CSRS switch centre resilience

to convert from BT European standards (E1, 2 Mbit/s, C7 signalling) to MCI USA standards (T1, 1.5 Mbit/s, S7 signalling). Connected to the AXE10 switch is a Northern Telecom (NT) trunk switch, the DMS250, which is used widely in the MCI network. On receipt of an incoming call, the DMS250 will connect to the ISN via the NT DMS100 switch for call validation. If successfully validated, the outgoing leg of the call will be set up and the connection to the ISN will be released (release trunk capability). It should be noted that the ISN is only used for call set-up, and as such has a relatively short call-holding time (20–25 seconds).

Calls remain routed through the DMS250 switch for the duration of the call. An MF receiver monitors each circuit for customer re-origination requests (MF signal ##). On receipt of a re-origination request, the DMS reconnects to the ISN.

The NT DMS100 switch is an automatic call distributor (ACD) and its function is to distribute the call to the appropriate speech circuit on the ISN, whether this be an automatic response unit or an operator console.

Figure 5 shows the failover routes provided between Crawley and Newcastle. These provide a greater level of resilience to the platform should operational problems restrict the capacity of the ISN on site. The failover routes are dimensioned to carry up to 20% of ISN traffic.

Figure 6 shows the ISN configuration at each switch site. There are two connections from the DMS100 to the ISN, one being for speech and the other for signalling. The signalling path connects to one of five DEC VAX front-end processors. The function of

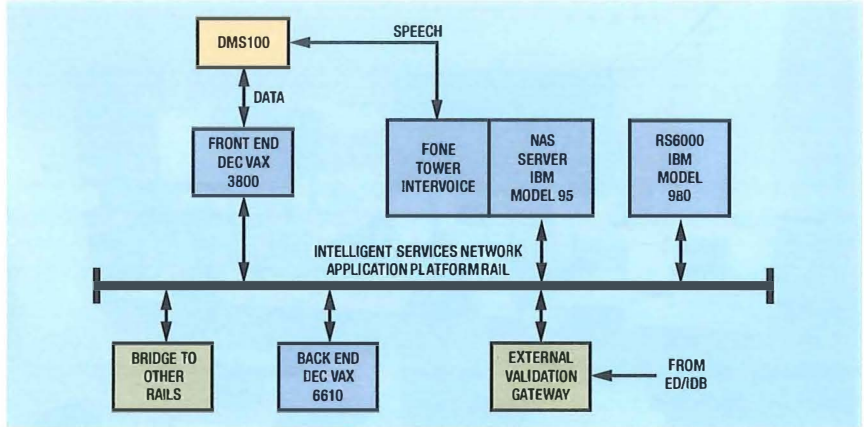


Figure 6—ISN switch sites

the front-end processor is to translate incoming signalling information from the DMS100 and pass this on to one of three DEC VAX back-end processors. Similarly, outgoing signalling information from the back-end is translated and passed on to the DMS100.

The back-end processor receives incoming signalling information from the DMS100, such as calling line identity and the number dialled to access the platform, and decides how to process the call. For a conventional Chargecard (144) call, the back-end passes the necessary information to a RS6000 (IBM Model 980) machine associated with the chosen incoming speech circuit. In the reverse direction, the back-end processor could receive a *ready state* signal from an operator terminal and pass this information via the front-end processor to the DMS100 switch.

Each RS6000 (five at each site) has call-processing application software as well as a database containing a full record of customer information such as account and PIN,

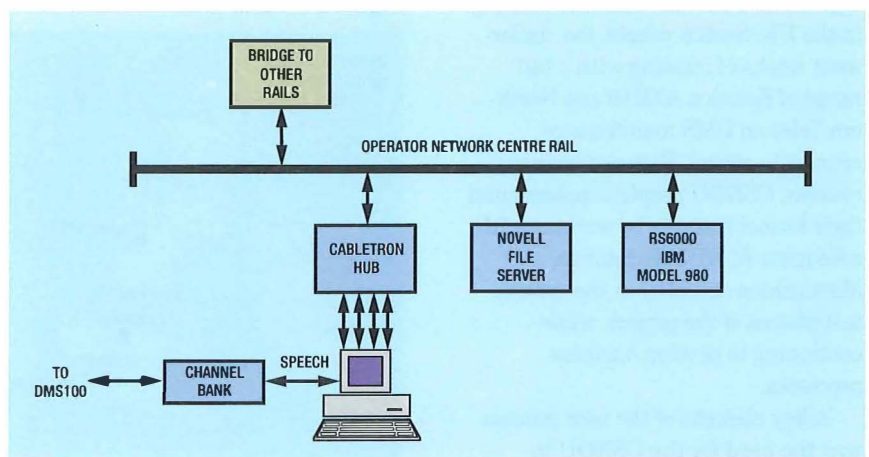
and administration data such as barred country codes.

The RS6000 controls three network audio servers (NAS) (IBM Model 95). Each NAS has an associated Fone Tower, each with two T1 links (48 circuits). The speech path from the DMS100 switch is connected to a Fone Tower circuit and the function of each NAS is to deliver the appropriate voice-prompt messages to the customer as well as monitor and pass on incoming MF signals to the RS6000 (account and PIN numbers, dialled digits etc.). An external validation gateway, again using an IBM RS6000 machine, enables access to 100 and international operators to validate Chargecard customer account and PIN.

Communication between the various devices is made over a local area network (LAN) based on thick-wire Ethernet technology. Communication to other sites is made over a bridged wide area network (WAN).

Figure 7 shows the ISN configuration at the operator centres and

Figure 7—ISN customer service centre



The technology deployed in the new network makes the CSNOU possibly unique within BT in that it combines telephony and computing operational management within one unit.

customer service centre. Again there are two connections from the DMS100 to the ISN, for speech and signalling respectively. The signalling path is provided over the bridged WAN to one of the switch sites via the back-end and front-end processors as described above. Communication between the various devices is made over a LAN based on thick-wire Ethernet technology. A Cabletron hub extends the LAN to the manual telephony operator console (MTOC), which is a 486 personal computer fitted with an 'Emerald Card'. This provides the speech connection direct from the DMS100 switch at one of the switch sites. There are two Novell file servers which hold operator console application software, and an RS6000 which holds a database containing a full record of customer and administration data. A customer making a conventional Chargecard (144) call is connected to a Fone Tower port as described above. If the customer fails to enter his/her account, PIN or destination number successfully (for example, does not have an MF telephone, or fumbles the entry) he/she will be connected to an operator console. An operator receiving a call from a customer is presented with dialogue scripts on the MTOC screen and prompts to progress the call. As an example, if the customer is having difficulty in using the Chargecard service, the operator is able to validate the customer details and connect the call. If the operator is unable to validate the call, the customer can be transferred to the customer service centre with appropriate notes.

It is also possible to connect calls directly to operator consoles; for example, when a customer dials 0800 345 144, the customer service centre helpline number, the digits are set up in data and the call is routed directly to a customer service centre MTOC.

Figure 8 shows that there are two billing messages generated for every call. The first, the billing data record (BDR), is produced by the RS6000 and contains information such as

calling line identity and number called. The second, the operator services record (OSR), is produced by the NT DMS250 switch and contains the time of call and call duration. The OSR is sent from the DMS250 switch to an adjunct processor (Stratus machine) which converts the record into a form that can be accepted by a mainframe application.

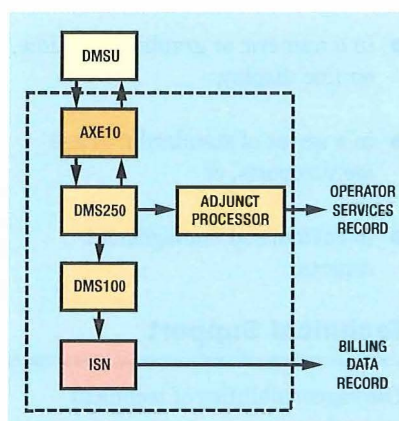
Both of the billing records are transmitted over the BT multi-protocol router network (MPRN) to a dedicated mainframe computer at Glasgow ISU where they are passed through a match/merge application to produce an enhanced operator services record (EOSR), which is forwarded to the BT mainframe billing engines for processing.

With the introduction of CSRS, all Chargecard customer details have been transferred to the customer services system (CSS). Changes to customer information are passed from CSS to the mainframe computer in Glasgow where it is received by a data distribution system (DDS) application. DDS downloads these changes to all ISN RS6000 machines via the BT MPRN network.

In-Service Roles and Responsibilities of the Cashless Services Network Operations Unit

The technology deployed in the new network makes the CSNOU possibly unique within BT in that it combines

Figure 8 – CSRS billing to mainframe



telephony and computing operational management within one unit.

Following the visits to the USA and discussions with Computing Services Operations colleagues, an organisation similar to that of BT Core Network Operations was developed for the CSNOU. This organisation with its recognisable disciplines of network management centre, technical support, data and capacity management and change control provides the foundation for the integration of different processes required for the real-time management of the platform.

An unknown factor at the development stage was whether this mix of American and British network management techniques for this new merged environment of computing and telephony would work. Early signs are that they do!

Network Management Centre

The network management centre (NMC) has the following responsibilities:

- 24-hour surveillance of the cashless services platform,
- first-line maintenance action,
- customer fault handling,
- tasking to field if on-site work is required, and
- management of all routine work.

Out of hours, a minimum of two people are on duty, one specialising in switch maintenance and the other in ISN maintenance.

Customer fault reports are passed to the CSNOU via CSS. The national activity management repairs system (NAMS) is used to scan the 29 incoming CSS queues for incoming faults. Tasking to field technicians at Crawley, Newcastle, Dover, Leicester and Liverpool is managed via the work manager (network) system (WM(N)).

A number of surveillance and support tools are employed to maintain visibility and control. As far as the switches are concerned, these tools are very similar to those used in the maintenance of PSTN System X and AXE10 exchanges.

ISN surveillance is more complex and a number of different systems are used. Some of the more important systems are described below.

VAX cluster console system

The VAX cluster console system (VCS) enables an operator physically to monitor alarms from the ARU RS6000s, front-ends and back-ends, and the Stratus AP. From a single terminal or VAX station display connected to the VCS host system, all console functions for front-end/back-end DEC VAX machines and Vitalink Bridges can be performed. In the case of the automated response units (ARUs) and Stratus adjunct processors, it is only possible to perform a monitoring function.

VCS scans incoming messages from the serviced nodes for specific text strings (error messages). The text strings are associated with an event record which may be allocated a priority number. The display is made up of node icons that alert the operator of scan events by a change of colour, the colour reflecting the priority of the events received.

Netview 6000

Netview 6000 provides visibility of the CSRS network via a graphical display. It monitors the network using the transmission control protocol/Internet protocol (TCP/IP) environment and graphically maps device locations onto a display screen. It dynamically displays status information, and by using filters and setting data thresholds, any key events/problems are detected and displayed. As an example of its use, all communication links (speech and data) between sites need to be converted from the MCI USA standards (T1, 1.5 Mbit/s) to BT European standards (E1, 2 Mbit/s). This conversion is achieved by using a

device called an *Orion*, produced by the USA company Larscom. The IP addresses of all these devices are built into the Netview 6000 system. Any network link failure will be reported by the Orion devices to Netview.

Network information distributed service network alarm monitor (NIDSMON)

NIDSMON is a stand-alone 486 DOS-based personal computer (PC) connected via the call-processing network (CPN). It connects to the operator/customer service centre LAN and receives error messages generated by the MTOCs. Error messages are presented on a scrolling screen. When a user-defined threshold is reached, an audible alarm is generated. The error message contains the site, console, service, server and the specific error. A file option facility exists, which enables the error messages to be written to disk for trending and further analysis. Filtering of error messages can be achieved by using the watch facility.

Call centre management information system

The call centre management information system (CCMIS) is a stand-alone Nortel product connected to the Crawley and Newcastle ACD switches via X.25 links. It gives management the ability to plan, manage, and monitor their agent requirements by collecting statistics on the performance of the network configuration.

CCMIS reports these statistics to management in one of three ways:

- in a numeric or graphic, real-time, on-line display,
- in a series of standard management reports, or
- in customised management reports.

Technical Support

The responsibilities of technical support are:

- to provide second-line technical support to monitoring centre team and field technicians,
- to own recovery following a major service outage (MSO),
- to investigate and report on MSOs, and
- to escalate technical to third-line support (MCI, vendors etc.).

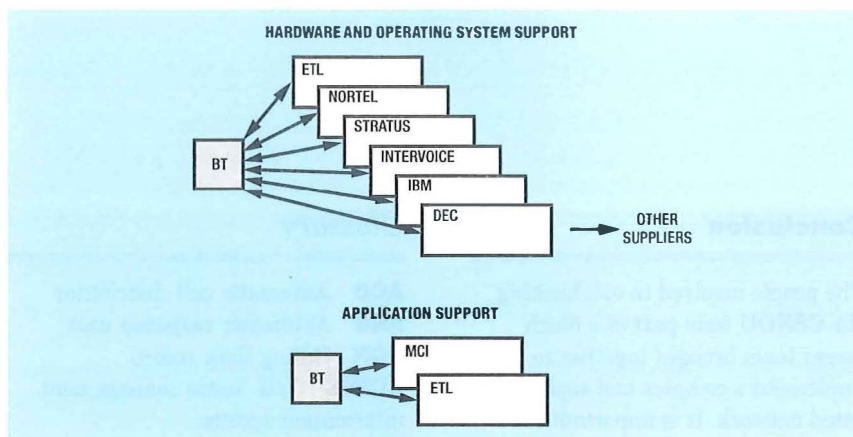
The wide range of switch, computer and communication equipment provides a considerable challenge for the technical support team where a greater in-depth knowledge of the systems is required. As a rule, individual members of the technical support team take a lead role in one or more specialist areas.

A support tool used extensively by the technical support group in problem identification and resolution is *Sniffer*. The system consists of a Sniffmaster console located in the CSNOU. This console consists of a standard PC connected to the core-processing network, running proprietary software. The console controls network monitoring and analysis tools known as *Sniffer servers*.

These servers are located at Crawley and Newcastle ISN sites. There are also servers at Liverpool customer service centre, and at Dover and Leicester operator centres. A server is made up from a standard computer with two interfaces and proprietary software. One interface supports communication with the console and the other uses software to capture data frames and statistics from the network. They are connected to both the Ethernet LANs and the token ring WANs.

CSNOU people use the Sniffmaster console to connect to the desired server. They can then use the analyser functions to capture network traffic for future analysis. This data can be filtered to save data fitting certain criteria. Alternatively, triggers can be set so that capture stops when the trigger event is

Figure 9—Third line support



detected. The captured data can then be displayed in various user-selected formats or filtered to show only frames that meet set criteria.

The Sniffer can be made to interact with other network management tools such as Netview 6000. Netview can then be used to signal an alarm when a trigger event is detected.

The CSNOU uses the Sniffer data to perform in-depth analysis of the data packets generated during the operation and interaction of the various components of the ISN platform. This enables the team to isolate the application or process that may be the cause of a problem.

From the outset of the project, the team determined that the third-party vendors should share the operational experience and, when necessary, responsibility for their products. Moreover, given the limited population of the hardware and software deployed, it was not cost effective to develop the in-depth expertise necessary if BT was to establish itself as best-in-class operators of this platform.

To achieve the required level of support, a series of agreements was reached with each of the suppliers, requiring all CSRS technical escalations to be made through the technical support team. Figure 9 shows the support arrangements in place with MCI and the other vendors.

With the exception of the AXE10 application software, MCI owns the application software on all the machines. Therefore, when a problem is diagnosed as being in the application software or if the problem cannot be clearly diagnosed as relating to operating system or hardware, escalations are made to MCI in the first instance. If, however, a problem is diagnosed as occurring in the operating system or hardware on any of the discrete platform elements, the escalation is made direct to the relevant vendor. In the case of some of the smaller suppliers, this facility is provided by DEC.

Where a complex problem requires contributions from several vendors, a virtual support team is established.

This team is led jointly by the technical support team and MCI.

Script and Data Management

The responsibilities of the script and data management team are:

- data build and configuration management of the cashless services platform,
- digitising and loading of scripts, and
- capacity management.

All data changes on the AXE10, DMS250 and DMS100 switches are made by this team. DMS100 data work is driven to a great extent by operator and customer service centre requests for change to the configuration of the ACD queues. In addition, all changes to configuration files on the ISN itself are made by the team. This is a critical area as a single line of data in a configuration file can influence the workings of the entire ISN platform.

Voice scripts are recorded in a studio and sent to Worthing where they are digitised and installed on the platform. This makes the deployment of new scripts considerably quicker and less expensive than on the 'old system'.

Change Control

Change control responsibilities are:

- management and control of software and hardware upgrades,
- management and control of high-risk work,
- data and configuration file audit,

- performance monitoring and reporting, and

- black-spot analysis.

The wide range of switch and computer systems and communications equipment means that there is a vast and disparate range of software (operating systems, application software and configuration files) in use. Tight control of change is essential to prevent conflict and mismatch. To manage such a platform and provide change control is an essential requirement requiring considerable attention.

To assist in the control of change, several audit tools and procedures have been developed.

Glasgow Command Centre

While the CSNOU has responsibility for the 'call processing' components of the platform, it should be noted that Computer Service Operations (CSO) has responsibility for the following:

- CSRS mainframe machines and applications,
- communication links between the ISN and the mainframes, and
- end-to-end integrity of customer data (CSS, DDS, RS6000).

The Glasgow command centre plays a lead role for CSO on CSRS. Overall it is essential that the working relationship between the two organisations is good. As such, the relationship forged between the CSNOU and the Glasgow command centre teams is key to the successful management of the CSRS platform.

Conclusion

The people involved in establishing the CSNOU were part of a much larger team brought together to implement a complex and sophisticated network. It is important to acknowledge and recognise the efforts of colleagues from the Head Office Project Team, Design and Build Directorate, Personal Communications and Computer Service Operations who together have delivered a network with a range of facilities comparable with any other in the world and with the flexibility to meet changing customer requirements and the challenge of competitors.

The CSNOU has proved that it is possible to manage the converging technologies of computing and telephony within one unit. It also clearly indicates that BT people are capable of speedily learning and adapting to new skills and working methods.

Glossary

ACD	Automatic call distributor
ARU	Automatic response unit
BDR	Billing data record
CCMIS	Call centre management information system
COU	Central operations unit
CPN	Call processing network
CSO	Computer Services Operations
CSDB	Cashless services database
CSNOU	Cashless services network operations unit
CSPU	Cashless services processor unit
CSRS	Cashless services replacement system
CSS	Customer services system
DDS	Data distribution system
DEC	Digital Equipment Company
DMSU	Digital main switching unit
EOSR	Enhanced operator services record
IP	Internet protocol
ISN	Intelligent services network
LAN	Local area network
MCI	Microwave Communications Incorporated
MF	Multi-frequency
MPRN	Multi-protocol router network
MTOC	Manual telephony operator console
NAMS	National activity management repairs system
NAS	Network audio server
NOU	Network operations unit
NIDSMON	Network information distributed service network alarm monitor
NT	Northern Telecom
PCM	Pulse code modulation
PSTN	Public switched telephone network
OSR	Operator services record
PIN	Personal identification number
TCP	Transmission control protocol
VCS	VAX cluster console system
WAN	Wide area network

Biography



Steve Ellett
BT Networks and
Systems

Steve Ellett joined the Post Office in 1966 as a Trainee Technician Apprentice. After completing a three-year apprenticeship he was assigned to circuit provision duties in the Brighton Telephone Area. In 1972, he entered management, carrying out PCM planning duties in the South East Regional Office. From 1979–1986, he was responsible for trunking and grading, TXE2 and TXE4 exchange design within the Regional Office, moving to System X data management in the latter part of this period. In 1986, Steve took responsibility for the System X support group in the newly formed South Downs District, then moved on to establish the District NOU in 1987. In 1991, he became the operations and surveillance centre manager in Worthing NOU. His latest career change in 1995 took him to the post of cashless services network operations manager within Worthing NOU.

Keith Beacham and Simon Barrington

CallMinder™

The Development of BT's New Telephone Answering Service

The CallMinder™ service, or 'answering machine in the network', is proving to be very popular with BT's customers. The introduction of the service broke new ground in being the first large-scale deployment of interactive voice platforms within BT's network. This article describes the service, the platform architecture and the development and deployment activities.

Introduction

CallMinder is one of BT's new Select Services. It combines the features of a normal telephone answering machine with the ability to answer calls on busy. The service is aimed at personal customers and single-line business customers.

Approximately 40% of all calls over BT's network either go unanswered or receive busy tone. By taking a message, CallMinder provides a valuable customer service and improves the call completion rate and hence network efficiency.

There are significant advantages for the user of the CallMinder service:

- no customer premises equipment required;
- simple access from home or remote telephone to retrieve messages or configure the service;
- simple and friendly dialogue with voice or TouchTone™† commands;
- no premium charged for usage: standard network tariffs for message recording—free access from home for message retrieval;
- the service takes messages both on no-reply and also on busy—this is a key differentiator from answering machines, particularly with the rapid introduction of data services such as fax and Internet, where a telephone line can be tied up for long periods, and yet CallMinder will continue to take messages for the user; and

- reliable and capacious message storage.

Initial human factors trials of CallMinder applications showed that usability was key to achieving market penetration for the service. For most users, CallMinder would represent their first experience of communicating with a machine.

Following a series of small-scale human factors trials, a marketing trial with 1000 customers commenced in Orpington in 1993. This trial was based on prototype hardware and produced valuable data on usability, usage patterns and service dimensioning.

The data was used to develop a business case for the national roll-out of the CallMinder service and refine the user interface still further. The subsequent roll-out programme included a pilot, and a limited launch within the London area prior to a national roll-out in 1995.

Main Features

The CallMinder application offers both a standard and an enhanced service. The standard service is targeted at residential customers and is described below. The enhanced service is targeted more at small businesses, providing greater mailbox size and greeting message length, and the simultaneous answering of a greater number of calls.

Call answering

In addition to answering calls on 'ring tone, no reply', CallMinder answers up to three calls simultaneously on 'busy'. This feature, not available from traditional answering

CallMinder™ is a registered trademark of British Telecommunications plc.

† TouchTone™, known technically as *dual-tone multiple frequency* (DTMF), refers to the tones sent when the keys are pressed on a tone dialling telephone.

A novel feature of CallMinder is that owners can retrieve their messages and change the features of the service by instructing the system using speech recognition.

machines, can be particularly effective for small businesses in managing their call flow and capturing enquiries.

Once the call is answered, the caller can be greeted with either a personal message from the CallMinder owner, or a pre-recorded BT greeting message.

When the caller leaves a message, they have the opportunity to listen to it and re-record it if they wish—very useful if they have said something they had not quite meant to! They can repeat this re-recording process up to four times, leaving a message of up to 5 minutes in length.

Message retrieval

CallMinder owners can retrieve their messages in one of two ways:

- message retrieval from their own telephone by dialling a short code access number (1571), and
- message retrieval from a remote telephone by dialling their own telephone number and entering a 4-digit PIN[†] number when the call is answered by CallMinder—this interrupts the CallMinder greeting message and switches the application into message retrieval mode.

In addition to being able to play, repeat and save or remove their messages, CallMinder owners can:

- review and change their personal greeting;
- change the 'ring tone, no reply' period before CallMinder answers their calls to one of four settings: zero, short, medium or long; corresponding to 0, 12, 21 or 30 seconds;
- change their PIN number (available from local message retrieval only);
- change to a 'fast track' message retrieval dialogue, targeted at experienced users; and

- change the addressee setting, which changes the caller dialogue to ask the name of the person who is calling.

Message waiting indication

CallMinder owners are informed that a message is waiting for them by two methods:

- changing an owner's dial tone to a distinctive 'stuttered' tone, and
- ringing the owner after messages have been taken on 'busy' and informing them that a message has been taken. The service will make up to four attempts to ring the owner, at 2, 5, 15 and 30 minutes after the message was taken.

Speech processing

A novel feature of CallMinder is that owners can retrieve their messages and change the features of the service by instructing the system using speech recognition. This provides a more natural interface than traditional TouchTone dialogues, and can be used on any telephone including the older 'pulse dial' types. To facilitate customer choice, TouchTone detection is provided in parallel with the speech recognition, and a caller may even use a mixture of the two in the same call.

The speech recognition used is speaker independent, allowing CallMinder to cater for a wide range of UK dialects without any user-specific training. The following words can be used:

'yes', 'no', 'yes please', 'no thank you';

'zero', 'short', 'medium', 'long';

'zero', 'oh', 'nought', 'one', 'two', 'three', 'four', 'five', 'six', 'seven', 'eight', 'nine'.

Recorded messages are stored at 8 kbit/s, using a state-of-the-art encoding algorithm developed at BT Laboratories. This algorithm provides high-quality recordings using a

bit rate which is significantly lower than the 32 kbit/s storage rate more commonly used. In subjective assessments against a range of analogue and digital answering machines, the 8 kbit/s encoding was found to provide higher quality recordings. The lower bit rate reduces the hard disk storage required for CallMinder and so lowers the platform cost.

Architecture

Overview

The CallMinder platform architecture is shown in Figure 1. The platform used is the interactive speech applications platform (ISAP)¹.

The platform is fully integrated into the network, connecting directly to the digital main switching units (DMSUs)* using BT standard C7 network signalling. The platforms are controlled remotely from existing network operations units (NOUs). Alarms are also extended remotely, via the BT network operations management system (NOMS1) alarm system. Customers are loaded automatically: from taking the order on BT's customer services system (CSS), through to the operations and maintenance centres (OMC) system and thence to the CallMinder platform. Service data is downloaded overnight to existing management information systems.

Interfaces

There are four primary interfaces to the platform, shown as A (network interface), B (data collection interface), C (operations interface) and D (alarm management interface) in Figure 1.

Network interface

Each ISAP is connected to two or three DMSUs. Each of the DMSUs connects via a single CCITT Signal-

[†] A 4-digit numeric password configurable by the owner.

* The backbone switching units in BT's UK network

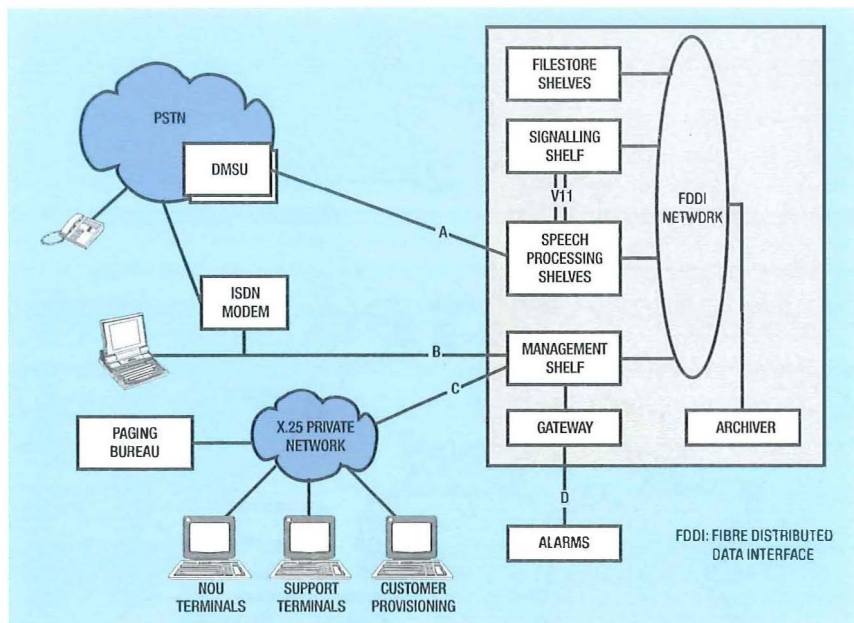


Figure 1 - CallMinder platform architecture

ling System No. 7 (C7) route, consisting of two links into the ISAP. Traffic can be routed to the platform via any of the attached DMSUs, and via either link of each DMSU-to-ISAP route. This therefore provides resilience against network routing, DMSU or platform failure, as traffic may route around such a problem.

Message deposit calls are routed to the platform by means of a divert which is set on the owner's line at their local exchange. The C7 information routed to the ISAP contains the caller's number, the owner's number (last diverted line identifier) and a service code to indicate to the ISAP whether the call diverted was on 'busy' or 'ring tone, no reply'.

The C7 nodal end-to-end data protocol is used to change the owner's 'ring tone, no reply' period and to apply and revoke changed dial tone to their line.

Data collection interface

Statistics collection is performed over an ISDN2 interface. Collection is undertaken overnight using standard file-transfer utilities.

Statistical data collected over the course of a day is downloaded to the Tinsley Park Computer Centre where it is loaded into an Oracle database. A management information system allows subsequent access to this data over BT's multi-protocol router network.

Operations interface

Each ISAP has an X.25 interface to BT's private data network. This

network allows closed user group access to the platform for system and service management purposes.

- *Automated customer provision* The (automated) customer provision process overcomes the problems of errors, cost and throughput associated with manual systems. When the customer's request for CallMinder is entered into the BT CSS, it is forwarded electronically to the appropriate BT OMC. The OMC software first creates the mailbox on the ISAP. If this is successful, then the OMC invokes the administration-controlled diversion on the customer's local exchange. The OMC interface to the ISAP is via X.25 over BT's private network. The interface can support up to 1000 customer transactions per day on each ISAP and can support:

- customer provision,
- customer audits,
- customer cessation,
- change to customer's directory number, and
- alteration of customer's CallMinder product.

- *System management* System management is menu driven using a character-based system. There are five levels of support:

-system management, allowing:

access to error logs, ability to reboot/shutdown system components (for example, speech shelves), C7 maintenance actions (for example, restart link), and alarm management;

-security, allowing:

account management, X.25 management, and password management;

-support, allowing:

platform operating system access, and access to system and service management and data collection commands;

-data collection, allowing:

access to platform statistics (for example, traffic levels);

-service management, allowing:

access to customer provision cessation facilities, and service statistics (for example, number of customers, usage).

Alarm management interface

Alarms are forwarded to the NOMS1 alarm-reporting system via a gateway system or via an alarm-collection unit. These standard BT alarm boxes collect alarms in the form of signals on input wires (0 = active, 1 = inactive) and send the outputs via a modem to BT's NOMS1 system. Alarms are prioritised into one of four levels: FATAL, ERROR, WARNING or INFORMATION.

Alarms can be viewed and reset via the system management menus.

Dimensioning

The CallMinder system has been designed to be easily extended. The initial installation was dimensioned to meet the first year traffic forecasts.

Figure 2—CallMinder UK sites

Individual platforms can be extended in capacity, or additional platforms can be installed in current or new sites.

Network dimensioning

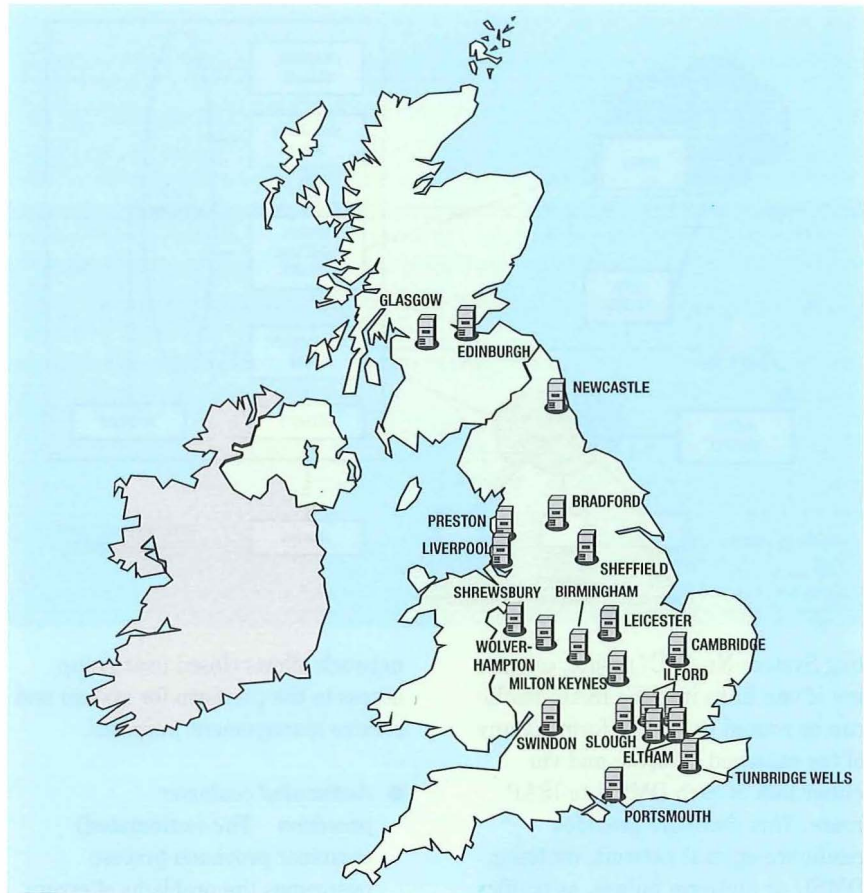
Based on the forecast traffic, and assuming an even distribution across the UK population, a network costing analysis showed that the optimum configuration would be for each CallMinder platform to serve three DMSUs, except in London where the optimum configuration would be four platforms serving the eight DMSUs. Thus the total number of platforms required was 21, spread nationally and located at 19 sites (see Figure 2).

Platform dimensioning

The forecast take-up of the service required the total initial installed capacity to support approximately 900 000 users. This capacity was spread evenly between the sites with the four London ISAPs dimensioned 25% larger than the other installations.

Platform dimensioning took account of many factors, including:

- number of mailbox owners;
- message storage capacity;
- call arrival rate;
- typical call duration;
- distribution of call types—message storage, message retrieval, user configuration;
- typical use of speech resources (for example, usage of speech recogniser, speech encoder and decoder);
- outgoing calling rate (for example, for message waiting indication calls and setting stutter dial tone);
- service provisioning rate (for example, number of new users per hour, day);
- typical data collection usage; and



- typical system and service management usage.

This information was derived from the business-case forecasts and from data from the marketing trials.

Reliability and availability were identified as key market requirements, and therefore the CallMinder platforms have been designed to be resilient to individual component failure. Some of the resilient features included are:

- disk mirroring for all data storage;
- $N + 1$ redundancy for speech processing components;
- hot standby of network signalling components;
- support for dynamic signalling link re-routing;
- uninterruptible power supplies, and $N + 1$ redundancy of power supply components;
- dual attached optical-fibre ring as platform local area network; and
- multiple and diverse network and operational connections.

Table 1 Initial Installation Configurations

Shelf type	Number of Shelves	
	London	Other
Speech Processing	6	5
Filestore	4	3
Management	1	1
C7	1	1
Archiver	1	1

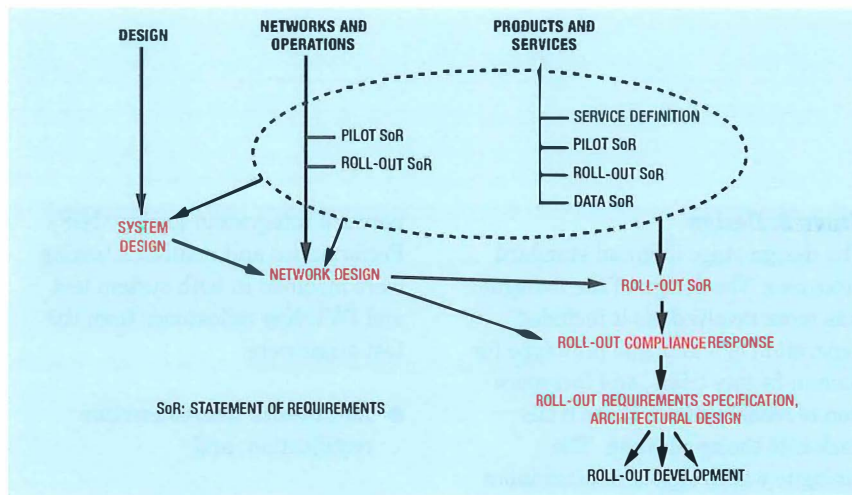
The ISAP architecture is described in reference 1. For information, the configurations used for the initial installations are as shown in Table 1.

Managing The Development

Introduction

After the evaluation of tenders for voice platforms from leading international suppliers, the ISAP bid by Ericsson Ltd. was chosen on the grounds of cost, functionality and future enhancements. The ISAP had originally been developed by BT Laboratories and subsequently licensed to Ericsson Ltd, who then acted as the prime contractor for delivery of the CallMinder service. BT Laboratories acted as a subcontractor for the delivery of the application on a fixed-price basis.

Figure 3—Requirements map



The CallMinder roll-out development was planned in three phases:

- a single platform pilot service to prove the platform, operational feasibility and marketing case;
- a limited roll-out to the four London sites to gain installation and operational experience (the service would not be officially launched at this stage); and
- a full roll-out to all the national sites, and a national launch.

Pilot service

The objective of the CallMinder pilot service was to prove the new ISAP platform technology, and trial the network integration and operational processes proposed. The service was based on a single platform sited at Ealing, with 1000 customers transferred from the 1993 Orpington market trial.

The pilot development followed a similar life cycle to that described below, although simplified to reflect the single site trial nature of the delivery.

A key element of the pilot development was the transfer of ISAP manufacture and hardware design and support knowledge to Ericsson Ltd.

In order to reduce launch timescales, the pilot development overlapped the roll-out development to some extent, requiring careful management of the two projects.

The pilot development commenced in June 1993, with the service launched on schedule in March 1994. The platform remained in service until deployment of the London roll-out service in October 1994. As well as providing an operational test bed, the pilot service was also used to fine-tune the marketing campaign for national launch, such as investigating customer price sensitivity.

Roll-out delivery structure

The roll-out of the CallMinder service was dependent on the delivery of a number of new components:

- ISAP Build 4.0—the core platform hardware and software build,
- Visage Build 3.2—the service creation tool for dialogue development,
- CallMinder application,
- OMC Build Q—customer provisioning software,
- CSS release 23—customer order handling software,
- AXE10 NEP—Ericsson local exchange software, and
- System X 4900 Build—GPT† local and main exchange software.

In addition, new customer service and operational practices needed to be developed across BT.

The integration of such a large number of new components and processes carried considerable risk, which required careful management. Both CSS and OMC changes were delivered through existing programmes, although their activities were coordinated with the CallMinder application development through an umbrella programme.

The Networks and Systems area within BT acted as the delivery agents and managed contracts let to GPT (for the System X switch changes) and Ericsson Ltd. As well as acting as prime contractor for the delivery of the CallMinder platform and application, Ericsson was also

† GEC Plessey Telecommunications—suppliers of the System X main and local network switches to BT.

responsible for AXE10 switch changes. Ericsson subcontracted the platform software and CallMinder application development to BT Laboratories. The platform hardware design and manufacture were transferred to Ericsson as part of the pilot delivery, although BT Laboratories provided design support as necessary.

The application development contract was a fixed-price fixed-timescale contract, involving various parts of BT and Ericsson in a relationship that, in practice, worked very well.

Application development

The Callminder application software runs on the ISAP, and manages the customer dialogue, the storage and retrieval of messages, customer provisioning, statistics and element management. There are several components to the application, running on different processors in the platform and managing the platform interfaces as described earlier in this article. The application development project was combined for the London and national roll-out deliveries, and split into seven stages:

Stage 0: Project initiation and support

Stage 1: Definition

The requirements capture process derived a single statement of requirements, compliance response and requirements specification from eight separate and often conflicting source documents, as shown in Figure 3. Close teamworking across BT and with Ericsson was necessary to achieve this.

Stage 2: Design

The design stage followed standard processes. The design of the dialogue was more involved, as it included generation of a dialogue prototype for human factors trials, and incorporation of feedback from those trials back into the application. The dialogue was designed in accordance with the BT Voice Applications Style Guide².

Stage 3: Implementation

The implementation stage included module code and test, and integration and acceptance of application and ISAP software into a stable system prior to commencement of the test stage.

Stage 4: Test

The test stage was the culmination of verification, validation and testing activities spread over the entire life cycle of the project, which included:

- requirements testability analysis,
- requirements tracking,
- design reviews,
- documentation reviews,
- code walk-throughs,
- test harnesses,
- call generator (specifically commissioned for this project),
- signalling test bed, and
- commissioning of a full-sized reference model.

The test stage was very complex and had to meet requirements including system testing, documentation validation, network interconnect validation testing (IVT), network interworking testing (IWT), system integration testing and formal development acceptance testing. The IWT activity included the integration of the ISAP with OMC, CSS, AXE10 and System X elements in the

network integration facility (NIF). Performance and resilience testing were included in both system test and IWT. Key milestones from the test stage were:

- BT network interconnection certification, and
- BT formal development acceptance.

Without interconnection certification, the system would have been prohibited from connection into BT's network.

Stage 5: London roll-out release

The approved platform and application software was released to Ericsson for distribution to CallMinder London sites. In parallel with the software development, Ericsson had been manufacturing, installing and commissioning ISAPs in each of the sites around the country.

One of the main risks to successful roll-out was identified as training. To reduce the risk, a joint BT and Ericsson team, in addition to the planned training events, acted both as installation and commissioning agents, and as first-line support to the novice operations teams. This proved critical to the success of the project.

Stage 6: National roll-out test and release

The national roll-out software was essentially the same as that for London release, but needed to be validated against different AXE10, System X and OMC builds.

The national release followed the same process as for London roll-out. The special joint team again proved invaluable as they toured the country installing and commissioning the systems and supporting the local maintenance teams.

Roll-out

The CallMinder roll-out development was delivered on schedule, and in fact installation was achieved slightly

earlier than planned. After a series of trial transfers to ensure no disruption to service, customers were transferred from the Ealing pilot system to the four London roll-out sites in October 1994. These sites were then upgraded to national roll-out software with all national roll-out sites operational by mid-December 1994.

Before the service could be launched, reliable operation and operational and customer services needed to be proven on all sites. This process highlighted the difficulties of the fast deployment of such a complex and advanced system. Some of these roll-out issues were:

- the scale of the roll-out, with a total of 21 platforms to be simultaneously installed and launched from 19 sites nationwide, severely stretching management, reporting and support processes;
- minor operational and network data configuration differences across the country, causing service problems which had not been encountered during the testing phase or London roll-out,
- a large-scale training programme which had to be fitted around existing commitments, and
- the roll-out and subsequent launch, often requiring the resolution of conflicting priorities across the business.

One of the critical areas for service launch was the customer order taking and provisioning process. Because of the large number of people and systems, much time had to be spent bringing this process on line.

Launch

After national roll-out, and prior to the service launch which was delayed because of operational difficulties and the discovery of a faulty batch of hard

disks (used for message storage on the ISAP platforms), some experience of platform operation was provided by 1000 BT people being recruited to take up the service. The service was formally launched nationwide in May 1995. Customer take-up nationwide has been extremely rapid, and has so exceeded expectations that the service is currently sold out. Expansion plans are currently in progress.

User feedback has been extremely positive, with customers appreciative of the service's ability to answer calls on 'busy' and its user friendliness.

The consumer magazine *Which* carried out a survey of answering machines in July 1996 and reported that 'BT's (CallMinder) answering service beat all standalone machines for sound quality, ease of use It was the best on test and is good value for money.'

The Future

The immediate future for CallMinder is the extension of the current platforms to cope with demand.

In the medium term, enhancements are planned:

- to allow better feature interworking with other Select Services, particularly call diversion and call waiting;
- to allow message retrieval on message waiting indication calls; and
- to allow different personal greetings for 'ring tone no reply' and 'busy'.

Other enhancements are likely to see different product sets of CallMinder aimed at particular market sectors such as students and small businesses.

Additional features being considered include

- fax store and forward,
- memo facility,

- message waiting indication to use ring back when free,
- call screening, and
- paging message-waiting indication.

Conclusion

The development of the CallMinder service has provided many technical, commercial, operational, management and marketing challenges. The successful delivery of CallMinder into service is a tribute to the many teams across BT and its suppliers that have worked together to overcome those challenges.

With high service take-up rates and good customer feedback, the future of the CallMinder service is looking very bright.

References

- 1 ROSE, KEVIN; and HUGHES, Pat. The BT Speech Applications Platform. *Br. Telecommun. Eng.*, Jan. 1995, 13, p. 304.
- 2 Voice Application Style Guide. BT internal publication. A summary of this guide has been published as BT publication 'Now You're Talking... Voice Services', Ref. No. 367, 1994.

Biographies



Keith Beacham
BT Networks and
Systems

Keith Beacham received an M.A. in Electrical and Electronic Engineering from Cambridge in 1985. Since then he has worked in a number of BT Divisions on voice and messaging products, and systems integration. He was the project manager for the CallMinder application development, and currently manages the Voice Services unit at BT Laboratories.



Simon Barrington
BT Networks and
Systems

Simon Barrington received a B.Sc. in Physics at The University of Wales in 1986. Since then, he has worked at BT Laboratories on the development of interactive voice systems. He was the team leader and system design authority for the CallMinder application development, and is currently project manager for the ACE automated customer handling service developments.

MeterLink™

BT's New Telemetry Service for Meter Reading

MeterLink™ is a new BT service to read utility meters (water, electricity, gas). The BT network has been modified to give telemetry capability without affecting normal telephone service by the use of 'no-ring calls'. Accessed through the telemetry platform, no-ring calls allow data to be communicated to and from customers' premises while the line is otherwise not in use. The main focus of this article is on the telemetry platform, which uses the interactive speech applications platform (ISAP), and the telemetry interface unit, which terminates telemetry calls at a customer's meter location.

Introduction

MeterLink™ is a new BT service for reading utility meters (electric, gas, water) over the public switched telephone network (PSTN). It is the first of a number of potential applications of network telemetry¹, transparently sharing telephone lines to collect and deliver relatively small amounts of data without affecting normal telephone calls. This article highlights the network elements of the underlying BT telemetry service and the role of the interactive speech applications platform (ISAP) as a PSTN calling engine and speech platform for the service.

This service is aimed at the 70 or so privatised utility companies, who are experiencing new competitive and regulatory pressures. Electricity consumers of over 1 MW on a site (about 5000 sites nationwide) have been able to choose their supplier of both electricity and metering since 1990. From April 1994, this was extended to consumers of over 100 kW/site (about 50 000 nationwide) and by 1998 it is expected that electricity will become an open market with all consumers free to choose their supplier.

A similar situation is emerging in the gas market, where large gas consumers are already purchasing gas via the British Gas transport network. Competition in water supply is less developed and the regulatory focus to date has been on improving the quality of service while containing costs. However, the water industry is under considerable political pressure to introduce metering to domestic customers, which has heightened interest in telemetry.

For the utilities, remote meter reading is only attractive if it can save costs over the current mainly manual methods. Efficient processes have been developed whereby consumers take some readings and others are estimated so that the cost of manual reading is low. The utilities are, however, coming under regulatory pressure to read more accurately and more often. MeterLink has been designed to be cost-competitive with manual reading—where more actual readings are required, the costs become even more attractive.

For the utility regulator, the BT telemetry service offers a very effective means of facilitating competition. Suitably equipped meters can be read throughout the UK, unlike broadcast telemetry systems which are limited in geographic coverage. Security is built into the system so that the utilities can only read meters to which access rights are owned. This facility is intended to preclude anti-competitive behaviour.

For BT, telemetry is a way of using the network when it might otherwise be idle, especially overnight. MeterLink is but one of a number of services which will be developed using this platform.

Systems Architecture and Major Components

The technical requirements of the telemetry system, simply stated, are:

- to automatically originate large numbers of calls per day (there are 23 million households in the UK and many with two or three meters each to be read at least quarterly);

Figure 1—MeterLink systems architecture

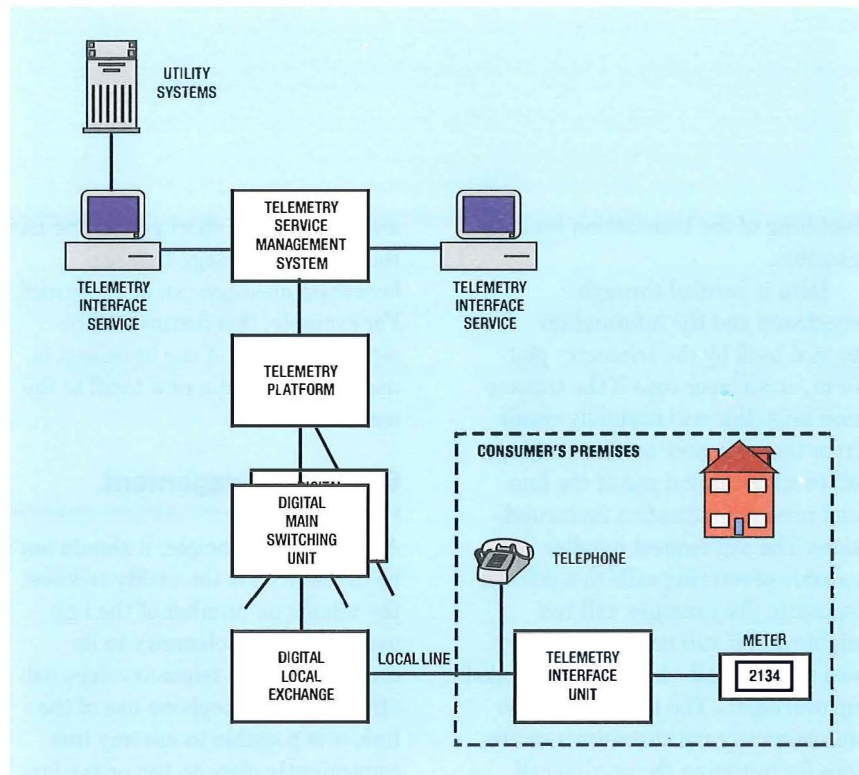
- to route these calls to the appropriate location, regardless of any network diversions the telephone user may require;
- to connect to and transfer data with a meter without ringing the consumer's telephone or interfering with normal telephone service ('no-ring calls');
- to route data back to the utility/MeterLink customer; and
- to charge for the service.

In addition, the utilities as well as BT are subject to regulatory controls and data must be segregated for data-protection and competition reasons. BT must also consider the impact on the network of originating many calls automatically and build in safeguards so that the service will not cause congestion for normal calls. In the future, it will be necessary to deliver telemetry over other networks including ISDN and other licensed operator networks.

These factors among others led to the systems architecture illustrated in Figure 1.

The functions of the systems elements in Figure 1 are given below.

- The telemetry interface system (TIS)² is a computer which is configurable to interface with the utility's existing systems and gives control of the service.
- The telemetry service management system holds data for the service and initiates call requests to the telemetry platform.
- Human assistance for users of the service is given by the telemetry service centre, which has terminals to monitor and operate the service.
- The telemetry platform is based on the interactive speech applications platform (ISAP) and is the calling engine for the whole telemetry



service. Multiple telemetry platforms will be deployed as the service grows which will allow for very large numbers of calls; different types will also allow for other transport technologies.

- Calls are routed into the network via a digital main switching unit (DMSU) which accepts CCITT SSC7 signalling and routes the calls from the telemetry platform. Two connections are used into the DMSU for resilience. Calls are then routed via digital local exchanges which have been modified to accept no-ring calls.
- No-ring calls cause tones to be sent to the customer's line but without the ringing which alerts the telephone customer to an incoming call. These tones instead activate a telemetry interface unit³ which answers the call and then interfaces data between meters connected to the interface unit and the telemetry platform.

System Functions

The MeterLink service features the following high-level functions.

Bulk read

Regular meter readings will be most efficiently made in bulk and the

utility has the facility of creating lists of meters which can be read on a regular cycle. Lists can be created on the telemetry interface system and downloaded to the telemetry service management system for regular reading—the customer will normally consult telemetry service centre people to ensure that the correct priority is attached and terms of the contract agreed.

Lists of meters to be read will normally be downloaded on to the telemetry platform only when needed and will have a time window attached within which all readings should be made. The call request handler can then prioritise calls in order to give the best match of demand to resources and begin making calls on the list starting at the specified time.

Each call is associated with low-level network routing data and passed through the call rate control before handling by the call-connection/data-handling process. This makes no-ring calls to the required telemetry interface unit and handles the resulting modem data transaction according to rules stored in the template for this kind of connection. Templates will have parameters passed and will translate data to a required format—a future upgrade will allow scripts to be run within the template so that intelligent

handling of the transaction becomes possible.

Data is verified through a checksum and the information passed back by the telemetry platform, or an error code if the transaction fails (this will normally result from the telemetry service backing off to allow normal use of the line and need not indicate a fault condition). The call request handler is capable of retrying calls to a defined sequence (for example, call two nights and if still no connection try one daytime call—useful for lines tied up overnight). The system can also automatically try alternative strategies for initiating the no-ring call including higher level wake-up tones and preceding the call with line reversal (which may cause bell tinkle and so reverts to normal mode if possible).

On-line reading

The system may also be required to make a single reading as soon as possible; for example, for home moves. The call sequence is exactly the same as above except that the highest priority is attached to the call. The call may still fail because the line is busy, but this information is passed back to the telemetry interface system so that the operator is informed of the reason why the meter cannot be read. It should also be noted that service management of the line during a house move has been considered in the design—changes in BT data will be coordinated so that contact with the meter is not lost in this transitional period.

Broadcast

A broadcast message can be initiated in much the same way as a bulk read; that is, a list of meters to be contacted and each sent the same message via a template. The essential difference is that calls must be made in rapid sequence and might be initiated from more than one telemetry platform in parallel. This emulates the facility of broadcast systems to contact a number of

meters in a very short period and has the added advantage that the broadcast message can be confirmed. For example, this feature will be especially useful if the broadcast is used to download a new tariff to the meter.

Service Management

As a general principle, it should not be necessary for the utility to know the telephone number of the line used to deliver telemetry to its consumer. Since telemetry does not affect normal telephone use of the line, it is possible to use any line conveniently close to the meter for telemetry. This line could give telephone service to someone quite unrelated to the meter user, or could be ex-directory. Utilities will have their own reference for consumers which must be mapped across to the appropriate telephone number by

interface to the telemetry service is specified (SIN252) to facilitate use of the service.

The service also has human operators who can help the utility to set up and use the service, via systems very similar to the telemetry interface system.

Utilities wishing to add meters to the service will be required to provide (electronically via the telemetry interface system, where possible) lists of consumer addresses for the MeterLink service which will be matched with BT data to ensure that there is a suitable line on the premises. Where a line is not present on the premises, it may be possible to use an adjacent line, in which case a site survey will be needed. However, the bulk of installations will be processed automatically through normal line records. The telemetry service management system will then return lists of consumers who are

telemetry does not affect normal telephone use of the line

the MeterLink service. The number mapping must then be maintained against telephone users/consumers moving house, changing network operator or simply going ex-directory.

The platform required for this function is the core of the service management system. Its primary function is that of holding consumer details and mapping these on to routing data for use by the network systems as required. Changes to consumer data are handled by links to BT's customer service systems (CSS), which have themselves been modified to handle telemetry. The telemetry service management system is also the interface to the BT telemetry service via the telemetry platform, issuing requests for calls to/from meters and handling the data.

Utilities will control the MeterLink service through a telemetry interface system. This interface is customisable by systems integrators to fit in with the utilities' operations. An open

pre-enrolled ready for installation of meters.

BT is not planning to provide and fit meters, but has instead provided automatic enrolment processes so that the fitters chosen by the utility can easily enrol meters themselves. This involves entering details about the specific meter and interface unit fitted on the premises, checked against the calling line identification of the line, so that the service can use the correct protocol and can check that these details match on later readings.

Once successfully enrolled and a check reading taken, utilities will be able to initiate meter readings through the TIS as and when required.

Telemetry Platform

Why choose ISAP as the platform?

The telemetry platform is the core calling engine for telemetry services.

Its requirements for PSTN use are that it:

- is capable of safe attachment to the core BT network (that is, without risk of degrading core network call transport);
- has many modem channels for parallel calls to give a high calling rate;
- has modems capable of upgrade for emerging customer requirements;
- is expandable, reproducible (multiple platforms will be needed to give a scalable service); and
- is capable of voice dialogue for enrolment (also permits other services).

Initial design studies in 1993 considered a number of options for this data application. These studies showed that the highly flexible digital signal processing (DSP) capabilities of the ISAP platform⁴ were ideal for PSTN telemetry. The platform also had many other strong points:

- each speech shelf can handle a large number of modems (108 V.23 modems), and many shelves can be incorporated into a single platform giving the capability of incremental growth as the service takes off;
- powerful multiprocessor UNIX environment, required for the many internal look-up and call-processing functions to proceed while maximising the use of DSP resources;
- reuse of core functionality of the ISAP to increase development efficiency and effectiveness (reduced risk over an entirely new platform);
- network hardened and certified (comprehensively tested to ensure that it cannot damage the network);

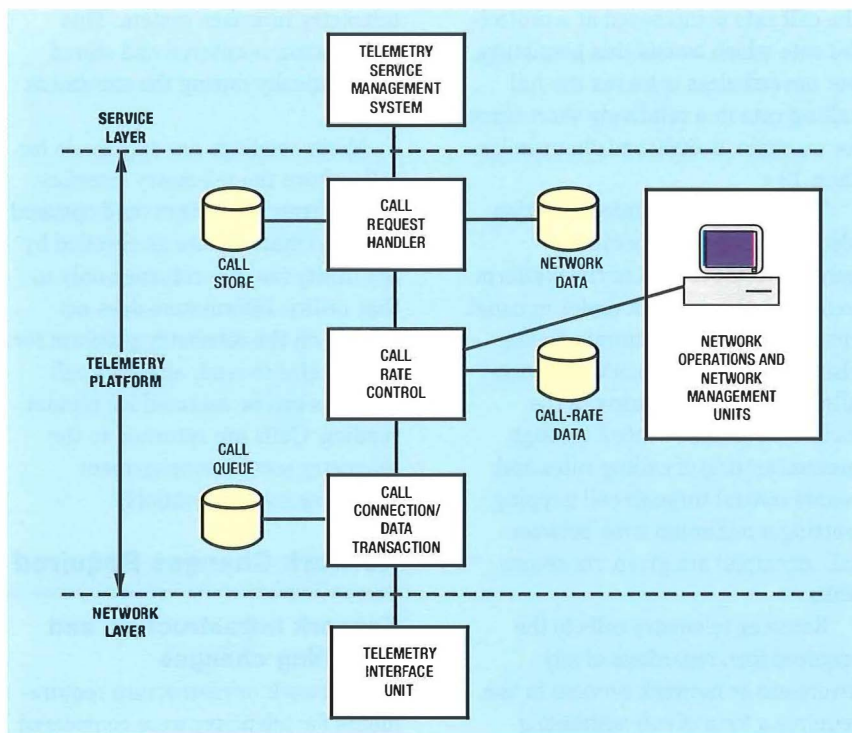


Figure 2 – Telemetry application architecture

- network management capability; and
- speech capability inherent, which gave the opportunity to automate the enrolment process and reduce the cost of meter installation.

The platform is also used for other applications, notably CallMinder, so that it will be familiar to operations personnel. These reasons led naturally to the choice of ISAP as the preferred platform for PSTN telemetry.

Internal operations of the telemetry platform

A simplified view of the internal operations of the telemetry platform is given in Figure 2.

Calls are received from the telemetry service management system and held in a queue. They will have different priorities depending on the service required but the system periodically increases the priority of calls so that the lowest priority calls will get handled after a time.

The generation of calls is carefully controlled through the calling rate

control function⁵. Currently, the telemetry platform is capable of originating 6000 calls/hour and this is expected to increase substantially as the software is optimised and resources added. Large batches of calls will become due at specific times of day or night at times specified by

The generation of calls is carefully controlled through the calling rate control function

the utilities and, unless managed by this function, would cause step increases in the numbers of calls originated in the network. As this is not a normal calling pattern, it would normally register as a fault condition, in which case protection algorithms in the network would prevent these calls from being forwarded. Instead,

the call rate is increased at a controlled rate which avoids this possibility, but nevertheless achieves the full calling rate in a relatively short time; for example, 0–6000 calls/hour in less than 13 s.

The calling rate control function also gives network operations personnel full control of the platform to protect the network under unusual circumstances; for example, faults elsewhere in the network or unusually high calling patterns in the vicinity. Both fine control through manual setting of calling rates and coarse control through call gapping (setting a minimum time between call attempts) are given via secure links.

Routeing telemetry calls to the required line, regardless of any diversions or network services in use, requires a form of sub-addressing within the digital local exchange. This requires a translation of the consumer's telephone number to a host network address. Because of the low-level nature of this address (to a specific concentrator and line card) the address length is variable and the number itself subject to change with engineering work on the network. A table mapping telephone numbers to host network addresses is held within the telemetry platform (network data in Figure 2) and maintained by network operations personnel.

Call connection/data transaction is via the bank of digital signal processing (DSP) modems which forms the heart of the telemetry platform⁶. Here a scheme of templates⁷ is employed to translate high-level metering instructions to the protocol required by a specific type of meter and telemetry interface unit^{8,9}. This allows the telemetry interface system to request calls in a common format regardless of the equipment installed. Information held on the telemetry service management system invokes the appropriate template run on the telemetry platform to communicate with the installed meter and telemetry interface unit and return results in a format suitable for use by the

telemetry interface system. This information is entered and stored automatically during the enrolment process.

Meter readings are only made for calls where the telemetry interface unit reference numbers (and optional security) match those as enrolled by the utility and are returned only to that utility. Information does not remain on the telemetry platform for anyone else to read, although call requests can be retained for regular reading. Calls are returned to the telemetry service management system as soon as available.

Network Changes Required

Network infrastructure and signalling changes

The network infrastructure requirements for telemetry were considered along with the wider upgrade of the UK network needed for calling line identification (CLI). CLI requires a limited signalling facility in all digital local exchanges in order to pass calling numbers to customers' telephones. In the UK implementation of CLI, this information is passed before any ringing is applied to line; that is, as a no-ring call. One benefit of this is that it can facilitate automatic routeing of the call on customers' premises in accordance with the CLI number.

CLI is an embedded feature in the digital local exchanges which support it, and required upgrade of all line cards to include hardware for no-ring calls. Accessing this feature for the telemetry service requires low-level access to this hardware and is via the host network address, as described earlier.

Several classes of no-ring call are available. The preferred type uses the line in entirely the idle state and the meter interface unit operates at a current level below that which will seize the line. This allows the most robust scheme for detecting that the telephone customer wants to initiate a call because looping the line is easily detected and the telemetry

service can back-off. However, there will be circumstances where more line current is required and a loop-answer mode is permitted where the meter interface unit may take normal telephone line current. In this case, both the telemetry platform and meter interface unit must continuously monitor line signal levels and drop the telemetry call if the telephone user lifts a handset.

Each no-ring call is initiated with a wake-up tone which indicates the class of service required (different tones are used for CLI). Experience in the US, with their simpler form of CLASS (customer local area signalling services) signalling, has shown that it is less robust and more prone to false triggering. The BT system uses a dual wake-up tone with careful selection of the frequencies to avoid duplication of signalling tones used worldwide. CLASS signalling also normally includes a line reversal, which is easy to detect and confirm on the tone. This, however, can give rise to bell-tinkle, normally masked when ringing starts after the CLI is received. In the case of telemetry, there is no following ringing and the line is likely to be used at night; the preferred mode of operation is therefore without line reversal to avoid bell tinkle where possible.

Before call set-up can be completed, the exchange must detect that the local connection has been successfully made. This is indicated by returning a DTMF† 'C' tone to the exchange, which can then complete the connection. Loop termination no-ring calls are of course detected by loop-answer as for normal calls. The no-ring call is then connected through to the telemetry platform.

† Dual-tone multi-frequency signalling, as used in modem telephone. A matrix of four by four tones is defined, used two at a time to give the normal keypad 0–9 and *, # keys. Four combinations are not used on the keypad but are available for CLASS and telemetry signalling, referred to as 'A' through 'D'.

The types of no-ring call are listed in Table 1.

Table 1 Types of No-Ring Call

No line reversal	Line reversal	
1	2	idle line
3	4	loop answer

Effects of MeterLink on the network

Telemetry services provide both an opportunity but also a potential problem for the UK PSTN. If the service is adopted widely for domestic metering, then very large call volumes are anticipated. Call volume will grow, especially if, as intended, other applications are added. These calls could swamp the PSTN if they are concentrated in time or location such as to exceed network capacity locally (the network as a whole has ample capacity for any foreseeable eventuality). The telemetry platform could also conceivably develop faults

Telemetry can use spare capacity overnight

(for example continuous looping through a batch to generate erroneously high call rates). For these reasons, a number of measures have to be included to protect the network against excess network loading.

The first level of control is to work with customers and set tariffs which encourage use of the network in less busy times. Figure 3 shows a typical calling pattern where the network is busy for large parts of the day and (increasingly) in the early evening when social calls are made and prices are lower.

Telemetry can use spare capacity overnight and for periods during the day—there is also always capacity for reasonable numbers of calls at any time for immediate meter reading.

Where it is necessary to place a large number of calls at the same time (for example, a broadcast to meters in a town), the application will scatter calls among the list of host network addresses to be used to avoid overloading single points in the network.

As the network evolves it will also be possible to scatter calls between telemetry platforms to avoid overloading single routes. In exceptional circumstances, for instance, television programme response, which can generate huge amounts of calls in a short period, it will be necessary for network operations experts to set call controls on the telemetry service to preserve the PSTN performance; a measure of control is also built in for automatic operation to cover unforeseen and fault conditions.

Managing these risks has been an important consideration from the beginning of this development. In order to ensure that the required high standards are met, all network-affecting elements of the overall system have been tested and the telemetry platform and application subjected to particularly thorough testing, including provocative testing (attempts to induce network-affecting faults by, for example, removing live links/cards). While this is a tough regime within which to develop applications, the contribution of testing, has been highly constructive and networks test specialists con-

tinue to support the service during the current phase of rapid evolution to help ensure that the network certification given to MeterLink is maintained.

Telemetry in the Consumer's Premises

Connection to the telephone line

In order to receive telemetry calls, a telemetry interface unit (TIU) must be installed where it has access to both a telephone line and the utility meter. Normally, the TIU will be connected to the consumer's own telephone line; however, since the service is transparent to telephone use, this is not necessary and any convenient and suitable line will do.

Access to the line is via the standard line socket allowing telemetry interface units to be supplied and fitted by any organisation which can meet BABT requirements. The equipment will therefore be the property of the installer, but has the potential to disrupt service to the telephone user should a fault develop. Connecting via a line socket gives the user the ability to disconnect the telemetry interface unit under these circumstances.

It should be noted that this facility to disconnect the telemetry interface unit did cause some concern to the design team since this might allow some consumers to try to frustrate

Figure 3—Calls in network (indicative)




Figure 4—BT telemetry interface unit

the meter reading service. However, advice from the utility experts indicated that a hard-wired connection from the line to the meter interface unit would be no more reliable since the type of customer who would frustrate the service is perfectly capable of cutting the wire.

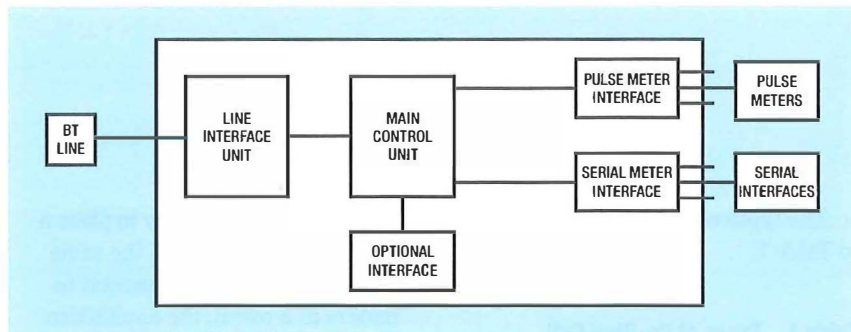
Telemetry interface unit

The telemetry interface unit (see Figure 4) has the function of accepting telemetry no-ring calls and converting meter readings into a form suitable for transmission over the network. The preferred mode of operation is the 'on-hook' mode where the TIU takes less than 2.5 mA in operation (less than 30 μ A on standby).

BT has developed a TIU¹⁰ as proof that line-powered operation is feasible and to expedite early roll-out; however, the service has built-in flexibility for other modem standards and message protocols. The BT TIU will be used as an example for the rest of this section.

The unit is connected to the standard telephone line jack unit. The low standby current taken means that it will not affect the use of other telephone equipment in the house and does not reduce the total ringer equivalence number (REN) available to the telephone customer (REN is a measure of the loading of the line by telephones especially for ringing—up to a REN of four is allowed per line and the telephone user should not fit telephones in excess of this total).

The line interface unit contains all telephone termination and powering functions and is required to meet the provisions of BS 6319, etc., for connection to the UK network. It also contains a modem capable of V.23 signalling, half duplex 1200 bit/s in both directions plus 390 Hz back-channel carrier (as used in the V.23 standard, required for parallel off-hook detection when the TIU is used in off-hook mode). The service requires an MF generator for signalling call set-up, and wake-up tone detection.



The wake-up tone has been specially selected to avoid false triggering. A guard tone of 1827 Hz is used for all calls, coupled with select tones of 520.6, 578.4 and 468 Hz for addressing up to three TIUs on the same line (more could be defined later). The select tone is cadenced 300 ms ON, 100 ms ON and 100 ms OFF to select the off-hook mode. This can be used if transmission proves unreliable for the on-hook mode and will be selected by the telemetry platform automatically if service on-hook is not reliable.

The main control unit contains a simple microprocessor and memory for supporting the message protocols required by the service. While the service is adaptable to any protocol, it is recommended that the RITA¹¹ protocol is used—this has been created in such a way that it makes efficient use of simple microprocessor memory mapping and is easy to monitor using simple protocol analysers.

Interface ports are provided for several pulse or serial interface meters. Pulse inputs simply count the number of contact closures, which gives a measure of incremental metering. The serial interface supports the IEC FLAG protocol, which is in wide use on electricity meters and has the benefits of standardisation and inherent safety isolation through the infrared interface used.

Other interfaces are provided for a hand-held terminal (to ease set-up on installation) and a display unit which can optionally give the customer confirmation of the meter reading taken.

Service Deployment

MeterLink is currently being trialled by major gas, regional electricity and water utilities. Even at the current

early stage of installation, many thousands of telemetry calls are made daily.

Currently, two telemetry platforms are deployed (for resilience) which gives sufficient capacity for the first four years of expected growth. The systems architecture has flexibility to increase ISAPs in line with business growth and share other ISAPs (for example, CallMinder) when these have spare capacity, say at night.

The BT-designed telemetry interface units have been BABT approved for limited production and nearly 200 have been tested in the Doncaster area. These units have proved reliable in use and have been of great value in testing the system under real operational conditions. Interface units have also been produced by GPT Ltd for major gas users and a number have been deployed to support trial MeterLink services. Other manufacturers have expressed an interest and are expected to supply interface units for telemetry.

Conclusions

This article has given a top-level description of the systems which are needed to deliver BT telemetry and the MeterLink service. The telemetry platform has featured as an example of an application of the interactive speech applications platform which is concerned with data as well as speech.

The ISAP platform has proved to be highly capable of configuration for this challenging application, generating large numbers of calls with the flexibility to change the data format for each call. The platform also made it easy to add other features such as enrolment.

This flexibility will be essential as the telemetry service evolves and

new applications are developed. For instance, telemetry will be able to communicate with vending machines to check on stock levels and notify the operator when these need restocking.

Acknowledgements

The telemetry service has required the development of highly complex systems in a very short time—driven by regulatory change.

We would like to thank our colleagues in BT Networks and Systems for their help in developing this ISAP application.

The service builds on an early trial of automatic meter reading which used equipment supplied by Schlumberger and was run in collaboration with Yorkshire Electric and Yorkshire Water companies. These customers have helped us a great deal in understanding the issues facing their industries and have been of enormous practical value in developing the service.

We have also enjoyed the support of many people across the whole BT group; in particular, the help of our colleagues in BT Network Operations has been essential in ensuring we can deliver a reliable and manageable service.

References

- 1 SIN236. Telemetry service description. (A copy can be obtained by calling 0800 318 601).
- 2 SIN252. Telemetry interface system. TSMS interface specification. (A copy can be obtained by calling 0800 318 601).
- 3 SIN237. Telemetry CPE design guide. (A copy can be obtained by calling 0800 318 601).
- 4 ROSE, K. R.; and HUGHES, P. M. The BT Speech Applications Platform. *Br. Telecommun. Eng.*, Jan. 1995, 13, p. 304.
- 5 TRICKER, D.; ROSS, G.; and DUFFY, M. ISAP Telemetry Application Architecture Design. BT Document Reference 0106-0305-3-00004 (CMC).
- 6 LINCOLN, S., and CLELAND, D. ISAP Telemetry Application Physical Design. BT Document Reference 0106-0390-3-00006 (CMC).
- 7 TRICKER, D. TIU Transaction Templates. BT Document Reference 0106-0305-3-00014 (CMC).
- 8 IEC 1107 Revision 1, 4th draft April 1994. Data Exchange for Meter Reading, Tariff and Load Control, Direct Local Data Exchange.
- 9 ISO 3309. 07-01 2nd edition. Data Communications—High Level Data Link Control Procedures—Frame Structure. 1979.
- 10 NOLDE, K. E. Technical Specification for a Telemetry Interface Unit (TIU) for Connection to the PSTN. BT Document Reference 0106-0325-3-00002 (CMC).
- 11 NOLDE, K. E. Telemetry Interface Unit (TIU) Message Protocol (RITA Protocol). BT Document Reference 0106-0325-3-00004 (CMC).

Biographies



Kevin Welsby
BT Networks and
Systems

Kevin Welsby joined BT in 1967 as an apprentice. Having received a B.Sc. in Electrical Engineering from Aston University in 1978, he joined BT Laboratories to work on the development of telephone testing standards and techniques. Later he moved to the development of telephone

instruments. He started design and specification work for telemetry in 1992, and now leads the team delivering the network systems (telemetry platform and interface unit).



Jeremy Adams
BT Networks and
Systems

Jeremy Adams received a joint honours degree in Philosophy and Psychology from Bristol University in 1970 and joined BT Laboratories in 1971. He spent nine years as BTL's technical editor and then worked on technical documentation for the Westminster switched-star cable TV system, before becoming involved in local loop architecture evolution studies. He has been involved in telemetry since 1990 and currently works closely with BT's telemetry product line as its trials manager.



Jeff Deslandes
BT Networks and
Systems

Jeff Deslandes joined BT in 1973 as an apprentice. He moved to BT Laboratories in 1982 after gaining a bachelors degree at Liverpool Polytechnic. He went on to gain a masters degree in 1987 and a Ph.D. in 1991, both from the University of Essex. He has worked on analogue line interfaces for digital exchanges and digital transmission techniques. Since 1992, he has worked on solutions designs and is the system design authority for the telemetry system overall. He is a Chartered Engineer and Member of the IEE.

Philip Johnson, Andrew Catchpole and Laurie Booton

Computer Telephony Integration

The Meridian 1 PBX

CTI-enabled services can deliver productivity gains for the smallest business to the largest multinational. The Meridian 1 PBX range is a key offering in BT's portfolio, and provides a sophisticated and flexible CTI platform which can grow to support a company's business needs today and in the future.

Introduction

Previous articles in this series have introduced the concept of computer telephony integration (CTI)¹ and described the CTI capabilities of the Meridian Norstar PBX². In this article, CTI on the Meridian 1 PBX range is presented. For background information, the features of the range are briefly described. The CTI aspects of the PBX are then examined in greater detail, essentially treating the PBX as a platform to deliver CTI-enabled services. Later articles will describe typical CTI applications which can be enabled.

The Meridian 1 PBX Range

The Nortel Meridian 1 is one of the world's leading PBXs, and is marketed in the UK by BT. Over 5600 Meridian 1 PBXs have been sold in the UK in the last five years, many of which are CTI enabled.

The Meridian 1 is a truly international PBX, and is now found in over 90 countries worldwide. As such, it is fundamental to BT's strategy of providing a complete business solution. Some of the more important attributes of the PBX are:

- **Comprehensive size range** The Meridian 1 is available from the 'small' option 11, supporting up to 288 extensions, to the large option 81 with up to 10 000 extensions (see Figure 1). The modularity of the system provides an upgradable solution with hardware and software common across the range.

- **Hybrid PBX/ACD** The system can be configured to provide advanced PBX features, such as integrated voicemail, and a highly sophisticated automatic call distribution (ACD) system with fully integrated customised routing, management statistics and networking. Additionally, multi-tenant and multi-customer installations are possible.
- **Numerous network interfaces** including DASS2, DPNSS, QSig and I.421.
- **Large terminal range** supports analogue telephones and a range of digital feature phones, with 'key and lamp' simulation and I.420 to the desk.
- **Strong CTI capabilities** including the Meridian Communications Adapter (MCA) for desktop-based applications and the Meridian Link interface for PBX-based applications.
- **High availability** reliable hardware, coupled with BT's ServiceCare maintenance, ensures dependable operations.

BT Support for Meridian 1

The national Meridian operations centre (NMOC) in Birmingham provides system support 24 hours a day, 365 days a year, with faults being reported automatically using remote access to customer's equipment (RACE). With RACE, many



Figure 1—Meridian 1 Option 11 (left), and Option 81 (right)

potential problems can be resolved by the NMOC before customers are even aware of them.

In the event of a complex problem, the NMOC can reproduce the fault in its laboratory, and consult experts worldwide to seek a solution.

The NMOC forms part of BT's ServiceCare maintenance, a 'one-stop shop' to provide support for all BT-supplied services.

PBX-Based Integration

A functional diagram of the Meridian 1 PBX is shown in Figure 2. To enable the Meridian 1 with PBX-based CTI capabilities, a module called *Link* is required. Link is a

UNIX-based module which communicates with the PBX's internal bus to support industry-standard TCP/IP and X.25 connections to a host computer.

The host computer must run PBX driver software that can communicate with the PBX using the Link messaging protocol. In addition, software known as an *application programming interface* (API) is also required to convert messages to and from Link into messages suitable for application creation. API and driver software packages available for the Meridian 1 include Genesys T-Server, Dialogic CT-Connect, IBM CallPath and Novell's TSAPI.

Link also provides operations, administration and maintenance

capabilities, including link control, monitoring, recording, status, statistics and message filtering.

Meridian Link Message Set

Link commands are categorised as general call-management instructions, routing instructions and voice processing instructions.


General call-management instructions

There are software packages for inbound and outbound calls.

Inbound calls

The inbound calls package allows an application to monitor the progress of

Figure 2—Functional diagram of the Meridian 1 PBX



an incoming call. The MONITOR DEVICE instruction can be used to see where the ACD routes the call, to check if the telephone rings and how long it takes to answer. The instruction can be used within the call centre environment to monitor the performance of an agent, or to provide complex management statistics of the overall call-centre operation. Statistics applications such as this are so important that a standard internal CTI package, called *Meridian MAX*, can be purchased for the purpose. This is described later in more detail.

Outbound calls

A typical use of CTI for placing outbound calls is within a debt collection department. Here, a list of debtors is compiled by the CTI application, which will then place calls automatically. CTI is then used to display account details to the collecting agent when the call is answered.

Routing instructions

The Meridian 1 supports ROUTE REQUEST and ROUTE SELECT instructions. These allow the external host to decide how a call is routed, rather than the traditional method of allowing the ACD to decide and only informing the host via the CTI interface. With the host as the master, it is then possible to write sophisticated call routing applications using dialled number information service (DNIS) and calling-line identity (CLI) information. For example, a call can be routed depending on which number was dialled (for example, sales, servicing, or emergency numbers) or on which client is calling, to provide a personalised service. Again, this is a common call-centre requirement and has been made available as a separate package called *customer controlled routing (CCR)*, which is discussed later.

Voice processing instructions

These instructions allow automatic interaction with a caller, by playing

Desktop integration

Desktop integration is the linking of a computer application to a single telephone, or line, via a physical interface **at the desktop**. The application's view of the telephony world is limited to that of the telephone itself; this view is known as *first-party CTI*.

PBX-based integration

Also known as *third-party CTI*, this approach uses a common CTI channel between the host computer system and the PBX. Applications can then access CTI functionality via the host—thus there is no need for a physical connection at the desktop between the application and the telephone.

Call centre

A call centre is an office which is designed to manage large volumes of telephone calls, which have a predictable and uniform content³. These offices are in direct contact with customers, and are used widely by banks, building societies and insurance companies for handling customer service. Call delivery to the centres is often from Freefone and LoCall numbers. BT's own customer services operations form the largest call centre in the UK.

Automatic call distribution (ACD)

ACD systems are fundamental to call centre operations, and allow agents to be placed into functional groups. Calls may then be directed to these groups with the ACD managing the distribution of calls to meet business needs. Call statistics are collected and used by the ACD administrator to adjust the system configuration.

It is often the case that call centres are physically distributed. This can be for many reasons such as enhanced disaster protection through multiple sites, or to a 'follow the sun' policy where 24 hour service is offered via multiple call centres worldwide. The Meridian 1 actively supports call centre distribution via the network ACD (NACD) facility. This allows calls to be queued against multiple groups of agents anywhere on the network.

voice announcements, collecting MF digits and recording messages.

Common applications are front-end systems for call centres ('For sales, press 1; for servicing, press 2...'), and home banking applications, where a database can be queried to supply a customer with requested information.

Once again, for simple dialogue, that is, 'press 1 for servicing...', the Meridian Mail package can easily be configured to provide this functionality. For more complex home banking type applications, an interactive voice response (IVR) is available to provide this.

Integrated CTI Applications

Meridian 1 can fulfil many business needs via its pre-packaged CTI

applications. These can offer significant cost savings when compared to bespoke developments or the integration of third-party applications. These applications are listed below.

Customer controlled routing (CCR)

The Meridian CCR package allows complex routing of incoming calls. It provides a high-level scripting language using 'IF...THEN...' functionality and mathematical formulae to decide how a call is routed. For instance, in the credit card business, DNIS could be used to move gold card customers to the front of the service queue. With Meridian CCR, routing parameters may be altered in real-time. An example script is shown in Figure 3.

```

GOTO Platinum_Callers IF DNIS = platinum_dnis
GOTO Gold_Callers IF DNIS = gold_dnis
GOTO Regular_Callers
    QUEUE TO cust_svc WITH PRIORITY 1
    QUEUE TO special_svc WITH PRIORITY 1
    GIVE RAN platinum_ran
    GIVE MUSIC soft_music
    QUIT
SECTION Gold_Callers
    QUEUE TO cust_svc WITH PRIORITY 2
    GIVE RAN gold_ran
    GIVE MUSIC soft_music
...etc...

```

Figure 3—Example CCR script

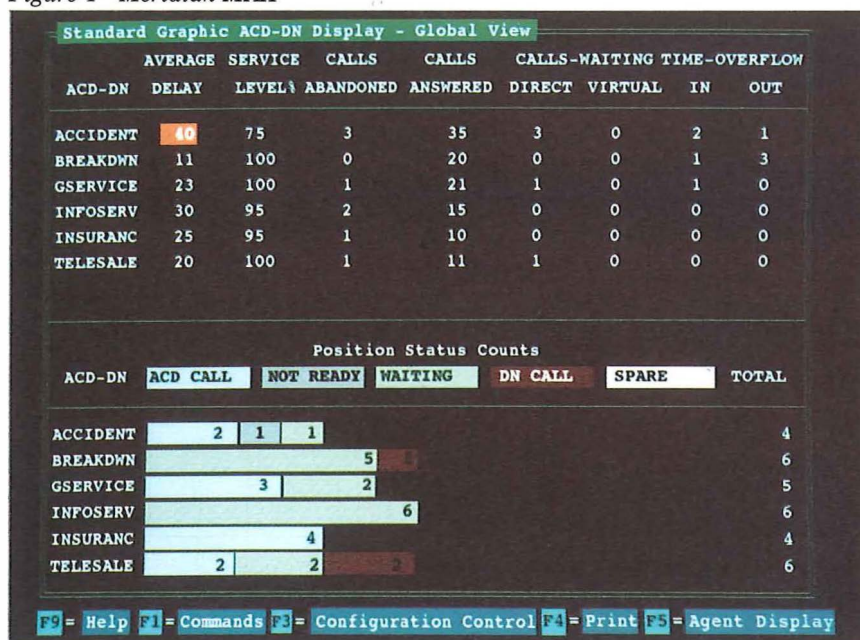
In this case, there are three types of caller—regular card holders, gold card holders and platinum card holders. The type of caller is determined using DNIS information. ACD queues have four priority levels. Platinum card callers are placed in the customer service queue with the highest priority of 1. If the call is not answered immediately, it is also placed in the backup service queue, labelled 'special_svc'. The call is then left in both queues, and the caller is played a recorded announcement (RAN) followed by music until the call is answered. Similarly, gold card

callers are placed in the customer service queue, but with a lower priority of 2. They are not placed in the backup queue, and hear a different RAN and then music.

MAX

Meridian MAX is a statistics and management reporting system which is used to set service levels and monitor performance in call-centre environments. A screen shot of MAX is shown in Figure 4. The performance of individual agents can be monitored to provide information to human resources departments to

Figure 4—Meridian MAX



calculate agent bonuses. Additionally, Meridian QMAX (and MAX-caster) will be available to provide forecasts of call volumes and patterns. These forecasts can then be used for employee scheduling and to determine staffing numbers and trunks required to meet desired service levels.

Mail

Meridian mail is a highly sophisticated integrated voicemail system.

Standard features include:

- *Call answering* used to leave messages when the dialled person does not answer, or is engaged.
- *Voice messaging* for composing messages which may then be sent directly to another user's mailbox. Distribution lists may also be compiled for multiple recipients.
- *Express messaging* for fast messaging without the need to log-on to Meridian Mail or call the person directly.

Calling the sender, dialling by name, through dialling and remote notification (via pager or telephone) are also possible. Special features are available for use in the hotel and hospital environments, which can be linked to a property management system (PMS). As an option, mail can also be networked using a centralised or distributed architecture to cater for large organisations. The other options available include voice menus, voice forms, and fax on demand.

Interactive voice response (IVR)

The Meridian IVR toolkit is an enhancement to mail. It provides an advanced object-orientated design environment to create IVR applications and an interface to back-end data systems. IVR systems are used for automated transactions with customers; for example, bank balance enquiries.

Desktop Integration

The MCA is a data board which can be installed inside the Meridian 1 digital telephone range to provide desktop-based CTI (see Figure 5). The MCA provides:

- synchronous data capability up to 64 kbit/s, and
- asynchronous data up to 19.2 kbit/s.

These capabilities can be used to support videoconferencing, local area network (LAN) bridging, Group 4 fax, and file transfer.

Desktop CTI Applications

The Meridian 1 has a range of desktop applications available that use CTI technology to improve the way in which the telephone, and indeed all media, is processed by the human user. The CTI applications described below all work by physically connecting the desktop computer to the MCA card installed in Meridian 1 digital telephones (see Figure 6). As well as using the MCA interface for high-speed data communications, it can also be used to support desktop CTI applications. These desktop applications, known as the *VISIT* range, include the following but may not all be available from BT:

- VISIT Voice,
- VISIT Video, and
- VISIT FastCall.

The *VISIT* range of products are designed to be fully integrated with the Meridian 1 telephone system to make the most of its powerful capabilities. Each of the *VISIT* products is described in more detail below:

VISIT Voice

VISIT Voice is a personal call management application that allows users to control their telephone by using a simple-to-use Windows



Figure 5 – M2616 digital feature phone with MCA card

interface. Desktop directories can be set up to save time in looking up and manually dialling numbers.

The CLI service allows users to know who is calling before they pick up the telephone. A technique called *screen-popping* allows users to have vital information (such as the caller's name and address and previous contact history) available instantly on their computer screen.

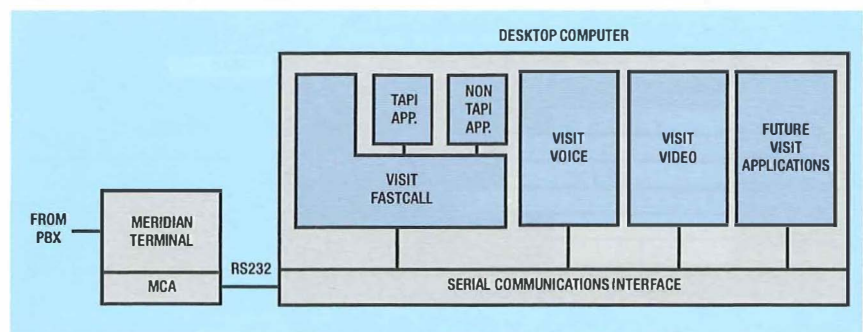
The call log feature captures calling traffic information for all calls to help track daily telephone activity.

VISIT Video

VISIT Video incorporates all the features of *VISIT Voice* and *Messenger* but adds full-colour videoconferencing capabilities. Collaboration with distant colleagues to work on documents and file transfers at high speed is all possible.

The *VISIT Video* system can support both IBM-compatible PCs and Macintosh computers. It is also compatible with both public telephone networks (such as ISDN) and private telephone networks.

Figure 6 – Desktop CTI applications using the Meridian MCA interface



VISIT FastCall

VISIT FastCall has similar features to VISIT Voice in that it provides call management functionality by controlling the telephone using the Windows interface or keyboard function keys. FastCall also has screen-popping so that users know who is calling before they pick up the telephone.

One of the most interesting features is the ability to CTI-enable applications that do not otherwise have CTI capabilities. FastCall provides a link between legacy applications and the CTI technology; thus non-CTI applications can now easily become CTI enabled.

VISIT FastCall offers a unique CTI solution that does not require any application changes. Instead of requiring an expensive custom application, VISIT FastCall is able to integrate with any Microsoft Windows application using a simple rules-based system and macro language. An application, such as FastCall, that provides this CTI-enabling capability is often referred to as *CTI middleware*.

FastCall connects to a Meridian 1 digital telephone through the MCA card and uses Microsoft's TAPI interface. This allows it to interwork with existing Windows applications that use the TAPI interface. TAPI, which was jointly developed with Intel, is now recognised as a de facto standard for CTI applications.

VISIT FastCall offers an affordable CTI solution for call centres and office environments alike. FastCall provides everything a company's telephone agents and knowledge workers need to operate in a call centre operation right from their desktops. Small call centres, even as small as a single agent, can enjoy the benefits previously only available to larger, more expensive call centres.

TAPI Applications

Microsoft's telephony API (TAPI) primarily allows Microsoft Windows

applications to manipulate telephony devices on the desktop; that is, telephones and lines. TAPI, which was developed in partnership with Intel, is generally seen as a low-end desktop CTI enabler. Although Meridian 1 terminals do not directly support TAPI-compatible applications, these can be connected when used in conjunction with VISIT FastCall. This potentially allows any number of off-the-shelf TAPI compatible applications to be connected to the Meridian 1 PBX. The number and diversity of TAPI-enabled applications is rising steadily.

Conclusion

The aim of this article is to show the strategic nature of the Meridian 1 as an internationally available CTI platform.

It offers customers the ability to solve a business problem rapidly via pre-packaged CTI applications, but retains the flexibility to interface to customer's applications via its external CTI interfaces. These external CTI interfaces cater for the more commodity applications, via its desktop interface, while the large bespoke applications are catered for by Meridian Link.

International availability offers multinational companies great economies of scale by allowing the use of identical PBX infrastructure worldwide.

Meridian is continually under enhancement. For example, in a similar way that local ACD operations may network to form one virtual operation, local Meridian Links may be networked to provide one virtual CTI environment. One virtual ACD operation, supported by one virtual CTI environment, internationally available, is a powerful future indeed.

Acknowledgements

The authors would like to thank Peter Wignall, BT product manager, and Geoff Batchelder, Nortel market-

ing manager, for their help in completing this article.

All trademarks acknowledged.


References

- 1 CATCHPOLE, A., CROOK, G., and CHESTERMAN, D. Introduction to Computer Telephony Integration. *Br. Telecommun. Eng.*, July 1995, **14**, p. 98.
- 2 CATCHPOLE, A. Computer Telephony Integration—The Meridian Norstar. *Br. Telecommun. Eng.*, Oct. 1995, **14**, p. 209.
- 3 BONNER, M. Call Centres—Doing Better Business by Telephone. *Br. Telecommun. Eng.*, July 1994, **13**, p. 128.

Glossary

- ACD** Automatic call distribution
API Application programming interface
CCR Customer controlled routing
CLI Calling line identity
CTI Computer telephony integration
DASS2 Digital Access Signalling System No.2
DNIS Dialed number information service
DPNSS Digital private network signalling system
IVR Integrated voice response
LAN Local area network
MCA Meridian communications adapter
MF Multi-frequency
NMOC National Meridian operations centre
PC Personal computer
PMS Property management system
QSig Q-signalling
RAN Recorded announcement
RACE Remote access to customers equipment
TAPI Telephony application programming interface
TCP/IP Transmission control protocol/Internet protocol
TSAPI Telephony services application programming interface

Biographies



Philip Johnson
BT Networks and
Systems

Philip Johnson joined BT Laboratories as an apprentice in 1985. After graduating from the University of Nottingham with a B.Eng. in Electronic Engineering, he joined the Network Intelligence Engineering Centre working on universal mobile telecommunications service (UMTS) and intelligent network (IN) structured multimedia services. He currently manages the Meridian 1 CTI testbed within the Customer Premises Equipment and Peripheral Intelligence group.



Laurie Booton
BT Networks and
Systems

Laurie Booton is a CTI specialist at BT Laboratories. He originally joined Post Office International Telephones as an apprentice in 1974. He soon realised that computers would be the dominant factor in telecommunications and received an award to study at Sussex University. He graduated with a first-class honours degree in Computer Science in 1983 and then moved to BT Laboratories. There he joined the team developing Monarch PBX software, where he specialised in multi-processor operating systems and diagnostics. He also worked on Mitel's SX2000 in Kanata, Canada. After that he set about installing the first Meridian in BT Laboratories and one of the first CTI-enabled Meridians in the country. This took the team in the direction of more commercially-based CTI research. He now leads a project responsible for downstreaming the knowledge gained in the course of this research to the sales and marketing areas of the business. He has registered a patent on CTI controlled teleworking of call centre agents. He is a member of the Institute of Electrical Engineers and a Chartered Engineer.



Andrew Catchpole
BT Networks and
Systems

Andrew Catchpole joined BT as an apprentice in 1981 in the Norwich telephone area. In 1989, he transferred to BT Laboratories and is now in the Network Intelligence Engineering Centre at BT Laboratories where he works in the Customer Premises Equipment and Peripheral Intelligence group. His main work is investigating CTI technology for the desktop and the Meridian Norstar PBX. He is one of the first students to be awarded the Martlesham M.Sc. degree from the University of London; his M.Sc. project was a video PBX telephone system using computer telephony integration.

Allan Drew and Valerie Jones

Driving on Customer Satisfaction

A Re-engineering Case Study in the Retail Line of Business

This article is a case study of re-engineering carried out in a customer-facing environment. Problems encountered with the traditional service delivery process are described, together with the solutions used. The key to improved performance was the CaseTeam – an alliance of engineering and order-entry skills. The project achieved significant improvements in cycle time and delivery date failures. Customer satisfaction increased markedly, and more unexpectedly, the job satisfaction and morale of the people operating the new process was significantly enhanced.

Introduction

Excellent service is increasingly becoming a key competitive advantage as BT establishes itself as a world-class supplier of telecommunications. Within BT's National Business Communications (NBC), there is an evolving line-of-business structure where sales and service are aligned. This structure is moving to a better understanding of customers' needs and has been driven by some key factors:

- continuing drive for efficiency savings,
- increasingly complex customer requirements,
- dispersed communications, and
- the need to combat the increasing competitive threat.

After Project Sovereign, BT created the Business and Personal

organisational structures for customer-facing activities. Figure 1 shows the different customer interface arrangements. Personal Communications has responsibility for BT's residential customers. Business Communications was further divided (segmented) into the following:

- volume business customers (now known as *Business Connections*), covering small businesses with typically 1–5 lines;
- major customers, who have service managed accounts and a national presence within the UK; and
- global customers, who are service managed but with a strong non-UK presence.

ServiceCentres give a single point of contact which is unique to each of the major business and global customers for all of their service requirements. The ServiceCentres forge strong links with both sales

Figure 1 – ServiceCentres and customer segments

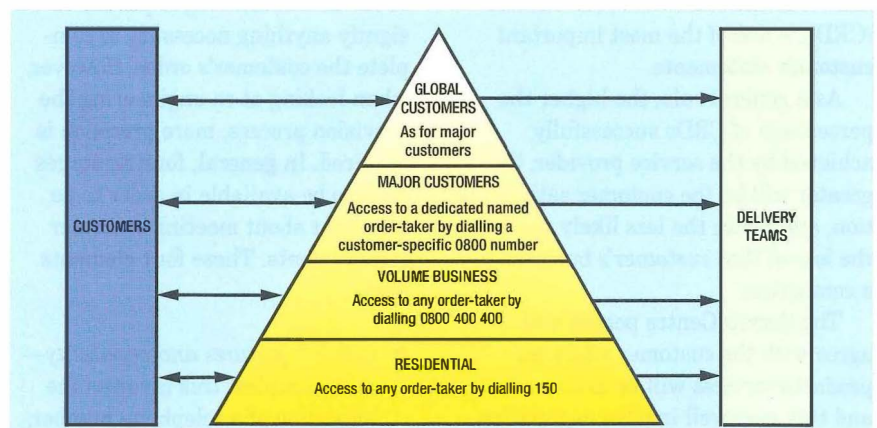



Figure 2—Existing process flow



personnel and customers to enable proactive management of customers' requirements from pre-sale to service delivery.

This case study describes a new way of managing delivery of products and services to major customers within the retail line of business.

Background to the Provision Process

To obtain a full understanding of this project to re-engineer service delivery within ServiceCentres some background on the present provision process within BT is necessary.

The current methods of providing products and services to customers run along traditional lines. Figure 2 shows the activities within the process for a typical product (providing an exchange connection). Firstly, the customer places an order with the ServiceCentre. This order may be the outcome of a protracted selling effort involving the account manager and the account team. By then it is likely that the customer has expectations around both the service functionality and the delivery lead time. As far as timing is concerned, BT works traditionally to a set of standard lead times; for example, six working days to provide a business exchange line, 10 working days to provide a small switch, and 20 days to provide a wideband (MegaStream) pipe. However, the customer will always be asked a preferred date for having the product or service delivered. In a customer-first culture, this date, known as the *customer required date* (CRD), is one of the most important customer statements.

As a general rule, the higher the percentage of CRDs successfully achieved by the service provider, the greater will be the customer satisfaction, and hence the less likely will be the loss of that customer's business to a competitor.

The ServiceCentre person will agree with the customer when the products/services will be delivered and this may well involve negotiation

around BT's capability compared with the CRD. The standard lead time approach is useful as typically it will not be known precisely at this stage whether the resources required for the customer's order are all available. However, with standard lead times the delivery commitments have been cast in such a way that resources not available when the order was taken (for example, line plant) can be provided within the standard lead time. When the date is finally agreed with the customer, it becomes known as the *customer confirmed date* (CCD). The CCD is important as all performance measures, such as cycle time, are based around the CCD. In some cases, the CCD is the same as the CRD, in other cases it may be later.

The drawback is that if the CRD lies within the standard lead time it is difficult to check if all the resources are available and hence whether or not a firm commitment can be given to the CRD. From a customer satisfaction viewpoint, the worst outcome is to agree an early delivery with the customer and then fail because of shortage of some crucial resource. In these circumstances, it is much preferable to agree dates with customers which may be later but where there is a high degree of confidence in delivery. The golden rule is 'don't make promises you don't know you can keep'.

The term *resources* has been used so far as a rather vague phrase to signify anything necessary to complete the customer's order. However, when looking at re-engineering the provision process, more precision is required. In general, four resources have to be available in order to be confident about meeting customer requirements. These four elements are:

- *network features and capability*—at its simplest this involves the allocation of a telephone number,

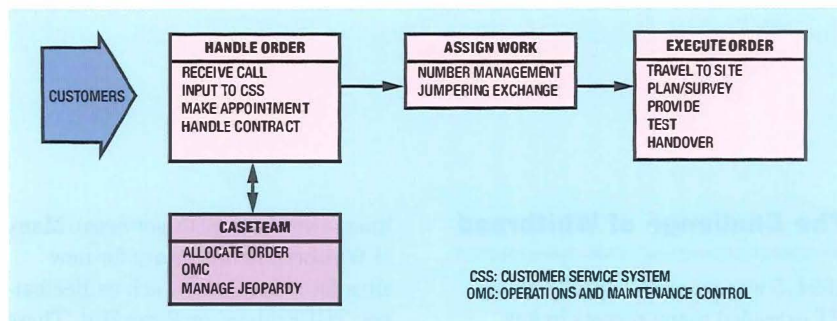
but increasingly involves more complex allocation of network elements such as Featurenet capacity or wideband (for example, switched multi-megabit data service (SMDS)) capability;

- *access capacity*—predominantly delivered on copper but increasingly via fibre or electronic methods;
- *labour capacity*—the customer's order may require a visit by a BT engineer to the customer's premises (this visit may be scheduled for the day agreed for service delivery or some days beforehand if the work is significant and complex); and
- *customer premises equipment (CPE) availability*—the customer's order may involve CPE which needs to be ordered directly from the manufacturer (for example, Norstar and Meridian customer switches).

By knowing the lead time of each of these components and building these lead times together in the appropriate way (not necessarily additive since many of these allocations can proceed in parallel), a date by which BT is almost 100% sure of delivery can be quoted to the customer.

The opportunity which can be lost with the standard lead time approach is that in many cases the resources are available more or less immediately to meet the customer's need. If the process was designed such that availability information about exchange equipment, line plant, labour and CPE was accessible to the ServiceCentre person when the order was placed, it might well be possible to meet the CRD, with a high degree of confidence, in cycle times much shorter than the standard lead times. Hence customer satisfaction would increase and competitive threat

Figure 3—Interaction with CaseTeam



reduce. This argues for a comprehensive expert system capability which can be called up to support the ServiceCentre. Unfortunately, such an expert system does not exist within BT, although CSS (customer service system) and COSMOSS (customer-oriented system for the management of special services) system developments are building in that direction. However, to achieve the benefits of such a system in practice, the re-engineering project utilised what became known as the *CaseTeam* concept.

CaseTeams

As BT steps up to the challenge of delivering world-class service excellence, CaseTeam working is the catalyst to achieving customer satisfaction and efficiency savings through end-to-end managed service delivery. CaseTeam working collocates people with engineering and service skills and enables BT to handle all aspects of service delivery from receipt of order, to installation and hand-over to the customer. Figure 3 shows how the process has been modified with the introduction of CaseTeam working.

Historically, orders from customers have come to ServiceCentres, which have then managed these orders across many organisational boundaries. This meant that order handling was serial and dependent on each part of the organisations fulfilling their commitment by a certain time. However, none of these organisations was responsible for the end-to-end delivery of the customer requirements. Instead, each had its own set of functional targets to meet. (An analogy would be of one department doing its part of the order and then throwing the job over the wall to the next department.) This serial activity can lead to problems not being identified until late in the process with an increased risk of failing to meet the date agreed with the customer. These problems are magnified if there is no clear control

of the end-to-end process. CaseTeam working brings all the office-based functions together and enables order requirements to be handled in parallel, thus reducing time and removing organisational boundaries. Further, by checking resource availability at the time of order-taking, the CaseTeam minimises the possibility of making promises to customers which cannot be delivered.

Job Management

In the traditional provision processes, the major mechanism of process control is through the job management function. As described above, one of the problems with serial processes is that no one functional operator within the process has end-to-end responsibility for the job. Each person carries out certain activities, and the hope is that when all the activities have been completed, the job is completed and the customer requirements have been delivered. Most people who work in a customer service or an assembly line production environment will realise that this happy state is never achieved 100% of the time. A certain percentage of jobs are roadblocked and action has to be taken. Either the roadblock is cleared and the job comes back on track, or the customer is advised that the agreed date will not be met and appropriate action is taken to fix a new date which must then be delivered. In the provision process, job management fulfils this function. By using jeopardy and failure flags on every activity and on every job, the jobs where roadblocks are occurring can be identified and action taken. Job management traditionally supports the ServiceCentre to deliver the commitments made to the customer. However, it is ultimately a cost of failure, and if better processes and systems could eliminate or prevent

roadblocks, then there would be significant cost savings.

Geographic Spread

Major customer ServiceCentres present some specific problems in the delivery of products and services to customers. Although all the orders are placed centrally at the ServiceCentre, the work required is likely to be spread throughout the country. Thus, for example, checking resource availability becomes difficult since CSS is organised around 27 databases. The relevant databases must be accessed, and the resources booked as appropriate to meet the customer agreed dates. Furthermore, any order queries or problems occurring must find their way back to the owning ServiceCentre so that problems can be resolved with the customers' order-placing point. This complexity leads to difficulties when monitoring progress, since job management ends up trawling numerous activity queues around the 27 databases in order to find the jobs they are involved with.

Traditionally, organisation of engineering office functions is based on the geographic location of the service. With CaseTeam working, all job management aspects of the orders are handled in the ServiceCentre irrespective of the geographical location. This requirement starts to define the systems modifications necessary to deliver what in effect is a seamless database across the UK (and potentially globally).

National Business Communications has already launched a £60M investment in ServiceCentres focused on state-of-the-art centres with world-class standards of accommodation and systems technology. This concept is already in use by Bell Atlantic who is setting the standards for world-class benchmarking.

The Challenge of Whitbread

1994–5 was an excellent year where BT exceeded many targets in key areas. However, many customers' perception was that BT was not meeting all of their requirements, especially around CRDs. A major customer, Whitbread plc, was concerned that BT remained focused on internal measures and, using its own measures, believed that, at best, BT delivered only 75% of their orders by CRD. This led to Whitbread scoring BT extremely low on satisfaction with service delivery—satisfaction measured on surveys prior to the start of the project was only 4 out of 10.

Allan Drew of Process Control and Measurement, and Mike Langston of customer Service Management and Strategy, agreed the project proposals for CaseTeam working to resolve these problems. Valerie Jones of Process Control and Measurement launched the trial of CaseTeam working for ServiceCentres in the Retail Sector in February 1995.

The first customer to join the trial was Whitbread plc. Three engineers joined the Whitbread ServiceCentre team bringing with them knowledge of CaseTeam working from earlier pilots in BT's Breakout project.

By collocating engineering and service people, all aspects of service delivery can be handled from receipt of order, to installation and handover to the customer. Job progression of all office-based processes empowered the team to drive the customer's order from end-to-end, providing a central coordination of activity. This resulted in reduced double handling, enhanced productivity and created room for proactive customer management.

Implementing CaseTeams

Initially, the CaseTeam focused solely on Whitbread plc. Whitbread places an average of 600 orders with BT each month and a large percentage of these are for services within the public switched network (exchange

lines, switches, ISDN services). Many of Whitbread's orders are for new sites for subsidiaries such as Beefeaters, TGI Fridays, or Pizza Hut. These are often complex orders on green-field sites and demand a high level of coordination to achieve success.

By July 1995, the operation expanded to cover all 20 retail customers at Milton Keynes, handling provision of all products and services.

Inflexible systems and systems access were, and still are, major issues. These have initially been resolved through manual work-arounds which have enabled the definition of system developments delivering more elegant solutions especially in the area of UK-wide queue management.

CaseTeam working has moved the job management function to the point where BT receives the order, and ensuring technical expertise is on hand. The team can identify potential difficulties and resolve them before jobs fall into jeopardy or failure. Bringing in technical skills at the customer interface enables clear interpretation of information resulting in better quality of orders. With better quality orders, follow-up enquiries from customers are reduced, allowing the ServiceCentre people more time to focus on the customer's new requirements rather than chasing failures. By building the customer's confidence in BT, existing business will be protected and the way opened up for additional revenue growth.

Benefits of CaseTeams

CaseTeam working delivers benefits in quality of service. Also, BT people much prefer the new way of working, and job satisfaction and morale rises. People have more accountability, more discretion and the ability to solve problems as a team. Sound working relationships have been forged between engineering and service people. Service excellence is a major determinant of improving

customer satisfaction and CaseTeam results are consistently best in BT.

Measured results delivered with CaseTeam working are impressive:

- Satisfaction of BT people (as measured against a selection of before and after results from BT's CARE survey against the line index) has risen from 45.8% to 95%.
- BT's ability to meet the customer's required date **as measured by the customer** has risen from 75% to a consistent 98%. This has led to customer satisfaction of service delivery (as measured in a major customer survey) rising from 4/10 to 8/10. The communications manager at Whitbread plc, Norman Smith, acknowledged that problems from his customers attributed to poor BT service have all but vanished as a result of this initiative.

Additionally there have been cost savings as a result of the reduced need for job management. Surveys before CaseTeam implementation showed that ServiceCentre personnel spent 20% of their time dealing with problems and negotiating new commitments with other parts of BT. Putting in a CaseTeam has eliminated much of the cause of this failure and the engineering expertise handles any remaining problems in a more efficient way. As a result, the 20% of ServiceCentre effort can be replaced by a CaseTeam which is only half the size. Overall, a saving of 10% of ServiceCentre staffing is possible.

BT Account Manager Lester Potter believes that integrating the CaseTeam function into the line-of-business team is an overwhelming success. For Lester this means a greatly improved provision of service that is responsive and efficient, and a seamless operation with all parts of the business at our fingertips.

Martyn Cooper, customer team manager, says that the CaseTeam concept has been a triumph of

innovation and a total success in customer satisfaction for Whitbread plc. This achievement is even more remarkable considering the expansion and decentralisation of the Whitbread account.

The trial of CaseTeam working ended in August 1995. CaseTeam working continues to be business as usual for the retail customers at Milton Keynes.

Next Steps

With such a strong case built on improved customer and people satisfaction and underpinned by cost savings, National Business Communications was very keen to get implementation moving across the lines of business. However, organisational change delayed progress for some months, but by June 1996 the new sales/service organisation under Tom Johnson and John Wheeler had settled down to the point where further progress became possible. Agreement was reached between National Business Communications and Process Control and Measurement, Networks and Systems, (the post-Breakout inheritors of the project) that full-scale implementation in the further two ServiceCentres supporting the retail line of business would commence immediately. Full CaseTeam working would be in operation by August 1996. After four months, the benefits would be rechecked and if the scalability was proven then implementation across all 10 lines of business would be strongly recommended.

Biographies



Allan Drew
BT Networks and
Systems

Allan Drew has spent most of his 26 year career with BT in the customer service area. He has held senior management positions in West End District and then in Business Communications within London. With the formation of Global Customer Service in 1992, he ran job management nationally for Business Communications. In 1994, he joined the Breakout project to lead the provision re-engineering stream which developed the CaseTeam concept. Along with field engineering multi-skilling, these concepts proved to be powerful re-engineering tools and already have several key applications within the service delivery arena in BT. After Breakout, Allan took the work into Process Control and Measurement, BT Networks and Systems, where radical re-engineering proposals for service delivery are being devised.



Valerie Jones
BT Networks and
Systems

Valerie Jones joined BT from school in March 1977 as a Clerical Assistant. Much of her career path has been spent in sales/service including Special Services Sales and ServiceCentres. She worked as a ServiceCentre manager on both global and national customer accounts from April 1990 to January 1995. She joined Process Control and Measurement in February 1995, bringing service, sales and customer-facing skills. She trialled CaseTeam working in the Retail and Leisure Sector from February to December 1995, and is now implementing CaseTeam working in the Retail Sector.

Shared Spaces

'Cooperation in groups is the catalyst that'll let all the new thinking about virtual corporations actually happen. It works with Doom.... it doesn't work at all with Windows '95. The most important part of Doom-like user interfaces – the element that will undoubtedly shave its legs, put on a suit and get an MBA – is the shared space.'

'... in the future everyone will have fifteen Megabytes of fame'

When Andy Warhol measured fame against a unit of time, how could he have known about the future possibilities of a medium called *shared spaces*?

Telephony provides universal access to people on a one-to-one basis. Networked computers routinely link a person to a remote application or to distant data. By combining and extending this functionality, shared spaces are quietly reinventing the telecommunications business.

Shared spaces are a new communication medium. They are three-dimensional spaces where people are represented by characters or *avatars*. The avatars can move around, converse (either voice or text) and interact in a common context of information and applications. Shared spaces are causing us to think again about the role of computers and how, with telephony, they can mediate conversation and social interaction.

Advances in standards, network capacity and connectivity, and in the

processing power of terminal equipment are opening exciting new avenues for shared spaces, and in this article we reflect on some of the early opportunities and challenges that this will afford.

'A phone call is a shared virtual environment with self-authored content and no graphical support'

When the telephone was first invented, it was thought that it would be ideal for listening to opera from afar. In fact, simple conversation became the main use. People have made their own 'content', and the growth of telephony has not been limited by the growth of opera.

High-performance personal computers, the growth of global networks and graphical standards such as VRML² (Virtual Reality Modelling Language) are bringing distributed virtual reality (VR) as a **capability**, to the mass market. Rather than content-rich applications such as shopping or real estate (the VR equivalents of opera) the authors believe that three-dimensional spaces will be most successful for building virtual on-line communities.

In today's shared spaces (Figures 1–5), avatars can express simple emotions, gesture and perhaps move their legs in a walking motion. Some spaces provide autonomic responses such as blinking, so that avatars become more believable characters. Elements of groupware and conferencing can be combined in a shared space, within which avatars can support turn-taking behaviour and other social constructs of group interaction. Users can also express themselves by customising their avatars – 'fashion' in cyberspace.

Adoption and trials of shared spaces will be advanced by commercial browsers from major companies, with

Figure 1 – Avatars in a teleconferencing shared space³




Figure 2—Avatars in an audio shared space⁴

support for three-dimensional worlds and voice conferences embedded in documents.

'An age of networked communities'


In the information age, we are witnessing the end of geography and the growth of networked communities^{8,9}. There are forces of fragmentation and cohesion within any society. We rarely live beside those with whom we work; given a choice of over 200 channels, we may not often watch the same broadcast TV. Yet it is good to have something in common with our 'neighbours'.

Computer systems are effective at supporting newsgroups and bulletin boards. Sometimes called *chat spaces*, these allow interest groups to form and die out as needed. A subject may be so specialist that it will attract only 10 people from around the world, and yet this will be enough to sustain a valued discussion. Finding those 10 people in the same physical neighbourhood is highly improbable. Networked computers help us to 'meet' people with a shared interest and then, having established a common context, we can engage in conversation.

Newsgroups support a vast diversity of topics—so vast that people only have time to read or contribute to a tiny fraction. The topic provides a common context for a conversation, and participants assume that their fellows are there by choice and because the topic is of interest to them. Sustaining this vast and lively topic base is only possible through networked computer systems.

Shared spaces as multimedia environments can be equally cohesive. In shared spaces we meet and converse with people we want to from anywhere in the world, based on desire not circumstances. The common context is provided by features in the three-dimensional world, and by the way in which it is marketed.

Chat spaces are normally text based; that is, messages are typed to



other participants. This has advantages of being able to reflect on what we say before it is posted, or being able to read different viewpoints rather than having to listen to a dominant voice in a group. At the same time text has limitations. Voice is more spontaneous, with live emotional richness, and will be an important element in broadening the market for shared spaces.

'People are content'

Newsgroups are fluid, self referring, diverse, and are used by over four million people. A further subtle aspect of their popularity is that by posting contributions, a topic is created by the participants. As in a telephone call, customers provide the

content. Contributors can say 'I wrote that'. Shared spaces content is also framed in a persistent context, and may itself contribute to the evolution of the space. Others can browse, admire or criticise. Contributors can say 'I built that'.

Today's multimedia applications are insatiable consumers of new content, but what is missing from the interactive multimedia revolution is other people. Imagine a town centre without people. It could be filled with features to add excitement and 'value for money', but people would not go there unless it offered new content every time. The behaviour of groups of people is important content in itself.

We need to build spaces with enough interest for people to feel comfortable and which allow them to

Figure 3—Full-body avatars in a text chat community⁵



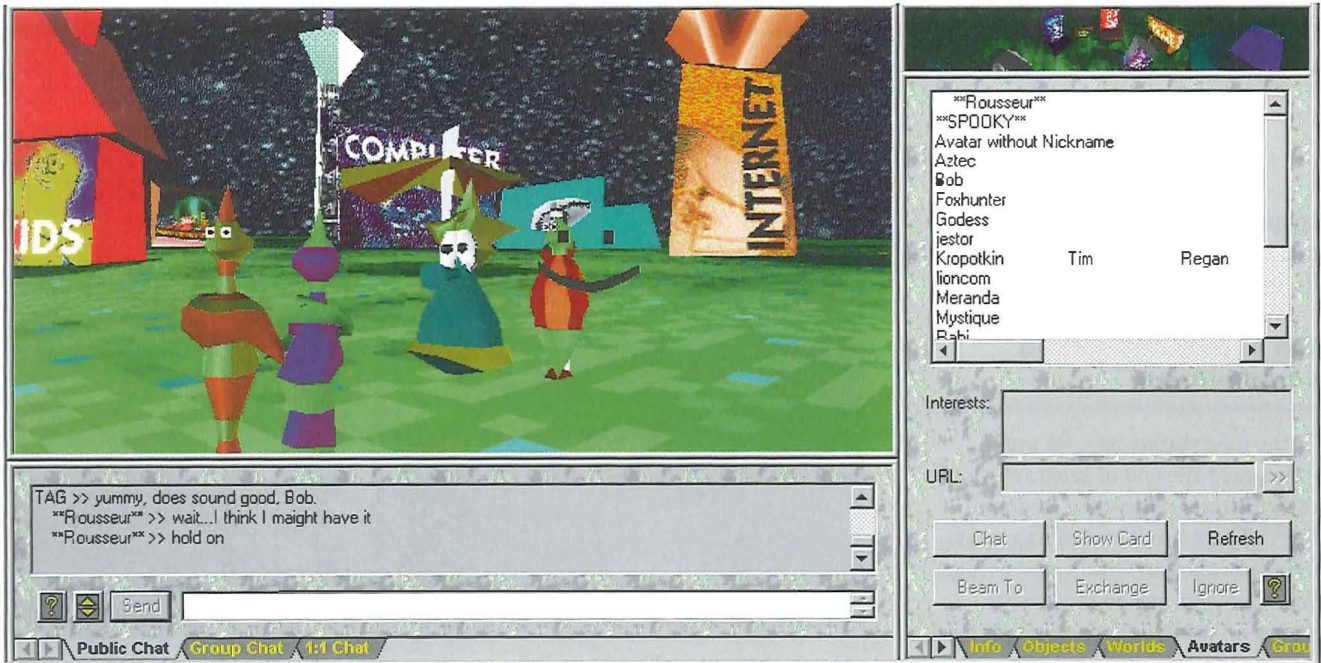


Figure 4 – Text chat shared space with additional environments control⁶

express their own personality and appearance. We also want services which are self-supporting rather than creating a hungry ‘content monster’. Imagine if we had to **design** every telephone conversation to ensure that it was interesting!

Drawing from the newsgroup experience, one solution is to provide a common context directly in the design of the three-dimensional world. Titling one world ‘Premier league football’, and another ‘Bridge club for beginners who wish to learn German’, is a first step. Designing features linked to the theme of the space is another. For example, we might provide a locker room, turnstiles and a gallery of famous players in the football world. Beyond this conventional thinking there are deeper ways to engage participants, particularly if we allow them to build for themselves. Some physical places like certain bars or clubs are popular mainly due to the mutual compatibility of the clientele. Others like our homes are full of personal constructs and act to tell our story.

‘Just seeing you there reminds me how important serendipity is to communication’

If we remove people from a shared physical environment and connect them with telephony and messaging, we have not given them everything they have lost in the transition. For shared spaces to be a satisfying alternative to real-world environments they must provide an exciting range of task-oriented and exploratory features, accessible through a simple interface. Although in real life we can simultaneously walk, talk and absorb information through our senses, in a VR space this is not so easy. Mastery of navigation and authoring must be sufficiently autonomous for the user to concentrate primarily on the task or activity in hand.

Moreover, large organisations have a host of common stories around duplication of effort, missed opportunities and how a chance corridor meeting made that all-important connection. Unplanned meetings and

random events are very important to the functioning of social and business groups. Yet we persist in the belief that the traditional office, fixed in space and time, lends itself to communication between its inhabitants.

Personal profiles and user agents may actively promote appropriate connections and groupings within an electronic community, but an effective shared space will also support chance encounters and will recognise the importance of serendipity in communication.

We are designing shared spaces where the useful social conventions of meetings can be employed to improve communication¹⁰. For business, we are building spaces which support distributed teams who need to work together across time and distance. In social communication, we are designing spaces which allow people to talk together after ‘chance’ meetings—spaces where a family can tell stories, co-present at a ‘kitchen table’.

‘Shared spaces turn surfers into settlers’

To attract people back to a location on the World Wide Web (WWW) again and again is both commercially and socially exciting. One method which has previously been shown to work in a two-dimensional graphical space is to let users **build** their own spaces and make their own rules for how to govern them¹¹.

Figure 5 – Avatar emotion and navigation control panel⁷



shared spaces

90 minutes

PRESENT ...

story - Graham Walker
drawn - Matt Polaine

One evening ...

Home at last! What a day!
Can't wait to see my mates!

As a fanatical fan of Martlesham United Matt always went 'on-line' via the 'Football Forum'!

Using the common gateway to on-line services ...

CONNECTED!

ID accepted ...

Matt selects his new avatar strip - which is automatically debited from his account.

Excellent! I'll get that!

NEW!

Hi Laurence, Graham!

Yo Matt! I haven't seen you for ages mate!

Have you seen the latest strip in the MU shop?

Oh no! I missed that goal ... I'll replay it at the club!

GOAL! ...Shiplee scores again!!

BOXES.

Graham and I are going to our private box - see you there!

In the private 'box' ...

GALLERY

Don't my photo's of the last match look great!

I want to win more on-line credits with our Karaoke night.

Yeah! The bigger the crowd we pull, the more credits!

We could also do 'Name That Goal' from the Golden Oldies vid screen in our club?

But Laurence had 'jumped'. He had gone ...

LYRICS

... 'match interactive' against The Sydney Settlers in 'VReal Football'!

Laurence goes in hard...

Pweeep!

WHAM!

BRADLEY "OFF" connection terminates in 10 secs

HOLOCARD

His new 'tackle mode' does not go unnoticed by the 'VRef'!

Laurence, outside the 'Football Forum' ...

Hmm ... Cricket next ...

To be continued ...

Developing the skills to create and sustain communities in shared spaces will be one of the most interesting marketing challenges of the next five years.

If we combine shared spaces with an emphasis on participant authoring, we will get network settlers—people putting a permanent, persistent presence of themselves on the network. They will start to build, to put down roots and form relationships in the network. The new Wild West is on the Internet and the indigenous surfers are being joined by settlers.

Developing the skills to create and sustain communities in shared spaces will be one of the most interesting marketing challenges of the next five years. Traditionally, mass communication is a one-to-many process whereby a firm transmits content through a medium to a large group of consumers. Shared spaces are a new medium, providing a many-to-many environment in which consumers can interact. In this mediated model, the primary relationships are not between sender (company) and receiver (consumer), but are rather with the shared space within which they take place. We already think of the WWW, a set of hyperlinked and largely two-dimensional pages, as a **place**:

'In this new medium, information or content is not merely transmitted from a sender to a receiver, but instead, mediated environments are created by participants and then experienced.'

Hoffman and Novak 1995¹²

Three-dimensional environments will serve to strengthen this metaphor. Creating an avatar for a fancy dress ball, planning a party with a friend, building a garden-patio to provide a context for the guests and mailing out the invitations are all possible in today's shared spaces. In the future, companies might create an educational theme park as a fun-site and employ a cast (hosts) to entertain visitors. For many of the visitors (for example a school party), interaction among the group and creating their own megabytes of fame will be a major element in enjoying the experience.

'Shared spaces will turn the network into a place, with all the opportunities and problems that will bring'

Persistent context, participant authoring, and evolving electronic communities are powerful aspects of a shared space. Over time, customers will develop a strong affinity with a **place**, in which they may have invested both material and emotional effort. Successful service management of this place is a major element of the overall challenge in marketing shared spaces.

The more of the management of shared spaces services we can move into the metaphor of the space itself, the more engaging and real the interaction becomes. Aspects of community governance such as membership privileges and censorship could be partly handled by the participants, and any external intervention should be aligned with the values and custom of the space.

Shared spaces also present interesting opportunities in charging models, an area of service management where Internet services are particularly ill-defined. The existence of stable communities provides greater scope to segment the advertising market. Moreover, in a shared space we can know where someone has looked. We can put adverts in the distance and bill the advertiser based on how close the customer gets to the advert. We might charge less for oblique views onto information.

'You can be anyone you want to be in a shared space, but sometimes we need to know exactly who we are talking to'

Trust is another important service management attribute of communication channels. Meeting a stranger in a shared space and developing a relationship is like meeting in an aeroplane or train. There is normally an exchange of pleasantries (conversation rather than communication)

and if the early hurdles are cleared a person may move onto deeper topics. Today it is common for people to have 'met by e-mail'. Enduring relationships can develop remotely, as we have seen from letters, amateur radio and e-mail. Although we may have heard of couples who agreed to marry on the basis of such correspondence, most strong relationships are between people who meet in real life.

In an age of networked communities, society will learn how to form, judge and maintain relationships with a higher level of remote interaction. This will happen more quickly and smoothly if there are services available which help to identify shared-space participants in a trusted manner. In a shared space, there is considerable opportunity for masquerade. We have to develop systems that can authenticate not only those to whom we are entrusting payment, but anyone with whom we establish more than a casual dialogue.

In a further twist on the notion of identity, there are applications where people may be accountable and yet anonymous. In role-playing games, the participants often hide aspects of their true identity, such as name or gender. Trust in this shared space means knowing that the avatar 'Kropotkin' is the same 'Kropotkin' to whom we talked yesterday!

Conclusion

Advances in technology and standards are promoting the creation and distribution of three-dimensional environments as a routine element of on-line information services. The BT Laboratories Portal is an early example of three-dimensional VRML interface which provides intuitive and exploratory access to complex information spaces². In this article, we have seen how such interfaces can evolve into shared spaces, bringing the creativity and spontaneity of human behaviour into an otherwise isolated information environment.

Realising the full potential of shared spaces to support evolving and

self-sustaining on-line communities will require a balance of technical and design skills. Turning the network into a place will raise new issues in areas of service management such as governance and trust. The challenge is to deliver scalable services, built around a simple persistent framework and aligned with the needs and values of a loyal customer base. As we move into an age of network **settlers**, there will be great rewards for those that can deliver this vision.

Synchronous, serendipitous, and social, shared spaces are a powerful new communication medium.

Acknowledgements

Our ideas and understanding of shared spaces have developed through discussions with other members of the shared spaces project team at BT Laboratories. '90 minutes' was illustrated by Matthew Polaine.

References

- 1 RUPERT, GOODWINS, *Wired*, May 1996, p. 38.
- 2 WALKER, GRAHAM, *et al.* Interactive Visualisation and Virtual Environments on the Internet. *Br. Telecommun. Eng.*, April 1996, **15**, p. 91.
- 3 GREENHALGH, CHRIS. MASSIVE Teleconferencing System. <http://www.crg.cs.nott.ac.uk/~cmg/massive.html>
- 4 OnLive! <http://www.onlive.com/>
- 5 Worlds Inc. AlphaWorld. <http://www.worlds.net/alphaworld/>
- 6 Black Sun Interactive. <http://www.blacksun.de/>
- 7 Sony Virtual Society. <http://sonypic.com/vs/>
- 8 O'BRIEN, RICHARD. Global Financial Integration: The End of Geography. Published by the Council on

Foreign Relations, ISBN:0876091230, Jan 1992.

- 9 HOCKLEY, ANDREW. The Information Needs of Network Communities. *Br. Telecommun. Eng.*, April 1996, **15**, p. 85.
- 10 BT Labs Shared Spaces. <http://virtualbusiness.labs.bt.com/SharedSpaces/>
- 11 MORNINGSTAR. CHIP; and FARMER, F. RANDALL. The Lessons of Lucasfilm's Habitat. Published in *Cyberspace: First Steps*. Michael Benedikt (ed.), MIT Press, Cambridge, Mass, 1990 and <http://www.communities.com/habitat.html>
- 12 HOFFMAN, DONNA. L.; and NOVAK, THOMAS. P. A New Marketing Paradigm for Electronic Commerce. Project 2000 Working Paper, Feb. 1996 and <http://www2000.ogsm.vanderbilt.edu/novak/new.marketing.paradigm.html>

Biographies



Laurence Bradley
BT Networks and Systems

Laurence Bradley manages the Virtual Business Campaign at BT Laboratories. The campaign considers the future of businesses and the communication services they will need. It aims to free business from the constraints of time, geography and a fixed composition of assets. He is also the marketing manager of the BT Laboratories shop, an experimental trading platform on the Internet. Before BT, he worked for the Hughes Aircraft Company. There he was European business development manager for automatic identification products in the airline industry. Prior to this, he established a new production line for advanced, hybrid-electronics packaging. He has a B.Sc.

in Electronics and a Ph.D. in Optical Telecommunications from Glasgow University and an M.B.A. from Heriot Watt University.
E-mail: lbradley@bt-sys.bt.co.uk



Graham Walker
BT Networks and Systems

Graham Walker manages the shared spaces project at BT Laboratories. The project team is exploring early opportunities for commercial shared spaces, and is working collaboratively to advance the underpinning technologies. He joined BT as a sponsored student and, after graduating from Oxford University in 1986, spent six years researching into coherent optical transmission systems. This work resulted in numerous publications, and the award of a Ph.D. from Cambridge University in 1992. More recently, he has been leading a group within Advanced Applications and Technologies, working on information visualisation and shared information environments. He has an M.B.A. from Cranfield School of Management and is a Member of the IEE.
E-mail: gwalker@bt-sys.bt.co.uk



Andrew McGrath
BT Networks and Systems

Andrew McGrath manages the applications work on the shared spaces project. Since leaving college in 1991, he has worked in the Human Factors Unit at BT Laboratories. For the last three years he has worked on three-dimensional interfaces, including management of the Portal development, a world-leading VRML site. He has a B.A. in Three Dimensional Design from Glasgow School of Art and an M.A. in Product Design from Manchester Polytechnic.
E-mail: andy.mcgrath@bt-sys.bt.co.uk

Alex Pentland, Rosalind Picard and Pattie Maes

Smart Rooms, Desks and Clothes

Toward Seamlessly Networked Living

The MIT Media Laboratory, Cambridge, Boston, USA, is developing smart networked environments that can help people in their homes, offices, and when mobile. The research is aimed at giving rooms, desks, and clothes the perceptual, affective, and learning intelligence needed to become active helpers.

The Opportunity

Inanimate things are coming to life. However, these stirrings are not Shelley's Frankenstein or the humanoid robots dreamed of in artificial intelligence laboratories. This new awakening is more like Walt Disney: the simple objects that surround us are gradually gaining network connections, computational powers, sensors, and actuators. Desks and doors, TVs and telephones, cars and trains, eyeglasses and shoes, and even the shirts on our backs...all are changing from static, inanimate objects into adaptive, reactive systems that are more useful and efficient.

Imagine a house (or a city!) that always knows where your kids are, and tells you when they might be getting in trouble. Or an office that knows when you are in the middle of an important conversation, and shields you from interruptions. Or a car that knows when you are sleepy, checks the network to find a nearby open coffee shop, and suggests that you stop there for a while. Or glasses that can capture the face of a person you just met, query the network for his/her name, and whisper it in your ear.

These examples are not a long-range fantasy: we are building research prototypes today. The key to changing inanimate objects like offices, houses, cars, or glasses into smart, active helpmates is to teach them how to pay attention to us the way another person (or even a dog!) would. That way they can adapt their

behaviour to us, rather than the other way around.

Making things smart and attentive

How can common objects become smart and attentive? After all, artificial intelligence researchers have been trying to build intelligent machines for 30 years, with little success outside of abstract mathematical domains or carefully controlled factory environments.

We believe that the main problem has been that current computers are dumb, deaf and blind. They are dumb because they are static. They do only what they were programmed to do, and do not learn to adapt to changing circumstances. They are deaf and blind because they mostly experience the world around them through a slow serial line to a keyboard and mouse. Even 'multi-media computers', which can handle signals like sound and image, do so mainly as a transport device that knows nothing of the signal's content.

If you imagine raising a learning-impaired child in a closed, dark, soundproof box with only a telegraph connection to the outside world, you can quickly realise how difficult it is for today's computers to become intelligent and helpful. They exist in a world that is almost completely disconnected from ours, so how can they understand and help us?

We believe that there are three keys to making the objects around us smart and helpful:

- *First*, they need to share our perceptual environment before they can be really helpful. They need to be situated in the same world that we are; they need to know much more than just the text of our words. They also need to know who we are, see our expressions and gestures, and hear the tone and emphasis of our voice.
- *Second*, they need to learn. They need to watch what we do in each situation, and learn our preferences and habits. Filtering and matchmaking agents need to know your likes and dislikes; remembrance agents need to know what you've seen and heard before; and negotiating agents need to know your goals and values.
- *Third*, they have to share knowledge with each other. Without communications between these learning, perceptive agents, they can get only a narrow view of the world, and consequently will be limited in what they can learn and what they can do. To be smart, efficient, and effective helpmates, they have to share knowledge, computational resources, and communication outlets.

Each of these elements is critical. Without knowing the situation (for example, perception), it is impossible to know what action to take. Without learning about the user's habits and preferences, it is impossible to be helpful in any but the most general way. And without communication, it is impossible to search for, integrate, or deliver required information.

Experimental Test Beds

To conduct research on making houses, desks, cars, and clothes more helpful and attentive, we have created a series of experimental test beds at both the Media Laboratory and at BT Laboratories. These test

beds are 'smart environments' that help users to navigate the World Wide Web, to teleconference in virtual worlds, to find multimedia information, and to interact with artificial life agents. They accomplish this by using new technologies like software agents for negotiation, memory augmentation, computer vision for face recognition and head/hand/eye tracking.

The test beds can be divided into three main types: smart rooms, smart desks, and smart clothes. The idea of a smart room is a little like having a butler; that is, a passive observer who usually stands quietly in the corner but who is constantly looking for opportunities to help, and who knows your preferences so well that they can act without detailed instructions. A smart desk is similar, but aimed at the office work environment rather than at living spaces in general; it is perhaps more like a good secretary. Smart clothes, on the other hand, are more like a personal assistant. That is, they are like a person who travels with you, seeing and hearing everything that you do, and who tries to anticipate your needs and generally smooth your way. Some of our prototype smart clothes are shown in Figure 1.

These smart clothes allow users to be in constant video, sound, and text communication with office, each other, and the World Wide Web.

Smart rooms, desks, and clothes all have embedded computers. Unlike today's computers, however, they are also instrumented with sensors (mainly cameras and microphones, but also biosensors for heart rate and muscle action), which allow the computer to see, hear, and interpret users' actions. They are networked together by IR ethernet, wire ethernet, and ISDN.

People in a smart room or at a smart desk can control programs, browse multimedia information, and experience shared virtual environments without keyboards, special sensors, or special goggles. People in smart clothes can obtain personalised information about their environment, such as the names of people they meet or directions to the next meeting. Moreover, smart clothes can replace today's laptop computers, cellular phones, personal stereos, appointment books, and watches. They have the potential to become the universal consumer electronic aid.

Our current smart clothes prototypes use off-the-shelf head-mounted displays (HMDs) to provide privacy

Figure 1 – The cyborgs of the MIT Media Laboratory. These 'smart clothes' allow users to be in constant video, sound, and text communication with office, each other, and the World Wide Web





Figure 2—Steve Mann wearing a head-mounted display (HMD) and video-rate radio link

and convenience, see Figure 2. Their CPUs are small and unobtrusive, and alternative input devices have been developed to utilise these machines in just about any context, see Figure 3. They use IR and wireless ethernet to stay in contact with the entire World Wide Web at all times.

The goal is that eventually such devices will be so small and light that they will be worn constantly, much as eyeglasses and clothing are today, thus providing access to computing power at all times. Today's smart clothes are not yet inconspicuous, particularly the HMDs, so they project a rather cyber punk look, see Figure 1. However the coming of continuous computing is not far off; two of the Media Lab cyborgs, Thad Starner at the extreme left and Steve Mann at the extreme right, already wear their devices all day every day.

Current Demonstrators

No wires, keyboards, or mice

In current virtual environments, users must use goggles and gloves, or keyboards and mice. In a smart room or at a smart desk, the position of the user's head, hands, and body are tracked by camera. This allows users to control an application directly with their body movements, without




Figure 3—The twiddler input device

needing any special hardware. This interface has been used to navigate in networked three-dimensional virtual game environments, and as an interface to various visualisation applications. (See Figure 4.)

Playing with artificial creatures

The cameras, microphones and other sensors of a smart room or smart

desk can be the eyes and ears for artificial life creatures. For instance, in the artificial life interactive video environment (ALIVE), the smart room's description of the user is used to composite a video model of the user into a networked virtual reality populated with computer-generated artificial life forms, as illustrated in Figure 5. Information about the user's gestures, sounds, and position

Figure 4—Playing in a multiperson virtual game at MIT; the other participant is at BT in Martlesham and they are connected by an ISDN line



Figure 5—Playing with a virtual dog; the dog is on a computer at BT; the person is at MIT

are used by the artificial life forms to make decisions about how to interact with the user. The human users and virtual creatures can be anywhere on the network, and still appear in the same virtual space.

Talking with deaf people

American sign language (ASL) is a sophisticated set of hand gestures that allow deaf people to communicate more naturally. By combining the smart desk's ability to track the user's hands and head by camera with hidden Markov modelling (HMM), technology developed for speech recognition, we have been able to build a system for reading sign language with a 99.2% accurate classification of a 40-word subset of real-time ASL, Figure 6. The ability to recognise ASL in real time opens the possibility of interfaces for deaf people that will match the speech recognition systems now being introduced for the hearing.

The cyborgs are coming

In a smart room, the cameras and microphones passively watch people move around. However, when we build them into a person's clothes, the computer's view moves from a passive third person to an active first-person vantage point. This means that smart clothes can be more intimately and actively involved in the user's activi-

Figure 6—Real-time reading of American sign language



ties, making them potentially a real personal (digital) assistant.

Augmenting human memory is a major application for smart clothes. For instance, when we built a camera into eyeglasses, we found we could use our face-recognition software to help remember the names of people. When we met people, the camera would send pictures of their faces over the net, our home computer would recognise them, and finally our smart clothes would whisper their names in our ear.

Similarly, we found that a computer built into your clothes can automatically remind you of important facts that are related to your current conversation. For instance, if you have a discussion about 'the Megadeal contract', your smart clothes can automatically project Megadeal's finances onto the display built into your glasses.

The Future

We now see two main research directions for development of smart environments. The first is development of completely seamless communications. The computers in my smart clothes should automatically coordinate with the computers in my home, car, and office. They should share not only text, but also sound and video. This is not just a matter of automatic call forwarding or teleconference forwarding; it is seamless services. For instance, if the kids are playing a video game or watching a movie when it is time to drive into town, then the

game or movie should migrate to the display in the back seat of the car. Great strides are already being made in this area at BT Laboratories.

The second research direction is making smart environments more intelligent; this direction is not as well defined as the first. We believe that to build really useful and helpful environments we need to develop three different types of intelligence. These are:

- *Perceptual intelligence*: hearing, seeing, and categorising situations. In journalism classes, they teach that the most important questions a newspaper article must address are: who, what, when, where, and why. This is exactly perceptual intelligence—you cannot evaluate a situation and determine a course of action unless you can answer these perceptual questions. Perceptual intelligence for smart environments is now becoming possible: microphones and cameras are now common in our offices, homes, and even in cars and on our person; moreover, we are beginning to see practical systems for recognising words, faces, and expressions.
- *Affective intelligence*: understanding user's emotional and evaluative responses and their subjective preferences. People are not cold, uniform automatons. They have preferences, likes and dislikes.

They get sleepy, confused, interested and excited.

Taking account of these affective states is critical for an agent that is trying to be helpful. For instance, people cannot learn or perform well when sleepy or confused. They will not return to a service that does not conform to their likes and dislikes, but will become almost addicted about a service that makes them feel interested and excited. As a result, any effort to build a smart, helpful environment must be sensitive to, and make use of, information about a user's effective state.

- *Cognitive intelligence*: symbolic and inferential knowledge, such as used by negotiating, matchmaking, or filtering software agents. People have very limited ability to find, assess, and react to information. Consequently, communication is always a primary problem in every organisation. We spend huge amounts of time answering questions like: who else has worked on this problem? Is there anyone who wants to buy this old computer? Who would be interested in using this new technology? Consequently, further development of software agents to help with these problems is a critical need.

We have recently begun serious research efforts in each of these areas, and already have made some promising advances. Much more, of course, remains to be done. To see the current state of our perceptual, affective, and cognitive intelligence research please visit us at the following World Wide Web sites:

<http://ive.www.media.mit.edu/projects/ive>

<http://wearables.www.media.mit.edu/projects/wearables>

<http://agents.www.media.mit.edu/groups/agents>

Biographies



Alex Pentland
MIT Media
Laboratory,
Cambridge, Boston,
USA

Alex Pentland received his Ph.D. from the Massachusetts Institute of Technology in 1982 and began work at SRI International's Artificial Intelligence Centre. He was appointed Industrial Lecturer in Stanford Universities' Computer Science department in 1983. In 1987, he returned to MIT and is currently Head of the Perceptual Computing Section of the Media Laboratory, a group that includes over 50 researchers in computer vision, graphics, speech, music and human-machine interaction. He has researched into artificial intelligence, machine vision, human vision and computer graphics, and has published more than 180 scientific articles in these areas.



Rosalind Picard
MIT Media
Laboratory,
Cambridge, Boston,
USA

Rosalind Picard earned a B.E.E. from the Georgia Institute of Technology and was named an NSF Graduate Fellow in 1984. She earned the M.S. and Sc.D. in Electronic Engineering and Computer Science from the Massachusetts Institute of Technology in 1986 and 1991 respectively. She was a member of the technical staff at AT&T Bell Laboratories from 1984-1987. In 1991, she was appointed Assistant Professor at the MIT Media Laboratory, in 1992 was awarded the NEC Development Chair in Computers and Communications and in 1995 was promoted to Associate Professor. She is one of the pioneer in algorithms

for content-based video and image retrieval. Her research interests include textile and pattern modelling, learning, video understanding and affective computing.



Patti Maes
MIT Media
Laboratory,
Cambridge, Boston,
USA

Patti Maes is an Associate Professor at MIT's Media Laboratory. She holds the Sony Career Development Chair. Previously, she was a visiting Professor and a research scientist at the MIT Artificial Intelligence Laboratory. She has a Bachelor's degree and Ph.D. degree in Computer Science from the Vrije Universiteit Brussel in Belgium. She is one of the pioneers of a new research area called software agents; that is, semi-intelligent computer programs which assist a user with the overload of information and the complexity of the on-line world. She is the editor of three books and is an editorial board member and reviewer for numerous professional journals and conferences. She is a founder of Agents, Inc., in Boston, Massachusetts, one of the first companies to commercialise software agent technology.

The Face of Talking Machines in a Multimedia World

In an increasingly multimedia-oriented world, people will come to expect services which provide both sound and vision. Services which traditionally require human-to-human discourse will be replaced and extended through the use of multimedia telephony. Computers will need to communicate in a way that is natural and pleasant for their human users. This can be achieved by the use of a synthetically generated person (synthetic persona) which moves and talks in a lifelike manner. This article briefly overviews how BT is combining two important technologies—low-cost facial-image generation and text-to-speech synthesis—to provide the first steps in the development of such multimedia interfaces.

Introduction

Communication is about the transfer of information. For human beings, the most natural form of communication is to have a conversation with someone. In any conversation, there are a large number of information-bearing signals passing between the speakers, through speech and gestures. Body and face movements are commonly used to supplement speech and may contain information about the emotional state or real intention of the speaker. Consequently, wholly natural information exchange can be achieved only through a medium which provides both sound and vision. This statement holds whether we are talking about human-human discourse or human-machine discourse. In other words, people like to talk and they prefer to see who they are talking to! The current generation of automated information services is difficult to use and intimidating to some people. The customer feels forced into a restricted and inflexible dialogue with a 'faceless' machine.

As technology progresses, synthetic personae will appear in almost every walk of life—education, entertainment, information and marketing, to name but a few. For example, imagine an activity such as booking a holiday. In the future, this may be done from the comfort of your own home via a video telephone or your TV! When deciding on a particular holiday, the customer, could interactively view different locations and ask questions of a synthetic persona, acting as their travel agent.

In addition, the ability to present a particular visual image has an

impact on the emerging area of video conferencing. In such applications, a synthetic persona could be used in place of real video images, either through choice, or when, for whatever reason, a real-time picture was not available. The business person would then have the ability to present the image they wanted at a meeting!

As stated, the eventual aim of a research team at BT Laboratories is to develop a synthetic persona, which behaves in a way we, as humans, see as pleasant and natural. This is an ambitious aim, and it will be some years before anything approaching truly natural communication with a synthetic persona is possible. However, the first steps to this goal have been taken, and soon we will start to see services which provide an inkling of what is to come.

The rest of this article provides a brief overview of the current research in this area. In particular, it concentrates on two technologies which are key to the development of a synthetic persona: face image generation and text-to-speech synthesis.

Face Image Generation

The MAXIS head¹, under development within the centre for human communication (CHC) at BT Laboratories, uses one of the original methods of computer animation—a so-called *wireframe model*, shown in Figure 1, which is really a data structure contained in the computer's memory.

In the MAXIS system, a technique known as *texture mapping* is used to project an actual colour image of a person's face onto a three-dimensional wireframe approximating the shape of the person's head (shown in




Figure 1 – Wire-frame model

Figure 2 – Head with texture mapped wireframe

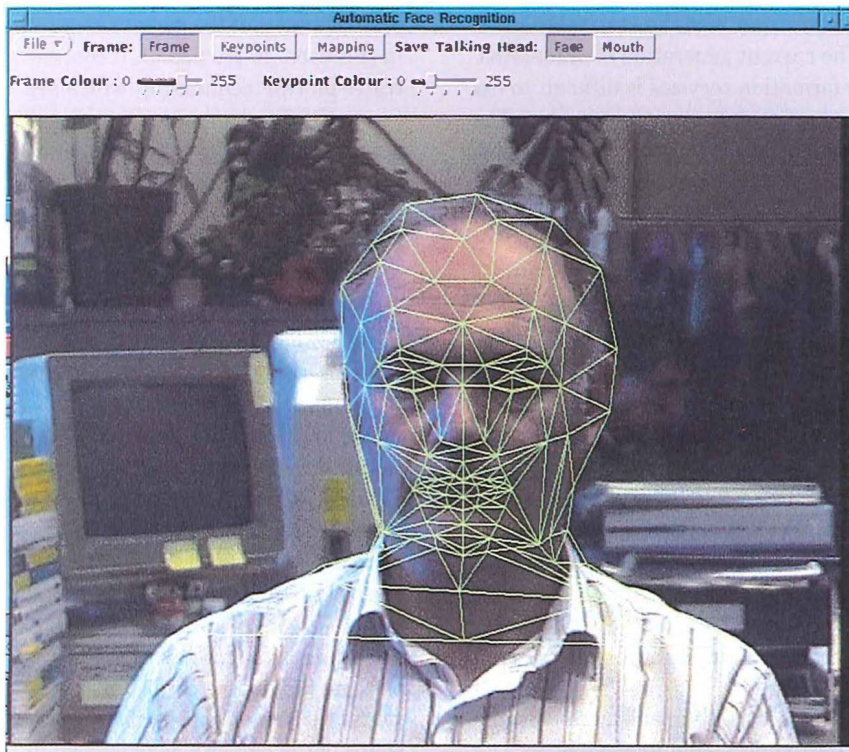


Figure 2). By moving the texture-mapped wireframe, an effect is obtained which has the appearance of seeing the actual person's head in motion.

The wireframe is controlled through a set of *action units* (AUs). An action unit is an index into a pre-defined set of wireframe shapes which describe how the head is to be modified when it is applied. These units may be sent singly or in combination to produce an overall effect. Any number may be defined and groups can be ascribed to specific head gestures. For example, action units are defined to control blinking, eyebrow movement and head rotation. In addition, groups may be used to provide emotional cues such as happiness, sadness and amusement etc. One of the most important groups are those used to control the tongue, lips and jaw.

From the definition provided above, it is apparent that an action unit can be very simple in concept, representing a basic facial movement, such as rounding the lips, or very complicated representing all the movements needed to produce a particular expression. Examples of these are shown in Figure 3. Figure 3(a) shows the face before any modification, while Figure 3(b), shows the result of applying the simple lip rounding to the image. Figures 3(c) and 3(d) show the results of applying the anger and sadness action units to the image respectively.

Generating visual speech

For the head to appear to speak, action units must change in response to the synthetic speech produced by the text-to-speech synthesiser². A text-to-speech synthesiser, such as the BT Laureate system, takes unrestricted text and converts it into synthetic speech. Unfortunately, the process of assigning action units to speech is no simple matter. Speech has evolved over the course of tens of thousands of years, and in the process changed human physiology, such that, today, we are capable of producing a




Figure 3—Application of action units

Table 1 SAMPA (speech assessment methodology phonetic alphabet) for British English²

p	pear	t	tear	k	king
b	bear	d	dear	g	gear
f	fear	T	thing	s	sing
v	very	D	this	z	zing
S	sheer	l	leer	j	year
Z	treasure	w	wear	r	rear
m	men	n	near	N	wing
tS	cheer	=n	button	h	hear
dZ	jeer	=l	bottle		
@	ago	{	bat	E	bet
l	bit	Q	cod	U	good
V	bud	3	bird	A	bard
i	bead	0	bore	u	boot
@U	zero	al	pie	aU	cow
E@	hair	el	pay	l@	peer
Ol	boy	U@	Ruhr		
"	Primary stress	\$	Syllable boundary		
'	Secondary stress	#	Word boundary		

bewildering array of complex sounds; these sounds are produced using the speech articulators—the jaw, lips, teeth, tongue, velum and larynx. When we watch someone speaking, we are seeing an intricate interplay of these articulators.

A moment's introspection will reveal that a sound such as 't' in 'tea' differs in lip shape to the 't' in 'two'. These, sometimes dramatic, differences in facial movements due to context are the visual correlate of the speech effect known as *coarticulation*. Any method of mouth-shape generation attempting to produce realistic mouth shapes must have the ability to appropriately model these complex movements. Coarticulation and its affect on the design and collection of a set of action units is considered later.

A language may be viewed as being composed of a discrete set of abstract symbols, the phonemes of that language. The set of phonemes in a language is defined as the minimum number of symbols needed to describe every possible word in that language. In standard British English there are some 46 such phonemes (see Table 1).

A phoneme is not a sound; it may be viewed as the label for a set of sounds which when spoken as part of a word do not change the meaning of that word. There are an unimaginably large number of sounds used in a language.

The problem is simply this: we, as humans, do not speak in discrete units—speech is produced as a continuous flow of articulatory movements which, upon introspection, can be 'written' as a discrete set of symbols. This flow from one articulated sound to another is coarticulation. In other words, it is a term used to describe the effect of local articulation on a given sound. For expediency, coarticulation is often described as being the effect of context on the production of a phoneme.

When an utterance is to be synthesised, text is presented to the Laureate text-to-speech system. Laureate

produces synthetic speech and a data structure which contains knowledge gained during the synthesis of the utterance. In particular, this data structure contains phoneme type and duration information. This information forms the input to the mouth shape generation component of the MAXIS head, where action units are associated with each of the phonemes (the effects of coarticulation are modelled at this stage). Once a sequence of action units has been generated, it, and the synthetic speech, are sent to the head player to be presented to the user. The next section describes in more detail the process of text-to-speech synthesis.

Text-to-Speech Synthesis

As stated above, text-to-speech synthesis is the process of converting plain text into synthetic speech. Early synthesisers produced the speech sounds by modelling the human vocal tract, but with the development of speech signal processing techniques it is now possible to synthesise speech from a database of phonemes in various contexts. BT's Laureate system^{3,4} is based on this latter technique and can produce speech that is recognisably that of a particular individual. This makes Laureate ideal for combining with the synthetic persona technology described earlier where both the image and speech could be based on the same person.

The process of text-to-speech conversion can be broken down into three main parts, as shown in Figure 4. The first activity is text analysis, where the underlying language structure is determined and a string of phonemes is produced in place of the text. Next come the parallel processes of prosody synthesis and speech sound selection. Prosody synthesis generates the intonation and rhythm of the passage to be synthesised, while speech sound selection involves choosing speech fragments from a recorded speech database. Finally, the speech waveform is generated by smoothly joining the selected frag-

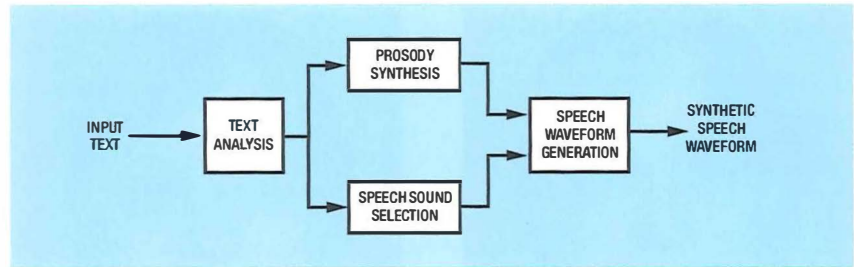


Figure 4 – The process of text-to-speech conversion

ments of recorded speech and imposing the pitch and rhythm onto it using speech modification algorithms.

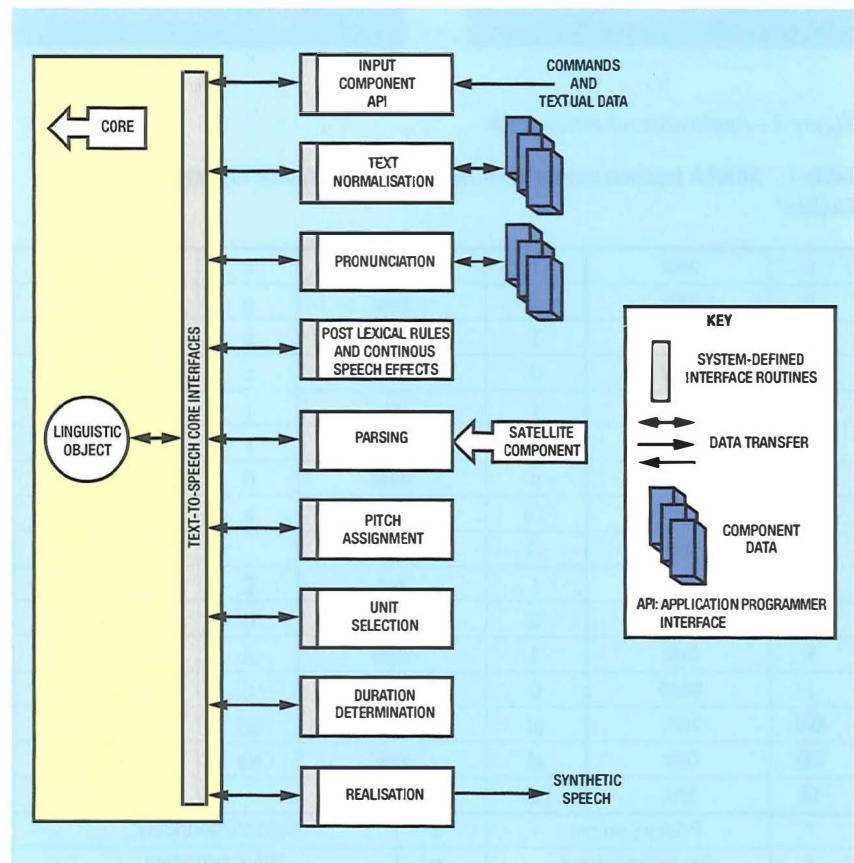
The structure of Laureate can be viewed as consisting of three main parts as shown in Figure 5: a core block, which is the skeleton of the system, a set of system components, which actually convert the text into synthetic speech, and some language specific data.

Looking at the components in more detail, consider the text 'Does the girl live at no. 156 St. Margarets st.?'

The second component, text normalisation, has a number of tasks. First it has to decide where the sentence finishes. Is it after 'no.', 'St.', 'st.', or '?'? Next the abbreviations are expanded where 'no.' becomes 'number', 'St.' becomes 'saint', and 'st.' becomes 'street'. The number 156 is also expanded to 'one hundred and fifty six'. This component also expands acronyms, times, dates and amounts of money.

The pronunciation component produces a string of phonemes for the

Figure 5 – The Laureate system structure



sentence using a combination of dictionaries and rules.

The parsing component assigns parts of speech to the words. This is required to deal with words that are spelt the same but pronounced differently; for example, 'I *live* in a house.' 'I went to a *live* concert.' It also breaks the sentences up into sayable blocks.

Pitch assignment has the task of working out a symbolic representation of the intonation to be applied to the speech.

The unit selection component selects fragments of recorded speech, from a specially prepared speech database. These fragments of speech will be joined together in the realisation component.

The duration component takes information from the parser and computes the duration of the segments of the speech to give it an appropriate rhythm.

The speech produced by the Laureate system retains many of the characteristics of the original speaker.

Once all the analysis components have completed their tasks, the realisation component performs the signal processing to join smoothly the fragments of speech, and impose the intonation and duration on the resulting speech signal.

This flexible structure has been modified to accommodate the information needed to generate action units used in the control of BT's MAXIS talking head. When combined with MAXIS, the Laureate system contains an extra component. This component sends information such as phoneme type and duration to the MAXIS action unit generation component.

Voices, accents, styles and emotions

The Laureate technology and architecture permits new voices, accents, and languages to be incorporated into the system with minimal changes.

New voices are introduced by recording the voice of the chosen speaker reading a specially devised script, which has a good coverage of the sounds of the language. Many regional British English accents may be included simply by recording a speaker with the desired accent. A different language would require changes in areas such as text normalisation and parsing, but the core system would not change.

The speech produced by the Laureate system retains many of the characteristics of the original speaker. The most noticeable of these is the basic voice quality, but also some aspects of accent are retained, and thus when combined with the head model of the same person make a very compelling human-machine interface. Text-to-speech synthesis systems have been around for a long time and most people will be familiar with the robotic sounding speech that many of them

produce. Today with the development of systems such as Laureate, which are based on recordings of real human speech, the level of naturalness has noticeably improved. In the near future, it is likely that synthetic speech will be produced that is indistinguishable from a real person. This will depend on the development of text-to-speech systems that can be made to speak in a style appropriate to the situation and the text being read out. To do this, sophisticated understanding components will need to be invented, which actually 'know' what it is being talked about.

Even with a fully developed understanding of the text to be spoken, techniques are still to be developed that can effectively mimic the myriad of different styles and emotions present in natural speech. Some speech styles and emotions may be more readily modelled than others.

In the current generation of text-to-speech systems we have an effective tool for modifying the basic pitch and rhythm of language, but only a very crude ability to modify the quality of the basic sounds stored in the database of units. It remains to be seen how many styles and emotions can be effectively modelled by pitch and rhythm alone.

Conclusion

This article has introduced the topic of synthetic persona development. In particular, it has concentrated on techniques for developing talking heads which can be used in existing or near-future systems. The aim is to produce technology which can be continually improved as knowledge and technology progress, while providing an effective tool for BT, both now and in the future. Synthetic persona technology will only progress through a combined effort in three disciplines: language understanding, spoken language synthesis and image generation. The challenge currently facing developers of this technology is how best to combine these different fields into a consistent and effective user interface.

References


- 1 WELSH W. J. Model-Based Coding of Moving Images at Very Low Bit-Rates. Picture Coding Symposium (PCS 87), Stockholm, Sweden, June 1987.
- 2 BREEN, A. P. Speech Synthesis Models: A Review. *IEE Electronics and Communication Journal*, Feb. 1992, 4(1).
- 3 PAGE, J. H.; and BREEN, A. P. The Laureate Text-to-Speech System—Architecture and Applications. *BT Technol. J.*, Jan 1996, 14(1), p. 57.
- 4 EDGINGTON, M., *et al.* Overview of Current Text-to-Speech Techniques: Part 1 and Part 2. *BT Technol. J.*, Jan. 1996, 14(1), p. 68 and 84.

Biographies



Andrew Breen
BT Networks and
Systems

Andrew Breen has a B.Sc. degree in Physics with Computing Physics from University College Swansea, an M.Sc. (Eng.) in Electrical Engineering from Liverpool University and a Ph.D. in Speech Science from University College London. After completing his first degree, he was employed in the Department of Electronics and Electrical Engineering at Liverpool University, where he worked for three years in the area of speech coding and digital filter design. After leaving Liverpool University, he was employed in the Department of Phonetics and Linguistics at University College London where he worked for 10 years on continuous speech recognition before moving into the area of text-to-speech synthesis. He is project leader of the language component of the Laureate text-to-speech system and project manager of the synthetic persona project. His current research interests include speech synthesis, speech prosody and the representation of linguistic and para-linguistic information in a synthetic persona. He is a member of the IEE.



Julian Page
BT Networks and
Systems

Julian Page graduated from Warwick University in 1972 with a B.Sc. degree in Engineering Science. He then joined BT and worked in a group designing submarine coaxial cables, several of which were installed across the Atlantic. In 1979, he spent a short period in a group responsible for the design of digital PABXs. In 1980, he moved again, on promotion, to a group responsible for the design of high-speed voice-band modems where he was involved in the design of several successful modem products, as well as having responsibility for the BT datapump project, a key component in the BT Fifth Generation range of modems. He is currently leading a team undertaking research and development on speech synthesis and analysis techniques, the main output of which has been the Laureate text-to-speech synthesiser.



Bill Welsh
BT Networks and
Systems

Bill Welsh graduated from Liverpool University with a B.Eng. degree in Electronics in 1980. He has worked on image coding and processing at BT Laboratories since 1984 and gained a Ph.D. in this area from Essex University in 1991. He was one of the early pioneers of model-based image coding, presenting papers on the subject at many international conferences and colloquia. This work led to the development of the synthetic persona in collaboration with the Speech Synthesis group at BT Laboratories.

A Model for the Project Management NVQ

This article traces the development of a National Vocational Qualification in project management from its inception through to delivery. It starts with an appreciation of the NVQ concept and background, identifies the mandate from the Department for Education and Employment (DFEE) to develop such an NVQ and describes the methods and processes used to deliver it. The role of BT in the development of this NVQ is explained, and the article concludes with the management process to be put in place for the ongoing marketing and delivery of the NVQ.

Introduction

Project management is as old as humanity, since by definition any management activity that introduces a new objective or causes change and has a definitive start and finish time, is a project. More specifically, a project is a unique set of coordinated activities, with defined start and finishing points, undertaken by an individual or organisation to meet specific objectives within defined time, cost and quality parameters. The project manager is the individual or body with responsibility for managing the project to achieve specific objectives. It will typically pass through the project life cycle as defined by inception, feasibility, definition, development, installation, pilot, roll-out and closure.

Formal project management, as we know it today, started from the development of network techniques in the late 1950s, yet formal qualifications in project management are generally lacking. Rather than being stand-alone qualifications, they will usually form part of general management qualifications or university degree modules, where this competence is specifically required—and will be manifested at various levels of intensity and scope.

It will be of interest, therefore, to the project management community that opportunities are being developed for those many and various skills to be recognised in the job that is actually done in the workplace. The vehicle for this will be the project management National Vocational Qualification (NVQ) at levels 4 and 5.

BT has supported the development of such an NVQ through a

project management exemplar. The BT representative (a project management practitioner), together with similar representatives from other sectors of industry, has been involved in all stages of its development since January 1995. The result is a draft document built on the experience of project management practitioners, field trialled across industry and amended ready for submission for accreditation.

Basic Appreciation of the Vocational Qualification

It would be pointless launching into the detail of the exemplar project itself without a basic appreciation of the concept of the vocational qualification (VQ) on which it is based.

An NVQ is a qualification as awarded by the National Council for Vocational Qualifications (NCVQ).

National means that it is a scheme supported by the Department for Education and Employment (DFEE) across all sectors of industry in the UK.

Vocational means that it is work-based—it does not depend upon exams or tests, although they may help, since it is about collecting evidence that proves you do your job competently.

Qualification means that, once sufficient evidence has been gained for the requisite units, the candidate can be awarded a certificate. Unit accreditation can also be recognised.

A Scottish Vocational Qualification (SVQ) is the equivalent qualification in Scotland and will be awarded by the Scottish Vocational Education Council (SCOTVEC).

NVQs and SVQs are available in five levels, from level 1, which covers

the competence and application of knowledge in basic predictable tasks rising with increasing responsibility, to level 5, the highest achievable.

A level 5 NVQ would suggest that the candidate has a high level of personal responsibility, autonomy and accountability over a wide range of activities. A level 4 would suggest a slightly narrower responsibility level. The project management NVQ was considered appropriate to levels 4 and 5 only.

An NVQ consists of several units covering a particular competency—in this case project management. Each unit is composed of a number of *elements*, which detail the area of activity to which the standard applies—all of which have to be completed to achieve an award in that unit. Each of the elements contains a group of *performance criteria* which specify exactly what level of performance a candidate will need to achieve, associated with a number of *range statements* which give a range of situations and contexts which need apply. It is the *evidence* which goes to make up the candidate's portfolio and will include *product evidence* (what the assessor can look at), *process evidence* (what the assessor can observe), and *knowledge evidence* (what the candidate can tell the assessor).

Background to NVQ Development

A previous article¹ gave an outline of the British Government's initiative in 1986 regarding vocational qualifications.

Competence-based systems, comparable with NVQs and SVQs, are at varying stages of development across Europe (Finland, Greece, Portugal, Spain, the Netherlands and Italy) and internationally (Australia, USA, New Zealand, South Africa and Mexico).

Occupational standards, defined by industry, form the basis of all NVQs and SVQs. To set these standards, and advise on the

requirements of the qualification, industry sectors have formed consortia of employment interests. These bodies are called *industrial lead bodies*. There are some 160 such organisations, each directed to set standards across a particular occupational field. Some are focused on a single occupational field, others such as the Management Charter Initiative (MCI) cover a very broad field. Within the engineering sector several Lead Bodies are involved in this work. For technical, professional and managerial occupations in engineering, there are four bodies directly concerned, under the umbrella of the Engineering Occupation Standards Group (EOSG) as identified below. BT has representatives on many lead bodies and Kents Hill is an accredited assessment centre for many awarding bodies.

Nearly 1 000 000 NVQs have now been awarded in the UK and some 80 per cent of all those in employment now have access to an NVQ or SVQ that relates to their job. The residual 20 per cent are mostly in technical, professional and managerial areas, where the development process is more complex. No sector of industry or the professions has backed away from this initiative even though it is based on voluntary commitment. BT is committed to having 50 per cent of its people working towards an NVQ by January 1997, and NVQs available to everybody by March 1997.

The Development of NVQs in Project Management

The DFEE gave approval to the Engineering Occupation Standards Group (EOSG) to run an exemplar to develop a generic project management qualification. Exemplars are currently being developed for other NVQs, but only one in project management, and the mandate was to concentrate on the highest levels 4 and 5 as follows:

Objectives of the exemplar

For the EOSG to 'design and trial pilot NVQ/SVQ at levels 4 and 5 in engineering project management'

Scope of the exemplar

Use representatives from the standing conferences, education and industry, in a voluntary capacity, to produce NVQs/SVQs which will:

- appeal to a wide population;
- not conflict with existing qualifications; and
- fulfil industry needs where no qualifications exist.

Deliverables

- A 'family' of pilot NVQs/SVQs in Engineering Project Management at levels 4/5, trialled by industry and revised accordingly.
- A potential awarding body partnership prepared to take it forward to accreditation.

Methodology

Project Management: The project will be project managed by a project manager and the appropriate expert guidance will be provided by a consultant. Both individuals will be approved by the EOSG and their names notified to the Employment Department.

Development Work: The development work will be carried forward by working groups dispersed throughout Britain and Northern Ireland, and made up of field practitioners.

Setting up the Process

The EOSG is an umbrella organisation representing all sectors of the engineering industry through various industry standing conferences (ISCs)—construction (CISC), service (ESSC), manufacturing (SCEM) and extraction and processing (SCEP). The EOSG project manager invited representatives from various sectors to a meeting at

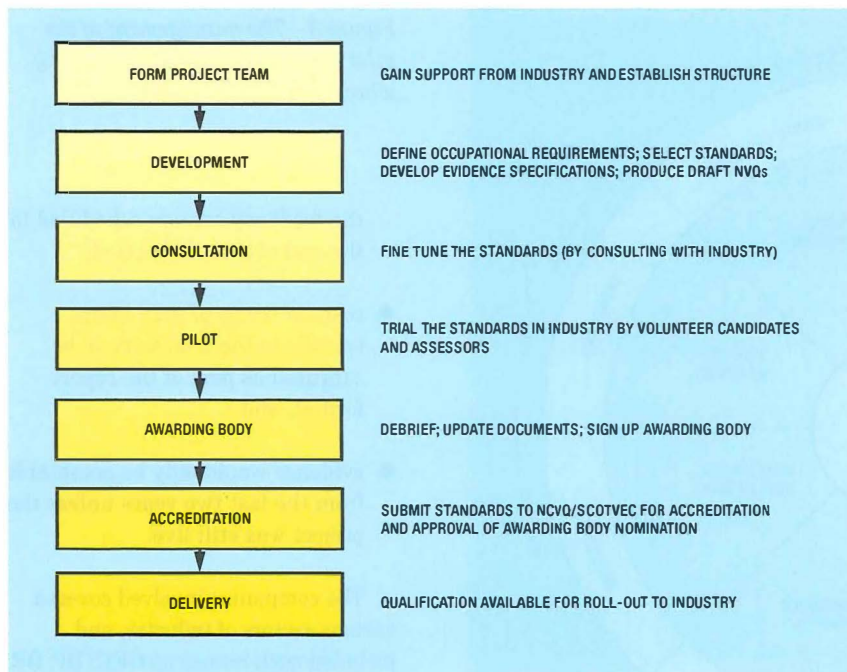


Figure 1 – Setting up the NVQ process: development to delivery

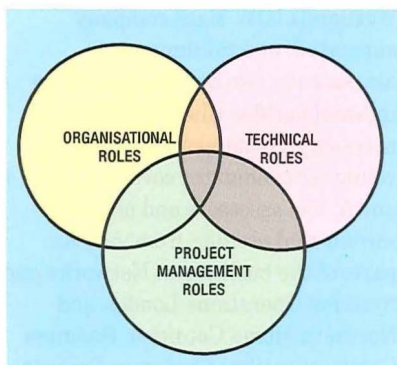


Figure 2 – The overlap of roles within project management

the Engineering Council in January 1995. At this first meeting, 15 representatives at senior level in industry, including those from the four ISCs, the Management Charter Initiative, and academia discussed the proposal and determined the way forward.

It was decided that a steering committee (to be known as the *validating group*) would provide overall direction of the project. At the working level, regional workshops (six throughout the country) were planned to take on the development tasks between them (Figure 1) and report back to the validating group. In reality, owing to the high level of experience required of the participants, and conflicting priorities on industry at the time, it was difficult to resource all of the regional groups as planned, and a decision was taken to resource the

whole development from a central body of people. A new development group was formed, therefore, with people drawn from a wide area and from several sectors of industry. The regional workshops, as they existed, acted as a sounding board and reported back to the validating group. This structure was maintained for the duration of the development and worked well, with practitioners from industry giving their time generously and at nil cost to the project. BT was represented at all levels and made a substantial contribution to the draft document put out for consultation during the period August to November 1995.

Identifying the Standards

Project management cannot be wholly separated from the functions of organisational management or from the technical roles in engineering design and development. There are overlaps (Figure 2).

The process by which occupational competence in project management has been defined started with an analysis of, and agreement on, the central functions that must be performed. Desk research into definitions of project management devised in the UK and internationally were also examined.

Against this functional profile, occupational standards relating to these functions were selected from

material developed across the engineering sector which match the identified requirements. Over 36 occupational standards were identified from existing approved sources for level 5 and 30 for level 4. These standards—elements (of competence)—when grouped together, form the **essential** units comprising the proposed qualification; nine units for level 5, and seven units for level 4 as follows (units for level 4 in italics):

- Proposing and agreeing objectives
- *Developing and agreeing specifications*
- *Managing contracts*
- *Developing plans*
- *Controlling risk*
- *Organising coordinated implementation*
- *Resourcing the plans*
- *Controlling implementation*
- Developing improvements

The need for additional units to meet the needs of some sectors of industry and individual companies was foreseen and would be subject to debate at a later stage.

Consultancy on the Standards

A series of five consultation workshops were held nationwide to start the consultation process, and the draft documents using the standards defined above were sent to over 150 companies, institutions, training bodies and industry lead bodies. Some 2000 comments were received in the four weeks of the consultation phase, with no company or individual suggesting any fundamental change to the standards. The revised pilot trial editions of the standards

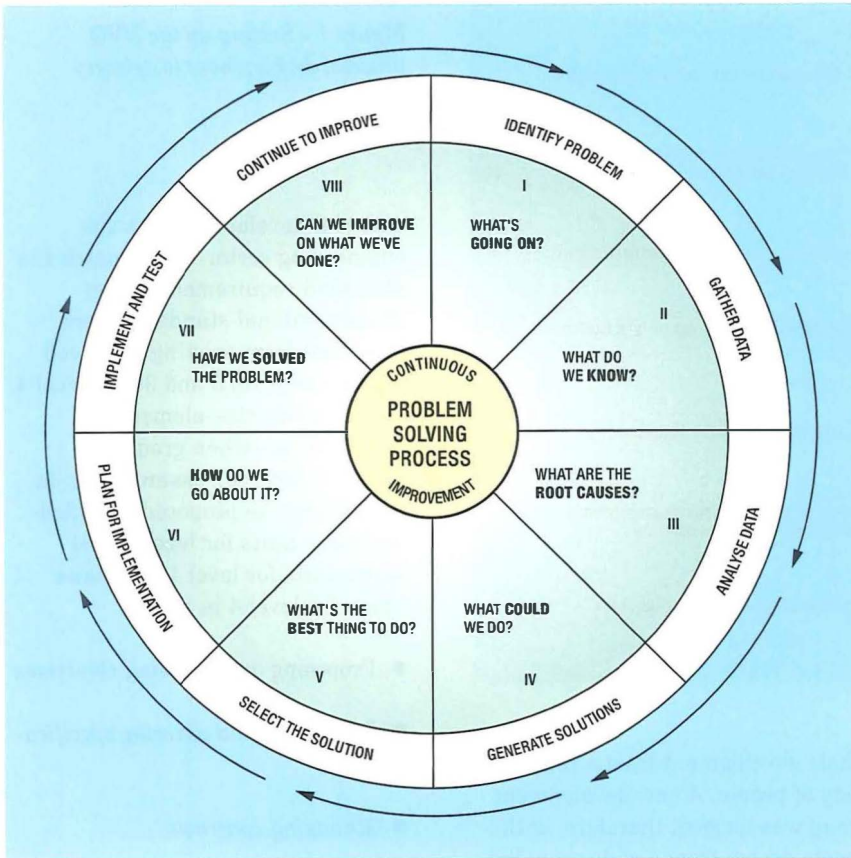


Figure 3 – The management of the pilot trial used the problem solving wheel

were used as the basis for candidate/ assessor workshops held at IBM offices in London on 5 December 1995 and at Rolls Royce Aerospace in Derby on 8 December, to launch the trial. Over 45 people attended the London session and 30 attended the Derby. All volunteers were briefed on the pilot arrangements and time-scales, and training was given in the skills required for the exercise.

Pilot Trial—Management

The pilot started officially on the 5 December, and was to run until the 29 February 1996. Participating in the pilot were 55 candidates, 36 assessors and 10 verifiers from 16 organisations.

The stages of the project were clearly identifiable from the sectors of the problem solving wheel (Figure 3), enabling a quality product to emerge from the months of data gathering, analysis, trial, decision making and implementation phases. The process of feedback and continuous improvement would of course continue even after the launch of the new qualification.

Although candidates and assessors were experienced project managers, they were largely new to

the NVQ environment, and it was necessary to lay down some ground rules for assessment, which were agreed as follows at the launch workshops:

- candidates were to familiarise themselves with the requirements of the trial;
- a project plan was to be agreed between candidate and assessor within the first two weeks of the trial, and be managed throughout;
- candidates were to select key projects as being the most productive basis for the trial;
- the methods of presenting evidence and the indexing required were to be agreed;
- the internal verifiers were to satisfy themselves on the above;
- the internal assessors, on a second visit, were to assess what happened, how it happened, and determine if there was a level playing field;
- the internal assessors, on a third visit, were to gather material for

the feedback session scheduled for the end of the trial period;

- costs in terms of man-hours specific to the trial were to be captured as part of the report format; and
- evidence would only be acceptable from the last two years unless the project was still live.

The companies involved covered various sectors of industry, and included such names as GEC, BP, BR, BNFL, Rolls Royce Aerospace, IBM, Courtaulds, Thames Water and Westlands IOW. Each company supported, at a minimum, one assessor per two candidates and an internal verifier where this was necessary. BT supported eight volunteer candidates covering levels 4 and 5, five assessors and one assessor/internal verifier, from various parts of the business in Networks and Systems Operations London and Northern Home Counties, Business Communications Customer Projects, and the BT Project Group. Each candidate undertook to complete two units each, selected on the basis of those which fitted best with their normal work function. It was necessary to cover all the units for the two levels of qualification; the units candidates undertook were agreed at the workshops. Throughout the trial, assessors ensured that standards were maintained at a high level and the internal verifier ensured conformity and interfaced with the EOSG project manager.

Reports were prepared by candidates, assessors and verifier as follows:

- candidates reported against performance and knowledge evidence, commented on gathering and assessing evidence and clarity of standards;
- assessors gave qualitative comments on the assessment process, documents and guidance state-

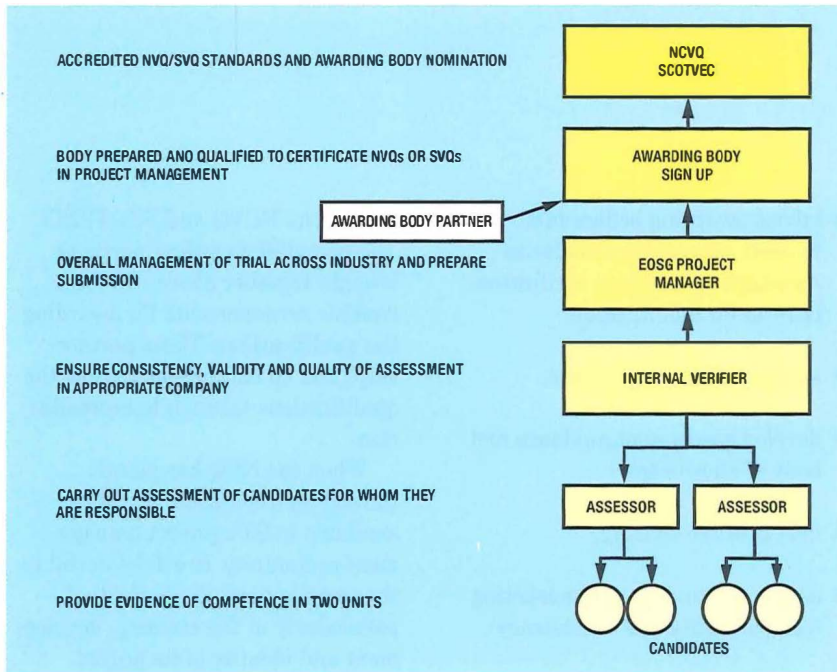


Figure 4—The pilot assessment and accreditation system

ments, as well as noting general progress of the candidate; and

- verifiers gave further comment on the assessment process; availability of evidence, rigour of assessment, assessment of knowledge and consistency, as well as providing qualitative comments on the documents and guidance statements.

This documentation was in addition to that expected to support an NVQ portfolio, and provided a platform for the EOSG to capture best practice for the qualification by making any necessary changes to the NVQ standards and structure, before submitting for accreditation and delivery.

All those taking part were enthusiastic, both from the point of view of contributing to a new and necessary national qualification, and in the expectation of being eventually certificated for those units in the interests of personal development.

Pilot Trial in BT and Feedback

The pilot required dedication and hard work from all the participants, who still had all the normal demands on their time. There was, additionally, the learning curve on the NVQ concept itself. Volunteer candidates and assessors were generally not familiar with NVQs, and the assessors

were not initially accredited with the Training and Development Lead Body (TDLB) D32/D33 assessor units. The *learning by doing* made a slow start but gained momentum after the first month.

With candidates and assessors dispersed around London and the Northern Home Counties, there was a conscious need for the internal verifier to arrange BT-based assessor/candidate meetings which would enable the sharing of ideas and experiences and determine improvements. One was held just after the start of the trial, and the other just before the finish. This was in addition to the visits to candidates and assessors in the workplace.

Candidates worked hard to complete the necessary units and provide sufficient evidence to support the performance criteria. There was some movement in the original volunteer candidates, with some unable to continue due to conflicting pressures, but there was no lack of interest in any individual. The fast tracked three months of the trial, which included a major holiday period over Christmas, was demanding indeed. The need for evidence from a normal work function to cover the performance criteria and the range of all necessary elements was sometimes difficult to establish, and was a concept new to the candidates. BT people tend to do their job in the best way they know how, and would not necessarily be able to support every

action with a piece of evidence. Such evidence, however, may well be required as part of the standards, and this issue required special attention. NVQ terminology—witness testimony, simulation, accreditation of prior achievement and learning (APA/L), special assessment requirements (SAR), and underpinning knowledge were certainly terms more familiar at the end of the trial than at the beginning.

The trial was completed on time, on 29 February 1996, and at the final meeting with all participants, it was clear that there were no major problems with the standards, but certain items needed attention before launch, and these were supported by evidence in the documentation to the EOSG project manager. A consensus view, from all participants at the final meeting, was presented to the EOSG. This consensus is summarised as follows:

- it was a useful qualification, with credibility, and gave more focus to the project manager on topics including quality, safety and general documentation;
- there were no perceived gaps in the standards,
- the standards were, however, generally too wordy and difficult to interpret;
- there was some duplication in the elements and overall the number could be reduced;
- the demarcation of levels were unclear, and a level 5 was more appropriate than a level 4 for the person actually managing projects to the appropriate standard; and
- emphasis was generally lacking on the client role, and there needed to be more customer focus.

It was unanimously agreed, however, that such a qualification was very necessary and timely.

Early indications in the feedback from other companies pointed to a clarity of view and some very useful comment. Despite the diverse interests of the participating companies, it was apparent that there was scope to fine tune the standards while maintaining the generic nature of the qualification.

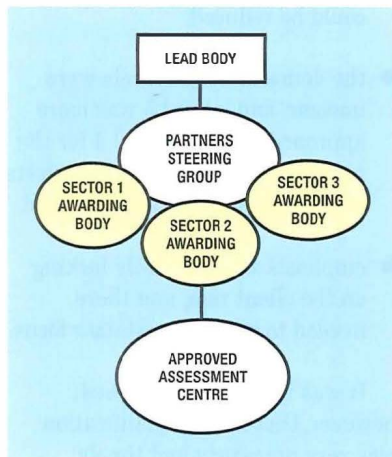
Preparing The Awarding Body/Partnerships

The awarding body, like an examining body for a written qualification, would certificate the candidates completed and verified NVQ/SVQ units. The awarding body(s) may be ably assisted by professional organisation(s) in the delivery of the qualification.

A meeting with the potential awarding bodies and partners convened after completion of the trial, agreed the value in working together. Partners would be able to offer help to members in their professional areas and assist in assessment guidance, particularly in knowledge and comparability of application. A steering group would provide the platform to move forward, as shown in Figure 5, and the following terms of reference were agreed:

- be agents for the EOSG standards (monitor and advise on necessary revisions),

Figure 5—Proposed awarding body partnership arrangement



- advise awarding bodies on all matters concerning standards/structures/assess and verification criteria for qualification,
- set qualification structure,
- develop assessment guidance and body of knowledge,
- host practice sharing,
- assist in branding and marketing (comparability and consistency),
- maintain a glossary of terms,
- provide a source of expertise,
- identify and facilitate research and delivery, and
- be a self-help group—at minimal cost.

Next Steps

The full appreciation of the views of participating candidates and assessors will be achieved at a seminar later in 1996. Apart from a review of the pilot trial itself, there will be a discussion on the perceived benefits of the NVQ/SVQ to employers, candidates and assessors. Such feedback will give important material to the champions of the NVQ in the marketing and delivery of the qualification itself. Some early thoughts on this are given in Figure 6.

The EOSG will consider all the written feedback from participating companies and the verbal feedback from the final review to update the qualification documentation to the final version. It is inevitable that there will be some movement in standards, change of emphasis, and fine tuning generally when all the feedback is assessed. This will be considered by the development group before the final version is presented to the validation group for acceptance that the qualifications are suitable to go forward for accredita-

tion by the NCVQ and SCOTVEC. The potential awarding partners brought together above will forge credible arrangements for awarding the qualifications. These partnerships will be encouraged to take the qualifications through to accreditation.

When the NVQ has passed through accreditation, it will be available to BT's project management community. It will be useful to the company and the individual, particularly in the training, development and identity of its project managers across the business

BT will have made, and be seen to have made, a substantial contribution to industry in the development and trial of a national qualification, available to all project managers across all sectors of industry.

Conclusion

The EOSG project management exemplar was living proof of the cooperation of industry in developing a national qualification, and this was achieved in very tight timescales and comparatively very low cost. In using practitioners from different companies across different sectors of industry, best practice was brought to bear, and this resulted in a quality, fit-for-purpose qualification of real value to industry. This must be the way forward for all new NVQ/SVQ qualifications, and there are many

Figure 6—Suggested reasons to invest in vocational qualifications

Individual:

- the aspirant (wants to get on)
- the defensive (seeks mobility, security, identity)

Employer:

- the 'good' employer (retention)
- give client confidence
- recruitment
- assessment and appraisal
- training and development

still to be developed in other areas of industry and at various levels.

It is still early days for the NVQ and its general acceptance across industry. It will serve to supplement rather than replace academic qualifications in the future, particularly at the higher levels, and make them more meaningful both to the company and the individual.

Acknowledgements

The project management exemplar was managed by Jim Watson—project manager for the Engineering Services Standing Conference (ESSC) under the auspices of the Engineering Occupation Standards Group (EOSG). Jim was assisted by Alistair Robertson, a consultant from Q-West. The success of the exemplar is a credit to their leadership and motivation and to all the companies who gave of their time and experience throughout the project life cycle—too many to name individually.

The BT participants in the trial as candidates or assessors were:

Rasik Patel, Terry Brown, Mick Tucker, Simon Redwood, Roger Beckley, Jim Place, Gerry Hewitt, Jim Williams, Dick Silk, Alan Matthews, Alec McIntock, Nigel Hellon and Ted Smith.

References

- 1 SMITH, H and KEMP, R. National Vocational Qualifications. *Br. Telecommun. Eng.*, July 1995, **14**, pp.170–174.

Biography



Michael Tuggey
BT Networks &
Systems Operations,
London and
International

Michael Tuggey is the Group Network (Projects) Enhancement Manager for Networks and Systems Operations, London and International. He spent many years involved in maintaining the switching network before moving into project management in 1988 on the Isle of Dogs fibre project, and has a continuing responsibility for core switch, transmission and customer-focused projects. His involvement in vocational qualifications originated from the nomination to attend the initial meeting of the NVQ exemplar, at the Engineering Council in January 1995, as BT's representative and project management practitioner. He additionally managed the pilot trial in BT and performed both an assessor and internal verifier role. He is an Incorporated Engineer, Fellow of the Institute of Electronic and Electrical Incorporated Engineers, and a Member of the Association of Project Managers.

National Vocational Qualifications for the Telecommunications Industry

BT has been involved in the development of National Vocational Qualifications for the telecommunications industry since 1989. The evolutionary nature of this development has resulted in changes to these qualifications and the way in which they are implemented. This article outlines some of these changes and provides details of the first of the next generation of NVQs available to candidates in the telecommunications industry.

Introduction

National Vocational Qualifications (NVQs) have been available for nearly 10 years, in one form or another. Quite a lot has been learnt by those who develop these qualifications and by those who use them. The Telecommunications Vocational Standards Council (TVSC), has been working closely with the telecommunications industry, and BT in particular, to develop NVQs appropriate to the industry. The TVSC has also been working with the industry to develop systems to support the assessment and certification of candidates, and the quality assurance arrangements required to maintain the integrity of the awards. This close partnership has produced NVQs for the telecommunications industry which set the national standards of competence in employment. These qualifications also encourage individual development.

The Nature of NVQs

Although NVQs have been discussed in many articles, it is helpful to begin with an overview of these qualifications. The TVSC develops and awards NVQs based on the principles below:

- NVQs are specifications of occupational competence;
- they are based on national standards set by industry,
- they are employment-based qualifications,

- they are used to assess performance in work roles,
- they are independent of the learning route taken, and
- they are open to all individuals having the potential to meet the standards.

In simple terms, a candidate provides evidence of their competency through normal work activities. This evidence is judged against the national standards. If the evidence confirms that the candidate's competency meets the standard, then the candidate is assessed as being occupationally competent.

Occupational Competence

Occupational competence is defined as:

'The ability to perform the activities within an occupation or function to the standards expected in industry'

Within this definition, competent performance has four components:

- *task or technical ability* required to carry out a task,
- *task management ability* needed to manage all the tasks and other components that combine in a work role,
- *contingency management ability* to cope with non-routine and unexpected situations (this ability is also needed in order to act positively when things go wrong), and

- *environment management ability* to perform the function within particular environments which have specific characteristics. Such characteristics may relate to health and safety, interrelationships and interfaces with other people and organisations, economic and social issues, and the natural environment.

Language of NVQs

The TVSC develops the standards of occupational competence in partnership with representatives from the telecommunications industry. These representatives are from large organisations, such as BT, as well as from much smaller businesses. The armed services are also consulted, as many occupations or trades within the services have direct equivalents in civilian organisations. TVSC and their industry and services partners meet to define the national standards. This is one reason for the style of wording of the qualification standards, which may not always be in the language of individual organisations. The various groups of representatives discuss and agree the wording of the standards, which often has to be a compromise acceptable to all parties.

The wording of the standards within NVQs has also to meet criteria set down by the National Council for Vocational Standards (NCVQ). Recent reports on the implementation and take-up of NVQs have suggested that improvements should be made to the wording in the qualifications. Therefore, the TVSC is looking at ways in which standards can be written to be more 'user friendly' while remaining national specifications for competent performance.

Some organisations have produced extra guidance material which provides interpretation of the standards into the environment of the particular organisation. This is acceptable and liaison between the organisation and the TVSC is encouraged to ensure the quality of the qualification standards is not compromised.

Levels of NVQs

NVQs are assigned one of five levels by the NCVQ. The level should reflect the characteristics of the work role for which the qualification was developed. Each level is dependent on:

- range of work activity,
- range of work contexts,
- responsibility and autonomy,
- interaction with others, and
- use of knowledge.

Qualifications in specific areas will not always exist at all levels, as work roles do not necessarily exist at all five levels. Also, candidates need not gain NVQs at the lower levels before achieving higher level qualifications. The achievement of any NVQ at any level depends on an individual's opportunity to generate sufficient valid and reliable evidence of competence to meet the national standard.

Approved Assessment Centres

Assessment of candidates must be through centres, approved by the body awarding the NVQs. Approval criteria relate to sufficient resources, management and administration systems and quality assurance arrangements. Centres are responsible for the registration of candidates and the administration of assessment, finance and certification. Candidates are not expected to go to their approved centre to be assessed, as it is effectively only an administration centre.

Candidate and Assessor Partnership

Assessment of candidates for NVQs is a totally open system; there are no hidden requirements that are kept from the candidate. Where the wording or any aspect of a qualification causes a candidate problems in

understanding what is required, then they should discuss this with their assessor. The TVSC requires assessment centres to make available all details of the qualification standards. Unlike earlier vocational qualifications, the NVQ system encourages a strong partnership between candidate and assessor. This is particularly relevant in the workplace, where objective feedback to employees on work performance is better than subjective feedback, or no feedback at all. A practical benefit of this partnership is that the employee/candidate knows more precisely the expectations of their supervisor or manager/assessor. The qualification standards can therefore be used in the appraisal process.

NVQs and Personal Development

As NVQs are constructed on a modular basis, a candidate's/employee's occupational development can be planned and measured using the units that comprise NVQs. Units of NVQs are, and should be seen as, qualifications in their own right. Several organisations, including BT, have identified the benefit of employees gaining units which, although not necessarily forming a complete NVQ, more accurately reflect the profile of competence required in particular work roles. However, it has been recognised that there is a certain reluctance to achieve a group of units which do not constitute a full NVQ award. The TVSC has therefore developed a second generation of NVQs for the telecommunications industry which provide more opportunity for individuals to gain a full NVQ.

Recent Developments in Telecommunications NVQs

In order to provide flexibility in the structure of NVQs to meet current and predicted patterns of occupations in the telecommunications industry, the TVSC proposed a

Table 1 Proposed framework of Telecommunications NVQs

LEVEL	PROPOSED NATIONAL VOCATIONAL QUALIFICATIONS (Indicative Titles Only)			
4	Designing Telecommunications Systems and Equipment	Planning Telecommunications Services	Providing a Telecommunications Service	Manufacturing Telecommunications Equipment
TVSC CODE	(DSN 4)	(PLN 4)	(SVC 4)	(MFG 4)
3	Designing Telecommunications Systems and Equipment	Planning Telecommunications Services	Providing a Telecommunications Service	Manufacturing Telecommunications Equipment
TVSC CODE	(DSN 3)	(PLN 3)	(SVC 3)	(MFG 3)
2			Providing a Telecommunications Service	Manufacturing Telecommunications Equipment
TVSC CODE			(SVC 2)	(MFG 2)
1			Providing a Telecommunications Service	Manufacturing Telecommunications Equipment
TVSC CODE			(SVC 1)	(MFG 1)

framework of qualifications that has mandatory and optional units. This framework has been accepted by the industry and the first NVQ of this type is now available. This framework is shown in Table 1. (It is not expected that there will be any telecommunications industry specific qualifications at level 5.)

The mandatory units within the qualifications specify the competences that are common to the most occupations in the telecommunications industry. They therefore give recognition to the transferable competences that support technical competence. These units will underpin all qualification structures. The option units relate more to the technical activities found in the telecommunications industry and provide the flexibility for individuals to select units most appropriate to their specific work role.

It is becoming an increasing requirement for employees within the telecommunications industry to be multiskilled. The option units proposed in these qualifications provide various routes for candidates to have these additional skills and knowledge assessed and recognised. The option framework also demonstrates to employers possible paths for development and progression for their workforce.

The first of these qualifications to be accredited by the NCVQ covers

the activities performed by employees who are providing a telecommunications service. These activities relate to installing, testing, commissioning, maintaining and operating systems, services and equipment. There are option routes for candidates working in different areas of the business; for example, switching, transmission, radio, power, and cable. This particular award is at Level 3 in the NCVQ National Framework. The title of it is *Providing a Telecommunications Service* and its full unit structure is shown in Table 2.

Candidates are required to achieve all the mandatory units and any four option units for a full award. It is expected that they will select four option units that reflect their own work role. For example, a candidate who installs and commissions switching equipment will probably require units: Mandatory+SVC3/SW/2, SVC3/SW/3, SVC3/SW/4 and SVC3/SW/5

A candidate who maintains radio equipment will probably need units: Mandatory+SVC3/RD/3, SVC3/RD/5, SVC3/RD/6 and SVC3/RD/7

It can be seen that the above are examples of 'vertical' groupings within the framework in Table 2. These are appropriate to occupational roles that relate to single skills in single equipment areas. 'Horizontal' groupings are also possible, and would apply to single

skills in several equipment areas. For example, a candidate who installs several types of equipment, may opt for units:

Mandatory+SVC3/SW/3, SVC3/TX/3, SVC3/RD/3 and SVC3/PW/3

Candidates who are multi-skilled, may achieve groupings that are not purely 'horizontal' or 'vertical'. There are in fact many combinations possible. This feature has been designed into the framework to allow the degree of flexibility required by the telecommunications industry.

Individual flexibility and development

Candidates are not limited to achieving only four option units. In fact this is where individuals can develop their occupational competence, by achieving more option units after their initial four. It is also possible that when four option units match the typical breadth of competence found in the industry, it may be decided to increase the number of option units required to achieve a full award. This is a benefit of vocational qualifications that are owned by industry. Industry controls the level and breadth of competence required by individuals to become recognised as being competent.

This qualification has also been designed to support the telecommunications industry's Modern Apprenticeship. Organisations and apprentices will be able to select

Table 2 Framework for the Qualification *Providing a Telecommunications Service*, at Level 3 (SVC3)

MANDATORY UNITS	OPTION UNITS				
	SWITCHING UNITS	TRANSMISSION UNITS	RADIO UNITS	LINKS (CABLES) UNITS	POWER UNITS
Provide and maintain customer care <i>MAN3/1</i>			Specify telecommunications radio systems and equipment which meets the needs of of the customer <i>SVC3/RD/1</i>		
Protect people, property and the environment <i>MAN3/2</i>	Install telecommunications switching systems and equipment <i>SVC3/SW/2</i>	Install telecommunications transmission systems and equipment <i>SVC3/TX/2</i>	Install telecommunications radio systems and equipment <i>SVC3/RD/2</i>	Direct the installation of telecommunications links and equipment <i>SVC3/LK/2</i>	Install telecommunications power systems and equipment <i>SVC3/PW/2</i>
Develop competence of self in the work role <i>MAN3/3</i>	Confirm operation of telecommunications switching systems and equipment meets performance specification <i>SVC3/SW/3</i>	Confirm operation of telecommunications transmission systems and equipment meets performance specification <i>SVC/TX/3/3</i>	Confirm operation of telecommunications radio systems and equipment meets performance specification <i>SVC3/RD/3</i>	Confirm operation of telecommunications links meets performance specification <i>SVC3/LK/3</i>	Commission telecommunications power systems and equipment <i>SVC3/PW/3</i>
Develop and maintain effective working relationships <i>MAN3/4</i>	Handover and demonstrate telecommunications switching systems and equipment to the customer <i>SVC3/SW/4</i>	Handover and demonstrate telecommunications transmission systems and equipment to the customer <i>SVC3/TX/4</i>	Handover and demonstrate telecommunications radio systems and equipment to the customer <i>SVC3/RD/4</i>	Handover and demonstrate telecommunications links installation to the customer <i>SVC3/LK/4</i>	Handover and demonstrate telecommunications power installation to the customer <i>SVC/PW3/4</i>
Maintain effective communications in the work place <i>MAN3/5</i>	Identify variances in performance of telecommunications switching systems and equipment <i>SVC3/SW/5</i>	Identify variances in performance of telecommunications transmission systems and equipment <i>SVC3/TX/5</i>	Identify variances in performance of telecommunications radio systems and equipment <i>SVC3/RD/5</i>	Identify variances in performance of telecommunications links and equipment <i>SVC3/LK/5</i>	Identify variances in performance of telecommunications power systems and equipment <i>SVC/PW3/5</i>
Contribute to the efficiency and effectiveness of the organisation <i>MAN3/6</i>	Identify causes of variances in performance of telecommunications switching systems and equipment <i>SVC3/SW/6</i>	Identify causes of variances in performance of telecommunications transmission systems and equipment <i>SVC3/TX/6</i>	Determine the action required to restore the specified performance of telecommunications radio systems and equipment <i>SVC3/RD/6</i>	Identify causes of variances in performance of telecommunications links and equipment <i>SVC3/LK/6</i>	Identify causes of variances in performance of telecommunications power systems and equipment <i>SVC3/PW/6</i>
Maintain the quality of service provision <i>MAN3/7</i>	Restore the specified performance of telecommunications switching systems and equipment <i>SVC3/SW/7</i>	Implement procedures to restore the specified performance of telecommunications transmission systems and equipment <i>SVC3/TX/7</i>	Restore the specified performance of telecommunications radio systems and equipment <i>SVC3/RD/7</i>	Restore the specified performance of telecommunications links and equipment <i>SVC3/LK/7</i>	Implement procedures to restore the specified performance of telecommunications power systems and equipment <i>SVC3/PW/7</i>
Prepare for telecommunications activities <i>MAN3/8</i>	Operate telecommunications switching systems and equipment <i>SVC3/SW/8</i>	Operate telecommunications transmission systems and equipment <i>SVC3/SW/8</i>	Operate telecommunications radio systems and equipment <i>SVC3/RD/8</i>		

TVSC Unit codes shown in italics

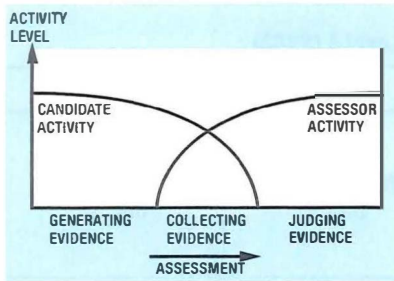


Figure 1—Candidate and assessor activities during assessment

those option units that match individual apprentice's training, development and experience. The progress of individuals can be monitored by their rate of achievement of the appropriate units. The nature of the apprenticeships may also be reflected, or indeed controlled, by the selection of the option units; that is, a generalist or a specialist apprenticeship.

Guidance material is being prepared to support this and all future qualifications developed by the TVSC. In addition to the standard for each element of competence, evidence requirements have been written which give details of what evidence is needed and how it should be collected. Approved centres will be provided with detailed information of the requirements of each element of competence. This information identifies typical sources of evidence that may be found in normal work activities and situations. Candidates will also be provided with guidance material.

A development from past practices will be that candidates will retain their own record of evidence and record of judgements. Guidance will be provided to candidates as to how evidence should be collected and recorded. Therefore, in addition to the NVQ certificate, candidates will have a personal record of the evidence they generated, collected and had judged to confirm their competence. It is believed this will enhance the value of the NVQ award. To support this, assessors are required to ensure sufficient detail of evidence is recorded.

Assessment of NVQ Candidates

What is assessment?

The definition of NVQ assessment that has used since NVQs began is:

'Generating, collecting and judging evidence.'

These three activities are shared between the candidate and the assessor as shown in Figure 1 which demonstrates a model where:

- the candidate generates evidence through normal work activities,
- the candidate and the assessor collect and record this evidence, and
- the assessor judges the evidence against the national standard.

The evidence considered is proof of the candidate's ability to meet the national standard and consists of performance evidence and knowledge evidence. This evidence should confirm that the candidate can carry out a specified activity and has the required knowledge and understanding that supports the activity. The standard against which the evidence is judged is the industry's specification for competent performance in a particular work role.

Collecting evidence

Evidence is collected using a number of methods. These are described below, with their respective uses:

Observation of candidate performance by the assessor

The assessor observes directly the candidate carrying out work activities. This provides performance evidence with the emphasis on how the candidate is carrying out the activity.

Inspection by the assessor of the outcome of candidate performance

The assessor inspects work done by the candidate after it has been carried out. This provides performance evidence with the emphasis on how well the activity was carried out.

While these two methods are used primarily to collect performance

evidence, some knowledge evidence will also be generated during performance. This may be explicit evidence—the candidate knows what to do; or it may be implicit evidence—does the candidate know why to do something? Any knowledge evidence which is implicit must be made explicit if it is required by the knowledge specification. This may be done by using guided discussion as described below:

Guided discussion with the candidate

The assessor asks the candidate questions in order to elicit and confirm the appropriate knowledge and understanding. This method is used to collect evidence across any part of the range statement for which performance evidence has not been collected. It can also be used to collect evidence of the candidate's competency for contingency situations. The guided discussion may take place before, during or after observation by the assessor of candidate performance. However, the discussion should be relevant to the activity specified and be based on the knowledge specification associated with the standard.

Guided discussion with others

The assessor may ask questions of other people who are able to speak authoritatively on the candidate's performance and knowledge. This source of evidence is used to enhance and supplement the evidence provided directly by the candidate. It is also used to collect evidence of the candidate's interaction with others; for example, customers, colleagues, supervisors.

These are the four main methods used to collect evidence of the candidate's competency. Several secondary methods that may also be used.

Simulation

These are situations arranged specifically for assessment. They are used to simulate workplace conditions

where, for example, safety, legislative, financial or contingency issues prevent assessment in the real work place or work role.

Case history

The candidate records evidence of work activities. These records must be validated by an appropriate person before being accepted as evidence. Case histories have limited use as they can only be a record of what was done, not how well it was done.

Accreditation of prior learning and achievement

This method of collecting evidence allows a candidate to offer any form of evidence of previously demonstrated competence. Although it may take many forms, it should always provide evidence of current competence as specified by the qualification standard.

Witness testimonies

These are statements, either verbal or written, provided by individuals who have experience of the candidate's work activities. The value of these testimonies will be determined by the assessor. This will depend on the status of the witness in terms of their occupational competence and understanding of the standard required.

It is unlikely and undesirable for the assessment of a candidate to use evidence collected through just one method. The TVSC believes the best assessments are those using a balanced mix of methods. Whatever methods are used in practice, the TVSC requires some evidence for each element of competence to be collected through direct observation of candidate performance. This requirement is set by NCVQ criteria for assessment.

Where to assess?

As NVQs are statements of occupational competence, based on industry's standards for competence in work, the best place for assessment to take place is the work place. However, it is recognised that this is not always possible owing to the nature of work

roles in the telecommunications industry. The TVSC therefore requires assessment to take place under workplace conditions. Assessment situations for NVQs should put the candidate in no more and no less than the conditions found in typical work environments. This means, for example, if typical work environments for the candidate have customers present, or require team working, then these conditions should apply in the assessment situations. For assessment to be acceptable to the TVSC, candidates should be exposed to the influences on their performance that are significant in their real work place.

What to assess?

Assessment of competence must not be for anything less than a complete element of competence. This means that it is not acceptable to make assessment decisions on some of the performance criteria, or some of the knowledge specification, or part of the range statement. Evidence, either of performance or knowledge, must be collected for all the performance and knowledge criteria, and the complete range before any judgement is made. Some methods of collecting and recording evidence encourage sub-element assessment. However, assessors should not follow this practice.

When to assess?

One of the original philosophies of NVQs was that they should set the standards for occupational competence in employment. Candidates should be assessed against these standards after providing evidence which is generated through normal work activities in the work place. When the standard set by industry has been met, then a certificate is awarded in recognition of an individual having demonstrated sustainable evidence of competent performance.

It would seem reasonable then that candidates should be assessed at a point in their development when they have the potential for sustainable or repeatable competent per-

formance. It would also seem reasonable that this point is reached after suitable vocational training, development and workplace experience.

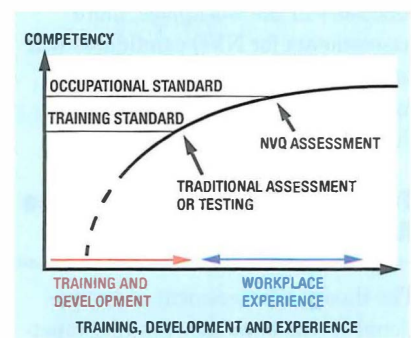
This is one of the most significant differences between NVQs and traditional course-based, examination-based and earlier vocational qualifications. Candidates for these qualifications were usually rewarded when they reached a prescribed standard of vocational education or training. Recognition was therefore given to reaching a training standard. Very few of these qualifications required any workplace experience. These qualifications therefore gave recognition to competence for employment, whereas NVQs give recognition to competence in employment—in the work role for which the qualification was developed—for reaching an occupational or operational standard.

Figure 2 indicates the point in an individual's occupational development at which NVQ assessment should take place.

Approach to assessment

Candidates and assessors should consider how to collect evidence which exists naturally. Early experiences of assessment usually involved looking from the qualification towards the generation of appropriate evidence. Candidates may have been required to generate evidence specifically for the qualification. Assessment will be more cost effective if collection is from workplace activities towards the requirements of

Figure 2—When to assess a candidate



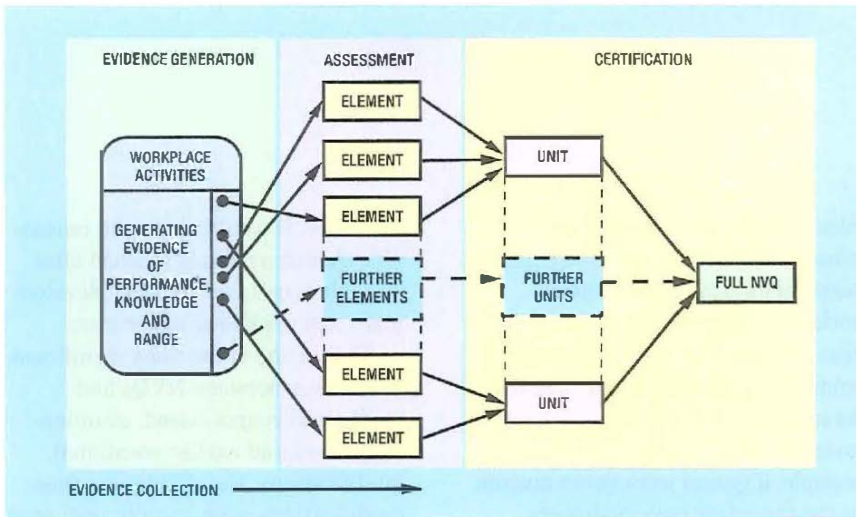


Figure 3—Approach to assessment

the qualification. Figure 3 represents this approach.

Assessment practices

Currently, assessment of NVQ candidates uses the full range of evidence-collection methods and takes place between a candidate reaching the training standard and reaching the occupational standard (Figure 2). The TVSC encourages the continued use of the full range of collection methods, while expecting more assessments to take place around the point when a candidate reaches the occupational standard.

There are at least two reasons why centres might assess candidates on reaching the training standard. Firstly, it has nearly always been the case that assessment for vocational qualifications has taken place at the end of a period of training. Certificates are awarded on completion of a course, or on passing an examination. Secondly, at present, some organisations have their accredited assessors in the training environment. The reason for this could be that NVQs are still seen as training tools, rather than assessment tools.

As familiarisation of the purpose of NVQs increases and with more assessors in the workplace, more assessments for NVQ candidates will take place in the workplace, after appropriate vocational training, development and experience.

Tensions Around Workplace Assessment

The theory of assessment of occupational competence has several compet-

ing requirements. The practice of assessment has to balance these requirements. Each organisation or approved assessment centre will find this balance. Figure 4 shows some of these competing requirements.

While attempting to provide maximum access, lowest cost, or highest credibility, each centre will have to consider any detrimental affect on the other two characteristics of assessment. The industry will be able to achieve the best balance by promoting and spreading good practice. External verifiers will play an important part in achieving this.

Telecommunications NVQs Within BT

BT has been involved with the TVSC since the Council became the lead body for the telecommunications industry in 1989. The current chairman of the TVSC is David Thomas from BT Management and Development Programmes, and the chairman of the TVSC Steering Group is Geoff Salter, from BT Personnel Policy. Representatives from BT have provided invaluable information and contribution to support the development of the occupational standards and NVQs. Also, BT has been involved in a number of pilot schemes for implementing NVQs. These pilots have provided both BT and the TVSC with useful information about the issues to be addressed when implementing NVQs in a large organisation.

The new generation of NVQs from the TVSC has, in part, been developed to satisfy BT's requirement for

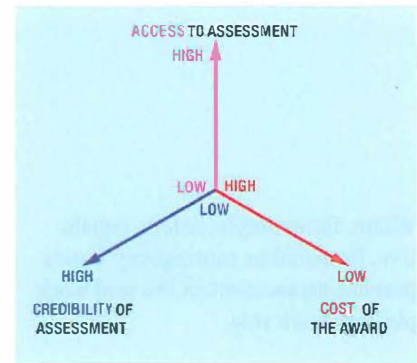


Figure 4—Potential tensions around workplace assessment

greater flexibility in a multi-skilled workforce. It is expected that the mandatory and option structure of awards will meet the needs of both BT as a business and individuals within that business. The TVSC is working closely with BT to establish sufficient approved centres through which BT people can be assessed as NVQ candidates. Establishing these centres and the network of assessors and internal verifiers requires significant coordination and training.

An interesting facet of NVQs is that candidates can be generating and recording evidence of their competency before being enrolled for an NVQ. When assessment is carried out, the assessor will consider all evidence offered by the candidate. The final judgement decision is whether the evidence provided is sufficient, valid, authentic and reliable to confirm the current competence of the candidate. Therefore, until BT has all the infrastructures in place to implement NVQs fully across the company, individual potential candidates can be recording what they do and collect reliable evidence on how well they do it. It must be emphasised that recording by candidates of their own evidence does not remove the requirement for assessors to observe directly some performance by the candidate. However, the more evidence available before formal assessment takes place, the more cost effective the assessment process may be.

Assessors are looking for quality of evidence, not quantity. Candidates should be selective about what is recorded and consider the requirements of the qualification standards when recording their evidence. The standards for all current TVSC NVQs

are available for candidates to identify what is required for the units that make up the awards.

Conclusion


This article provides background information on telecommunications NVQs available to telecommunications workers. The mandatory and option structure of the qualifications should provide greater flexibility for individual employees to select units that match their work role. The article also describes current thinking at TVSC on assessment practices. It can be seen that partnerships are encouraged between the candidate and the assessor for collecting and recording evidence of competency. This partnership should begin at the assessment planning stage when the candidate and the assessor agree what evidence is available, as well as what and how the remaining evidence is to be collected. There should be mix of collection methods used, with special emphasis on direct observation of the candidate.

Ideally, assessment should be carried out in the workplace, but assessment under workplace conditions is also acceptable. The timing of assessment is preferred to be after the candidate has had real workplace experience. Evidence should be collected during normal workplace activities. This may not always be possible and simulation may have to be used in certain situations. Assessors should be certain they have sufficient valid and reliable evidence for a complete element of competence before making any assessment decision.

NVQs and the occupational standards on which they are based have benefits for both individuals and organisations. Being employment based and assessed in the workplace, the individual generates evidence through natural work activities and is assessed in the normal working environment. Recognition is given to individuals who meet the national standard for competent performance,

without the need for possibly stressful examinations. For organisations, having employees with NVQs provides confidence that these employees have demonstrated competence to a national standard. But there is more to NVQs than the awarding of the certificate. NVQ candidates are required to record evidence of their competency. This requires individuals to think about what they actually do in their work roles and how well they are performing. This, in turn, helps employees to become more professional in their attitude and their performance. An NVQ workforce therefore develops into a professional workforce.

Biography



Alan Martin
Head of Standards
and Awards
Telecommunications
Vocational Standards
Council

Alan Martin has the responsibility for the development of occupational standards and vocational qualifications for the telecommunications industry. He began his telecommunications career in London with the Post Office, working on exchange installation. A move into the Research Department at Dollis Hill and then Martlesham Heath required a move from Strowger systems to electronic-based equipment. He spent three years in South Africa working with the South African Post Office Research Laboratories, and five years in Zimbabwe working at the Zimbabwe Posts and Telecommunications Corporation Training College and Philips Electrical. His interest and experience in vocational education and training began in Zimbabwe and has included positions as training manager and further education college lecturer. He has been with the TVSC since the Council was established as the lead body for telecommunications. The TVSC may be contacted at:
Blackfriars House
399 South Row
Central Milton Keynes
MK14 5DJ
Telephone: 01908 240120
E-mail: TVSC@wildnet.co.uk

In his column for July, Peter Cochrane, Head of Advanced Applications and Technologies, at BT Laboratories, Martlesham Heath, considers the evolution of the computer and how much of our lives we should let it run.

The Third Lobe


Following the Cambrian explosion of 500 million years ago, humankind emerged only during the last 2 million years or so. All the evidence suggests that our fundamental mental abilities have changed little over the last 10 000 years. The size of our cranium and its contents, our intelligence and our creativity have remained unchanged and are likely to stay that way. In contrast, the evolution of electronics, telecommunications and computers has seen a far faster process since the invention of the thermionic valve in 1915. Computer capabilities have roughly doubled every year since 1960 and look set to continue for a decade, probably two, and possibly three or more.

The different evolutionary rates of technology and biology are a function of entirely different processes. On the one hand, accidental mutation due to environmental and competitive threats invokes a slow and gradual change in biological systems where 'fit for purpose' is the order of the day. Once you are dominant, have reached an equilibrium state with the ecology, and you are the king of the jungle, why try and improve further? Also, biological systems can evolve into a cul-de-sac—once you are a human, you are extremely unlikely to evolve into a fish! In complete contrast, technological evolution is more direct, more focused and realised by successive generations standing on each other's shoulders. We improve, optimise, do more with less, cross-fertilise from different technologies, fight for market share in an environment of intense competition. Here chance mutation as an evolutionary process is far too slow, unguided, and uncertain. In technological evolution we target and design to solve a problem, meet a need, and create change in sustainable directions.

The exponential change in technology now sees rates of doubling anywhere between 12–24 months. We might confidently expect, therefore, that the computers on our desks will be 1000 times more powerful in 10 years, perhaps even 1 000 000 times in 20 years, and it is now just thinkable, that they may achieve 1 000 000 000 times the present-day capability in 30 years. All of this would mean a computing and storage capability roughly on a par with humankind by the year 2015 in the form of a leading edge super-computer that would then arrive on our desks some 10–15 years later. Some 10 years on we might be wearing this computer, or perhaps, it will be wearing us!

In a period of 100 years of electronics, from 1915 to 2015, we might see the 50 million year evolution of the vertebrates, including us, overtaken. The question arises: what are we going to do with all this computing power and intelligence? Well, given our most definitely limited brain power and analogue mode of processing and thinking, it could be that this digital capability will provide us with the natural extension to our biologically limited brain. After all, it could provide us with the next set of peripherals for seeing, hearing, touching and thinking, and we might even conceive it to be a new third lobe! Of course, there are those who favour tampering directly with our biological brain, through the use of drugs and/or genetic engineering. Some have even debated that the gene pool took action (2 million years ago) to create us for this purpose having recognised, consciously or otherwise, that a continuation of our species through random error and mutation was far too slow! Be this true or not, I favour the former argument over the latter. Modifying me is not an option!

Such a technological extension to our abilities might be the only way we are likely to keep pace with the technology itself and prevent our civilisation stalling. Future evolution and progress is unlikely to be limited by technology; more likely the limiting factor will be us and our inability to adapt and keep pace. Using the



technology to help us analyse, visualise and understand faster may be our only option. It will also allow us to escape many of the less rewarding activities and provide more time to interact and be creative. In many respects we are already there! Most raw material and production processes have been automated and denuded of human activity. From steel making, farming and bread making, car production and robots, to medical diagnosis and hip replacement surgery we see the human content progressively being reduced. The baker and car painter are now artificial intelligence systems—just software!

Much of our scientific and engineering effort with computers, systems and robotics is devoted to making machines that think and act like us. In many respects this seems to be a flawed approach—we have us! Why not create something that thinks and acts as a machine? Would it not enrich our environment to have another form of being that did and thought new things? Would it not be as enriching as another ethnic group, race or language? If it is true understanding and betterment that we seek, it could be that our present brain power, that has evolved to see us superior in the animal world of the hunter and the hunted, is too limiting. Biological evolution is far too slow a process for us to progress rapidly, but perhaps electronic enhancement offers us a new direction. For example, our abilities are enhanced significantly in the virtual environments of the computer. We have a demonstrated ability to consume and understand vast amounts of data when it is in a

MESSAGE FROM MARTLESHAM

pictographic and interactive form. It is always worth remembering that the written word is an unnatural artefact created by us for the storage and transportation of human thought. Our more natural abilities and interfaces lie in different directions!

Such arguments alarm some, but look back at the industrial revolution, the sacrifice of human life, working and living conditions, and think of today. How much better life is, how much healthier and happier people can be—if they choose to be so! Looking to the information age, have we made the transition well, but have we done all we can? Are we using the technology to our best advantage, or are we perpetuating the processes and established protocols of the past?

If Charles Dickens was to walk into a modern office, he would no doubt be amazed by the inventions and the technologies that would confront him. He would probably be amazed to see double entry book-

keeping making the migration from the ledger to the screen. In fact he might even be prompted to ask the question what are you doing? All you have done is taken the paper processes of my office and literally placed them on the screen. Surely this cannot be the right thing to do! Isn't a picture worth a thousand words? Perhaps a moving picture is worth volumes? Then, why perpetuate the old ways of recording and transmitting information with the written word? Is it not possible that you could find new ways of working, new ways of communicating and new ways of presenting information?


Having raised these questions, you would have to look Charles Dickens in the eye and admit that what has been done has perhaps been a little pedestrian and unimaginative. Perhaps if we had considered our innate visual ability for correlation, that primeval capability to spot an animal on the horizon, stalk and kill

it, or our ability to recognise the male or female friend or foe at a glance, it could have been used and exercised to greater advantage. Probably the single biggest improvement that could be made would involve the personal computer becoming truly intelligent and able to make the most basic of decisions across a broad front of our daily activities. People already run computer-based scheduling and diary systems that they implicitly trust, the whole financial system is run by a computer and even the stock market has an increasing element of artificial intelligence. The question is would we trust this third lobe to run a significant part of our lives? The truth is we will probably have no choice for as the technology continues its exponential growth the only way that we as humans will be able to keep up will be to invoke the technology itself in our aid. In time, this silicon artefact has to become part of us—we will have to trust it!

book reviews

High Speed Networking Technology: An Introductory Survey (third edition)

By Harry J. R. Dutton and Peter Lenhard



This is one of the *Redbook* series of publications, produced originally by the IBM technical support organisation for its engineers in the field to help them keep abreast of current and future technologies and applications. The approach taken is both informative and practical and the information

will benefit the novice and the reader with some knowledge of the subject.

Potential readers should not be discouraged by the title of the book, it is neither a survey nor is it aimed solely at those new to the area, although their needs are well catered for. They should not be put off by the apparent narrowness of the original target audience, the potential scope of the book is vast and it will prove to be of value to those seeking an introduction to the technologies available, to those who wish to improve their understanding of those technologies and to students wishing to reference specific high-speed networking issues.

Due to the breadth and magnitude of the subject, the book can devote only a limited amount of space to each part of the subject area. This is both its major strength and its weakness. The strength is that the book is written in an extremely readable form which most people will find informative and entertaining but the weakness in this approach is that, of necessity, each subject can be dealt with only to the level of detail

imposed by the constraints of one chapter. The IBM-centric nature of some of the product and internal standards used does not detract from the authors' efforts to produce an exceptionally readable book.

Each chapter stands alone as an inherently useful reference source while contributing to the authors' philosophy of producing a book which rarely overwhelms the reader with information while managing to drip feed a surprising amount on each pass. The book is first and foremost readable.

Chapters generally begin by placing the technology in context before outlining developments through to the present day. There is sufficient detail to interest those with some pre-knowledge of each subject area although the needs of the novice are never neglected. Readers with very considerable knowledge of any, or all, areas covered may find the depth of analysis restricted, although the target audience will be well served.

There are few books covering such a wide area of technology so well

presented or so approachable. I would recommend it to those new to the subject and to those who wish to improve or update their knowledge in the high-speed networking area. As a reference book, it is one that will prove invaluable to anyone in the fields of data communications or voice networking.

Published by Prentice Hall


xxxii+ 536pp. £26.95

ISBN 0-13-242421-5

Reviewed by Tom McGrane

Digital Telephony and Network Integration

By Keiser and Strange



Modern students are amazed that analogue telephony could have been made to work over intercontinental distances, showing that we have left the analogue world behind and nowadays the norm is quite clearly digital. Thus the inclusion of 'digital' in the title of any modern telecommunications book might be taken to imply that it is from an earlier age when such things needed saying. Actually 'digital telephony' in this title may well relate to the opening extensive coverage of analogue-to-digital voice-band coding techniques.

The author then goes on to cover the usual areas of multiplexing and transmission on various bearers (copper, fibre, microwave, satellite and mobile), switching, and network and service integration. 'Modern' subjects include asynchronous transfer mode and B-ISDN, SMDS, and network intelligence, but with such brevity that it is hard to endorse the back-cover claim that this book will be 'your

guide to digital telephone into the twenty-first century and beyond'. Overall the coverage is neither at the academic theoretical level, nor at the equipment description level, but lies somewhere between. Just because a technique is described does not mean that it is regularly in use. On the other hand some very specific descriptions of equipment structures are given; unfortunately for European readers, these examples are drawn entirely from North America, although mention is often made of European variants of standards. Here we have a useful general reference book for the North American environment but workers on this side of the Atlantic will need to supplement their understanding by other documents such as IBTE's own *Structured Information Programme*. The general technical coverage of basic principles is entirely adequate for the present day.

In summary, a useful reference book in a collection but not **the** one book to keep by your side.

Published by Chapman & Hall


xvii+ 669pp. £49.00

ISBN 0-412-09881-4

Reviewed by John Griffiths

Digital Broadcasting

By Paul Dambacher



This book reflects the digital age and its impact on traditional analogue broadcasting techniques. Coverage is broad indeed, from satellite receivers to video coding with numerous topics in-between. It probes into specific areas, but not from a mathematical viewpoint. Basic electronic knowledge is assumed, from which many of the systems are described.

There are two uses this book is suited to: background reading and ready-reference in the engineering business. The subject area is broken into a three-layer hierarchy. Major sections are conveniently divided into sub-sections grouping together related technologies. Dambacher has been effective in this process providing information quickly and easily.

The book starts with an introduction which describes the concept of digital as opposed to analogue broadcasting, with its advantages and new issues which are raised. This is followed by sound and TV broadcasting media covering terrestrial, satellite, cable and fibre transmission systems. Coverage of technological advances in audio broadcasting includes the compact disc, digital audio tape (DAT) and other digital recorders for the home, radio data system (RDS), digital satellite radio, terrestrial digital audio broadcasting (DAB) and audio baseband coding.

A major section covers advances in television. This starts with the progression of TV standards over time and these standards are then described and compared. A brief diversion is made into psycho-optics, leading into the ideas behind video coding for data reduction. This section is brought to an end with digital TV in Europe.

The book finally looks to the future with ideas that may become standard soon. This includes digital video broadcasting and narrowband DAB. The viewpoint is 'How it is done now has disadvantages—here's a solution'.

Numerous references are included throughout the book to allow the reader to delve further into areas of particular interest.

The book covers a wide variety of digital broadcasting topics and would be suited to engineers who are new to this field, or as a starting point for investigation of a particular topic. With such a large number of topics covered, the depth is limited but adequate for the purposes described here.

Published by

The Institution of Electrical Engineers

vii+336pp. £50.00

ISBN 0-85296-873-6

Reviewed by Chris Armson

'999' Video For Youngsters

BT has launched a new video aimed at ensuring youngsters know how to use the 999 service—and why they should not abuse it. BT commissioned the BBC television programme '999' to produce the 10 minute video presented by Michael Buerk. In it he explains that the service is a life-saver and that BT 999 operators handle around 22 million calls a year.

But, he says, nearly half these calls are false alarms—each one potentially distracting vital resources from the real emergencies. And most false calls, he points out, are made by children.

The video gives simple, basic advice on how to make these free calls so that children know what to do in an emergency. It also points out that the number and address of anyone making a 999 call is automatically displayed on the operator's screen. This makes getting help to an emergency even quicker; it also means hoaxers can be identified easily and their names passed to the police for possible action.

The video closes with the key words: 'Don't let that be you who is caught out—making a 999 call at the proper time in the proper way is something to be proud of—and you could be saving someone's life.'

There is also a Freefone number—0800 320500—for all enquiries about the video or any other BT education material and services.

BT is on the Ball for EURO 96 Tournament

For the Euro 96 soccer competition, BT assembled one of the most complex telecommunications solutions.

As official telecommunications supplier for the tournament, the company had over 400 staff working full time on the project to make sure the games were relayed, without a hitch, to an estimated 220 million fans worldwide. The project ranged from installing special data networks to providing media centres at the eight venues for the estimated 13 000 journalists.

Voice telephony services were available for all outgoing calls and

correspondents and media groups were issued with a special edition BT Chargecard. Calls were billed to a nominated credit or debit card. Cards were issued free and charges for calls made at BT's Chargecard rates. This solution allowed access to the BT service from any free desk in the Media Centre, plus cashless calling from virtually any telephone in the UK and overseas. BT Chargecards could also be used for fax transmissions.

Euro 96 football matches were broadcast to 150 countries, with an average television audience per match anticipated at 220 million, and a total cumulative audience of 6.9 billion.

To handle such broadcast capacity BT's Broadcast Services, Europe's leading supplier of global broadcast solutions, was contracted by FORTO (the joint BBC/ITV Sport host broadcaster for Euro 96) to provide a comprehensive range of fibre- and satellite-based transmission facilities, for coverage of the third largest sporting event in the world.

BT also supplied facilities for overseas broadcasters including Italy, France, Germany, Australia, BBC, ITV Sport and Reuters.

In total, BT Broadcast Services provided 50 visual circuits, each with a capacity of 140 Mbit/s bandwidth.

BT and MCI Introduce World's First Global Internet From Concert

BT and MCI have announced the world's first high-speed high-reliability global Internet backbone network.

Offered by Concert, the BT-MCI joint venture for global services, Concert InternetPlus will combine the substantial existing Internet networks of the two companies, initially into eight new regional Internet superhubs.

These will expand within a year to a total of 20 hubs in key locations around the world. The initial Concert regional superhubs will be located in Australia, Germany, Japan, The Netherlands, the UK and US.

In addition to MCI's extensive Internet presence in the US, the Concert InternetPlus network will be

offered this summer through BT's existing and planned European hubs including Belgium, France, Italy, Spain, Sweden and Switzerland. MCI and BT plan additional Internet hubs throughout Asia and the Americas.

The Concert InternetPlus services will provide Internet service providers, international carriers and businesses with an array of Internet and intranet transport, dedicated and dial-up access and value-added services available from more than 1200 locations in 70 countries. It will include the first-ever global Internet service performance guarantees, improved response times and greater availability.

The higher performance and reliability of the Concert InternetPlus services will allow businesses to use the Internet as their private global network for applications such as groupware, e-mail, messaging and electronic commerce.

Sir Peter Bonfield, Chief Executive of BT, said: 'Just as the BT-MCI alliance was the first to offer multinationals seamless global telecommunication services, today we are launching the first class of global Internet services.'

'Already we have achieved significant inroads into the European Internet market. This initiative will result in BT taking the lead in Europe in the same way MCI has taken the lead in the US Internet market.'

The new Concert InternetPlus network is expected to increase by 30 per cent the overall international capacity of today's Internet. BT and MCI together already offer direct Internet services to around 70 countries. MCI is the first US Internet provider to operate its network at 155 Mbit/s, providing more capacity than any other network of its kind. For the first time, this same level of performance and capacity will be available globally.

Gerald H. Taylor, president and chief operating officer of MCI, said: 'Concert InternetPlus now makes the Internet industrial-strength and a reliable credible substitute for private global networks. We expect that the global reach and performance of the

Concert InternetPlus network will attract many new distributors such as other leading phone companies and Internet service providers around the world.'

With the introduction of Concert InternetPlus, Concert has added NTT Data of Japan as a new distributor for the important Japanese market. NTT Data is Japan's leading information systems and computer networking company. With annual revenue of \$5 billion, NTT Data, creates value for its clients with *Professional Services*, its concept for the provision of total information services including strategic planning, systems planning, systems design and installation, systems maintenance and facility management. NTT Data also provides Internet services under its Intervia brand.

BT Internet Developments

BT has been heavily involved in the Internet, introducing a number of on-line services providing access and content.

BT Internet is a 'mass market' dial-up service aimed at residential and small business customers as well as new users on the Internet. The service emphasises ease of use, competitive subscription costs with no hidden charges, comprehensive customer support, and a jargon-busting approach suitable for the non-technical user.

BTnet, BT's Internet access service aimed at business customers and resellers, launched in December 1994, provides connectivity via a high-speed transatlantic link to internetMCI and to Europe through membership of EBONE (European Backbone). Since November 1995, Internet service providers, resellers and corporate users have been able to access BTnet and internetMCI via Concert Frame Relay Service. This service is now being used by around 50 Internet service providers and is available in Belgium, France, Germany, Netherlands, Spain, Switzerland and South Africa.

WebWorld, a BT Web hosting service, is currently on trial and due for launch this year. It is designed to

help customers establish a presence on the Internet and market their products and services.

In July, 1995, BT won a large international contract to provide DANTE a pan-European Internet transit network for the European research and academic community. DANTE's international network service to its customers is known as *EuropaNET* and is operated and managed by BT.

With this contract, BT has captured 20 per cent of the European Internet traffic market.

BT and Microsoft Announce New Small Business Initiative

BT and Microsoft have announced that they are jointly to market a range of computer networking and on-line services to small- and medium-sized businesses using BT's direct sales channels and a number of Microsoft accredited resellers, otherwise known as *Solution Providers*.

Under the agreement, BT has recruited Microsoft Solution Providers to promote the benefits of using local area networks (LANs) to link together PCs installed with leading Microsoft software such as Windows 95 and Windows Server NT.

BT and Microsoft believe that by combining their respective networking and software expertise with the support and integration skills of solution providers, smaller businesses will benefit from one point of contact for all their communications needs.

Companies with 10-50 personal computers (PCs) will be targeted initially and will either be looking to implement a networked computer system for the first time or considering upgrading their existing network.

Kathleen Flaherty, Marketing Director of BT Business Communications, said: 'Computer ownership among UK small businesses is high but not many customers have networks to link them together, allowing them to share resources and access the Internet and other on-line services.'

'BT and Microsoft will provide expert advice and assistance for customers who would like to know how best to integrate their computers and telecommunications services.'

'With the right computer network, a PC becomes the doorway to the office. A user can gain access whenever and from wherever they want and find the right information.'

Neil Holloway, Director of Organisational Customer Unit at Microsoft Limited, said: 'The small business market is a huge priority for Microsoft this year and we have now begun the process of looking at ways we can work to make our customers' lives easier. In general terms, small businesses do not employ specialised IT or technical staff and are ill-equipped to cope with the deployment of technology.'

'In order for a small business to take advantage of opportunities such as e-mail and the Internet they currently have to have a number of different relationships with commercial vendors and manage a complicated marriage between computers and communications.'

'The agreement with BT will make the purchase cycle simpler and generally provide a greater level of support through our Solution Providers at a time when it is critical. Best of all, this partnership will enable any business, however small, to use the latest technology to improve their communications and ultimately make them more competitive in their chosen market.'

The key to the agreement is the promotion of LAN access services to meet a number of business functions such as accounting, credit checking, access to shared applications and remote working, or teleworking. These packages include LAN-to-LAN connections and remote PC access to company LANs.

Also, customers will be given advice on how take advantage of the latest data communications products such as high-speed ISDN access to the Internet and other on-line services.

BT will continue to install and maintain separately those products and services covered by this agreement.

Africa ONE Will Transform Continent

A fibre-optic cable around Africa called *Africa ONE* will provide vastly improved and greatly expanded telecommunications capabilities—a critical element in Africa's future economic and business growth as well as to Africa's ongoing political and social progress. The 39 000 km cable will use state-of-the-art technologies, including optical amplifiers in terminal equipment and submarine repeaters and wavelength-division multiplexing capable of handling up to eight independent optical channels on a single fibre. Construction, carried out jointly by AT&T Submarine Systems, Inc. and Alcatel Submarine Networks, co-developers of the project, is set to begin in 1997 with full service expected in 1999.

First Web-Based Service from AT&T Business Network

AT&T has launched Lead Story, the first of a series of World Wide Web sites from AT&T Business Network designed to serve the needs of business managers, professionals and entrepreneurs. Lead Story provides busy professionals with an in-depth view of one major business or general-interest story per day, drawing from all the journalistic and background sources on the Web. Lead Story can be accessed through the URL: <http://www.leadstory.com>

Lead Story features links to news from a wide variety of Web sources, plus analysis, opinion, background materials and other information that help give professionals the broadest possible perspective on a major story. By organising and presenting this content in one place, Lead Story lets professionals spend valuable time digesting and acting on important information, rather than seeking it out. Examples of Lead Story topics include the federal budget show-down, Internet censorship and China's controversial military exercises.

By selecting links to only the best and most original material, and

augmenting it with summaries, information graphics and other helpful visuals, we are providing a complete, 'one-stop shopping' view of a major story,' said Michael E. Kolowich, president of AT&T New Media Services. 'With Lead Story, whether you give us as little as five minutes or as much as an hour, we'll give you the most comprehensive briefing available today on the Web. Any serious professional should bookmark Lead Story and check it every day—it's a big timesaver.' Lead Story topics are archived for several days after they appear. LeadStory is updated at noon Eastern time every business day, plus Saturday. Lead Story is a free site; registration is not required and there are no subscription fees charged.

Vodafone and Energis Networks Connect

Energis has announced its first interconnection agreement with a UK mobile telecommunications operator. The contract with Vodafone creates a direct link between Energis's 120 000 business lines covering many of the UK's top 1000 companies and over 2.45 million Vodafone customers. The agreement also provides cable companies and telecommunications service providers with a new alternative to using BT or Mercury interconnection points for accessing the Vodafone network. The first Energis/Vodafone switch is located in West London and plans are already in place to create a second switch in the North of England.

'We intend to provide our customers with direct interconnection from the Energis network to all the UK's major mobile communication services,' says Irene Cackett, Director of Marketing, at Energis. 'The Vodafone contract marks our first agreement in this strategy.'

Global Interoperative Broadband Network Project

France Telecom announced its participation in the first global

interoperative broadband network project with five of the seven G-7 countries—France, Germany, Italy, the US and Canada—beginning this year. The testbed will include implementing a high-speed research network to be used by research users of five countries. Among other applications, the international testbed will provide collaborative research and testing across continents, as well as sharing of feedback information between research users and operators on a global basis. The project will be the first global interoperative broadband network project to go online when it commences this year. Asynchronous transfer mode is a communications technology that promises to provide more cost-effective bandwidth and quality-of-service capabilities than currently available, a welcome development for the international research community, which is experiencing a growing shortfall in network reliability, predictability and available bandwidth. The ability of ATM to carry many data types over the same communications infrastructure implies reduced costs for the carrier, and therefore eventually to the end user.

The France Telecom testbed is currently in place and the pilot project is due to begin this year. Initial applications will include high-speed collaborative research between France, Germany, Italy, the US and Canada. The remaining G-7 countries—Japan and England—have expressed interest in joining the testbed to make it the first comprehensive G-7 high-speed testbed project.

Telecommunications Pricing From 1997

Don Cruickshank, Director General of Telecommunications, has announced proposals for the next control on BT's retail prices from August 1997 and the introduction into BT's licence, and the licences of other significant operators, of a fair trading condition prohibiting anti-competitive behaviour.

The proposals are contained in OFTEL's document: 'Pricing of Telecommunications Services from 1997'. Don Cruickshank, the Direc-

tor General, said: 'The main points of these proposals are:

- retail price control to 2001 only where consumer protection is required—low to medium spending residential customers (first 80% by spend) and small businesses. The control will cover only a quarter of BT's revenues;
- an X of 4.5% for those customers, better than the 2.7% they've actually had over the last 6 years.
- this will be the last retail price control;
- price controls on network charges (the input costs of competing operators)—the detail to be determined early in 1997; and
- vastly more pricing freedom for a still dominant BT—at both retail and network level—demanding more effective rules on anti-competitive behaviour to ensure BT uses this freedom fairly.

'On calls to mobile, I have announced that OFTEL will be investigating the levels of outpayments by BT to Vodafone and Cellnet for terminating calls on their networks. I consider these remain high, despite recent reductions and consequent reductions in BT's retail prices.

'The selection of a precise X for the purpose of the RPI-X formula for the price cap is, in the end, a judgement. It is necessary to assess where, within the range of values suggested by the financial modelling, the right balance is struck between price protection for vulnerable customers and the efficiency levels that are reasonable for the UK industry. It is also a duty of the Director General to take account of the need for BT to be able to finance the services which it is obliged to provide. A key consideration in arriving at the value of X has been the fact that the best protection for customers and the best spur to higher efficiency will be effective competition. In OFTEL's view an X of 4.5 per cent strikes the right balance. It gives residential customers a

substantially better deal than they have had before without making assumptions about market growth or BT's ability to improve its efficiency which would put the development of the industry at risk.'

Mobile Phone Contracts

OFTEL has supported the action being taken by the Office of Fair Trading (OFT) to improve the terms used in consumer contracts by the mobile phone industry. Don Cruickshank, the Director General of Telecommunications, said: 'OFTEL currently receives around 4000 complaints and queries a year from customers using mobile phones. The single biggest area of concern to these customers is the terms included in contracts they sign when buying mobile phone services. They are unhappy about a variety of contract terms; the length of time they are tied into the contract (often a minimum of a year or more); the lack of a 'cooling off' period once the contract is signed; the fees payable for disconnecting from a service, the small print, and some of it is very small, which may contain unpleasant surprises; and the absence of a full cash price alternative to the price of a subsidised handset.

'I fully support the action being taken by OFT to tackle these, and other poor contract practices, through their application of the Unfair Terms in Consumer Contracts Regulations 1994. I hope the mobile industry will respond positively to OFT to demonstrate its commitment to customers.

OFTEL is also urging the mobile sector to take measures to improve service quality and help customers.

'The key steps are:

- developing measures to allow customers to compare the performance of different mobile networks; and
- establishing codes of practice for mobile telecommunications service providers and mobile telecommunications dealers which:

—provide better, clearer information to customers at the time they purchase the service; and

—assure them of fair treatment if they have a complaint.

'The mobile telecommunications industry has grown rapidly in the past few years. By taking action now over customer concerns, it can help build customer confidence so its continual growth can be assured.'

MCI and the Internet

Building on its early association with the Internet, dating back to 1987, MCI is today one of the world's largest Internet-services providers. In late 1994, MCI entered the commercial Internet market with the launch of *internet.MCI*, a full suite of Internet and intranet services ranging from high-speed and nationwide dial-up access to the Internet, to software and turn-key services such as web hosting, intranet management, and security.

In March 1996, MCI announced MCI Internet 2000, a sweeping initiative that included: tripling the capacity of its internetMCI network; expanding its consumer Internet offerings; and providing businesses with more value-added Internet services.

Virgin Atlantic Airways Signs With Energis

Virgin Atlantic Airways, the airline of international entrepreneur Richard Branson, is signing a contract with service provider Energis. The three year contract, expected to be worth some £600 000 per year to the UK's newest national network service provider, will see the integration of Virgin Atlantic's central reservations system and internal and external communications within a virtual private network (VPN), provided and managed by Energis. The service, called *CustomNet*, will link Virgin Atlantic sites nationwide and its head office, with plans to bring more sites on-line as the company's communications needs evolve.

BRITISH TELECOMMUNICATIONS ENGINEERING

ISSN 0262-401X

Published in April, July, October and January by
The Institution of British Telecommunications
Engineers, Post Point G012, 2-12 Gresham Street,
London, EC2V 7AG.

Formerly *The Post Office Electrical Engineers'
Journal* Vols. 1-74: April 1908-January 1982.)

The Board of Editors is not responsible for any
statements made nor the opinions expressed in
any of the articles or correspondence in *British
Telecommunications Engineering* or the *Structured
Information Programme* unless any such state-
ment is made specifically by the Board.

© 1996: The Institution of British Telecommuni-
cations Engineers.

British Telecommunications Engineering and the
Structured Information Programme are printed in
Great Britain by Headley Brothers Ltd.

Subscriptions and Back Numbers

Annual subscriptions (including postage)—com-
panies, universities, libraries and other bodies
£40-00 UK, £45-00 overseas; private individuals
(paid by personal cheque) £23-00 UK, £28-00 over-
seas.

Single copies and back numbers—companies etc.
£10 including postage UK, £11-25 including post-
age overseas; private individuals (paid by per-
sonal cheque) £5-75 including postage UK, £7-00
including postage overseas.

Overseas customers should pay by a sterling
cheque drawn on a London bank, a sterling
Eurocheque or a sterling travellers cheque pay-
able in London.

Price to BT employees: £2-25 per copy.

Orders, by post only, should be addressed to *British
Telecommunications Engineering Journal* (Sales),
PP: G012, 2-12 Gresham Street, London EC2V 7AG.

Remittances for all items should be made payable
to 'BTE Journal' and should be crossed '& Co'.

Advertisements

All enquiries relating to advertisement space res-
ervations should be addressed to The Advertise-
ment Manager, *British Telecommunications
Engineering Journal*, PP: G012, 2-12 Gresham
Street, London EC2V 7AG. (Telephone: +44 171
356 8050. Fax: +44 171 356 7942.)

Communications

All communications should be addressed to the
Editorial Office, *British Telecommunications Engi-
neering Journal*, PP: G012, 2-12 Gresham Street,
London EC2V 7AG. (Telephone: +44 171 356 8050.
Fax: +44 171 356 7942.)

Copyright

The entire contents of *British Telecommunications
Engineering* and the *Structured Information Pro-
gramme* are covered by general copyright and
special permission is necessary for reprinting long
extracts, but editors are welcome to use not more
than one-third of any article, provided that credit
is given at the beginning or end, thus: 'From Brit-
ish Telecommunications Engineering'.

Authorisation to photocopy items for internal or
personal use, or the internal or personal use of
specific clients, is granted by The Institution of
British Telecommunications Engineers for users
registered with the Copyright Clearance Centre's
(CCC's) Transactional Reporting Service, provided
that the base fee of \$5-00 per copy is paid directly
to CCC, 27 Congress Street, Salem, MA 01970,
USA. For those organisations that have been
granted a photocopy license by CCC, a separate
system of payment has been arranged. Payment is
additionally required for copying of articles pub-
lished prior to 1978.

0262-401X/96 \$5-00 + .00



THE INSTITUTION OF BRITISH TELECOMMUNICATIONS ENGINEERS

IBTE Local Contacts

The following is a list of local contacts to whom enquiries about IBTE membership and services can be directed.

Zone / Centre	Contact	Telephone Number
London	Terry McCullough	(01707) 601151
Martlesham Heath	Peter Shaw	(01473) 224488
Midlands	Mark Broughton	(0116) 253 4476
North East	Geoff Jenkinson	(0113) 237 8163
North West	David White	(0161) 600 2468
Northern Home Counties	Terry Birdseye	(01702) 372557
Northern Ireland	Allan McDougall	(01232) 215301
Paris	Delphine Lalire	+33 1 46 67 27 89
Scotland	Raymond MacMillan	(01463) 665913
South West	Peter Shaw	(0117) 972 3216
Southern Home Counties	Nick Smith	(01227) 474823
Wales	Ferg Brown	(01952) 249886

IBTE International HelpLine — +44 171 356 8008

BRITISH TELECOMMUNICATIONS ENGINEERING

Journal Subscription Order (Not BT Employees)

PLEASE COMPLETE THE FORM IN BLOCK CAPITALS AND RETURN IT, WITH
THE APPROPRIATE SUBSCRIPTION PAYMENT, TO THE FOLLOWING ADDRESS:

The IBTE Administration Manager
Post Point G012
2-12 Gresham Street
London EC2V 7AG
United Kingdom

Annual subscription rates (four issues, including postage and packaging) are as follows:

Companies, universities, libraries and other bodies: **£40-00 UK; £45-00 Overseas**
Private individuals (to be paid for by personal cheque): **£23-00 UK; £28-00 Overseas**

Please supply four quarterly issues of *British Telecommunications Engineering*. I enclose
a cheque/postal order to cover the yearly subscription as given above.

(Cheques and postal orders, payable to 'BTE Journal', should be enclosed with the order.
Overseas customers should pay a sterling cheque drawn on a London bank, a sterling
Eurocheque or a sterling travellers cheque payable in London. Alternatively, customers
may wish to use a local subscription agent.)

Title (Mr/Mrs/Ms etc) _____ Initials _____ Surname _____

Mailing Address _____

Postcode (Zip Code) _____ Country _____

Signature _____

Date _____

Please state with which issue you wish your subscription to commence:

(April, July, October or January) _____

VAT Registration Number (zero rate): 243 5917 54

Photocopies of this form are acceptable.

Contents

Guest Editorial: Keeping Operations Ahead	113
The BT Operational Support Systems Architecture Framework	114
Enhancement of the Digital Core Switched Network	122
Cashless Services Network Operations Unit	128
CallMinder™ The Development of BT's New Telephone Answering Service	135
MeterLink™ BT's New Telemetry Service for Meter Reading	142
Computer Telephony Integration The Meridian 1 PBX	150
Driving on Customer Satisfaction A Re-engineering Case Study in the Retail Line of Business	157
Shared Spaces	162
Smart Rooms, Desks and Clothes Towards Seamlessly Networked Living	168
The Face of Talking Machines in a Multimedia World	173
A Model for the Project Management NVQ	179
National Vocational Qualifications for the Telecommunications Industry	186

