

VOLUME 18 PART 3 OCTOBER 1999

BRITISH TELECOMMUNICATIONS ENGINEERING

Included in this Issue

Alliance Engineering Symposium

Access in Focus



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



BRITISH TELECOMMUNICATIONS ENGINEERING

Contents

VOL 18 ■ PART 3 ■ OCTOBER 1999

Guest Editorial 1	74
Alliance Engineering Symposium Bruce Wiltshire	
Guest Editorial 2	75
Access in Focus Lawrence Bickers	
<hr/>	
<i>Alliance Engineering Symposium 1999</i>	
AT&T's 21st Century Network Architecture John Newell and Richard Roca	76
The Market-Driven Network John Helleur and Tim Wright	85
The Full Monty: PSTN-Equivalent IP Telephony for Spain Steve Brabner and Clive Dellard	92
Global Convergence—The Challenge of the 21st Century: Tele Danmark Experience Jean-Marc Salles	99
GPRS and the Mobile Data Revolution Peter Lisle, Howard Green and Peter Mason	107
UMTS Core Network Issues Thor Gunnar Eskedal, Arne Folkestad, Bjørn Harald Pedersen, Astrid Solem and Stein Svaet	113
Dense Wavelength-Division Multiplexing and Optical Networking Dietmar Garthe, Clemens Epple and Bernhard Edmaier	119
Interworking in the Future Broadband Access Network— A Question about IP and ATM Frode Beckmann Nilsen and Leif Aarthun Ims	125
UK Digital Terrestrial Television Ian Parke and Andrew Rayner	132
Managing IT Investment to Deliver Business Benefit Howard Ferguson and Graham Longmate	138
Extending the Boundaries of E-Business Transforming Back-Office Operations Between Telecommunications Operators David Freestone, Jonathan Mitchener, Alan Pengelly and Andy Childerhouse	142
Advanced Global Electronic Messaging Mark Fraser	149

Cover Picture

BT Corporate Picture Library;
a BT photograph

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. Roger Blake, C.ENG., M.I.MECH.E. Jason Chilton Andrew Dawson-Maddocks, B.SC., AM.I.E.E.

Martin Evans, B.SC., PH.D. David Greenop, B.SC., M.SC. Peter Howell, B.SC., C.ENG., M.I.E.E., M.C.I.B.S.E., M.I.O.M. Steve Jordan, B.SC., M.SC., C.ENG., M.I.E.E.

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. David White, M.I.P.D., M.B.I.M., L.I.A.A.

British Telecommunications Engineering Journal, Post Point 2D05, The Angel Centre, 403 St John Street, London EC1V 4PL ISSN 0262-401X

Access in Focus

Optimising Investment in the BT UK Access Network Ian Lawrence and Derek Cottrell	154
Investment Decision and Control System for Access Planning Integrated Business and Network Analysis—A Case Study Christine Stratton	160

Customer Service

Call Centres in BT UK Customer Service David Duxbury, Rob Backhouse, Mike Head, Graham Lloyd and John Pilkington	165
--	-----

Telecommunications in the 21st Century

Telepathy Jerry Bowskill, Craig Wisneski, Barry Crabtree, Peter Platt and Stuart Softysiak	174
Whither Video?—Between Sentiments and Signals Alwyn Lewis and Graham Cosier	182
Active Networks—Making the Next-Generation Internet Flexible Ian Marshall	188

Structured Process Improvement

Simulation Modelling in BT Operations Martin Hind	195
Statistics—An Indispensable Management Tool for Telecommunications Jonathan Malpass	201

External Engineering

Network Electrical Protection Bob Ruddock and Pete Whelan	210
---	-----

38th European Telecommunications Congress	219
Journal Awards 1998/99	221

Regular Features

pcochrane@btlabs	222
Book Reviews	223
Telecom Focus	224
Field Focus	230
Forthcoming Conferences	230
Information	231

Theme Editors

Access in Focus Lawrence Bickers	Customer Service Mark Jakeman	Telecommunications in the 21st Century Steve Sim
Structured Process Improvement Carole Jones		

Bruce Wiltshire

Alliance Engineering Symposium



I am delighted to introduce the selection of papers published in this edition of *British Telecommunications Engineering* from the Alliance Engineering Symposium—AES99—held in Oslo, Norway, 13–16 June 1999.

This prestigious event was co-hosted by BT and its alliance partner Telenor. We are very grateful to the Telenor team for their active support which resulted in an event hailed as the 'best ever'. We are particularly grateful to Dr. Ole Petter Håkonsen, Senior Executive Vice President Telenor, and Chris Earnshaw, BT Group Engineering and Technology Director, who acted as corporate sponsors for the event.

The key aim of the Alliance Engineering Symposium is to be the foremost international engineering event for the growing BT family of alliances. AES99 was the fourth annual event in the series with objectives as follows:

- to bring together the engineering community of the BT alliance family;
- to demonstrate technical leadership of the alliance engineering community; and
- to promote the importance of engineering to BT alliance commercial success.

In excess of 225 engineering specialists and senior managers attended the event from all parts of the BT family of alliances with some 25 countries represented. Over 60 leading-edge papers were presented covering the four conference themes (IP/ Networks; Mobility; Access and Convergence) with additional keynote presentations given by top managers from across the alliance family.

Feedback from the event has confirmed that AES99 met all the objectives described above and presented a rare opportunity for delegates to network with others to share experiences in pursuing exchange of technology, information and experiences.

The papers beginning on page 76 are a selection of the high-calibre presentations made at AES99—I am sure you will find them interesting and informative.

Bruce Wiltshire

Programme Manager, AES99

BT Group Technology

Lawrence Bickers

Access in Focus



It gives me great pleasure to introduce a series of articles from people in the access operational domain (begins on p. 154). System developments to assist with forecasting and modelling customer demand and guide network planning form the primary theme of this series. More specifically, the papers concentrate on developments where users have taken the initiative to design and build systems focused in meeting these needs. Articles that describe new access technology, its deployment in the network and the importance of reliability and asset assurance will also be presented.

The access network is of increasing importance to any communications company, and it is the part that is most visible to customers and the public at large. It is a huge and diverse structure that continuously evolves to meet demand for new residential and business services. Providing ubiquitous access to customers, the 'last mile' connection, in a deregulated and competitive environment is a key BT differentiator among the plethora of telecommunications operators and service providers in the UK.

BT's access network currently serves almost 30 million lines and has a replacement value of about £10 billion; a massive financial investment by any company's standards.

Network access, part of BT UK, is responsible for managing this asset efficiently and effectively with the main aim of delivering a quality service to BT's customers. There is increasing understanding of the business value and financial impact a well-managed access network can have. The operation of such a network is labour intensive with both a high current and capital expenditure. All products and services, together with their associated business practises and network components, must therefore be designed and evolved to reduce whole-life costs through increased productivity and network reliability.

Today's regulated and fiercely competitive telecommunications market demands investment in the access network to be targetted in a way that was never previously envisaged. In the past, planning the network focused primarily on technology and products. Today's business demands planning and investment decisions that take full account of customer requirements, network capability, regulation and service profitability.

Ongoing evolution of process and planning systems is employing technologies and practises to provide suitable decision tools. BT has a wealth of data sources from which to build a better understanding of its customer base and the network. Network and business planners need decision-support systems that bring together data from a variety of business systems. The vision and main driver is to ensure that customer demand for network connections is met in a cost-efficient manner while service provision times continue to be reduced. The target is a common platform architecture supporting automated data sharing between people and systems that will also support electronic commerce with suppliers.

Good information systems in the planning offices of today are essential. The increasing volumes of work, the

pressure for reducing lead times and the complexity of new product offerings mean there is a need to provide people with the right information and automated processes to plan and build the network, for today and tomorrow, at the lowest possible cost to meet customer demand. Increasingly, the use of information technology is coupled with people's experience and expertise to fully exploit information through innovative automation solutions. Some of these initiatives, practises and solutions, already successfully deployed in the business, will be described in the articles which will include IDACS, LPF, CONGO and MFS.

New developments in broadband copper access technologies and the massive growth in recent years of the Internet are also changing the map of the communications market. BT has already announced its investment programme to upgrade the access network for broadband services driven by the advent of asymmetric digital subscriber line (ADSL) technology. An article will describe the technology and the architecture designed to deploy it in the network. Another article will give an insight into the innovative work taking place to ensure that the reliability of the network is not simply maintained but increased to cope both with customer demand and new service requirements. The merging of traditionally separate industries, information technology, telecommunications and media entertainment, is creating new service opportunities that also seek to exploit further the network capabilities and the skills of BT's people. A typical example of this will be given in an article describing a digital terrestrial TV project.

I am certain you will find these articles to be informative and of great interest.

Lawrence Bickers

Manager, Access Infrastructure
BT Networks and Information Services

John Newell and Richard Roca

AT&T's 21st Century Network Architecture

During 1998, AT&T completed a study to determine the future direction of the AT&T network. This article describes the result of the study – AT&T's 21st century network architecture. Current technical and market trends are reviewed, together with AT&T's network evolution and architectural directions up to 2004.

Introduction

The telecommunications industry is undergoing tremendous change. Major industry trends include an increasingly competitive environment; the emergence of the Internet, which is beginning to compete with the public switched telephone network (PSTN); and the accompanying phenomenal growth of data networking traffic. In order to deal with these fundamental changes in a structured and orderly manner, a study was commissioned during 1998 within AT&T to outline the evolution of AT&T's network into the next century. Specifically, the study period projected the architecture of the AT&T network into the 2004 timeframe.

This article begins by reviewing the technical and market trends driving today's telecommunications industry and by illustrating how these trends are impacting AT&T's networking directions. Next the history of AT&T's network evolution is summarised, highlighting the important technical advances which have enabled today's intelligent voice and data networks. Addressing the timeframe to 2004, the study then focuses on the specific architectural directions which will be followed in evolving AT&T's network in the areas of:

- *information architecture*—the intelligent software-driven environment required for operations support as well as for value-added network applications;
- *transport and restoration*—the fabric of the long-haul broadband transport network along with the

infrastructure for interconnection and failure recovery;

- *switching and routing*—the evolution to Internet protocol (IP) packet technology as well as the integration with the evolving circuit-switched network; and
- *local and access*—the architecture and technology for local distribution within metropolitan areas and the access methods to the premises, including wireless, hybrid fibre-coaxial (HFC) and digital subscriber line (DSL) technologies.

Finally, the article gives an example of a network messaging application implemented under this architecture and an overall summary of the study.

Technology and Market Trends

Technology trends

Several aspects of networking technology are experiencing trends which, in turn, are having a significant effect on the way networks are being implemented. A significant trend in the computer industry is Moore's law, which states that the processing power of state-of-the-art computers is doubling every 18 months. This trend is driven by advances in the semiconductor industry, which enable the continuing ability to add more and more transistors functioning at higher and higher speeds onto the semiconductor devices which comprise the computer. However, there is a parallel trend in the telecommunications technology which is even more profound than

Moore's law. This trend is the simple fact that the amount of information which can be transported over a single strand of fibre optic transmission cable is doubling every 12 months. Even more remarkable, these increases can be realised by reusing fibre cable which is already installed and by simply upgrading the equipment which drives, amplifies and terminates the cable.

A second technology trend involves breaking the high bandwidth barriers that have traditionally existed in the local network which is still dominated by older technology only suited for low-bandwidth voice communications. Techniques for delivering high bandwidths to homes and businesses have been developed and are now starting to be deployed in earnest. Cable modems now allow two-way high-speed communications into homes which in the past had simply used the coaxial cable systems for one-way broadcast television entertainment. DSL techniques have enabled the reuse of traditional twisted pair wiring to achieve two-way data rates which rival coaxial cable in speed and error performance. Finally, wireless technologies have evolved to the point where they are becoming a viable alternative to 'wired' access to end-points to achieve high bandwidths to fixed and mobile end-points.

A third technology trend is the emergence of packet data networking, routing and switching technology. In the past, this technology had been applied to the traditional data communications environment in which the transmitted information was comprised of characters and images. For these information types, arrival rates were of no concern; the only requirements were that the information needed to arrive correctly and in the proper order. However, circuit switching, which provides a dedicated connection between two end points, was the technology of choice for conveying voice and video information. By

supplying a dedicated connection, circuit switching technology guaranteed the regular arrival rates required for voice and video intelligibility. Recently, however, it has been shown that acceptable performance for voice and video can be achieved in the packet environment by paying careful attention to the requirements of the information type being transported and assigning an appropriate priority in the communication process to the information type. By assigning such priorities, it is possible to meet the arrival rate requirements of voice and video and still remain in a totally packet environment. Also, the cost of packet switching is continuing to decrease and now rivals that of circuit switching. Assuming this cost trend continues, it will become economically viable to provide all wide-area communications on a packet data network and, at the same time, satisfy the needs of all the information types being transported.

A fourth technology trend is the increasing dominance of IP as the protocol of choice for data transmission. In the past, vendors of data communications equipment would compete for customers using proprietary data networking protocols which were unique to the providing company and would not necessarily work with the equipment provided by other companies. The Internet and IP have changed all that by using open and well-tested protocol standards which were specifically designed for interoperability between different networks and equipment from different vendors. In fact, the popularity of IP has become so great that it also is now being viewed as the protocol of choice for other information types beyond the character and image data which comprises most of the traffic conveyed on the Internet today.

A fifth technology trend is the use of layered software to implement computer-based applications. When the computer industry was in its infancy, it was not unusual for

applications to be designed and implemented with a totally new set of programs, which were unique to that application and which had little or no relationship to other applications. However, as the application development industry matured, it was realised that major efficiencies in cost and development speed could be realised by reusing software. The major enabler of software reuse is the development of software which performs a well-recognised and understood function or set of functions and provides a standard set of software interfaces to other programs used in the system. As this concept was further refined, the application development industry began to view such modular software as layers of functionality, with each layer having a well-defined set of interfaces to other software layers. Software constructed in this fashion is sometimes called *middleware* because it bridges the gap between computer resources like printers and networks (the bottom) and the applications (the top) being implemented using standard interfaces, now increasingly referred to as *application programming interfaces* (APIs). Middleware platforms have emerged as the preferred application development technology of choice for reasons of cost efficiency and development speed.

Networking trends—migration to the Internet

A second set of trends involves network evolution and specifically the fact that new networks are utilising and being driven by Internet technology and components. The public Internet itself has experienced a phenomenal growth in importance and usage. Emerging from almost total obscurity five years ago, there are 150 million users of the public Internet worldwide in 1999, and the number of Internet users is doubling every year. Initial usage of the Internet involved information exchange and electronic mail, but electronic commerce on the

Internet has already become a major business direction. Internet businesses which sell travel, books, music, clothing, automobiles and consumer electronics have already attracted millions of customers and the growth rates of these businesses routinely exceed 100 per cent per year. Recognising that such electronic commerce is just beginning to take advantage of the inherent creativity enabled by the Internet, we can expect a steady growth in the number and scope of the businesses thriving on the Internet.

A direct consequence of the internet technology is the *intranet*, the new internet-based direction of corporate and enterprise networking. Increasingly, corporations are using internet technology to create their corporate networks. In the past, corporate data networks tended to be very application specific because they were expensive to build and typically driven from large mainframe computers. Now, the notion of the intranet has taken over. An intranet is simply a corporate network which is implemented using IP and internet technology. It might not even connect to the public Internet at all for security reasons. However, the benefits of using internet technology are compelling. First, the technology is well developed and mature for delivering character and image information, and the data networks implemented with this technology tend to be more stable than past networks using proprietary technology. Second, the open and competitive nature of internet networking technology tends to make internet-based network implementations much more cost-effective than other alternatives. Finally, intranets are multipurpose utility networks which can be used for a variety of functions. It is not unusual for intranets to mature and change in the way they are used. Typically, intranets are used initially for e-mail and information exchange. For most enterprises, the intranet quickly becomes the dominant method of information

exchange and corporate web sites become the rule rather than the exception. Once this happens, people inside the enterprise start viewing the intranet as more than a fast way of exchanging mail and distributing documents, but rather as a corporate nervous system in which new innovative business applications can be created by linking information sources which had never been related in the past. For example, customer demographics contained in a billing database may now be used by the marketing organisation for regional sales campaigns.

Another new and fascinating further evolution of the intranet is also about to take place. It is becoming increasingly clear that IP networking technology can be used to transmit voice and video information in addition to the classical character and image data information types. Therefore, as IP voice technology matures, it will be a natural migration to include the corporate voice traffic in the intranet traffic flow. Enterprises will finally achieve integrated voice and data; not by using PBX or T-Multiplexer technology, but rather by using Internet technology.

Trends in voice and data traffic

The ever-increasing use of information technology in homes and businesses combined with the Internet growth described earlier are causing unprecedented growth of data networking traffic in wide-area networks. If we include fax as data traffic, it is generally acknowledged that, today, the amount of traffic supporting data communications is about equal to the traffic generated from voice conversations on the worldwide PSTN. However, the growth rates for voice and data traffic are substantially different. Worldwide, aggregated voice traffic tends to grow at a rate of roughly 5–10 per cent a year. While this is certainly a respectable growth rate for many industries, the growth rates

of data services are substantially larger. It is not unusual for service providers to report data traffic growth rates of 50 per cent a year or more for fast packet services such as frame relay and asynchronous transfer mode. Data traffic growth rates for Internet service providers (ISPs) tend to be even greater. Recently, ISPs are reporting a doubling of traffic within intervals ranging from 3–6 months. To put this in perspective, a doubling every three months translates to an annual traffic growth rate of 1600 per cent. Although it is difficult to predict what the sustained growth rate of data traffic will be over the next five years, it is possible to characterise the overall trend. Today, voice traffic and data traffic on the PSTN are about equal. Given the high growth rates of certain aspects of the data networking marketplace like the Internet and intranets, we can safely predict that data traffic will dominate our networks in the 21st century. The only question that remains is how dominant data traffic will become. However, to put that question in perspective, it is likely that multimedia communications comprising voice, video and image will emerge as the dominant information format flowing to network end-points in the 21st century (see 'Impending Inflection Points' later). When that happens, the distinctions we cite today between voice and data will blur and then finally disappear because they are no longer relevant.

Implications of the trends

The trends described above are related. Data traffic will grow because the combination of low-cost internet technology and technology insertion to increase fibre transmission capacity allows it to grow cost-effectively. The demand for Internet access and internet functionality fuels the market for this growth. However, the overall implications are clear. The telecommunications industry is changing from the voice-dominated PSTN to a high-growth

global data network offering a wide range of public and private information services.

20th Century Network Evolution

It is instructive to examine the recent evolution to today's voice and data networking. From an architectural perspective, it is possible to identify some technological turning points during this evolution. These turning, or inflection, points enabled much more rapid improvements in network performance and functionality and thereby give us some insight into the kinds of changes we can anticipate during the advent of the 21st century.

The early voice network

Direct distance dialling (DDD) was realised in the United States after World War II with a network of mechanical switches and transport provided by analogue cable and microwave radio systems. Signalling between switches was carried in-band and the switching network was arranged in a five-level hierarchy which was rigorously engineered and administered. This hierarchy provided relatively efficient use of expensive, long-haul transmission facilities, and accommodated other traffic engineering necessities such as overflow routing and load balancing. Customer-based switches, now called *private branch exchanges* (PBXs) were beginning to appear, and could be considered to be a sixth level in the switching hierarchy. Towards the end of this period, switching intelligence became based on computer technology, but switching matrixes and associated transmission systems remained analogue.

The intelligent voice network

Beginning in the 1970s, the intelligent voice network emerged by using computer technology to create stored-program-controlled switches. The intelligent call control components of these switches were then linked using

a packet data network which was out-of-band relative to the actual telephone call voice information. This packet data network, called the *common-channel signalling network*, allowed the computer-based switches to communicate with one another and to take advantage of processing and database resources embedded in the network. This intelligent network infrastructure allowed the creation of a host of value-added voice services in the emerging competitive telecommunications environment in the United States during the 1980s. The use of stored-program-controlled switches combined with common-channel signalling and network-based computers was a major inflection point in our network evolution.

Early data networks

Data networking began in earnest during the 1970s. In the US, business drove the data networking thrust by deploying modem-based networks which used IBM's Bisync and SNA architectures and AT&T private voice lines. Even though the PSTN transported most of the country's data, the service was transparent and was treated simply as private line voice. In short, there was little concept of an explicit data network architecture through this period. Also, the error rates on modem-derived data links over analogue facilities was high, usually no better than 1 bit in 100 000, even on a good circuit. Frequent data re-transmissions were required, forcing considerable overhead in the transport protocols and delays in the data transportation through the network.

In other parts of the world, public data networks achieved a significant amount of usage and importance. These networks utilised an open protocol called *X.25* and could serve the data communication needs of both large corporations and occasional dial-up users.

Fast packet networks

The introduction of fibre optic transmission led to significant

changes in both the quality and technology of voice and data networks. The quality of the digital transmission over fibre was orders of magnitude better than the quality achieved with the older analogue facilities using coaxial cable and microwave radio. This quality improvement was reflected in the use of much faster and efficient protocols for data transmission such as frame relay and asynchronous transport mode. Also, fibre quality removed constraints of distance. Voice telephone calls sounded as good across the country as they did across the street.

The use of high-quality fibre transmission for data transport was another networking inflection point, having a beneficial effect on voice and a revolutionary impact on data.

Internet technology

Client/server networks began during the 1980s as a natural consequence of networking personal computers (PCs) using local area networks (LANs). The need for shared servers became apparent for common functions such as file storage and backup, printing and wide area communications. The first technology for connecting remote LANs used the concept of bridging to make a big virtual LAN out of two geographically separated LANs. Routing technology gradually replaced bridging because it allowed the interconnection of disparate LAN types (for example, token ring and Ethernet) over fast data communications services like frame relay, offering an unprecedented degree of flexibility to the data communications marketplace. The early 1990s saw competition for the routing protocol of choice, with many vendors vying to make their proprietary protocol the de facto standard. However, the adoption of the open and pragmatically simple Internet protocol (IP) and the associated network rules established by the Internet Engineering Task Force (IETF) has already led to an explo-

sion in data networking demands for virtual private enterprise networks as well as for the public Internet.

The astounding growth of the Internet accompanied by the popularity of the open standard internet technology is an important inflection point and a dominant driver as we move into the 21st century.

Impending inflection points

By examining the past and extrapolating the confluence of technologies into the future, we can expect at least two more inflections in the networking domain over the next several years which will have a profound impact on the way we live and work. The first significant change will be the rapid migration to multimedia communications, as broadband communications finally becomes available to homes and small businesses. Fibre bandwidth continues to double every year. Therefore, backbone network providers have the ability to provide ever-increasing capacity by inserting new equipment into the installed fibre network plant. The real bandwidth bottlenecks involve the local access environment which historically has been controlled by telecommunications monopolies with no competitive incentive to improve the speed or quality of the communications. With the advent of deregulation throughout much of the world combined with emerging new access technologies, we can expect that the bandwidth bottleneck will be rapidly removed. We are rapidly approaching the day when high-speed communications to our homes and business premises will become commonplace and our day-to-day information exchange will be increasingly conveyed electronically as high-quality real-time two-way voice, video and image composite communications.

The second inflection point we can expect is the proliferation of value-added network-based applications. Presently, the Internet is enabling a transformation in the way commerce is conducted. The Internet has

already demonstrated that most of the traditional commodities like clothes, automobiles, travel reservations, and books can be bought and sold electronically. Moreover, new forms of commerce like day-trading of stocks and bonds and individual auctions are being invented and introduced. The Internet is a powerful demonstration of the power of computer-based value-added applications which can be introduced on top of a ubiquitous and user-friendly network. We can expect an acceleration of such applications, especially as we realise the multimedia two-way environment.

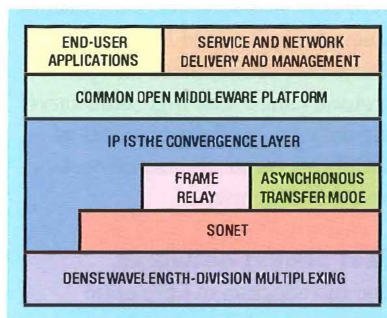
21st Century Network Architecture and Evolution

Before we can talk about the specifics of the architectural evolution in the 21st century, we need to first establish the architectural models which capture the architectural concepts, which, in turn, provide the structure for the evolution. Traditionally, communication architectures are described from two perspectives which are mutually complementary and necessary to achieve a complete understanding of the network operation. These perspectives are the *logical* and *physical* views, respectively.

The logical architecture

The logical architecture (see Figure 1) is intended to illustrate the various functional layers of the network, ranging from the basic

Figure 1—21st century network logical architecture



transport to value-added applications, network management and support systems at the top. At the very centre of the logical architecture is the IP layer which is the common meeting ground for all of the information types conveyed by the network. We expect that voice, data, video and image will ultimately be represented at this layer. However, we believe it is still important to offer the more traditional packet data services of frame relay and asynchronous transfer mode (ATM). Although the SONET multiplexing hierarchy is used extensively for traffic aggregation, we expect some situations in which IP traffic loads will rapidly grow to the point where it will make economic sense to put the IP traffic directly on top of the dense wavelength division multiplexing (DWDM) transport network, bypassing intervening layers for multiplexing, grooming and prioritisation. IP on top of DWDM will be especially attractive once the IP routing infrastructure evolves to the point where appropriate priorities, called *grades of service*, can be assigned to the demands of the information type being transported.

An open common middleware platform is essential for the efficient use of the network resources beneath for both applications and management functions. Certain functions recur for almost all applications (like authentication and usage recording) and network services (like provisioning and billing). These common functions are embedded in the middleware layer and are done in a consistent and efficient manner for all applications and services.

The physical architecture

The physical architecture (see Figure 2) relates to the location and quantity of major network components such as switching and routing offices, traffic aggregation points and end-user locations. The view shown here depicts a large national network such as the United States. Note that some notions of the logical architec-

Figure 2—21st century network physical architecture

ture are also contained in this view, including the hierarchy of transport and restoration/switching and routing/applications. Note that at the top of the hierarchy is the notion of a common information platform which is accessible to all of the intelligent application elements, including those located on the customers' premises. The availability of such a platform will become essential to the implementation of network-based applications, as described in the section entitled 'Impending Inflection Points'.

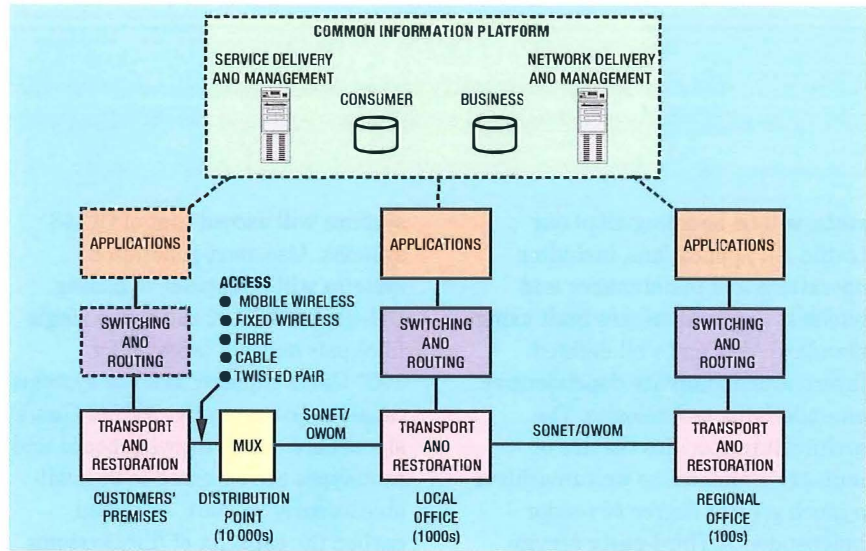
Evolution of the major network areas

Now that we have established the reference models for the logical and physical network architecture, the remainder of this article focuses on the areas of information systems, transport and restoration, switching and routing, and local access and networking. In these four areas, we will outline the expected evolution through the period of 2004/2005.

Information systems

Referring again to Figures 1 and 2, at the top of the network hierarchy are the information systems required to manage and provision the network as well as provide the applications. In the past, network information systems architectures were aligned with specific network offers. Management functions were redundantly implemented in multiple systems with offer-specific logic embedded in each application. Customer data was stored in product-specific databases. Enterprise-wide processes were difficult to achieve and often required manual workarounds. A customer's experience was different across products and services. As our network services grow in kind, scope, complexity and customer importance, it is necessary that we revamp the way these services are managed as well as the way applications are implemented on the network infrastructure.

In defining the target information systems architecture for the 21st



century, it is necessary that it causes a migration from a product-specific environment to a customer-specific environment and provides integration to the customer across all products and services. All network domains must use the same platform. Network operations systems move from technology-specific to function/process-specific and will employ open standards to foster a multi-vendor environment for network elements. The migration to open standards is necessary for reasons of cost, efficiency and time-to-market. The target architecture consolidates business functions by moving offer-specific rules to databases, integrates enterprise-wide processes through workflow technology, bridges network platforms through function/process orientation, integrates customer experiences through an integrated information architecture, and standardises software architecture through a common service operating environment.

A key element of our target information systems architecture is the notion of a common information repository which drives the information system functions. The information repository can be viewed as a two-dimensional matrix with three classes of data (customer, business rules and network resources) within three ordered layers (customer-specific, logical implementation-specific and physical implementation-specific.) The information architecture supports internal network operations and maintenance, service-related business functions such as ordering and

provisioning, and value-added applications, such as messaging and multimedia conferencing. The centerpiece is a single, logical customer database which supports all service delivery and management functions across all products/services.

Process-driven applications will be integrated through workflow technology, providing the ability to bridge all service platforms and applications. Workflow also enables flow through automation which seamlessly traverses functional boundaries (for example, 'order-to-provision-to-cash') across all offers. At the industry level, process automation also facilitates a more complex set of arrangements and interfaces with other providers by integrating the supply chain.

In the target information systems architecture, on-line product catalogues are provided which contain data that describes the configuration rules for each component, service and offer. This catalogue separates business rules from applications, and enables dynamic construction of new offers without requiring coding changes in the application. In terms of application design, the usually hard-coded logic is replaced with applications dynamically driven by external rules. The result is cross-product/service applications (ordering, billing, etc.) which support all variations of products/services.

A major consideration in achieving a simplified network infrastructure is to establish a well-defined and rigorously-enforced interface between our systems architecture and our packet network which, in the end

state, will be handling all of our traffic. All applications, including operations and maintenance and business applications, are built using standard APIs and well-defined layers which minimise dependencies on underlying technologies. The architecture specifies the use of industry standards so we can achieve a much greater degree of vendor independence. Third-party service providers will need to provide applications to our network customers using a standard set of application programming interfaces (APIs).

Transport and restoration

Referring to Figure 2, the transport functions provided at the local and regional offices are very similar. The fundamental interconnection speed of today's AT&T long-distance network is the DS3 rate of 45 Mbit/s. Restoration in this network is performed at the DS3 level using electronic cross-connects and the FASTAR restoration system. FASTAR is a centralised real-time restoration platform which quickly calculates the optimal network configuration after a facility failure and restores the network transport functionality based on that calculation. In addition, SONET/ring capability has been deployed which provides a very fast restoration capability for ring-routed facilities. The restoration in our local network is achieved via SONET rings.

DWDM technology, which provides the frequency multiplexing of multiple optical wavelengths on a fibre pair, has evolved quite rapidly in the last few years. Today, DWDM systems are capable of multiplexing a significant number of OC-48 (2.5 Gbit/s) channels or OC-192 (10 Gbit/s) channels on a fibre pair. In terms of the state of DWDM technology evolution today, about four times more wavelengths can be multiplexed onto an OC-48 system than an OC-192 system, resulting in the overall capacity payoff for a single fibre pair to be roughly the same for both speeds. We expect that ultimately the yield from the OC-192

systems will exceed that of OC-48 systems. Also, next generation systems will be capable of mixing OC-48 and OC-192 rates on a single fibre pair and are targeted for 1000 Gbit/s capacity. DWDM systems capable of carrying OC-768 (40 Gbit/s) signals are on the drawing board and prototypes are expected to be available as early as 2001. As stated earlier, the capacity of fibre systems is doubling every year and is expected to keep up with the traffic growth described by inserting new DWDM technology into the existing fibre base on a regular basis.

As the AT&T network evolves and the local environment grows in importance, we will implement a tiered hierarchy for the transport architecture, with the highest speeds occurring in the long-haul network. We will migrate our basic connection and restoration speed from the DS3 rate to the OC-48 rate within the planning period. It is likely that optical cross-connects at speeds of OC-48 and OC-192 will become a reality within this same period. Also, within this period, we will migrate from a centralised to a distributed restoration infrastructure and, by using ring rapid recovery techniques, will migrate to failure restoration speeds of 100–200 ms.

Switching and routing

Presently, AT&T offers both voice and data services. The four primary kinds of data services include: frame relay, ATM, IP and private line. IP services include the consumer ISP service called *WorldNet® Service* and the business transport and electronic commerce and web hosting services. Backbone routers provide connection to the public Internet. Dial platforms are connected to access routers, and provide dial-up access to the IP network. Higher rate, private line services connect directly to the transport network. The current voice network includes both local (wireline and wireless) and long-distance switches. The local switches support local access and nodal PBX trunks, as

well as trunks to other carriers and networks. The local switches support intelligent network services. The long-distance switches support nodal PBX trunks, intelligent network services and connections to other networks. Dial-up access to the IP network is provided from the local voice switches via the dial platform. The voice switches share a common transport network with the data services.

Interworking between the voice and data networks has been introduced with the offering of the IP voice transport services. These services use IP gateways in the domestic US network and selected international locations to convert voice traffic from the local voice switch to IP voice. The IP voice is then carried across the IP network to the far-end, where the call egresses to a voice switch via another gateway.

For our target switching and routing architecture, we plan to introduce high-capacity backbones for the frame relay, ATM and IP services using quality of service (QoS) capable packet switches and routers. Multi-protocol label switching (MPLS) is an IP-based signalling and routing capability which promises to greatly improve QoS and scalability of IP networks. Service-specific edge vehicles will be used, both to aggregate traffic for transport to the backbone and to provide local switching. Concentrators may also be used to aggregate traffic in locations that are too small to justify an edge switch. These concentrators could be shared across services. Voice switches will be interconnected via a packet-based backbone.

We intend to provide some of our intelligent network services which we presently offer in our intelligent circuit-switched voice network in the new emerging IP environment. As described earlier, the major enabler for the original intelligent network in the late 1970s was the combination of stored-program-controlled switches combined with a reliable out-of-band signalling network and

feature-rich network-based computers. We expect to see workable solutions for voice in the evolving IP environment in which the intelligence function is provided by a combination of gatekeepers (intelligent entities within an IP network) and service control points (intelligent entities within the traditional voice network). These solutions will include interworking, whereby intelligent multimedia calls which originate on one network can terminate on the other network.

Local networking and access

The access and local networks consist of various access technologies, a common transport structure, switching and routing mechanisms and the methodology to hand off large volumes of traffic to the regional/global tandem network. In addition, the local networks must support interconnection with other local networks, as well as other regional/global networks. The fundamental building block for the networks is a local fibre infrastructure, which links together various distribution points. The local network provides either access to other networks such as Internet, data, and long-distance or local switching for intranetwork communications or access to intelligent devices to perform service functions such as messaging services. Most traffic from the local environment remains in the local environment, so a significant amount of local switching and routing will take place. In the target architecture for local networking, these distribution points will serve a given geography, with the most cost-effective access technologies providing connectivity from residential and business customers to local offices. Target access technologies include coaxial cable, fibre, narrow and broadband wireless, as well as twisted pair carrying broadband DSL (see section on 'Technology Trends').

In the target architecture, access to consumers requires the use of

multiple technologies, with coaxial cable access playing an important, if not dominant, role. In specific situations where tethered access is problematic, fixed wireless will play a critical role. Fibre will continually extend into multiple dwelling units such as large apartment buildings. As stated above, DSL technology rounds off the portfolio of broadband access alternatives.

Within the planning period, we expect businesses to continue to demand greater bandwidth; and fibre will remain the dominant access vehicle. Broadband radio will provide a mechanism for providing broadband access to businesses on a short-notice or emergency basis when fibre alternatives do not exist. The interface to many business customers (and ultimately residential customers) will be through multiservices access multiplexers (MSIAMS) which provide a combination of traditional telecommunication interfaces like DS1 and DS3 to PBXs and non-traditional interfaces like Ethernet-based 100 Base T interfaces to data centre local area networks. MSIAMS will perform adaptation of many services, converting them into a common packet IP format for transport, routing and switching, while maintaining the required QoS. Within the planning period, we expect MSIAM technology to be integrated with broadband radio technology, offering an alternative to local fibre distribution.

The wireless mobility access network is evolving to higher bandwidths and convergence to a worldwide standard. Within the planning period, we expect the network evolution path to be towards an improved time-division channel structure that will first support 45 kbit/s IP data services and thereafter 384 kbit/s IP connections. It is expected that general packet radio services (GPRS) and 200 kHz enhanced data rate GSM evolution (EDGE) will be important technologies in this evolution.

An Application Using the 21st Century Network Architecture

Today, network providers offer an array of messaging services which are targeted towards specific market segments. While these services add customer value and meet revenue targets, they typically do not share common platforms, they require different customer databases, they do not share application software and they do not interoperate. As these messaging services increase in importance, customers are becoming increasingly dissatisfied with the need to check e-mail, office voice mail, wireless voice mail and answering machines, while sometimes still having to respond to paging systems. Also, receiving bills from multiple sources, even from the same provider, can be an added annoyance.

However, if we follow the directions established by our 21st century network architecture, we can make major improvements in the quality of the service we provide, while gaining substantial operating efficiency and economies. Using this architecture, we will implement our messaging services on a common IP platform which utilises a standard computing environment and set of APIs. Because this platform is open and standardised, the actual development of these applications will be quicker and less expensive than today's applications because many of the required functions are already programmed into the platform. Moreover, once the development is complete, we can share application resources such as directories across a number of market segments. Best of all, the individual messaging environments can interoperate. For example, 'road warriors' can retrieve IP messages while travelling in their cars. This can be accomplished because each individual is identified as a single customer in the information services platform and is recognised in the same way across all of the messaging services. Because there is inherent

integration among the messaging services, it is a straightforward feature to process a road warrior's text messages through a high-quality, text-to-speech converter in real time. Real competitive advantages can be realised by following the architectural directions spelled out in this article.

Summary and Conclusions

In summary, the 21st century network architecture is pointing us towards a global multi-provider service environment interconnected by open standards which will be based on fibre transport and packet data switching and routing technologies. This network environment will be accessed through a variety of access technologies with ever increasing emphasis on mobility and high capacity. A sophisticated signalling structure will link the various network elements and provide the nervous system for the integrated open layered information system infrastructure. In addition, this same signalling structure will enable value-added applications, network management and operations support. Finally, the 21st century network environment will provide an increasingly rich set of multimedia-based value-added services.

Acknowledgements

The authors would like to acknowledge the contributions of over 50 AT&T associates to this study. Special acknowledgements go to Alan Jones, Dan Sheinbein, Ken Shulman and Girish Varma who led their organisations in providing support to this study and who made numerous direct contributions themselves.

Biographies



John A. Newell
AT&T Laboratories

John A. Newell is currently Technology Planning Vice President for AT&T Labs in New Jersey. In this role, he has been instrumental in coordinating AT&T's 21st century network architecture project. He graduated from Drexel University in 1965 with a BSEE degree and from MIT in 1966 with an MSEE degree.



Richard T. Roca
AT&T Laboratories

Richard T. Roca is Vice President, IP Services Planning and Development for AT&T Labs, the research and development arm of AT&T. In this role, he is responsible for planning and developing business and consumer services based on IP technology. He is also responsible for maintaining AT&T's strength in certain key technologies and for ensuring that these technologies are effectively applied to the needs of AT&T's customers. He has also served as the Chief Technical Officer for AT&T Solutions. He began his career with AT&T in 1966. He joined Bell Laboratories after receiving a BS degree in engineering from Lehigh University. He attended MIT under Bell Labs sponsorship, where he received both SM and ScD degrees in 1967 and 1972.

John Helleur and Tim Wright

The Market-Driven Network

The Market-Driven Network is BT UK's strategic network initiative to enable it to meet the major challenges of a very volatile and competitive market, using new technologies and existing assets in the most commercially effective manner.

The Challenge

Given that all network operators have the same technologies with which to fight for market success, what will differentiate the winners from the losers? For a large and profitable incumbent, like BT in the United Kingdom, this question is even more difficult to answer than for a new entrant, since it has to play defence and attack with a mixture of old and new technologies.

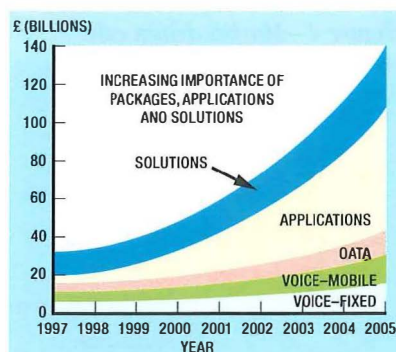
How do we make sense of the two management paradigms of milking the existing public switched telephone network (PSTN) 'cash cow' while investing to win in the new market opportunities now rapidly opening up?

Picking the next exciting technology is fun, but how can BT ensure cost-effective asset lives, and where and when does it take each key migration step with all the implications of disruption to customers, large money and skills investment added to potential shortening of existing asset lives. BT's answer in the UK is to adopt the focus of the *Market-Driven Network* (MDN).

The Background

In the UK, BT spends some £1.8 billion per annum of capital on the network, and keeps reducing costs

Figure 1 – UK market forecast revenues



and people numbers year on year to give productivity gains of the order of 10% per annum. The new build focus has been on green-field and brown-field sites and providing second lines, but now the deployment logic established over many years has to be re-assessed, stimulated by the growing reality of the datawave and the increasingly serious threat to BT's fixed networks presented by cable and mobile.

In the UK we are seeing the exponential growth kicking in for both Internet access and the take-up of mobile telephones. Narrowband traffic growth is largely data dial-up traffic (BT is adding data ports at a tremendous rate), mobile telephone usage in the UK is increasing at over 25% per annum and BT's revenues grew by 19% in quarter 4 of 1999.

Despite major productivity gains throughout the 1990s and increasing profitability, it is difficult to match the competition in pace and unit pricing, particularly in the new data products and with new, young customers. Why is this, and how can BT deal with it?

The Market

Several market analyses have been carried out on the UK market (Figures 1 and 2) and broadly the

Figure 2 – Total traffic volumes in the UK market

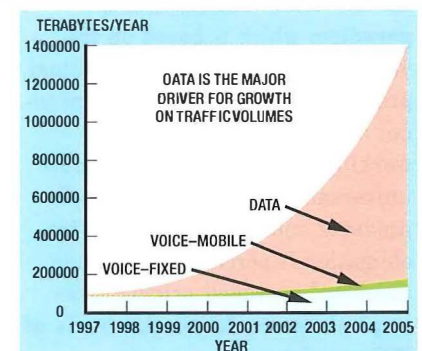


Figure 3 – Customer segmentation map

view is that the total market revenue will double to over £60 billion per annum by 2004, and that traffic will grow by approximately 10 times over the same period, with a massive shift to wide-band data/ Internet protocol (IP) products.

This exponential growth is being fuelled by the drive towards e-business, based on companies using intranets and extranets combined with mobile and home-working packages to enable flexible workforces and value chains.

The consumer growth is again fuelled by both IP and mobile products, but with the focus on entertainment; however, it is important to recognise the overlap of business and consumer demand with the growth of home-working.

Figures 1 and 2 anticipate a massive reduction in prices over the five years, led and supported by traffic growth reducing unit costs, which reduce unit prices. In turn this will stimulate increased usage and more traffic; a virtuous circle—if managed correctly.

What is clear are the following commercial imperatives:

- fast growth;
- packages and solutions: data/IP; multiple voice; video; mobility; and virtual private networks (VPNs);
- uncertainty and flexibility; and
- best performance and unit costs.

Focus on Segmentation

BT has been working with a paradigm which is based on the viability of a portfolio of individual products, each seen to be independent and profitable. BT is also working with the presumption of universal service delivery for all products and services, where this obligation is perceived to include universal averaged pricing. **Perhaps these could be the basis of BT's problems!**

SEGMENT GEO-TYPE	BUSINESS LARGE	BUSINESS MEDIUM	BUSINESS SMALL	S3	S2	S1/SOHO	PREM	POTS+ IP	POTS+ TV	POTS	TOTAL
RURAL											
URBAN											
METRO											
CITY OF LONDON											
TOTALS											

The first principle of MDN is to always think like a new entrant in order to compete; the market has to be segmented and each segment has to be fought appropriately. BT's approach is to view the 'battle-map' based on a matrix of customer segments and geographic segments (Figure 3), the latter because the delivery unit costs vary widely by geo-type.

The customer segments tell us about customer needs, indicating what products and packages are required to satisfy those needs; this also gives the size of pipe needed to deliver the package.

The geography is a rough indicator of traffic density and distance, and therefore gives a view of the delivery costs, plus the most cost-effective technology to use for each geo-type.

The map also shows where the competitors are impacting or are likely to impact; it helps to assess where the best profits are to be made, and it helps to show what BT's real market share is in each cell; that is, which customers BT has and which customers the competitors have!

BT can now treat the UK as a set of market cells (Figure 4), combining the needs of a customer segment with

delivery capability and costs of the geo-type, and recognising the competition for that cell. It now has a view of where to build to deliver packages which satisfy customers, and it can be even more accurate with investments (Figure 5).

Market modelling is used to assess market futures to a postcode or zip code, and BT is beginning to match this with network capabilities at the distribution point (DP). This gives a very accurate view of the market and the network build required. It also enables the company to be more accurate about where to spend network improvement resources.

Finally, this micro-market approach will enable the operation of a very accurate 'sell from stock' process, with the opportunity to target selling to the postcode.

The fundamental point to recognise is that satisfying this set of United Kingdom market cells is not just about the datawave or mobility, and it is not just about investing in IP technologies, it gives a framework to deliver value-adding packages from old or new assets in a commercially effective way.

Figure 4 – Market driven cell

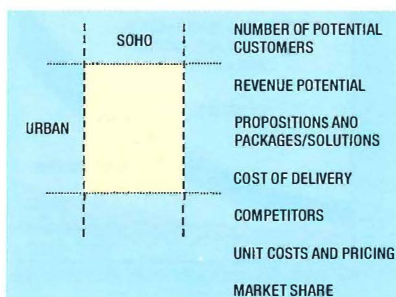
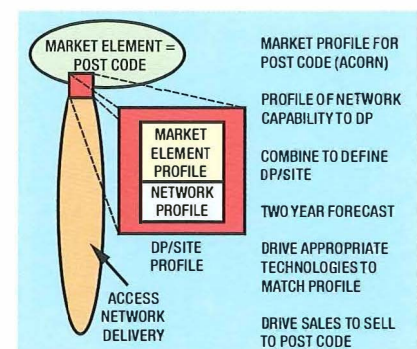


Figure 5 – Micro-market driven distribution point



MDN Commercial Principles

The MDN deals with the changing market and the network economic challenges head-on by applying a number of commercial principles:

- Build towards profitable customer segments and geographies:

– *Customer segmentation* The translation of customer segment needs into ‘propositions’ of packages, made up of multiple voice, data/IP, video and mobile/mobility. Aiming these to the customer segment by a translation to postcodes.

– *Geographic segmentation* Unit costs are determined by traffic density and distance of delivery; customer demographics also have a significant effect, with some 80% of traffic by revenue being derived from BT’s top 10% of nodes.

– *De-averaged pricing* The unit cost of rural access is around 10 times metropolitan access, and the cost of rural transmission is five times metro and inter-city transmission (per kilometre). Competitors currently, therefore, exploit BT’s geographically averaged prices, building against BT in cheap/dense geographies and renting from BT elsewhere. De-averaging is essential.

- Build capacity for anticipated market share, then sell the built capacity:

– *Build big* ‘Fat’, high-capacity networks with good utilisation yield, low unit costs. New technologies offer much higher capacity—typically five times higher—for only twice the cost. Building big with new technology makes sense when steered by good, confident market analysis based on the segment map, micro marketing and a set of linked models.

– *‘Build and sell’* BT is implementing positive selling in areas where network capacity already exists or immediately following a planned build, in order to reach competitive utilisation levels quickly and obtain best returns on the assets. This reduces costs by some 30% over incremental build.

- Aggressively use the new technologies to reduce unit costs:

– *Strategic, multi-service platforms* The MDN assumes a move to multi-service platforms, including access, data and intelligence platforms, to drive down unit costs as more products share the same platform, spread the risks and reduce time-to-market for new products.

– *Rapid focused roll out* New technologies offer significant reductions in unit costs (> 3:1). BT has to operate with a mix of old and new unit costs, and therefore must move to new technologies as early as possible for key customer segments and geographies, and must consider disaggregated pricing across old and new platforms.

- Build for flexibility and timing of investment:

– *Investment payback* Equipment build will minimise fixed costs (racking etc.) in favour of variable costs (for example, line cards). More appropriate assessments of asset lives are also being undertaken to better match market and technology volatility. Again this will be affected by segmentation.

– *Scalability* The volatility of the market arising from customer fashions, price competition, competitor activity and new technology means that volume demands cannot be forecast from past trends, except in the rela-

tively near-term. Designs, therefore, need to be highly scalable to cater cost effectively for very wide ranges of growth.

- Design to the best competitive affordability targets:

– *Cost-driven design* Traditionally, BT has designed new or enhanced platforms, estimated the costs (and unit costs) implied and, by adding a ‘reasonable’ profit margin, used this information to determine product pricing. With increasingly aggressive competition and falling prices, BT now has to design to meet competitive costs in each segment.

- Keep unit costs low by driving up utilisation:

– *Pricing and utilisation* Prices will be set to optimise utilisation and win highest traffic volumes by encouraging the use of spare capacity (pipe, bandwidth, time), via BT retail arms and other operators. Higher utilisation improves unit costs and protects margins.

– *Arbitrage* Tariffing must be based on a sound understanding of where real costs lie in order to minimise arbitrage opportunities. Any significant divergence between cost base and wholesale price provides opportunities for arbitrage.

– *Risk and reward sharing with strategic suppliers* Improved value-adding relationships are being built with key strategic suppliers, including risk sharing, joint evolution planning, joint research and development, and use of their design skills.

- Offer the best customer experience processes with maximum automation:

– *Service differentiation* The customer experience, a key

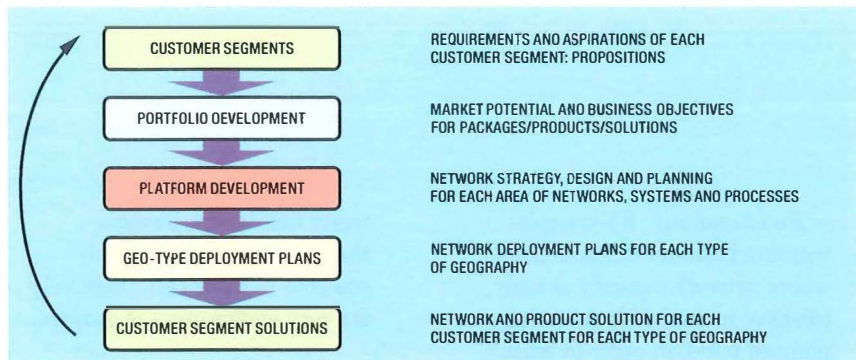


Figure 6—The business logic

differentiator, will arise from the use of world-best processes and operational support systems.

- Maximise the commercial value of existing assets:

—‘Sweat the assets’ The existing network will continue to supply the bulk of the revenues for some time. This will be for conservative customer segments and in difficult geographies where their commercial viability will remain for several years.

The Business Logic

Underpinning the commercial principles is the relentless business logic as represented in Figure 6. It starts with an understanding of the requirements and aspirations of customer segments, translated into propositions of multiple products. The product portfolio must be developed to meet these package requirements, and drive through to network/platform requirements for each customer segment and for each geo-type. Platform requirements are split across engineering domains, access, bulk transmission, switching, intelligence, processes, support systems etc, which must be developed in a coherent manner via a consistent set of evolution plans, migration steps and milestone targets for the delivery programmes. The most appropriate technology is used for each geo-type and route. Modelling is then used to assess the most commercially-effective rollout strategy.

The Market Proposition

The market propositions for the MDN are primarily business led. The aim is to provide:

- an intranet solution for leading-edge customers, namely large corporate businesses and their mobile workers;
- an extranet solution linking those customers with their customers and suppliers, including small/medium enterprises and people working from home;
- a data-ready and video-ready platform for recreational and educational use;
- a mobility and universal mobile telecommunications system (UMTS) ready network, offering seamless serving and seamless mobility irrespective of wireless, tetherless or wireless access;
- a solutions-ready network based on application programming interfaces (APIs) and interworking with customers’ equipment; and
- low-cost wholesale bandwidth for all big customers, including other

operators, service providers and BT service platforms.

The Products

Formal definitions have yet to be fully developed for MDN products. The products will be categorised in terms of their layer in a value chain (Figure 7) and based on their capabilities, not the technology. Each product will have retail and wholesale variants. The value chain nature of the products reinforces the move away from product stovepipes in which, for example, the access solution is specific to a network application. Products at the IP services layer and below will become commoditised increasingly and subject to fierce price competition.

- *Managed access* This provides connectivity between customer sites and the BT serving sites with the aim of supporting all services on a single access link. It is characterised primarily by the access bandwidth offered.
- *Managed transmission* This provides high-capacity, fixed-bit-rate pipes between network nodes and to a small number of very large customer sites. It is characterised,

Figure 7—The value-chain layers

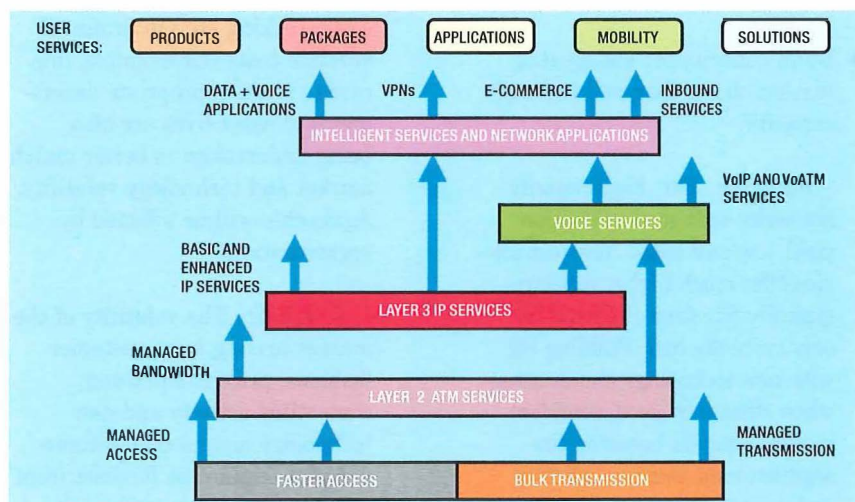


Figure 8—The basic network components

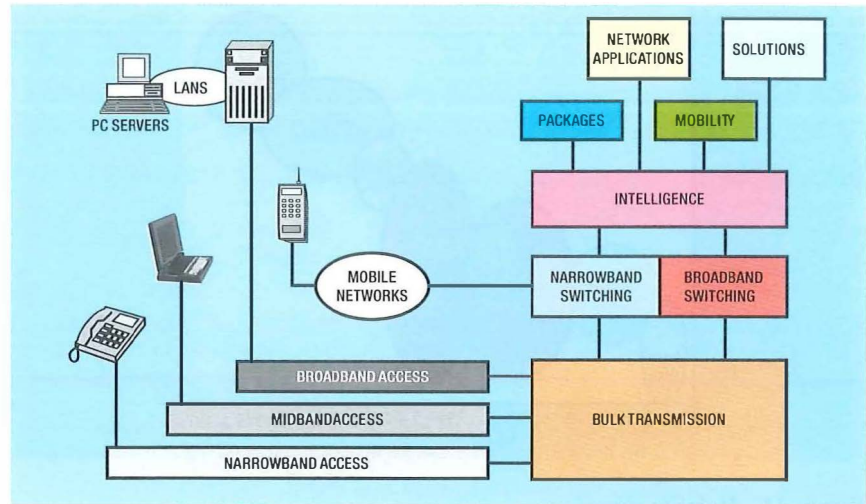
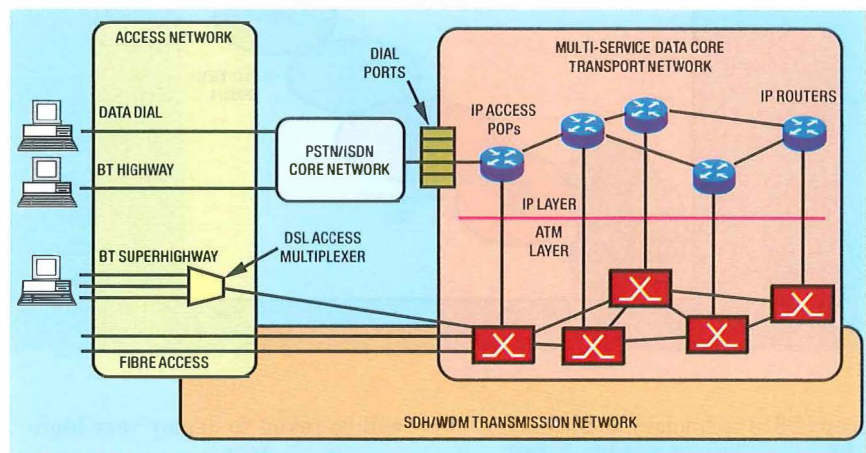


Figure 9—The network design in support of MDN



primarily, by the bandwidth offered, which typically will be at 155 Mbit/s and multiples thereof with the likelihood of optical channels being offered in the longer term.

- **Managed bandwidth** This provides connectivity at a finer level of bandwidth granularity which may be fixed on connection set up, or variable between network equipment, or between network equipment and customers' terminals, or even terminal to terminal. The primary managed bandwidth products are expected to be asynchronous transfer mode (ATM) private virtual circuits, characterised by bandwidth and quality of service, which will steadily displace existing leased-line products.

- **IP services** These will be the main focus of the data business since it is confidently expected that IP will become the universal end-to-end Layer 3 of choice supporting BT, third-party service provider and end-user applications. They fall into two types:

- basic IP service—simple transfer of IP packets, with no extras; and

- enhanced IP service—transfer of IP packets with support for flexible virtual private networks, security, user roaming, and guaranteed quality of service. Each type will provide name management, address management, routing, resource management, and accounting.

- **Voice services** The migration of voice services away from traditional 64 kbit/s circuit switched networks to IP- or ATM-based networks remains a complex commercial issue. Crucial to this migration will be the evolution of customers' terminal equipment and the need to keep the cus-

tomers happy during the transition. Existing PSTN voice will be supported on ATM switch fabrics as voice over ATM, while new voice services, particularly those emerging from the Internet and intranet arenas, will be supported on IP as voice-over-IP; and will evolve into industrial strength integrated voice/data applications.

- **Intelligent services and network applications** Real value add comes from the integration of basic and enhanced IP services, voice services and network applications, including mobility applications. These enable virtual private networks, e-commerce applications and other network applications such as inbound services.

The Network

The basic network components are shown in Figure 8. The components

and their interconnection have strong parallels with the product value chain discussed earlier.

Increasingly we will see open interfaces between the components and yet specified in ways which permit functional integration and thereby reduction in numbers of physical equipments and their interconnection. In particular, we expect to see the integration of bulk transmission functionality and switching functionality at various transport layers including ATM and IP. The access network components will be independent of the end-user products that they support.

To date, the main network design focus in support of the MDN has been on the data transport aspects; the network design is shown in Figure 9.

At its core is a scalable multi-service (broadband) switching and routing platform (MSP) with points of presence at the top 100 BT sites, giving direct access, either over fibre or copper with digital subscriber

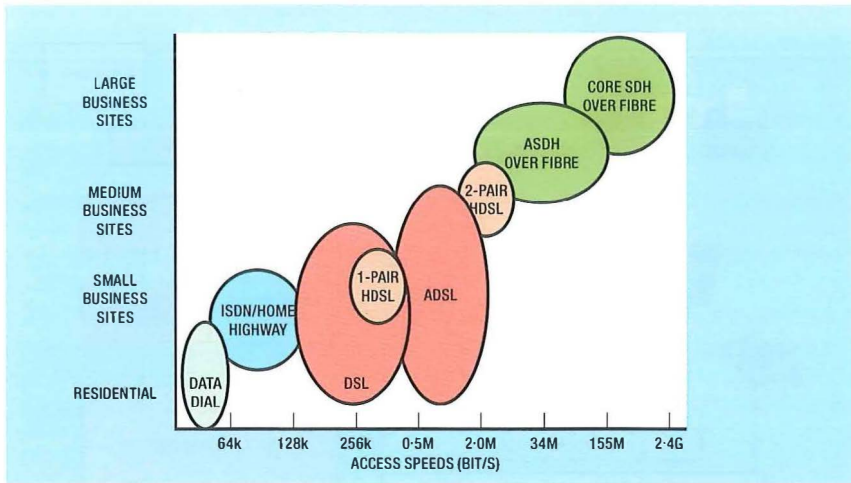
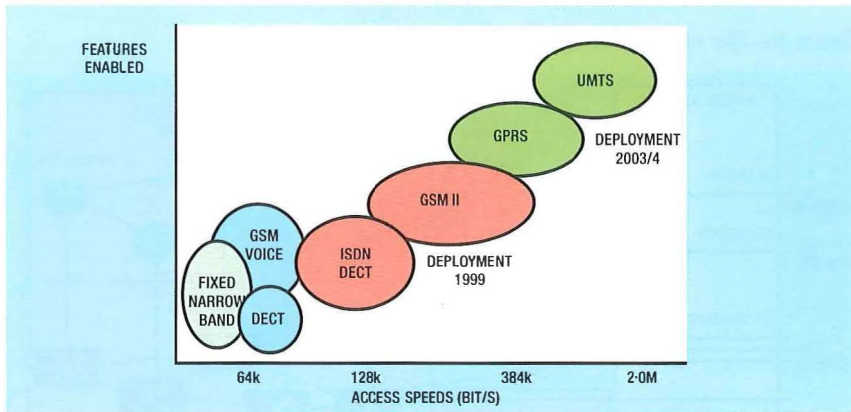


Figure 10—A full fixed-access armoury

Figure 11—Mobility access



loop (DSL) technology, to about 80% of business sites. In addition, DSL access multiplexing functionality will be deployed at an additional 1500 sites to permit high-speed access to 98% of business sites and homes in key competitive areas. For customer sites beyond the reach of these 1600 serving sites, BT Highway and PSTN dial access will continue to be available; using a full armoury of fixed and radio access technologies (see Figures 10 and 11).

The solution includes dial ports for the support of dial IP access from the PSTN/ISDN and gateways to support voice service interworking between PSTN and IP domains. All of this will be underpinned by the lowest cost and highest capacity transmission network employing synchronous digital hierarchy/wavelength division multiplexing (SDH/WDM) over fibre.

Many of the piece parts for this are in place, including the initial phases of the MSP with ATM edge switches and high-capacity IP core routers and preparations for large-scale DSL (mid-band) roll-out. As the roll-out proceeds, opportunities

will be taken to deploy very-high-capacity multi-function equipment with high-bit-rate interfaces. Use of such equipment will lead to physically simpler core networks with fewer equipments in tandem, fewer interconnecting cables and interfaces, less floor space, power consumption, considerable simplification of network management and generally much better economics.

The datawave is already hitting the narrowband network which will require enhanced switching capacity until the MSP fully takes the data growth. Next generation switches (NGS) are being deployed from next year to replace the current PSTN transit switches. These NGS transit switches will be based on an ATM switching fabric supporting PSTN voice directly as voice over ATM as well as offering the possibility of supporting broadband services.

The MDN will require some key enhancements to network intelligence, which will have to operate over several networks with either wireline or wireless accesses to support user mobility, data and voice networks to support data plus voice applications,

and service inter-working between BT's and other operators' networks. Finally, changes to the structure of network intelligence will be required to accommodate users with PCs and other intelligent terminals where the application is distributed between the terminal and the network-based intelligence.

The Geographic Focus

BT's network in the UK spans a large variety of geographies, each of which have different physical cost characteristics, different customer-demand profiles, different population densities and different competitors. The tenet of geographical segmentation is fundamental to the MDN; the following geographies have been defined to underpin this tenet:

- **Metropolitan**—covers major central business districts, including the City of London:
 - includes the largest 50 nodes;
 - mainly large business customers;
 - accounts for 30–40% of total revenues; and
 - intensely competitive, but net growth in connections overall.
- **Urban**—covers the remainder of the major conurbations and large cities:
 - includes the next 350 nodes;
 - some business customers, but mainly residential;
 - accounts for 30–40% of total revenues; and
 - highly competitive and negative growth in connections overall.
- **Cable areas**—remaining active CATV areas:

- includes the next 1350 nodes;
- mainly small business customers and residential;
- accounts for 10–20% of total revenues; and
- increasingly competitive and negative growth in connections overall.
- *Small towns*—remaining towns in rural areas:
 - includes the next 500 nodes;
 - is mainly residential plus minimal businesses;
 - accounts for 5–10% of total revenues; and
 - has some growth in connections.
- *Rural*—remaining country areas, very small towns and villages:
 - includes the remaining 4250 nodes;
 - is mainly residential plus minimal businesses;
 - accounts for 2–5% of total revenues; and
 - has minimal growth in connections.

This article has outlined how customer segmentation can be translated into geographic prioritisation using postcodes. Equally, information can be gathered on the network capabilities at the matching distribution point, such as sparing and utilisation, network performance including fault rates, vulnerability to weather and interventions, type and age of equipment, quality of data, etc. By coupling this network profile with the micro-market profile, an accurate two-year investment plan will

be created. Network planning systems will then drive the roll-out of appropriate technologies to the distribution point and matching sales to the postcode.

MDN Status in BT UK

The MDN has recently been adopted as the network strategy for the UK. It is now being implemented as the primary programme of transformation, changing the key business processes and skills to match the new approach. During 1999/2000 most of these changes will be put in place, driven by commercial targets and price changes to encourage the new commercial behaviours. To guide decision makers, a number of enabling activities are underway to deliver commercial targets, a network blueprint, network product plans, evolution plans, and programme objectives.

Acknowledgements

The development of the MDN thinking over the past 20 months has been a team effort on the part of many people in *networkBT*, Technology Directorate and Advanced Communications Engineering units; the authors wish to acknowledge their ongoing contributions.

Biographies



John Helleur
Networks and
Information Services
BT UK

John Helleur was appointed Head of *networkCommercial* in April 1999. He is responsible for the business and commercial focus and targeting of *networkBT*, the network operator in the UK. His team covers business

planning, market segmentation, network products, network strategy and evolution, network performance and utilisation, and network economics. He has been a director of the BT software joint venture in India, Mahindra-BT, for the past four years. He gained a B.Eng. in Electronic Engineering and an M.Sc. in Computer Science; he has worked for BT for 29 years, and his previous roles have included network strategy, OSS architecture, network management systems development, and switch software development.



Tim Wright
Networks and
Information Services
BT UK

Tim Wright works in the *networkCommercial* unit of *networkBT* where he is responsible for network evolution. He gained a B.Sc. (Eng) in Electrical Engineering from Imperial College, University of London, in 1972 and joined the then Post Office in that year. His work has mainly been in the transmission engineering areas, including a 10-year spell on international standards where he was deeply involved in speech coding and SDH standards. More recently, he has worked on network strategy and evolution planning. He is a Chartered Engineer and a member of the IEE.

Steve Brabner and Clive Dellard

The Full Monty: PSTN-Equivalent IP Telephony for Spain

BT Spain has embarked upon a world-leading project to implement a fully PSTN-equivalent network over an Internet protocol (IP) infrastructure shared with data traffic. The significance of this project to the economics and technology of the global telecommunications business is fundamental. This article describes the background to the project, the architecture of the chosen vendor, Nortel Networks, and some of the challenges being faced by the design team led by BT's Alliance Engineering.

Business Drivers

The economics of creating a national telecommunications network to compete with a long-established incumbent are harsh. Although local access is perhaps the most difficult problem to crack, the core network infrastructure must be highly cost-effective and capable of supporting multiple services. Traditional voice switches are notoriously expensive, and for this reason it has long been an objective of network designers to combine voice and data networks onto a single data network infrastructure—the so-called *datawave*. This has gradually been coming to fruition with the use of shared physical bearers and transport networks, but BT Spain took the bold decision two years ago to provide basic business voice services over their Nortel Passport Frame Relay network using voice-over-frame-relay technology, under the *Multivoz* banner.

The success of this service and the liberalisation of the Spanish telecommunications market have since led to the necessity to offer a wider range of voice services over a network capable of meeting the requirements laid down for interconnection with Telefonica and the 20 or so other aspiring network licensees in Spain. From these beginnings, the challenge was issued—was Internet protocol technology (IP) now sufficiently mature to provide voice and data services in a single infrastructure that could rival the dedicated networks and enable a rapid evolu-

tion into the multimedia services of the future?

Theory said that the cost benefits of integrating a broad mix of packetised voice and data streams should be significant. Given the falling capital costs of data network elements, and the future prospects for innovative services, it was an intriguing prospect, and the target set was to achieve the lowest cost-structure in the marketplace—a minimum 50% reduction in costs over conventional circuit-switched voice solutions.

In 1998, the big questions were whether IP vendors could replicate the rich diversity of basic and supplementary services routinely offered by modern intelligent networks, and whether IP—a protocol designed for non-real time 'best efforts' applications—could deliver the quality and reliability expected from network providers. To answer these questions, a request for proposal (RFP) was issued to the industry in June. The responses were sufficiently encouraging for a high-level design team to be established under the leadership of BT's Alliance Engineering which produced detailed designs and cost modelling for a business case towards the end of the year. This has ultimately proved to be the most ambitious project to implement PSTN-equivalence over any IP network worldwide.

BT Spain and Spain are particularly suited to a project of this type. Early experience with supporting voice-over-frame services has led to an understanding of the particular technical and customer issues

associated with this type of solution. Spain is a technologically advanced country, and BT Spain in particular is proud to have the opportunity to be a pioneer in the field. The fact that BT Spain is wholly owned by BT was also a factor in that the BT Board could develop this technology in a relatively self-contained network. Another important factor was that BT Spain did not have a large existing investment in a conventional voice network and so migration issues would be lessened. BT Spain was also in the process of deploying a country-wide synchronous digital hierarchy (SDH) network which would give the project the luxury of plentiful owned bandwidth for the first year or more. These positive factors came together like a fortuitous conjunction of the planets, and the business case was approved by BT Spain and BT in November 1998.

Service Requirements

The list of service aspirations from BT Spain Marketing was long and challenging:

- Multivoz replacement (basic business voice, virtual private networks (VPNs));
- indirect access;
- inbound (for example, freephone, shared rate, personal numbering, televoting);
- wholesale/reseller;
- value-added services (for example, fraud control, conferencing, voice messaging, fax store and forward);
- card;
- basic and advanced virtual private networks (VPNs);
- multimedia;
- direct access (ISDN primary rate interface); and

- GSM and international PSTN interworking.

The initial focus was on telephone-to-telephone services based on classic circuit-switched network techniques, with the IP telephony network providing a transit capability between other licensed operators (OLOs) and to directly-connected customer PBXs.

In addition, to satisfy the Spanish regulator for interconnection to OLOs, the IP telephony network must behave identically to a conventional voice network, with features such as local-number portability, emergency services access and Signalling System No. 7 (SS7) transparency. Voice quality had to be indistinguishable from the PSTN, and regulatory certification and initial service launch had to be achieved before the end of 1999.

Vendor Selection and Project Organisation

An Alliance Engineering-led team of voice, data, integration and systems specialists began work on a detailed specification, and the invitation to tender (ITT) was issued to eight vendors in November 1998. The technical schedule was necessarily of unprecedented complexity, combining the demands of both voice and data worlds, and the team was quietly proud of the strong reaction from vendors ranging from 'the ITT from hell' to the more positive vendor who described it as a perfect product description which rendered their marketing department redundant.

From those vendors who chose to bid, a shortlist of three was produced, and meetings and demonstrations arranged. Within a month, Nortel had become the clear leader, the recommendation being finally endorsed by BT Spain in February 1999. Common themes from the leading vendors were the difficulty in carrying their service control experience across from voice to data products, and their unpreparedness

for the regulatory demands of SS7 transparency and number portability. For some established data vendors, the learning curve associated with the world of carrier-class voice services, signalling systems and network availability was clearly too steep, given the available time.

Although we knew that we were in the vanguard of those attempting the 'Full Monty' of true PSTN-equivalence over IP, it transpired during discussions with Nortel that we actually had the earliest timescales of **any** customer and that the European version of their products—needed to meet BT and Spanish requirements—would be more advanced than the version for North America, where the bulk of the development had been undertaken. To our surprise we have therefore found ourselves at the front of the worldwide pack, and naturally many of BT's other worldwide partners have been keen to understand the project and follow our progress.

A joint Nortel/BT team has been established since March and is working closely together with the common objective of beginning trials before the end of 1999 to synchronise with the availability of the underlying SDH and IP transport networks. Separate design, systems, testing and implementation teams have been formed, each having joint BT/Nortel leadership, under a shared project management team.

In addition to Nortel test facilities in North America and the UK, the project has established a further network at BT's Adastral Park, Martlesham, as a primary test site for integration, security, management and subjective voice quality; this network will transfer to Spain once the service has been successfully launched there.

Network Design Overview

The chosen supplier

From the beginning, the adjudication team knew that, due to the immaturity of the technology, a

solution meeting all of the requirements would not be found. BT Spain's primary requirement was to provide a carrier-grade PSTN-equivalent network capable of supporting traditional voice services. The introduction of IP access types and multimedia services was not a day-one requirement although the marketing desire was for these to be launched shortly afterwards. The adjudication team recognised that, in order to provide a carrier-grade network, proprietary or pre-standard interfaces would have to be accepted, as currently-published standards have a number of deficiencies in this area. Having said this, future commitment to supporting standard interfaces when available, particularly ETSI TIPHON, was a major condition on all suppliers.

Carrier-grade is an often-used but not well-defined term. For the tender adjudication, the team concentrated on the following areas:

- *generic network capabilities*—the ability to support existing call types (for example, voice, fax and voice-band data) and services (for example, direct and indirect access, inbound services);
- *resilience and robustness*—better than 99.99% reliability, hot changeover to standby components and overload protection are some examples;
- *QoS and performance*—matching the performance requirements associated with voice and signalling in a traditional carrier switched network;
- *signalling and interconnect*—the ability to interface to existing networks with SS7 and to transparently carry SS7 information across the network;
- *scalability*—gateway devices capable of scaling to large time-division multiplex (TDM) routes;

- *monitoring of network resources*—the ability of both network management systems and, where appropriate, the call control entity, to monitor the resources in the network;
- *routing capabilities*—ability to perform standard routing functions, such as E.164 routing, digit manipulation, alternate route selection, and point of entry routing and restrictions; and
- *accurate event recording*—for billing records and secure storage and transfer of records to off-line systems.

In the final analysis the adjudication team felt that the Nortel Networks proposal best met these requirements in the required timescales. The solution has been designed from the start with carrier-grade services in mind. The solution offers traditional trunk network functionality and the support of ISDN-connected PBXs including a good range of supplementary services. The call control is designed such that additional access signalling types, such as H.323, can be added to provide multimedia services in the future.

Another key element in the adjudication was the support of

services by elements within the IP telephony network environment. The Nortel solution will provide almost all of the required services during the early phases of the project. As well as customer-focused services, a vital issue within Spain is the introduction of local-number portability which will also be provided in the initial release.

The services are phased over a number of software releases thus enabling their gradual introduction on the IP telephony network in a controlled manner. This does, however, bring its own problems as by the time the network is launched a number of services will already have been introduced on the BT Spain interim network—discussed later in this article.

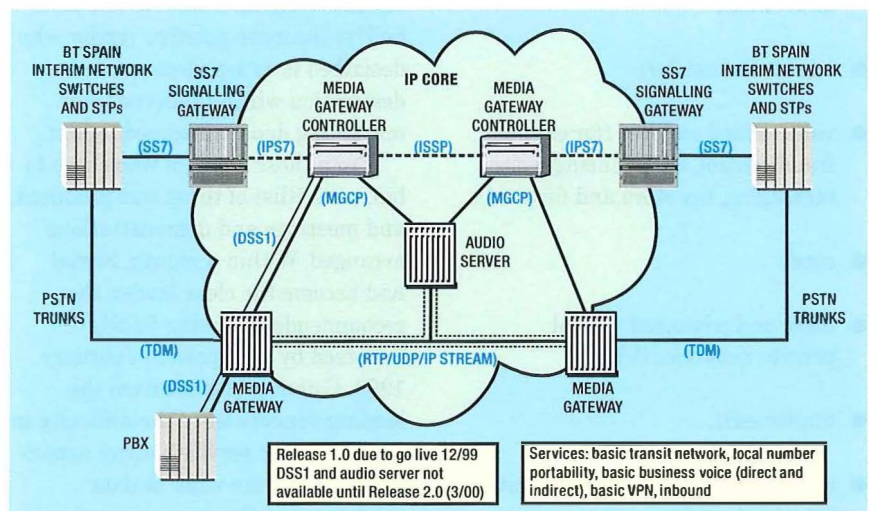
Proposed network design—initial phase

This section of the article describes the network design to be adopted for the first two software releases. Figure 1 shows the IP telephony network in the initial phase of the network. The IP telephony components and interfaces shown in the Figure 1 are described below:

Media gateway

The media gateway function provides the transfer of media streams between the IP domain and the switched-circuit network (SCN)

Figure 1—BT Spain IP telephony network—Phase 1 (4Q1999–1Q2000)



domain. The gateway terminates the protocols in each domain and provides conversion between the bearer channels on the SCN side and data packets on the IP side. This function includes any speech processing required, such as, voice compression, silence suppression and echo cancellation.

The Nortel CVX1800 product provides this function. Originally a remote access server (RAS) product, the CVX1800 has been re-engineered to meet the demands of IP telephony. During the first phase of the network, each gateway will provide 36 E1 interfaces and 1152 digital speech processors (DSPs). In order to provide optimum voice quality, G.711 encoding will be used without silence suppression. Speech will be encoded into IP packets every 20 ms and transmitted over the IP network using the Internet Engineering Task Force (IETF) real-time protocol (RTP). Three 100BaseT Ethernet interfaces are provided to connect into the core IP network. The gateway also provides the ability, under the management of the call control, to switch 64 kbit/s data calls between SCN circuits without routing the call via the IP network.

The CVX1800 communicates with the call control entity via a pre-standard form of the media gateway control protocol (MGCP). A modified version of MGCP has now been adopted by the International Telecommunications Union (ITU) and IETF, and is known as *H.248*.

Media gateway controller

The media gateway controller (MGC) provides the call control and connection control functions. The controller receives and sends call control messaging from the SCN and other elements within the IP telephony network. It communicates with the media gateways in order to set-up and take down connections, and connect tones and announcements.

The Nortel IPConnect Call Engine (ICE) performs this function. ICE is a software package that runs on

commercial high-availability Sun server pairs providing load sharing across a distributed network. Protection against failure is provided using redundancy of each ICE across two physical computing platforms. Checkpoints are provided between the primary and secondary instances of an ICE so that, under failure conditions, established calls are not lost.

The ICE receives integrated services user part (ISUP) signalling from the signalling gateway and Digital Signalling System No. 1 (DSS1) signalling from the media gateway. For BT Spain, the Spanish variants of these signalling systems are supported. The ICE is able to interpret, modify and act on the messages in order to set up and take down connections across the IP telephony network and transmit the appropriate signalling on the terminating side of the call.

Call control signalling between ICEs is provided by a Nortel proprietary interface known as the *Inter Soft Switch Protocol* (ISSP). This protocol transparently carries ISUP and DSS1 signalling. Each ICE involved in a call generates a call detail record (CDR) and all records for a call will contain a unique call identity to enable records to be correlated off-line. CDRs are stored on the hard disks of the Sun machines and will be transferred to the billing mediation processor via the file transfer protocol.

The network supports a number of services during the initial phase by distributing the intelligence across the ICEs. The ICE runs an intelligent network (IN) like call model and accesses on-board service logic to support the services. The following services will be supported:

- indirect access service,
- direct access service,
- inbound services,
- virtual private network service, and

- local number portability.

ICE provisioning is performed from a centralised data distributor. Network, service and customer data is entered into the distributor either via its graphical user interface (GUI) or via a bulk download interface and then distributed to each ICE.

Signalling gateway

The signalling gateway provides the transfer of signalling streams between the IP domain and the SCN domain. The gateway provides a means to transfer the call control protocol to the media gateway controller over the IP network.

The SS7 signalling gateway is based on Nortel Network's broadband STP signalling server technology and is known as the *universal signalling processor* (USP).

The USP supports 88 signalling links on the SCN side and four 100BaseT Ethernet interfaces on the IP network side. On the SCN side the gateway terminates MTP2 and MTP3. The gateway has an IP path established with each ICE instance; two paths are provided to each ICE. The gateway identifies the destination ICE for a message by examining the originating and destination point codes. The ISUP call control messages are transported to the appropriate ICE via a signalling gateway to ICE protocol. This protocol provides a reliable transport mechanism using a pre-standard version of reliable user datagram protocol (R-UDP) and a management and maintenance protocol that provides similar functionality to SS7 MTP3. The signalling gateway has its own GUI-based management system.

Audio server

Tones and announcements can be provided by the media gateway; however, to provide greater flexibility and capacity an audio server will be introduced. The IPConnect Audio Server will also work in conjunction with the ICE to provide an intelligent

Figure 2—Support systems for the IP telephony network

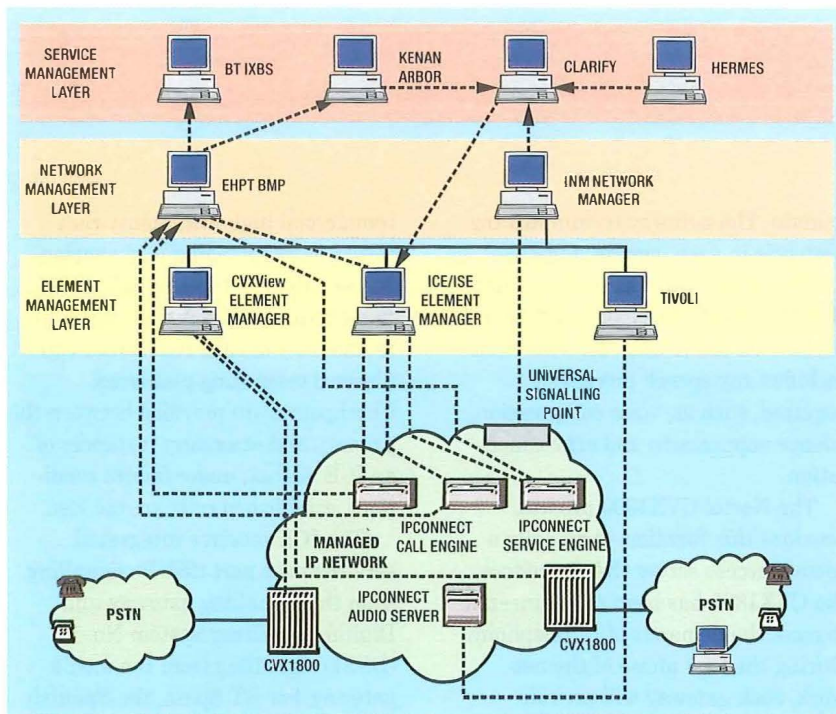
peripheral function. The audio server will be controlled using an MGCP interface from an ICE. The audio server will be used to provide announcements in the four languages used in Spain based on customer preference.

Support systems

Essential to any network are the support systems. Each IP telephony element has its own management systems, and BT Spain already has a number of systems to support the existing network and services. BT Spain will use the Nortel integrated network management (INM) system to integrate the various element managers into a single system and to provide interfaces between legacy systems and the new network. Figure 2 shows the main systems and the interfaces to the network.

Proposed network design—second phase

Figure 3 shows the network at the second phase. There are two main additions to the network at the second phase of the project. Firstly, an application server, known as the *IPConnect Service Engine* (ISE), is added. The ISE will enable more advanced services to be deployed in the network, such as that involving user interaction and call event monitoring. Secondly, the support of additional access types will be added; these will take the form of customer-located gateways, PC clients and soft PBXs. These new access types will use the ITU H.323 standard signalling protocols to set up calls, and hence the H.323 gatekeeper functionality will be added to the network. It is also possible during this phase that the network will also support clients using the IETF session initiated protocol (SIP). During this phase these new access types will have access to the services provided by the network, but it is also envisaged that they will add a new, multimedia, dimension to those services. The design and marketing teams are currently working to define the new services for this phase of the project.



Design Challenges

The introduction of this new technology and the nature of the underlying IP transport network present several significant challenges to the design team; the main areas are highlighted below.

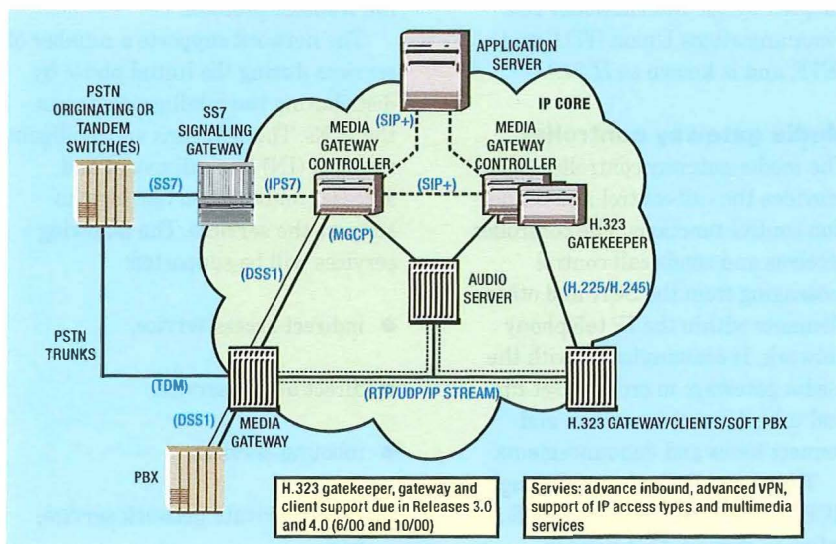
Quality of service and impacts on the IP network

Probably the biggest challenge of all is to reduce delay and delay variation to that of the existing public network. The variable delays introduced by the gateway as a result of encoding/decoding and jitter buffers plus the delay added by the IP core network will take up a significant portion of

the 150 ms delay budget allowed for by ITU recommendations. The design target is that 95% of national calls will experience less than 125 ms end-to-end delay. Keeping this delay to a minimum is critical, and initial testing has indicated that this target will be met.

The IP core network is being provided by Cisco Systems. The CVX1800 provides the ability to set the type-of-service byte in the IP header for a given type of data stream and this will be used by the routers to form the basis of an elementary quality of service (QoS) policy within the IP core. This will become an important issue as the volume of data traffic on the IP core network increases.

Figure 3—BT Spain IP telephony network—phase 2 (3Q2000–4Q2000)



In the longer-term, more sophisticated QoS mechanisms are being standardised and tested in the labs. Diff Serv and MPLS are the two most favoured, although RSVP offers the intriguing prospect of a connection-orientated mode for IP which would enable bandwidth and levels of service to be guaranteed for the duration of a session, something that would be very suited to voice services. However, the consequent processing overhead in the routers is high.

Although logically independent from the underlying IP network, the IP telephony service imposes heavy demands upon it. Delay and delay variance are issues that are well understood, but the impact of a traffic mix in which large data packets can severely delay small voice packets has yet to be understood in a real-life environment, especially on the lower-speed access circuits. IP was never designed for real-time applications and the typical reaction of a router to congestion is to simply discard packets. For data applications, lost packets can be re-transmitted, but for voice there is no time to do this, so the lost packets will be evident as a short break in transmission. The resulting requirement is for a network that avoids congestion situations, or implements QoS. Because of the plentiful availability of bandwidth in the new BT Spain transport network, these problems can be temporarily avoided by the time-honoured practice of increasing the bandwidth. However, the aim of this project is to demonstrate that IP telephony and data services can co-exist economically on a common IP network, and so this can only be a short-term solution.

A perhaps less obvious issue is the fact that the incoming speech is broken down into short bursts to be packetised and sent across the IP network. For example, 20 ms packets in a fully loaded gateway supporting 20 E1s can generate 75 000 packets per second. The resulting voice traffic

profile to be supported by the IP network could be very different from normal data traffic with the packet throughput of the router becoming the potential bottleneck rather than the bandwidth.

Speech quality

BT Spain's experience with voice-over-frame shows that speech quality is the most important issue for users. The choice of codecs and the quality of their inherent algorithms contribute much to the quality of the final connection—the sensitivity of the codec to background noise being a particularly difficult problem to resolve. With the delay inherent in IP-telephony solutions, echo is another potentially serious issue. Echo cancellation is a highly specialised field, and the characteristic delay variation, or jitter, in this type of network adds further challenges.

Control over signal levels is another area of novelty for the data network designer contemplating IP telephony, which has long been solved in the speech world. When end-to-end connections include mobile and international networks, the issues of achieving good-quality voice connections are magnified still further.

Service support

Services and functionality in existing public networks have been built up steadily over many years, and there is an expectation that the next-generation networks will very quickly provide the same services as well as more enhanced services. Although this will undoubtedly be true in the future, these are early days, and many functions taken for granted today do not necessarily yet exist in the new networks and will take time to develop. The market requirements for multimedia services are still in their infancy. These services will enter a new realm of true integration between network and customer equipment not normally seen in voice networks.

Resilience and robustness

The IP telephony network can be considered as an exploded switch architecture with the elements separated by an IP network. This gives rise to a number of issues relating to the resilience of the elements, the handling of overload and congestion in the elements and network management. While redundancy is built into the network elements, the means to allow operational monitoring of the network is essential. The real-time packet-loss-sensitive nature of voice means that the IP core network must be tightly managed. A mechanism to dynamically inform the call control of problems in the IP core network does not yet exist but should be the long-term goal.

Migration

BT Spain's business needs have driven the implementation of an interim network that will support some services prior to the launch of IP telephony. This network is based on small programmable voice switches interconnected via a DMS-100 international gateway. Migration of customers and services between two networks is never an easy task. Initially the IP telephony network will provide a transit capability for the existing BT Spain network including number portability. Once this is in place the other services will be migrated as they become available on the new network. This task will be a major challenge to the network, systems and operational teams, and is a key factor in the future success of the solution in other BT alliance partners where there are established voice networks.

Figure 4 shows how the two parts of the network will initially be interconnected.

Once customers and services have been migrated onto the IP telephony network and the CVX1800 supports both dial access for Internet service providers and IP telephony functionality, the interconnect switches can be

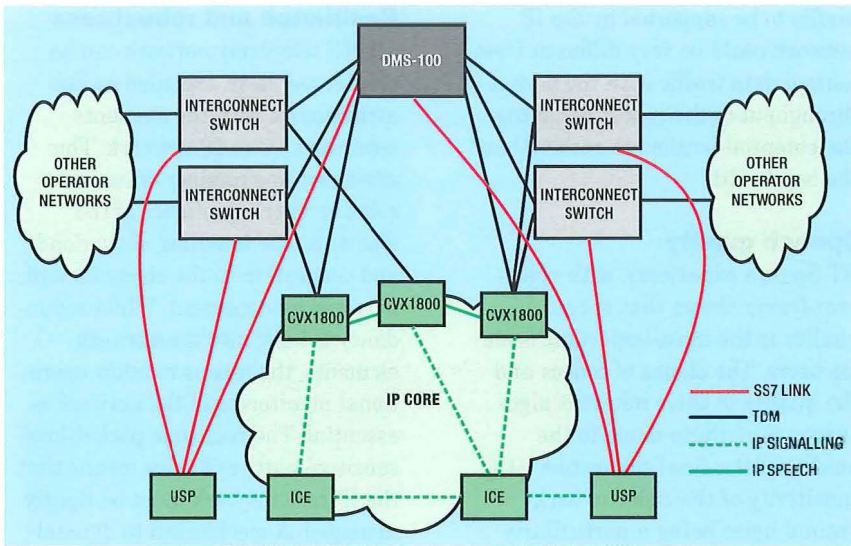


Figure 4—Interconnect between the BT Spain interim and IP telephony networks

recovered. The DMS-100 will remain as an international gateway switch.

Data calls

With the high quality of the G.711 codecs, there is confidence that fax and modem traffic up to 33 kbit/s can be carried in-band by automatically disabling the echo-cancellors and modifying the jitter buffers. However, 64 kbit/s data calls and ISDN traffic pose a greater problem, and the current approach is to switch these at the gateways to a parallel TDM network under the control of the MGC.

Conclusions

Organisationally, this hybrid technology necessitates a re-think of the traditional divisions between data and voice engineering, and the tender exercise highlighted the problems that companies are having in breaking down internal barriers. The alternative is to embark upon a rapid program of re-training or acquisition.

The immaturity of the IP standards is revealed by the stringent demands of real-time carrier-grade services and it remains to be seen if the protocol can be enhanced rapidly enough and far enough to meet the expectations being placed on it. The runaway success of IP in the Internet

does not necessarily mean that it can become the universal protocol for all services. Only time will tell if the enthusiasm of its proponents is justified.

Nevertheless, if this project proves that the full range of voice services can be economically provided over an IP network that is simultaneously carrying a broad mix of data traffic, then the movement towards the datawave will have taken a huge step forward. The prospect for the not-too-distant future is common service control for voice, data and multimedia services. This opens the door to an environment in which traffic can flow between users under the control of service intelligence residing at a remote location, and radically new multimedia services can be implemented by software upgrade.

Biographies



Steve Brabner
Networks and Information Services, BT UK

Steve Brabner joined BT in 1975 where he has spent most of his career in the development and implementation of data networks and applications. He was responsible for the creation of the Cohort 500 directory product range and Inter-View, the BT internal directory system. Between 1994 and 1998 he was seconded to Concert in the USA where he was responsible for non-voice services strategy and design, in particular the Frame Forward project to re-architect the Concert Frame Relay Service. Since returning to the UK he has been a team leader in Alliance Engineering in Advanced Communications Engineering (ACE) managing the delivery of IP services to BT's alliance partners.



Clive Dellard
Networks and Information Services, BT UK

Clive Dellard joined BT as a Trainee Technician (Apprentice) in 1973. He joined the System X Trunk Network System Design Authority team in 1986 and following this he ran the Network Coherence model in London before moving to the Advanced Service Unit design team. There he played a key role in the design and implementation of the FeatureNet network. Since 1997 he has been a technical consultant providing network and service design support to BT's joint ventures including those in Switzerland, New Zealand and Holland. Now, as a member of Mobility and Network Services in Advanced Communications Engineering, he is leading a number of aspects of the design for the BT Spain IP telephony project.

Jean-Marc Salles

Global Convergence— The Challenge of the 21st Century: Tele Danmark Experience

The rapid growth in mobile and Internet traffic is driving forward global convergence in many spheres, such as: technology, industry structures, regulatory frameworks and the market itself. This article reviews the trends from Tele Danmark's viewpoint, and outlines their Duet product: a fixed-mobile service with common billing, which has proved to be a market success.

Convergence of Technologies

The telecommunications market (equipment and services) is estimated to be worth in excess of US\$ 600 billion per year, to have a 7% annual growth rate and to represent on average 4% of GDP. To support this effort and the telecommunications industry in general, it is also usually assumed that the total standardisation effort is approximately 0.5% of the market value (US\$ 3 billion per year).

Most markets in the world are increasingly dominated by a small number of large, global players, often by consortia of former market leaders. This is particularly true for the telecommunications market, and it is not expected to change in the foreseeable future as indicated by recent mergers. Economies of scale are increasingly being perceived as only really possible and efficient in markets with global dimensions.

Fixed-mobile

Some would say that fixed-mobile convergence is no more than call-forwarding backed by a cutting-edge billing system. Others that it spells the death of copper loops as the means of termination in the PSTN.

Which of the two scenarios—if either—should be backed? Probably the cautious one. Cellular companies which predict the demise of fixed-line

telephony tend to end up looking very silly indeed. They ignore the fact that increasing amounts of information and entertainment are being sent down the fixed 'pipe' to the home. It is fanciful—and it will still be fanciful in five years' time—to pretend that wireless terminals will usurp this role. In fact, wireless is a very effective conduit for the delivery of voice telephony, but quite nonsensical when it comes to almost all other applications.

It is not hard to spot where competition is going to hurt. Traditional PTTs are seeing their fixed-line revenues slashed by the ingress of competition. Mobile services and the Internet provide the healthiest streams of revenue for the old-style telecommunications operators. The dream scenario is that cellular usage and Internet access will converge. The nightmare scenario for traditional telecommunications operators is losing their foothold in both markets. For operators not owning the local-loop infrastructure, offering converged fixed-mobile services is an uphill struggle.

Therein lies the contradiction and conflict behind the easy and eye-popping concept of fixed-mobile convergence. Incumbents see this convergence as a safeguard for copper loop revenues. Competitors see it as the way to break the stranglehold of traditional telecommunications operators local loop dominance. Telia, for example, owns a fixed and mobile infrastructure

throughout the whole Nordic region and has already announced its commitment to fixed-mobile convergence by reintegrating its cellular arm back into the parent company in 1996 to counter for this threat.

On the regulatory front, this is bound to lead to confusion. It is hard enough to order an incumbent PTO to unbundle the local loop and provide cost-based interconnection—a prerequisite for any fixed-mobile offering. Wireless operators, too, should be prepared to open their networks for resale. They are not because the margins on which they have grown up are now firmly entrenched in the industry's mindset.

In short, the technological platform required to provide an integrated fixed and mobile service is composed of a mobile licence, an intelligent network (IN) infrastructure which recognises home zones, and a powerful billing system. The second and third of these requirements can nowadays be bought off the shelf. The first one is a regulatory issue difficult to resolve without introducing discrimination per se, and the experience of issuing licences for GSM-900 and GSM-1800 does not presage intelligible regulation.

Third-generation (3G) and Universal Mobile Telecommunications System (UMTS) licensing offer regulators and politicians another chance to repeat complex licencing practices. As long as UMTS expectations are not too high, this will not make much difference. If fixed-mobile convergence actually means anything, the industry and its customers would be better off with full competition without any regulation.

Telecommunications and IT

One of the main drivers of growth in the telecommunications market is cellular mobile communication. It now accounts for 20% of the overall telecommunications market. Significantly, in 1997, the mobile services market was almost three times the size of the international telecommunications services market. Yet, even that

figure understates the impact of mobile communication because it underestimates the stimulation created for the fixed network. In other words, calls between fixed and mobile networks are generating considerable revenues for public telecommunications operators (PTOs).

Internet access is also generating large new revenue streams for PTOs in traditional and new service areas. For PTOs with measured local service, the Internet is boosting revenues through calls of longer duration. Demand for second residential lines, for higher-speed access services, for ISDN, for leased lines (local, national and international) have all increased at unprecedented rates due to the Internet. PTOs are also entering new markets, pioneered by ISPs, such as Internet access, hosting content for the World Wide Web and in managing private intranets.

The demand for Internet services provides an important example of how, where demand is buoyant and investment opportunities are unrestricted by regulation, technologies can rapidly evolve and diffuse and markets can grow. Internet deployment in the United States has led to over 40 major backbone networks, over 4500 ISPs, and 23% of households connected to the Internet. It is predicted that the Internet access market in the United States will increase from US\$6 billion in 1997 to US\$38 billion by 2002. Although bandwidth is growing, it cannot always keep up with demand growth. At the same time, however, while the question of bandwidth shortage is being discussed in the United States, it has been estimated that about 60% of installed fibre optic network capacity is not being used but is being held in reserve by the large telecommunication operators.

Interchangeability between communication networks, as well as features of complementarity, implies a policy allowing bandwidth expansion and technological innovation in infrastructures to take place in a competitive environment which is

commercially driven. The need is less policy influence, more emphasis on ensuring that obstacles to demand are removed, that diffusion of terminal equipment and services can take place rapidly and that convergence is facilitated. Many of the following 'technology drivers' are still developing fast:

- *Digital wireless access technologies—not only radio* Wireless access technologies must interwork with other access mechanisms, including cable modems, digital subscriber line (DSL), integrated service digital network (ISDN) and premises wiring ('smart' house or office). Direct opto-electronic access is also foreseen shortly. Technologies are being brought to bear on the voice and data networks which can bring about the convergence of these services onto single high-speed access and trunk networks. Voice over DSL (VoDSL) is emerging as a service likely to see deployment in 1999. Service providers are now trying to identify their VoDSL strategies.
- *Data compression techniques* As the demand for all forms of multimedia keeps growing, these techniques will increasingly be used to overcome the natural limitations of transport networks; for example, JPEG2000 will allow greater compression with higher image quality.
- *Optical switching* Optical switching cannot yet replace electronic switching. The need for opto-electronic conversion for switching and transmission purposes is growing and, in the near future, PCs will be directly connected to the optical stream.
- *Processing power* Trends of growth in processor power, memory size, etc. are expected to continue, enabling powerful multimedia applications to be brought to the desktop rather than workstations.

Figure 1 – The shift in the ICT value chain

- **Digital speech processing** The impact for users of digital speech processing will continue as long as bandwidth is a limitation and a resource to be paid for in proportion to usage.
- **Choice of technology** The choice of technology becomes tactical and strategic. Although Minitel was once seen as an adequate tool, it delayed the adoption of the Internet in France, which is now having to catch up.

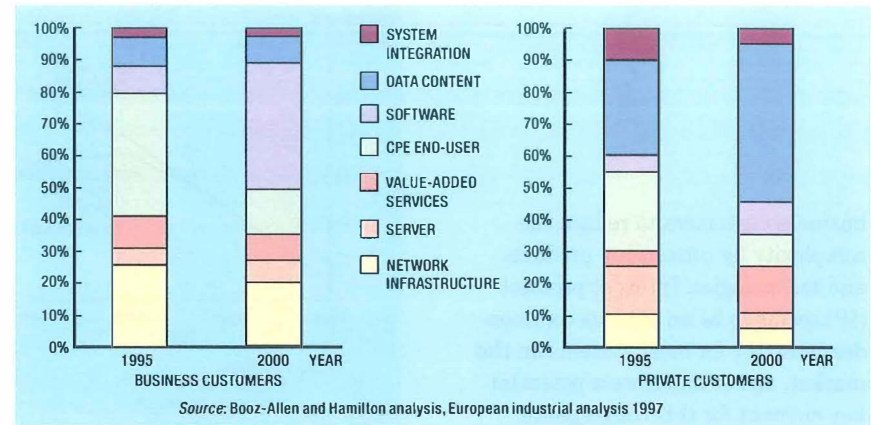
Converging (Politically Defined) Industry Structures

The traditional industries of the telecommunications sector have long been politically defined. Until the mid-1980s, for example, the supply industry for telephones and PBXs was largely dependent on the monopoly regimes that were in place in most European countries. As a result, there were neither physical nor functional convergence between the regulated telecommunications equipment (telephones, modems, PBXs, etc.) and the mostly deregulated computer and data equipment industries (computers, routers, etc.).

While new technologies have undoubtedly played a role in changing communication markets, many of these developments have been brought about by regulatory reform. The emergence of the Internet is a good example of a service enabled by liberalisation. In the early 1980s many public telecommunication operators introduced online information services and packet-switched data networks. Yet the idea, current at that time, of connecting the world's computer hosts and databases did not take off until the advent of the Internet and the World Wide Web.

Convergence in regulations

The historical regulatory barriers to the earlier emergence of a public Internet are worth recounting. The monopolies granted to PTOs often prevented other value-added suppliers



from offering information services. Even where value-added suppliers could provide service, they were deterred by PTOs, who would limit the supply of bandwidth available, charge high prices for leased lines, and restrict customers' freedom to attach the equipment of their choice to the PSTN and communicate with other users, etc.

The infrastructure and service provision—transport as well as application—must account for a better data accessibility. The convergence of technologies and industries leads to network-independent service platforms but increases the need for quality-of-service management. The need for standardising services may decrease, but standards are needed for the mechanisms that support them. As the tools and processes to customise services to user's individual needs and skills, services must however remain consistent and compatible, and therefore demand their own new standards. Future architectures are likely to exhibit a clear separation between the *service layer* where the access, operation and the service components are the three basic functions and the *connectivity network layer*. They are the basis for end-to-end service management. Finally, intellectual property rights (IPRs) are now embedded in products and services, and often declared too late or without licensing information. This can jeopardise or waste standardisation efforts.

Converging telecommunications and ICT

However, to achieve convergence is not simply a question of advanced technology. It is also, and not the least, the convergence of two radically different cultures, of two different worlds where the players have

different ways of thinking and often have opposing values. When a telephone is taken off hook, the user normally expects to hear dial tone. If there is no dial tone, the user would conclude that the line is unplugged or out of order. Normally, however, nobody is worried if a computer cannot connect to the Internet at the first attempt. The user would try again until a connection can be established and would tolerate the connection even if it is extremely slow.

The reason for accepting a lower quality of service on a data network is its great flexibility and the possibility of transferring large data files at low price. The tolerance threshold for a telephone user is much lower because the telephone network has for more than a century been optimised to almost 100% up-time and full reliability. A telephone user expects the network to be connectable and reliable at almost any cost.

From a market point of view, the relationship between end-user and service provider is also different in the telecommunications and data worlds. Telecommunications operators normally have contract-like conditions for service and the quality of service is clearly specified. Data network operators can normally only promise best effort. 90% of business users today have fully separated data and telecommunications networks. It is even not unusual to have several more or less incompatible networks (Ethernet, asynchronous transfer mode (ATM), routers, etc.). In addition, a clear shift of focus from terminal equipment to software for business customers and from equipment and infrastructure to data content and value-adding for private customer is identified (see Figure 1). It is absolutely necessary for most

business customers to reduce the complexity by converging products and technologies. Internet protocol (IP) seems to be an obvious common denominator for most systems on the market, and is therefore a potential key element for this convergence.

To be successful in providing converged communications services, service providers must be able to combine the best of both worlds: up-time, connectivity, reliability and high quality of service at the much lower Internet rates. For example, after a successful merging with Northern Telecom and Bay Networks, Nortel Networks has completely re-structured its sales, marketing and support functions into one organisation to compete more efficiently against network operators like 3Com and Cisco and manufacturers like Lucent, Alcatel and Siemens. This new strategy is called *Unified Networks*.

Converging industry structures

The changing structure and growth of international undersea cable capacity gives a good indication of the rapid pace of change and the new operating environment faced by PTOs. This is, in particular, noticeable to and from the United States. In 1996, the number of active circuits in use between the United States and the rest of the world increased by 50% from the previous year. This growth was not, of course, primarily driven by PSTN traffic.

By way of contrast, the number of active circuits allocated to other uses, such as Internet backbones and refile, grew by 207%. As a result, some 59% of the active circuits between the United States and the rest of the world were used for the International Message Telephone Service (IMTS) at the end of 1996, compared to 80% a year earlier. In a single year the number of active circuits used for IMTS between the United States and the Netherlands decreased from 64.5% to 35.4%. The growth in intercontinental capacity, stimulated

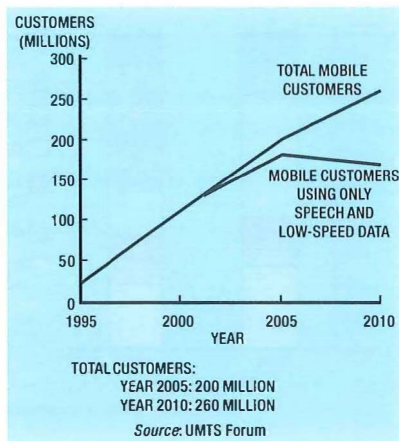


Figure 2—Mobile customer growth in the European Union

by Internet traffic, has now already resulted in capacity allocated to Internet traffic exceeding voice capacity on some routes (for example, between Japan and the United States).

The new balance of international transmission capacity in use reflects the change in underlying industry structures. Revenues have increased for those PTOs that adapt quickly to the new environment. For example, the size of the international leased line market between the United States and the rest of the world increased by 17% in 1995, by some 28% in 1996 and by a further 29% in 1997. Yet, at a time when the number of cable and terrestrial circuits active between the United States and the rest of the OECD area increased by 75%, the number of active satellite circuits decreased by 8%. This undoubtedly reflects a judgement by the market on the benefits of different technologies for high-volume traffic routes. It might also reflect the fact that the market structures evolving for undersea and terrestrial cable have adapted more rapidly than the traditional structures surrounding the provision of satellite capacity.

Market Convergence

Figure 2 shows the continuing growth in the number of mobile customers in the 15 EU countries. By 2010, it is expected that around 65% of mobile customers will expect to use mobile communications for speech or low-speed data including electronic commerce, while about 35% of them will demand the ability to access multimedia information via high-

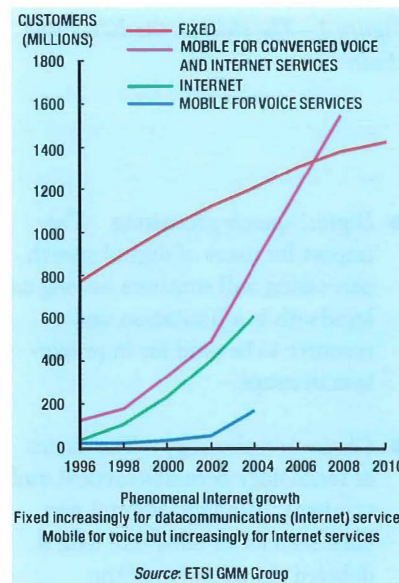


Figure 3—Global customer growth

speed data access. The take-off and growth of high-speed mobile data services will depend on the rate of introduction of 3G (UMTS) networks and the cost of high-feature user terminals.

Figure 3 is derived from multiple sources and was presented at the last General Assembly of the European Telecommunications Standards Institute (ETSI) to describe the expected growth in mobile and Internet services. Although the source figures differ somewhat, the trends are quite clear: that continued strong growth in mobile, Internet and mobile Internet, with an eventual ‘cross-over’ between fixed and mobile can be foreseen. Mobile figures have been consistently underestimated.

Mobile terminals are becoming consumer products, and many of the rules of the world’s consumer markets are leading to the same consumer expectations for the telecommunications market. For developing countries it is cheaper to bypass the fixed access network by installing wireless networks, whereby ‘instant’ deployment of services becomes feasible.

Profile mobility—the example of Duét

The Scandinavian countries with their high fixed and cellular penetration rates have become the pioneers of fixed-mobile convergence. The typical converged service either being trialled or already commercially offered works around a similar

Figure 4—Basic Duét concept

principle: the re-routing to a mobile of incoming calls to a fixed telephone number, and the re-routing to a fixed or mobile voice mailbox if the mobile is busy or non-reachable. A combined bill is then produced for fixed and mobile.

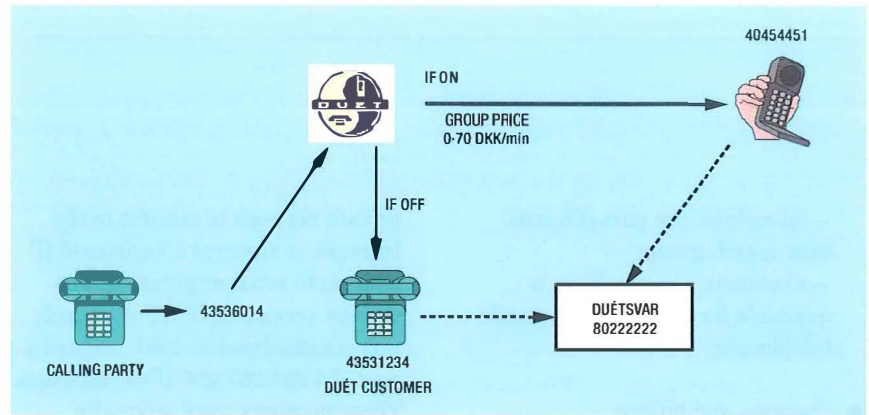
While this may appear a logical combination of fixed and mobile, a significant regulatory and strategic undertow has threatened to bring the launch of such services to a grinding halt. This has been particularly the case in Denmark, where Tele Danmark—the dominant Danish operator—has launched a fixed-mobile service branded *Duét*. This service relies on intelligent networks to combine its fixed and mobile network, and bills its subscribers for the combined service—with a discount. Tele Danmark has launched the service simultaneously in Belgium, where it has a large stake in the dominant national operator, Belgacom.

However, both in Denmark and in Belgium, the service came under the scrutiny of the European Commission in June 1998 when Mobilix—the newest entrant in the Danish cellular market—alleged that Tele Danmark was pursuing a predatory pricing policy. Mobilix, which is also aiming to launch fixed-mobile converged services through a link with the Danish railways infrastructure claims that Tele Danmark has an unfair advantage through its ownership of Denmark's local loop. Mobilix has called for the unbundling of the local loop and for fixed and mobile services to be billed separately.

Basically, Duét is a call-routing system. It can be activated when the mobile is on or off (see Figure 4).

The basic Duét service is being enhanced with additional features and by making the activation/de-activation procedures more user-friendly. Two major enhancements are the *Duét Family* service and the *Duét Business* service.

Duét Family (Figure 5) provides basic Duét with the following additional features:



- multi-user for up to five mobiles, as part of a common search group;
 - availability both on PSTN and on ISDN2 (basic access);
 - a common voice mailbox for the fixed and mobile telephones;
 - an individual voice mailbox per mobile for calls directly to the mobile telephones;
 - a total subscription discount per mobile of DKK 400; and
 - calls from the fixed network, mobiles and DuétSvar priced at a special Duét rate.
- Duét Business (Figure 6) provides in addition to the features of Duét Family the following supplementary features:
- call forwarding, if the mobile phone is switched on, to the following numbers:
 - direct dial in (DDI) numbers only (PSTN/ISDN2/ISDN30);

Figure 5—The Duét Family service offering


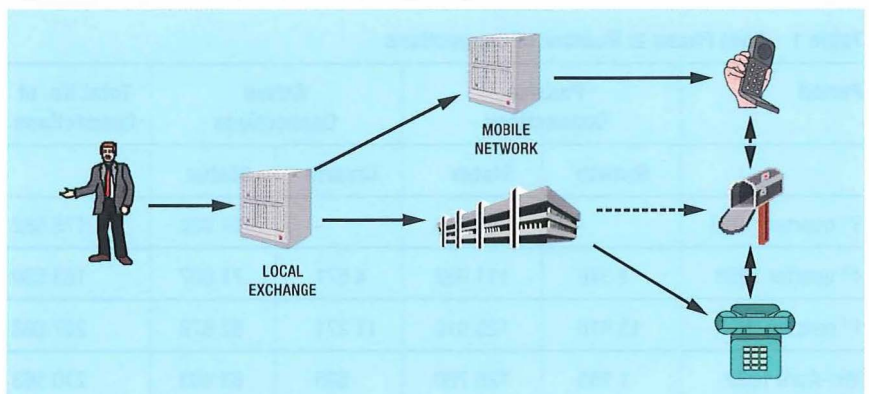


Figure 6—The Duét Business service offering



- all mobiles are part of a common search group;
- a common voice mailbox is available for the fixed and mobile telephones;
- charging and billing:
 - single billing;
 - advanced IN supervision on calls;
 - normal charging of calls between fixed and mobile Duét telephones;
- Duét control:
 - one single access number (80 22 22);
 - only authorised access (controlled by A-number on national calls, by PIN-code on international calls);
- call transfer from the mobile telephone to any other number.

The introduction of Duét is considered to be a market success (see Table 1). By the end of 1998, almost 200 000 Duét customers were registered (almost 10% of the total number of fixed access lines and 25% of the number of cellular customers). The growth rate is currently almost 20% per year.

The addressing problem of the Internet

On a stand-alone private TCP/IP network, IP addresses can be assigned at random as long as each one is unique to that network. For that

private network to connect to the Internet, it requires a registered IP address to avoid duplication. The current version of IP (IPv4), which was standardised in 1981, created a pool of 4 294 967 296 IPv4 addresses. These numbers were originally assigned under the three classes A, B and C. However, as the Internet expanded, concern arose that the existing numbers would be exhausted and that the size of the global routing tables was in danger of growing faster than the capabilities of the underlying equipment.

To address these concerns, the Internet Engineering Task Force (IETF) introduced classless inter-domain routing (CIDR), a new IP-addressing scheme that replaces the older system. CIDR enables more efficient allocation of the IPv4 address space, allowing for the continued growth of the Internet until a new numbering system (IPv6) is deployed. It is projected that IPv6, which will create a virtually unlimited resource of IP numbers, will be increasingly used from the year 2000, and play a significant role by around 2003 to 2005. Before that time, it is envisaged that IPv4 allocations and assignments will be made bearing in mind the finite nature of the existing resource and the need to minimise growth in the size of the routing table.

Originally, the Internet Assigned Numbers Authority (IANA) allocated blocks of IP addresses directly. As a general rule, end-users receive IP assignments from their ISP. Notwithstanding, some ISPs receive

addresses from upstream or backbone ISPs, and some end-users receive allocations directly from a registry or directly from the IANA. In Europe, telecommunication carriers, or alliances between carriers, hold the majority of IP addresses. This is due to telecommunication carriers emerging as some of the largest ISPs in their own right, and by taking over the largest independent ISPs. Several university networks retain large IP address blocks and some government agencies have large allocations, such as the National Health Service in the United Kingdom.

UMTS

The original GMM model

The original ETSI Global Multimedia Mobility (GMM) model supports a multi-service environment and convergence aspects. It has been widely accepted (in ETSI, ITU and elsewhere) and is, for instance, the basic architecture for UMTS. The basic elements of the GMM model are:

- independence within the terminal equipment domain (both for fixed and mobile users);
- a multiple access domain (accommodating both for public and private network solutions); and
- a core domain containing different kinds of service networks (including network intelligence needed to support specific service offerings in the core).

From that perspective, the model encompasses fixed-mobile convergence (FMC), since UTRAN is defined to include the ability to connect to any type of feasible core network.

In addition to the three basic domains, the model contains a fourth domain related to applications, content and end-user services. Internet services using IP as

Table 1 Duét Phase 2: Number of Connections

Period	Passive Connections		Active Connections		Total No. of Connections
	Growth	Status	Growth	Status	
3 rd quarter 1998	-	108 646	-	67 036	175 682
4 th quarter 1998	3 346	111 992	4 571	71 607	183 599
1 st quarter 1999	13 018	125 010	11 271	82 878	207 888
Mid-April 1999	1 750	126 760	525	83 403	210 163

Figure 7—The refined GMM model

transport mechanism independent of the underlying infrastructure being increasingly used to support the application domain, this relation will be of less importance in the future. Users are not concerned with the infrastructure mechanisms supporting a particular service or application.

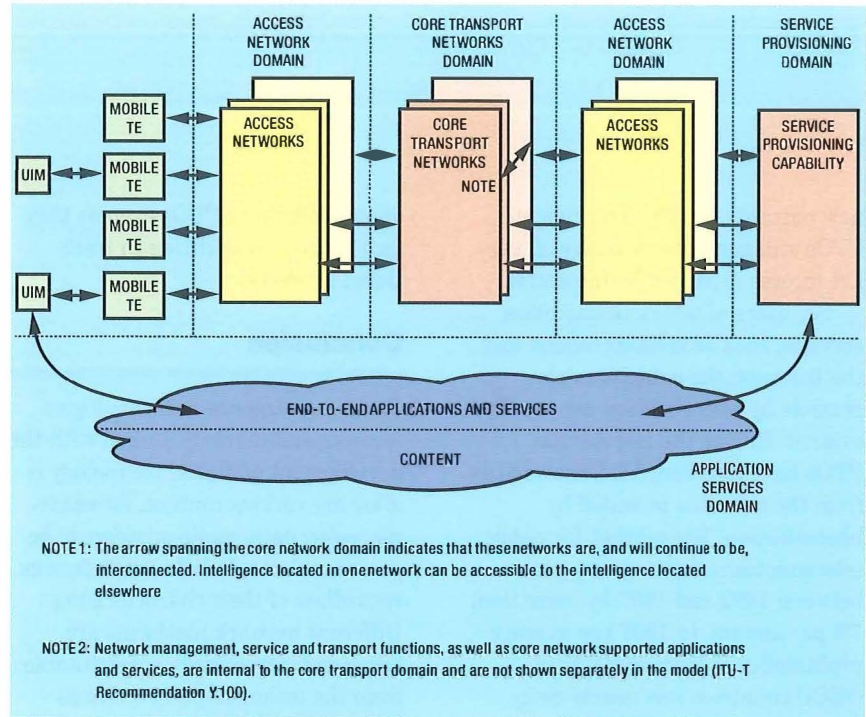
The refined GMM model

See Figure 7. Service provisioning, traditionally embedded in the terminal domain, needs increasingly to be recognised as such. The core network must thus be unfolded to reveal the role of the access domain in providing connection with the ‘far end’. The initiation of a communication need not any longer be from the left-hand side of the model. The far end may be a service, and the initiation of a communication could take place from the right-hand side of the model. Indeed, initiation could be from almost any part of the model where the necessary ‘intelligence’ is located. Various scenarios can be constructed, based on this model, such as communication between two service provider domains. A clear separation of the applications and services elements from the basic connectivity components appears. Network connectivity is transparent to the different types of information carried.

The ‘intelligence’ and the connectivity may still reside in the same equipment. The difference is that they need not, and that each can be substituted independently as networks and applications evolve at different rates. The decoupling of elements also allows independent service providers and content providers to offer their services without owning the infrastructure.

Convergence of Cultures

In 1997, there were 64 PTOs with revenues greater than US\$1 billion. This was an increase over 1995 when there were only 52 PTOs with the same characteristic. This number was already up from 42 PTOs in



1992. The growth in the number of PTOs with revenues greater than US\$1 billion is mainly due to liberalisation allowing new players to enter fixed and mobile markets. In 1992, there were only two cellular mobile companies independent of traditional PTOs, McCaw and Vodafone, with revenues over US\$1 billion. By 1995, there were four and by 1997, 10. McCaw was taken over by AT&T, but Air Touch was split off from Pacific Telesis and 360 Communications from Sprint to form separate companies.

Significantly, the concentration of revenues among the leading PTOs is decreasing. In 1992, the top 25 PTOs accounted for 92.4% of the revenues of PTOs with revenue over US\$1 billion. By 1995, the top 25 PTOs’ share of this total had dropped to 87.6%, while in 1997 their share was down to 84.4%. New entrants, particularly in cellular mobile communication, are partly responsible for this growth in the overall size of the market. In 1997 there were 25 PTOs with mobile revenues greater than US\$1 billion. NTT (DoCoMo) is by far the largest cellular mobile communications company in the world, with revenues of US\$22.8 billion in 1997.

Liberalisation

There is a significant new market entry in market areas where the liberal climate is adequate. Still,

there is an increasing number of mergers between leading PTOs across geographical boundaries. With the global networked economy, corporate culture is therefore also changing rapidly towards short-term ‘shareholder’ value—as pointed out in the original report of the ETSI GMM Group.

New entrants have captured 45% of the total market in mobile communication. In terms of the overall telecommunication market, they have gained a market share of just under 19%. With their small percentage of access lines relative to their share of mobile subscribers, most of their revenues are drawn from cellular mobile and long-distance services rather than from local access services. In 1997, new entrants provided less than 0.9% of the access lines. This number was up from 0.3% in 1995 but it is still primarily made up of new entrants in the United Kingdom.

In 1997 the assets of the telecommunications services sector surpassed 1 trillion dollars. The fixed assets, of which telecommunication networks make up the vast bulk, was around US\$700 billion. Noteworthy developments include the increasing share of capital expenditures and employment provided by new entrants. At the same time, the initial costs of entering telecommunication markets are considerable. This is evident in the overall loss made by

new entrants in 1997. Incumbent PTOs still generate virtually all the net income produced by the sector.

For users of telecommunication services, such as cellular mobile and the Internet, the benefits made possible by liberalisation are readily evident. During the last decade, PTOs have also benefited enormously from the stimulus provided by liberalisation. The market for public telecommunication services grew between 1992 and 1997 by more than 7% per annum. In 1997 the average profitability of leading PTOs in OECD countries was nearly twice that of the average for *Fortune Magazine's* largest global 500 companies. Revenues in the broadcasting market reached US\$145 billion in 1997. Growth in this sector has been more modest than in the telecommunication market, with an annual growth rate of 3% per annum for broadcasters since 1995.

Commercial versus technical

The new pricing structures were pioneered by Internet service providers (ISPs) and have enabled far greater access to information. Consider, for example, that in 1975 there were around one million interactive bibliographic searches performed in the United States. In December 1997, Internet users viewed an average of approximately 65 million Web pages per day from Yahoo! alone. By September 1998 this had grown to 144 million pages per day. Inktomi, a company operating Internet search engines, estimated that in December 1997 the World Wide Web contained 200 million pages which it was possible to index.

If developments in the communication sector have taken unexpected directions for PTOs, the outlook is, nevertheless, very promising. While new service suppliers are free to enter telecommunications markets, PTOs are also taking advantage of new business opportunities. The development of new services using communication networks can be expected to open new markets and

opportunities for PTOs even as they face greater competition in traditional markets.

Conclusion

Rapid convergence in technologies, services and markets, linked with the development of digital technology is allowing various content, for example, voice, data, audio or video, to be provided through different networks regardless of their characteristics. Different network platforms are becoming increasingly substitutable from the technical perspective as they attain the ability to carry essentially the same services. Taking advantage of this technical progress, a number of market participants are strategically expanding service provision beyond their traditional sectors through cross-platform and cross-product development. Data services provided via digital broadcasting platforms, on-line services combined with television via digital satellites and cable modems, and web-casting of various audio-visual services, are just some examples of this advancing phenomena.

The process of technological and service convergence in telecommunications, information technology, broadcasting and multimedia industries poses a number of significant regulatory challenges. The reformulation of specific regulations applicable to broadcasting networks and to telephone and data networks is difficult, but necessary. The major challenge to governments and regulators is to stop viewing the different communication network infrastructures as different sectors or markets. The speed of technological and market change is one element of this challenge because regulatory change moves at a slower pace than market requirements and cannot keep pace with market development. Other complications arise from the fact that the 'old' regulatory issues have not all been resolved given the continued—and probably inevitable—dominance of incumbent public

telecommunication operators in many markets, and that the transformation to competitive telecommunication markets is relatively recent in many cases. Thus, 'new' issues arising from convergence are likely to continue to coexist with current issues.

Biography



Jean-Marc Salles
Tele Danmark

Jean-Marc Salles was born in 1948 in Paris, France. After graduating in Polytechnics in 1971, he started his professional career as a consultant in on-line computer systems in 1973. He joined Jutland Telephone in 1976 as traffic engineer and network planner, becoming head of Network Development in 1983. After the merger of the five independent telephone companies into Tele Danmark in 1993, he joined the headquarters of Tele Danmark as deputy director of International Relations and Standardisation. He participated in the standardisation of ISDN, SDH, ATM, UMTS (previously FPLMTS) and others in ITU-T Study Group 13, later SG 18. He represents Tele Danmark on a number of international General Assemblies, Boards and Advisory Groups, including the Board of Governors of EURESCOM and the board of ETSI.

Jean-Marc Salles can be contacted at jsm@tdk.dk

Peter Lisle, Howard Green and Peter Mason

GPRS and the Mobile Data Revolution

The mobile data market is about to undergo a transformation as the primary enablers in network, terminal, and application developments come to fruition simultaneously. This technology push is paralleled by an equally potent market pull as more businesses and individuals become dependent on mobility to gain market advantage and assist in managing their busy lives. Entering the data world with mobility requires answers to some basic but demanding questions in the areas of security, addressing, and a clear evolutionary path for users as networks and terminals evolve. These issues are all addressed by the introduction of GPRS on GSM networks.

Introduction


The mobile data market is poised for explosive growth. In mature wireline networks, data services now employ similar amounts of overall network capacity to voice, and yet in mobile networks data drives only a few percent of network capacity. The reason for this mismatch is straightforward—too many barriers, too few enablers. All of this is about to change, as each and every barrier is overcome and the latent market demand propels forward enablers covering a wide spectrum of mobile data needs.

Latent market demand

Market demand for mobile data is surfacing like an epidemic across all market segments in the mobile marketplace (Figure 1):

- *corporates*—becoming dependent on e-mail, the corporate intranet, and e-commerce as key components of their business processes and yet with more and more itinerant workforces;
- *small/medium businesses*—needing flexibility, speed of

Figure 1—Expected growth trends in circuit switched and packet data users over BT Cellnet's mobile network



response, and access to wide sources of information to enable them to succeed in the dynamic business environment;

- *consumers*—growing in data awareness, and with the youth segment leading new trends of social communication.

Market forecasts from such reputable analysts as Ovum estimate at least a ten-fold growth in the market for mobile data over the next five years—a growth rate that will outstrip voice in the same period. So what are the barriers and how are they being overcome, and what are the enablers that will fuel this revolution in mobile data?

Barriers—the dismantling has begun

Life on the end of a 9.6 kbit/s circuit switched mobile data call is at best frustrating, and at worst a disaster. The speed is slow, the data session intolerant of a variable radio environment, and the cost to the user can be high. Examples of frustration are the inadvertent opening of a large e-mail attachment and the consequent inexorable download, or the loss of a data connection requiring a complete re-run of a log-on and authentication procedure. These barriers for circuit switched users are being addressed by Cellnet through the introduction of compression and spoofing capabilities:

- *Compression* takes the data stream and interrogates it to remove redundancy in enveloping overheads and the multi-layer protocol stack. Perceived user data rates of two to three times the

underlying bearer rate are achievable in this way.

- *Spoofing* monitors a data session between a mobile device and a host server and interacts with both devices to maintain that session even if the underlying transport service has been disrupted. It does this by emulating the 'chatter' that takes place between a client and server until such a time as the transport service is re-provided.

As a by-product, spoofing also provides the capability to suspend the transport service when no data flow is required between client and server, and re-establish it when data flow recommences. This serves to reduce radically the amount of circuit-switched airtime employed (and therefore charged for) for many interactive data services.

Terminals for mobile data—the choice widens

Historically, the options for mobile data terminals were restricted to a small number of specialist devices (for example, for workforce management or fleet management applications) or the general purpose, but somewhat expensive, laptop. The last two years has seen a dramatic increase in the variety of mobile data devices with prices now starting at \$300 for a modest personal digital assistant (PDA)/palmtop device and \$1500 buying quite a respectable laptop. The continuing barrier for users with these devices is the need to couple together the data device with a separate mobile device—wires, data interface cards, extra software to install—it all adds cost and deters users.

Mobile inside—the data terminal evolution

This year sees the start of a new era for mobile data—the widespread availability of 'mobile inside' devices for the global system for mobile communications (GSM). Trailblazers in the smartphone business such as the Nokia 9000 will be joined by

wireless application protocol (WAP) enabled mobile telephones providing access to the Web and some simple interactive data capabilities. In the PDA and laptop markets we will see a growing range of devices with built-in GSM mobile transceivers.

Applications—a key enabler

First-class mobile network devices and a plethora of data terminals on their own will not drive the use of mobile data services. The essential added ingredient is applications that are tailored to the mobile environment and provide real value to the end user. These are now becoming available from a number of sources, covering the following broad market segments:

- *horizontal business applications*: addressing a wide cross-section of business users with general-purpose products covering common needs such as e-mail, diary management, and access to company intranets and the World Wide Web;
- *vertical business applications*: addressing an individual customer application or a defined segment of the market with specific linkage between the mobile data service and the user's business operation (examples in this segment are workforce and fleet management systems, enterprise e-commerce systems, and applications with specific security or resilience requirements for utilities, healthcare, etc.);
- *consumer applications*: likely to be triggered by the advent of WAP, these will initially focus on web access and e-mail, but will branch out into a variety of personal information services, location-based services, gaming and e-commerce; and
- *telemetry applications*: machine-to-machine communication covering a variety of requirements such as stock control, alarms, meter reading, and remote control.

GPRS—The Marriage of Mobile and IP

While the introduction of compression and spoofing on circuit-switched data goes some way towards providing an enhanced-speed data service, it is essentially an attempt to offer a packet-switched look-alike service and has a number of limitations:

- Only one user can employ the circuit-switched bandwidth: in a packet environment a number of users can be statistically multiplexed into the same bandwidth.
- For the spoofing service to operate satisfactorily, the circuit must be maintained for a reasonable period (for example, 15 seconds) to avoid constant set-up and clear-down activity, introducing end-to-end data transmission delays. This places a limit on the cost saving experienced by the end user.
- The underlying bearer rate is not enhanced.

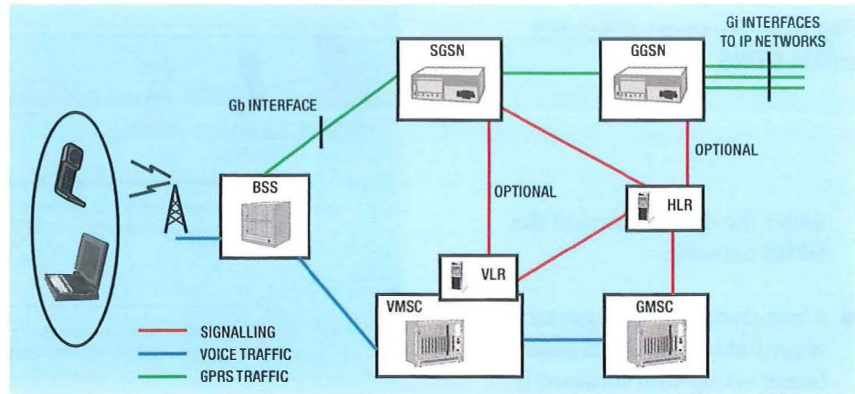
The high-speed circuit switched data and 14.4 kbit/s answers

High-speed circuit-switched data (HSCSD) and 14.4 kbit/s are features becoming available later this year. HSCSD allows up to eight channels to be multiplexed together while the 14.4 kbit/s service allows the per-channel bearer rate to be increased by reducing framing and error correction on the air interface. Between them, these techniques theoretically yield a total bearer rate in excess of 100 kbit/s. While these services address the bearer rate limitation, they do not address the multiplexing or cost limitations of circuit switched data. Furthermore, HSCSD compounds capacity problems of cellular networks by allocating multiple channels to single users.

The general packet radio service answer

The General Packet Radio Service (GPRS) is a feature becoming

Figure 2 – GPRS network architecture



available in late-1999/early-2000 which provides a fully-functional packet data service over the air interface. Various channel coding schemes can be implemented with varying levels of error correction, providing the end user with data rates of between 6 kbit/s and 14 k/bytes per channel. The user can be allocated up to eight channels, giving a theoretical bearer rate in excess of 100 kbit/s. In practice, early implementation of both network infrastructure and terminal equipment will restrict peak data rates to around 30 kbit/s.

Why is GPRS a perfect partner for GSM?

The radio component of a GSM network drives the largest single slice of capital investment, yet is the least well utilised. Typical average carrier utilisation, even on mature well-loaded networks, is less than 40% during the network busy hour. This relatively low utilisation is driven by the need to dimension cellular networks to deliver low levels of blocking, by the modularity of the equipment, the small trunking size of typical GSM cells and the requirement to manage large variances in network traffic across different geographical areas during a working day.

Support of voice and circuit-switched data services, with the need for instantaneous call set-up, high probability of successful call handover between cells, and well-bounded transmission delays, drives network dimensioning in a different way from packet-switched network dimensioning where queuing, variable transmission rates, and variable delays can be tolerated. The simultaneous combination of these two service types—one delay-intolerant fixed bandwidth, the other delay-tolerant variable bandwidth—enables radio network utilisation to be increased to well over 50% without impacting performance of the circuit switched service.

What do we need to make GPRS a reality?

The basic architectural concepts behind GPRS, depicted in Figure 2, comprise:

- Provision of a way of ‘tapping into’ the spare capacity in the radio network and managing the allocation of that capacity between voice users and data users. This typically requires software upgrades to GSM radio infrastructure, a data port addition into radio base station subsystem (BSS) equipment to gain access to the capacity, and a new range of customer terminal equipment.
- Provision of complementary Internet protocol (IP) routing devices to transport the packetised data, and to interact with the visited and home location registers (VLR and HLR) of a mobile network. These new routing devices are serving GPRS service nodes (SGSN) and gateway GPRS serving nodes (GGSN)—they undertake complementary roles to the visited and gateway mobile services switching centres (VMSC and GMSC) employed for circuit-switched calls.

The key to successful operation of the system in a multi-vendor, multi-operator environment is the implementation of standard interfaces between key components of the network, in particular:

- the interface between the radio infrastructure and core IP switching infrastructure—the interface denoted Gb; and
- the interface between the GPRS core infrastructure and other IP

networks such as Internet service providers and corporate IP networks—the interface denoted Gi.

In addition to this basic functionality and standardisation, GPRS networks must fulfil exacting requirements in the areas of security and interconnectivity.

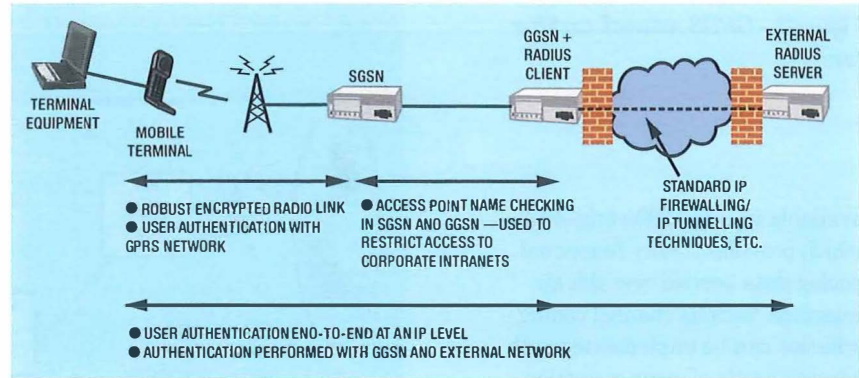
Security in GPRS networks

The term *security* within GPRS, as with any data network, encompasses many different aspects, but the basic requirements are clear: GPRS users (individuals, companies etc.) will naturally expect the data they transmit and receive to be protected against eavesdropping and tampering. Companies will require security to prevent unauthorised users from accessing their corporate Intranets. GPRS network operators will need to prevent unauthorised subscribers gaining access to their network.

GPRS itself has a number of built-in mechanisms to ensure a high level of access control and data security, such as:

- encrypted air interface using a new GPRS A5 security algorithm, which has proved to be extremely robust and secure in existing GSM networks;
- GPRS subscriber authentication and service request validation when attempting to connect to the GPRS network, which use existing GSM mechanisms;
- *restricted access point control* facility to ensure that only terminals authorised by an individual corporate are able to access that corporate’s network from the GPRS network—this is

Figure 3—Summary of security within GPRS



under the direct control of the GPRS network;

- a 'non-transparent' access technique linking the GPRS session/bearer set-up with standard IP access and authentication servers such as RADIUS (remote authentication dial-in user service)—this gives more direct control to the corporate when a user attempts to access the corporate's network; and
- barring options; for example, the ability to control international 'calls' both from the UK and when roaming.

Other aspects of security are outside the immediate scope of the GPRS standards but are just as important. Standard IP network techniques can be employed, such as firewalling for source/destination packet filtering, IP-in-IP tunnelling and encryption for secure transmission across the Internet. With many of these security control mechanisms, GPRS users and companies are heavily dependent on the GPRS network operator to implement the necessary controls, and this implies a certain level of trust between the customer and the operator. It is becoming increasingly important for the network operator to be able to reassure customers that their data is secure by demonstrating appropriate data security policy, compliant with recognised data security standards (for example, BS7799, Information Security Forum).

Figure 3 illustrates some of the typical security mechanisms and where they are used within GPRS for connecting to external IP networks.

Connectivity to external IP networks

It is important for GPRS operators to provide 'basic' IP network connectivity for a diverse range of scenarios in order to provide the necessary foundations for flexible mass market service and application developments. These scenarios include:

- *Internet access*
 - the GPRS operator serving as an Internet service provider;
 - the GPRS operator providing access to external Internet service providers (for example, CompuServe, AOL); and
- *'private' LAN access* to corporates and small/medium companies.

The key challenge for GPRS operators when offering external IP connectivity via GPRS is to balance between offering sufficient flexibility, features, control and coverage while minimising complexity and management overhead, for both the GPRS operator and customers alike. External IP connectivity is an area that is only standardised up to a certain level, with much of the detail left up to implementation.

Many aspects of external IP connectivity are not GPRS-specific and are associated with security, as introduced in the previous section. Other aspects that GPRS operators need to consider are IP addressing, and the provision of enough transport options to cater for a range of distance, quality, cost and bandwidth requirements.

IP addressing

One of the fundamental requirements for any IP networking is the allocation of an IP address to the end user so that packets of data can be successfully routed to and from their associated terminal/application, and GPRS is no exception. Having the necessary flexibility and getting the addressing right is also key to managing other aspects such as security, flexibility, and administration.

IP addresses can either be publicly registered or private. Public addresses are needed for use on the Internet, to avoid duplication

problems. With the growth of the Internet in recent years, public addresses are now in very short supply using the current IP version 4 standard. IP version 6 will provide sufficient addresses but is still some way off. Private addresses are plentiful but cannot be used direct on the Internet. Address translation techniques are available, but introduce their own set of issues.

IP addresses can also be either statically or dynamically allocated. Dynamic allocation potentially saves on IP address space, as these are only allocated when needed and are returned to a common 'pool' for others to use when free. However, this gives problems with mobile terminated services, as each GPRS user can have a different IP address every time, which makes tracking very difficult. There are techniques which can be employed, such as dynamic domain name systems (DNS), but these also introduce their own set of issues. Static addresses make mobile-terminated services easier but use up valuable address space and can give administration challenges.

The GPRS standards do not put any artificial limitations on IP addressing. A combination of sound implementation in GPRS node design can make life much easier, giving the necessary flexibility. Basic flexibility requirements would include the ability to handle potentially overlapping address ranges from different logical 'pools' for different customers from the same GPRS network support node.

Figure 4 illustrates a typical IP bearer set-up sequence, showing how the gateway GPRS support node (GGSN) dynamically allocates an IP address from different address 'pools', depending on the configuration and authentication requirements. It does

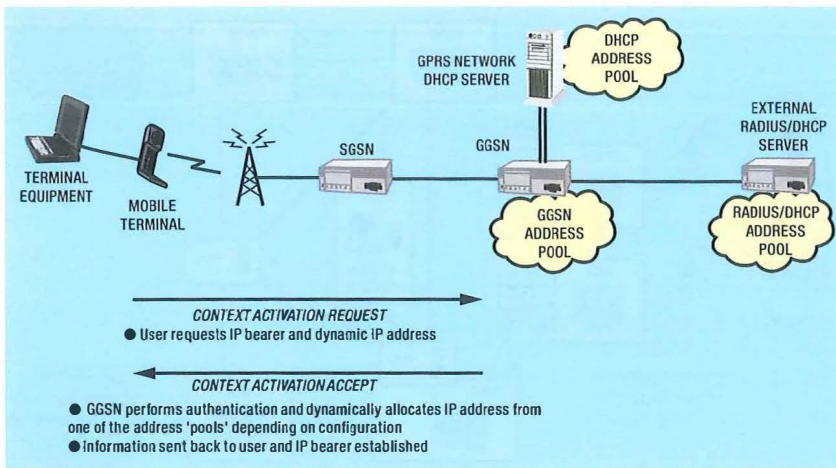


Figure 4—GPRS dynamic IP addressing

this by employing a dynamic host configuration protocol (DHCP) server associated either with the GPRS network or with an external IP network.

Roaming

As with current GSM networks, GPRS will provide the opportunity of roaming to other countries that offer GPRS. With the increased wireless data rate capabilities offered by GPRS, a customer may be half way around the world, but expect to have a connection back to his/her home local area network (LAN). GPRS will also allow operators to offer local Internet access, eliminating the need for an international 'leg' back to a home ISP.

Roaming for GPRS presents new challenges to the operator community such as:

- **Interconnect** To provide a successful high-quality GPRS roaming capability, a global IP backbone will have to be created and managed. Many options exist such as using the Internet, managed virtual private networks,

or dedicated private circuits. Quality-of-service (QoS) parameters (for example, loss of packets, delay, etc.) must be supported by the backbone in order to meet users' roaming requirements.

- **Access point resolution and domain name system** GPRS makes use of DNS (Figure 5) in order to resolve the IP routing for connecting to requested access points (corporate networks, ISPs etc.). Inter-network DNS message exchanges are required when roaming, and this potentially requires a single root DNS server which has knowledge of all roaming networks. A key question is 'who provides and manages such a server'. A decision, which has already been made for security reasons, is to keep the DNS structure private; that is, separate from the public Internet DNS structure.

How Does GPRS Evolve?

The consensus view of UMTS (Universal Mobile Telecommunica-

Figure 5—Domain name system for roaming


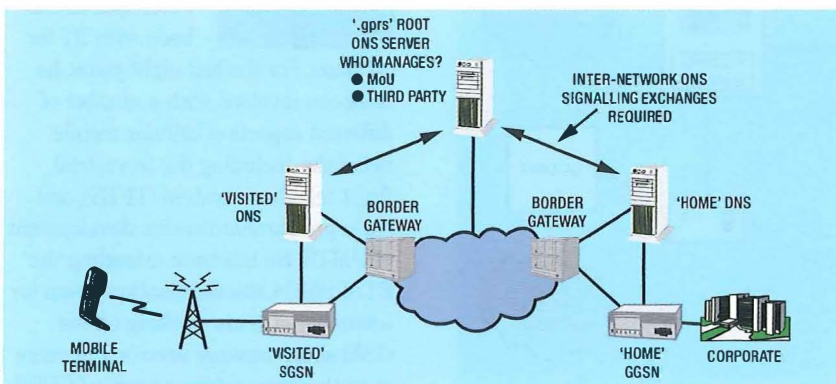


Figure 6—Predicted growth of multimedia services in Europe

tions Service) Forum members and industry consultants is that from around 2005 the primary component of mobile growth will come from users wanting access to multimedia services (Figure 6). Given the lifetime of existing GSM network infrastructure equipment, network operators need to see an evolutionary path through GPRS to multimedia networks such as UMTS.

The network requirements for multimedia services are very different from today's telecommunications networks:

- real-time components which are ideally suited to connection-oriented transmission require tightly defined transmission delay performance; and
- non-real-time components that are more effectively supported using connectionless transmission require defined peak and average throughput capabilities.

Current second-generation GSM supports these two requirements using 64 kbit/s mobile switching centres (MSCs) for connection-oriented traffic and GPRS service nodes for packet services. Although GPRS offers a number of QoS capabilities, it is unlikely to be practicable to support real-time services such as voice or video on the currently defined QoS values. GSM networks can address some of the QoS issues by introducing faster air interface capability through developments such as EDGE (enhanced data rates for global evolution). The real solution, however, lies in the introduction of third-generation UMTS systems that will enable much higher packet data rates to be offered with highly spectrum-efficient coding schemes. Commercial deployment of these systems is expected in 2002.

Figure 7—GSM/UMTS Phase 1 architecture

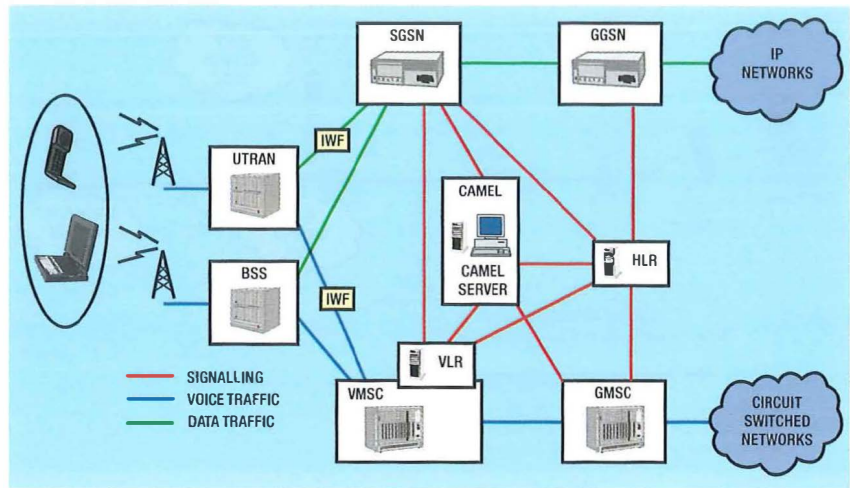
UMTS architecture—initial evolution from GPRS

The initial UMTS network will be largely based upon existing GSM/GPRS networks with the introduction of the UMTS radio interface (UTRAN) (Figure 7). Circuit-switched voice continues to be supported on evolved MSC platforms with packet services and multi-media supported on the GPRS platform. Interworking function units (IWF) enable the necessary protocol conversions between the UTRAN interface and the GSM and GPRS air interfaces. Interworking commences between the GPRS nodes and the GSM intelligent network (IN) capability (CAMEL).

Longer-term UMTS architecture

The longer-term industry view, which Cellnet supports, is that to support a multimedia capability in the most cost-effective way requires the move to an IP-based solution for all services. This requires the GPRS network to mirror the intelligence of existing MSC nodes, which will drive the further development of the standardised IN capability (CAMEL) to produce a common intelligence function for connection-oriented and packet-switched services (see Figure 8).

The rate at which operators progress from the short-term to long-term UMTS architecture will depend



on several key financial factors, centred on the value of their legacy investments in GSM technology and the comparative costs of interfacing second-generation and third-generation networks versus changeout of second-generation. With challenges still ahead on the completion of standards, both the supply industry and operators still have considerable work to do to optimise their migration plans.

programme manager responsible for all technical and commercial aspects of the development and launch of GPRS-based services.

Peter Lisle can be contacted at peter.lisle@btcellnet.co.uk

Biographies



Peter Lisle
BT Cellnet

Peter Lisle was with BT for 20 years in a variety of network strategy and planning roles, before joining BT Cellnet three years ago to lead its network design and planning department. He has recently taken on the role of the company



Howard Green
BT Cellnet

Howard Green has been with BT Cellnet for 14 years. For much of this time

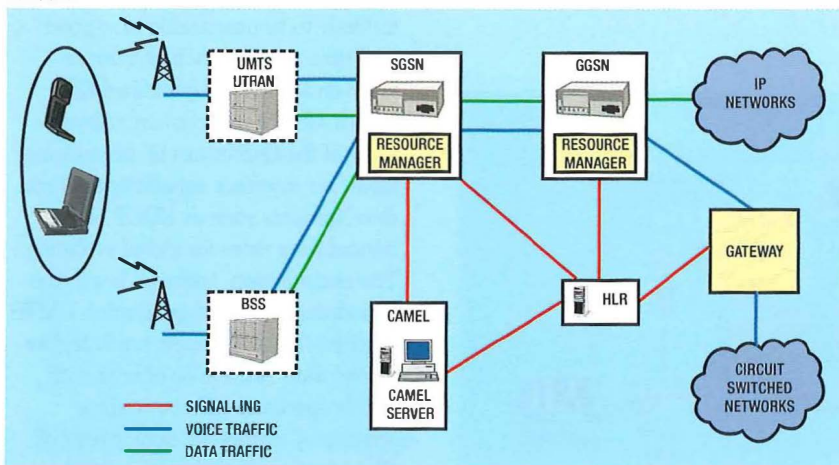
he has been involved with the development of services and features for both the TACS and GSM mobile networks. For the last three years he has been in the Technology Strategy group responsible for the strategic assessment of new network capabilities including GPRS, packet-based technologies and UMTS evolution.



Peter Mason
BT Networks and Information Services, BT UK

Peter Mason has been with BT for 16 years. For the last eight years, he has been involved with a number of different aspects of cellular mobile networks, including the terrestrial flight telephone system (TF'TS), and early pre-standardisation development of UMTS. He has been attending the ETSI SMG4 standardisation group for a number of years working on the GSM short message service and, more recently, interworking aspects of GPRS.

Figure 8—Longer-term UMTS IP network supporting voice, data and multimedia services



Thor Gunnar Eskedal, Arne Folkestad, Bjørn Harald Pedersen, Astrid Solem and Stein Svaet

UMTS Core Network Issues

Two phases for the development of the Universal Mobile Telecommunications System (UMTS) are planned. The first phase focuses on the radio access network, reusing the pre-UMTS core network. In the second phase the development of a new UMTS core network is foreseen. This network could be an evolution from existing networks or it could be designed as a new network. This article discusses three possible core network architectures for the second phase of UMTS, focusing on mobility management, routing and switching, and the ability to offer the required quality of service for multimedia services.

Introduction

The Universal Mobile Telecommunications System (UMTS) standardisation has been driven by the European Telecommunications Standards Institute (ETSI) since the beginning of 1991. The standardisation has recently been globalised by the formation of the third-generation partnership project (3GPP). A phased approach to the standardisation of UMTS has been taken. The first phase relies on standards to be issued by the end of 1999, while the second phase will be evolved in annual releases the following years. The first phase has focused on the UMTS terrestrial radio access network (UTRAN). It has been decided to reuse the GSM and GPRS (General Packet Radio System) core networks, serving as the core networks for circuit switched and packet switched services, respectively. In the second phase the development of a new UMTS core network is foreseen, see Figure 1. Implementation of the first phase is expected around 2002, and a full implementation of the second phase may be expected around 2005¹.

An important anchor in the UMTS work is the Iu reference point (see Figure 1). It forms the division line between the UMTS core network and possible access networks, like the UTRAN, satellite access networks and wireless local area networks (LANs). The UMTS core network provides infrastructure supporting network features and telecommunication services. This includes functionality such as management of user-location information, control of services, switching and routing, and transmission of signalling and user-generated information. The core network may also provide connections to other fixed telecommunication and data networks, which are called *partner* networks.

The UMTS core network for phase 1 will be based on the GSM and GPRS core networks. This implies a conceptual division into a circuit switched domain represented by the traditional GSM network subsystem and a packet switched domain represented by the GPRS core network. The two domains are also reflected at the Iu reference

Figure 1 – The interconnection of UTRAN and the UMTS core network with the pre-UMTS infrastructure².

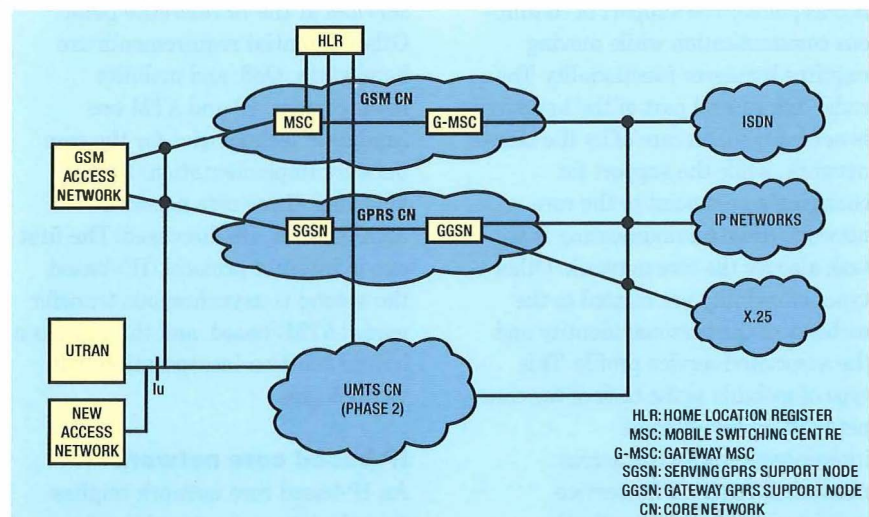
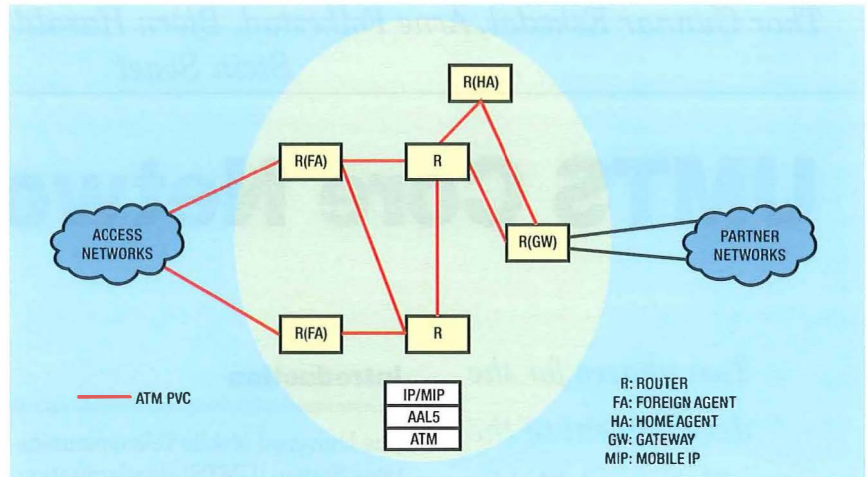


Figure 2 – The IP-based core network. The protocol stack is indicated in the figure

point. The UMTS standards for phase 1 require that the Iu efficiently supports dedicated circuits, best-effort packet services, real-time multimedia services, UMTS signalling and backward compatibility with the GSM signalling scheme³. UMTS phase 2 will be based on annual releases of specifications for the years 2000, 2001 and 2002. Regarding the core network in this phase, there are many open issues and options still under study within research and standardisation.

This article considers technological options for the core network regarding the demanding tasks of mobility management, routing and switching, provision of quality of service (QoS) and multimedia. The mobility-management functionality is, in general, distributed between the access network and the core network. The requirements for mobility management in the core network will therefore vary according to the type of access network, and the core network has to cope with this diversity of responsibility. Mobility support between different access networks will also challenge the mobility functionality in the core network, and demand traffic analysis regarding for instance optimal routing functions and call acceptance control. The problem of mobility management may be subdivided according to the type of mobility. Terminal mobility implies that the terminal may move between different access points. The support of continuous communication while moving requires handover functionality. The radio link control part of the handover is normally taken care of by the access network, while the support for changing access point to the core network while communicating is a task also for the core network. Other types of mobility are related to the mobility of the personal identity and the associated service profile. This type of mobility is the task of the core network. However, some interoperation with the access network is required for service mobility in order to cater for the



service capabilities of the access network. A major challenge for the UMTS phase 2 will therefore be to support future advanced multimedia services across different access networks.

The core network has to be designed to cope with the predicted evolution of bandwidth usage and QoS requirements. It must be able to support real-time and non-real-time services from a multitude of simultaneous users. This will influence the choice of transport and network technology. Also, the core network has to be able to map the QoS requirements from different access networks into QoS requirements in the core network. All the access networks will have their own characteristics which the core network has to interface.

UMTS Core Network Architectures

As mentioned above, the UMTS core network must support both circuit switched and packet switched services at the Iu reference point. Other essential requirements are bandwidth, QoS, and mobility management. IP and ATM are candidate technologies for the core network implementation. In the following, three core network architectures are discussed. The first one is Internet protocol (IP)-based, the second is asynchronous transfer mode (ATM)-based, and the third is a hybrid solution incorporating both technologies.

IP-based core network

An IP-based core network implies that the internal networking func-

tionality is implemented based on the IP protocol suite and using IP routers. The major motivation for considering this option is the global popularity and rapid development of the IP technology. Extensive research has been done lately to evolve the IP technology to be able to support real-time services with similar QoS as the circuit switched networks. If concepts like differentiated services (DiffServ)⁵ and integrated services (IntServ)⁶ can enhance IP to support real-time, high QoS-demanding applications, this can be both cost-effective and provide easy network management. Whether the IP technology will be able to fully meet the strict QoS demands coming from real-time applications remains to be seen. Figure 2 illustrates the IP-based UMTS core network. As shown in the figure, IP is used as the network technology. The core network may use the services from an ATM adaptation layer (AAL) on top of an ATM transport layer, as in the recommendations from ETSI². ATM is a packet switched transfer technology, based on short packets denoted *cells*. IP packets from the applications will be segmented into AAL5 frames which in turn are segmented into ATM cells, and forwarded on permanent virtual circuits (PVCs). The IP header is examined and the output link, that is, an ATM PVC, for the destination is selected by ordinary IP routing mechanisms.

Mobility management

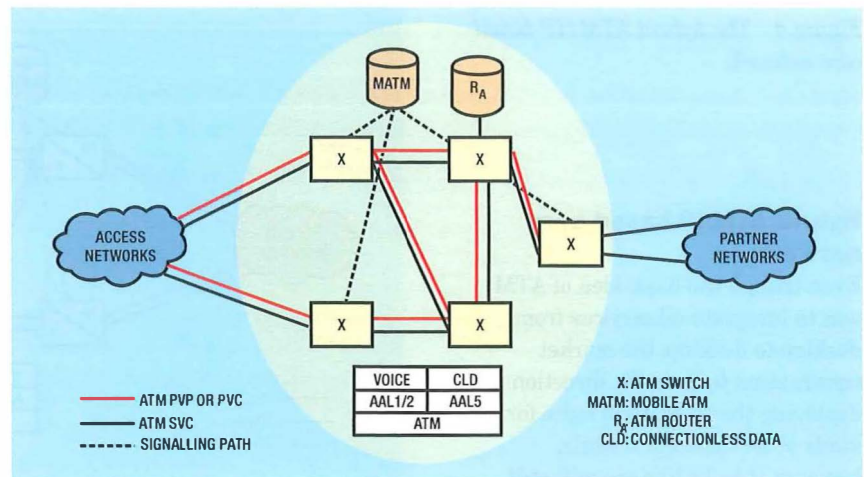
The mobility management in an IP-based core network may be based on the Mobile IP protocol⁴. Mobile IP is designed to support global mobility between different access points. This

Figure 3—The ATM-based core network

is done by means of home agents and foreign agents in addition to functionality to preserve security. The networks that allow mobile terminals to roam have a home agent, and the networks that accept visiting mobile terminals have a foreign agent. When a mobile terminal visits a foreign network, it registers at the foreign agent. The foreign agent sends a message with a care-of-address, often its own address, to the mobile terminal's home agent telling it the current location of the mobile terminal. The home agent uses this address to forward the packets coming to the mobile terminal's home network to the visited network. Mobile IP seems suited to handle mobility within the core network. In micro- and pico-cell environments, cellular IP⁷ may extend the IP technology. This protocol is under development with the goal to support fast handover across radio cells, and is being harmonised with Mobile IP to create a seamless end-to-end IP-based network.

ATM-based core network

ATM has been developed with the aim of handling all kinds of applications, both real-time and different kinds of non-real-time services with connection-oriented and connectionless service support. The connections provided can either be established on demand as switched virtual circuits (SVC), or pre-provisioned as permanent virtual circuits (PVC) or paths (PVP). Different kinds of applications are supported in ATM networks using various ATM transfer capabilities. Connection-oriented and connectionless services are provided in the ATM network by different ATM adaptation layers on top of the ATM infrastructure. Adaptation of connection-oriented real-time services to the ATM network is provided by the AAL1 or AAL2 adaptation layers. AAL1 has been available for some time, while the AAL2 protocol under development is foreseen to incorporate signalling



and switching of AAL2 connections. Connection-oriented non-real-time applications may be supported by AAL 3/4 or 5. The AAL3/4 and AAL5 protocols are also able to carry connectionless higher-layer protocols with no timing constraints, suitable to carry non-real-time data like IP packets.

In a UMTS core network, see Figure 3, all applications may be handled by the same ATM technology that supports switching and routing, and addressing. The ATM switching capability and control and management functions are used for both the connection-oriented and the connectionless data. For the connectionless data, specific connectionless servers are introduced deploying connectionless layer protocols; that is, the Connectionless Network Access Protocol and the Connectionless Network Interface Protocol for routing, addressing, QoS selection, and preservation of the PDU sequence integrity⁸. Packet-based applications segment the packets into AAL3/4 frames and hand them over to the ATM layer which forwards them towards the connectionless server. At the connectionless server the ATM cells are reassembled into AAL3/4 frames and given to the connectionless layer protocols which conduct the routing etc. The connectionless protocols therefore act as an 'ATM router' and only forward the connectionless traffic through pre-provisioned ATM-links that interconnect the connectionless part of the network.

Mobility management

In an ATM-based core network, the mobility solution could be based on

the Mobile ATM scheme being developed by the ATM Forum⁹. The Mobile ATM is based on the E.164 addressing scheme, and uses the PNNI (private network-network interface) hierarchical routing mechanism. Mobile ATM uses location servers handling rerouting and localisation of mobile terminals on the move. Also authentication and authorisation servers handling security, so called *authentication servers* (AUS), are implemented. Each mobile terminal is connected to a mobile enhanced ATM switch (EMAS). Each EMAS is connected to an authentication server (AUS) and a location server (LS) having up-to-date information of the whereabouts of the mobile terminals. If a mobile terminal makes an attachment onto another network, it registers at the current EMASs LS. A registration update message is sent to the mobile terminal's home switch, updating the current location of the mobile terminal in its LS. Regarding incoming calls to the mobile terminal at its home switch, the home switch queries its LS of the current attachment point of the mobile terminal and re-routes the connection. In a UMTS core network, Mobile ATM will offer both roaming between different access points and fast continuous mobility across radio domains. All types of services and applications, both connection-oriented and connectionless, will be handled. Mobile ATM will also be able to interface different access networks enabling coordinated handover. It is not clearly specified which mobility functionality will reside within the core network and which will reside within the different access networks.

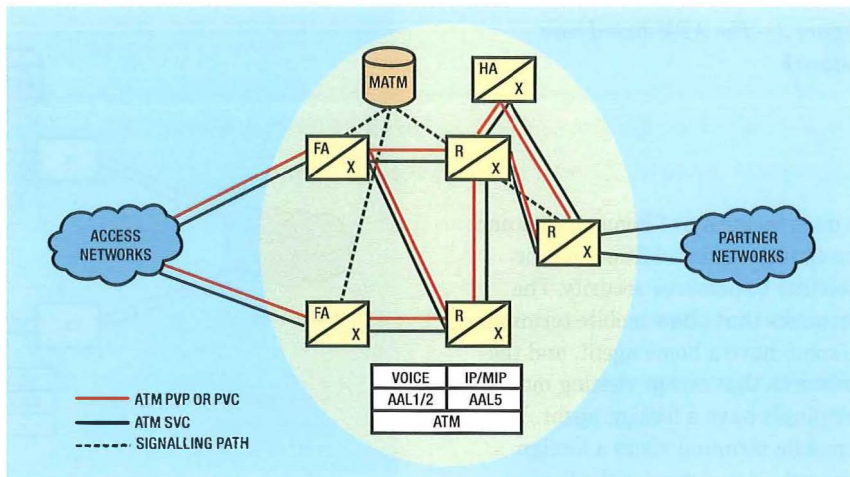
Figure 4—The hybrid ATM/IP-based core network.

Hybrid ATM/IP-based core network

Even though the basic idea of ATM was to integrate all services from desktop to desktop, the market trends seem to be in the direction of deploying the IP protocol suite for all kinds of services. For a while, however, it looks like one will still have to rely on the existing connection-oriented technologies for good-quality voice and video, and efficient packet transport for non-real-time data. With the enormous deployment growth of the IP protocol, it is foreseen that IP also in the near future will be the most important protocol for connectionless data services. An integration of real-time and non-real-time support deploying a common transport platform has been suggested as a solution for the future core network, see Figure 4. In this scheme, ATM is deployed as the common transport platform for both connection-oriented ATM services and connectionless IP services. This implies an overlaid network with different addressing schemes depending on whether connection-oriented or connectionless services are provided. Real-time services as voice and video may use the AAL1 or AAL2 adaptation protocols, and IP-based applications use the AAL5 adaptation protocols for connectionless services. This network structure implies a layered network with respect to routing, and each layer is not aware of the structure of the other network layer. Hence, optimal network structures are not easily obtained. The benefit obtained with this hybrid solution is that a common transport infrastructure based on ATM will give the possibility to utilise available bandwidth efficiently.

Mobility management

The hybrid ATM/IP-based core network discussed uses two different mobility management mechanisms, namely Mobile ATM and Mobile IP. In the same manner as GSM and GPRS today deploy the same overall



mobility management functionality for circuit switched and packet switched communication, the same integration could be possible with this hybrid core network solution. The GSM HLR could then be replaced by the Mobile ATM LS/AUS, and the GPRS mobility functionality could be replaced by Mobile IP. This introduces the possibility to have different visitor registrations for the connection-oriented and connectionless services. However, this may introduce problems not yet thoroughly investigated.

Discussion

There are pros and cons with each of the three core network solutions presented earlier. It is difficult to say which one is better; it depends on which perspectives are regarded the most important. The following sections will discuss some of these issues. This is not an exhaustive coverage—some issues are only briefly discussed and many issues remain open.

First, however, there is a need to clarify the circuit and packet switched conceptual domains. Circuit switching implies that a path is established by dedicated time-slots. Packet switching, on the other hand, is a store-and-forward service, where each packet contains an address or a label enabling the intermediate nodes to identify which next node to pass the packet on to. Contrary to circuit switching, which is inherently connection-oriented, packet switching may be connection-oriented or connectionless. Thus, the 64 kbit/s telephone circuit is a circuit switched transfer service, ATM is packet

switched and connection-oriented, and IP is packet switched and connectionless. For the next generation of UMTS networks, the solutions should not be related to the circuit or packet switched domains, but rather be discussed and evaluated related to the ability to support real-time and non-real-time services.

Mobility

The mobility management within the core network may be subdivided into the handling of terminal mobility and the handling of user and service mobility, where the latter may be regarded as part of the service platform, and thus outside the scope of this article.

Both Mobile IP and Mobile ATM address the terminal mobility problem. It is expected that either solution used exclusively is viable for the terminal mobility requirements in a macro cellular environment where the frequency of handovers is relatively low. For a micro cellular environment, further studies are needed to prove the feasibility. Cellular IP has been proposed as an extension of Mobile IP for fast handovers. However, this is currently an Internet draft only, and further development is still to be expected. In ATM, the same functionality will be provided by wireless ATM, developed as an integral part of Mobile ATM.

For the hybrid solution, both Mobile ATM and Mobile IP are used. The extra complexity related to separate location registrations for the two domains implies increased signalling load as well as more processing. The latter may be significant for the mobile terminal, which has to implement and run two

different protocols for terminal mobility management. In the initial phase, the mobility management of UMTS is based on GSM procedures, and regardless of the choice of future technology, the backward compatibility should be ensured.

Addressing and routing

An IP-based core network solution implies that IP addresses will be used both to identify the user and for routing purposes within the network. In the same manner, NSAP (network service access point) addressing, including the E.164 option, will be used as the addressing scheme within an ATM-based core network. In the hybrid solution, the user will be reached by different addresses for different types of services. The use of 'naming' may hide this fact from the users.

The Mobile IP scheme related to IPv4 (IP version 4) implies non-optimal routing of the packets, since the packets sent to the roaming user always are relayed by the home agent. Packets sent from the roaming user travel a direct route to the corresponding node. This implies that packets in opposite directions related to the same session follow different paths, and support of applications with delay synchronisation requirements between the directions will be complicated. Introduction of IPv6 will solve the problem of non-optimal routing. The solution is to redirect the session between the communicating end-points instead of using forwarding as in IPv4. However, it is foreseen that it will take some time before IPv6 will be deployed throughout the network.

Mobile ATM makes use of a signalling system. It is then possible to let the signalling network cater for the non-connection-oriented part of the mobility and service control. Through signalling, the current address of the mobile terminal may be obtained and the call may be routed directly to the correct destination.

Going from GSM and GPRS addressing and routing to ATM and

IP addressing and routing, the hybrid scheme inherits the problems from UMTS phase 1 with respect to the need to support two different addressing and routing solutions.

Multimedia

The UMTS core network shall support multimedia, implying simultaneous voice, video and data components within the same application. These components have different network requirements. Real-time applications such as voice and video put strict requirements on the network for instance in terms of delay. On the other hand, different data applications put strict requirements regarding an error-free connection. The future network solution should be able to meet both requirements in an efficient manner without introducing unnecessary complexity. The most straightforward solution is to use circuit switching for real-time services and packet switching for non-real-time connections. These transport mechanisms though have their disadvantages. The circuit switched technology is not really bandwidth efficient for most data-oriented applications. The packet switched technology, on the other hand, is bandwidth efficient but has limitations regarding QoS.

There are benefits obtained by choosing one single technical solution for phase 2 of UMTS. The main problem, however, is that there are no obvious multimedia network solutions. The main candidates are ATM and IP as discussed earlier. ATM has for a long time been considered as a candidate technology to be the network solution for both real-time and non-real-time applications. Utilising the flexibility of a packet-based solution, the key to supporting both real-time and data applications has been the choice of a short fixed-length packet size. However, this packet size is optimal for neither voice nor data. IP, as the other option for the future network technology, does not have sufficient support for real-time services today.

Evolution of today's IP-based networks is required to be able to support all kinds of applications including voice and video. Concepts such as differentiated services (DiffServ) and integrated services (IntServ) are developed to bring IP up to a level comparable with ATM regarding the possibility to support QoS. If these mechanisms turn out to be successful, IP would be able to support real-time services.

In the hybrid architecture different conceptual domains are chosen to realise different types of services. How multimedia services can be realised is an open issue. The use of two domains gives the flexibility to put the multimedia functionality on the technology supported by the partner networks. This flexibility may, however, be at the cost of building two separate networks to handle multimedia.

Migration aspects

Having introduced some possible implementations of a future UMTS core network, the question is how to migrate from existing networks. Two different types of UMTS network operators will probably dominate the UMTS market: the current mobile operators and the Internet service providers.

ETSI has decided to reuse the GSM and GPRS core networks for the first phase of UMTS. The GSM circuit switched backbone will be used for voice traffic, while data traffic will be supported by the GPRS packet switched infrastructure. The current mobile operators will follow the recommendations from ETSI in order to reuse as much as possible of their widely deployed cellular networks. This could be seen as the evolutionary approach to a new UMTS network. The result of this evolution is not necessarily given. An introduction of Mobile IP into the GPRS-based phase 1 network could result in a future IP-based core network. The failure of solving the problems of supporting the required QoS for multimedia applications in

such a network could pave the way for a hybrid solution. This may also be the best solution with respect to the backward compatibility requirements. Introduction of Mobile ATM in the UMTS phase 2 core network may be based on an evolution of the GSM HLR infrastructure. Work on interfacing GSM to Mobile ATM is being conducted within the ATM Forum. Other access networks are specified for a seamless interface to the ATM technology.

The Internet service providers, who do not have the legacy of GSM and GPRS, may choose to implement the core network functionality by adding mobility functionality to their IP networks. The IP may also be the basis for mobility management in new access networks, and UTRAN may face competition for instance from wireless LANs in campus areas. This is a revolutionary approach compared to the one likely to be chosen by present mobile operators. The result of this development could very well be the true mobile Internet consisting of an increasing number of fixed and mobile subnetworks with full mobility support.

Conclusions

The article has discussed three different solutions for future UMTS core networks in terms of mobility management, routing and switching, and QoS support, in particular for multimedia services. It is not easy to say which solution is better and which solution will gain commercial acceptance. Many issues are open, and this should stimulate researchers to work in this area.

References

- 1 ETSI, SMG. UMTS 30.00 v3.0.0, SMG Work Programme for UMTS, 1997.
- 2 ETSI, SMG. Draft Technical Report UMTS 23.20 v.1.5.0, Evolution of the GSM platform towards UMTS, 1999.

- 3 ETSI, SMG, Draft Technical Report UMTS 23.30 v.0.6.0, Iu Principles, Dec. 1998.
- 4 PERKINS, C. E. Mobile IP. *IEEE Communications Magazine*, May 1997, 35(5).
- 5 BERNET, Y., *et al.* A Framework for Differentiated Services. Internet Engineering Task Force, Internet Draft, Feb. 1999.
- 6 BRADEN, R.; CLARK, D.; and SHENKER, S. Integrated Services in the Internet Architecture: an Overview. Internet Engineering Task Force, RFC 1633, June 1994.
- 7 VALKO, A. Cellular IP. Internet Engineering Task Force, Internet Draft, Nov. 1998.
- 8 ITU-T, Recommendation I.364, Integrated Services Digital Network (ISDN), Overall network aspects and functions. Support of the broadband connectionless data bearer service by the B-ISDN. Nov. 1995.
- 9 ATM Forum, WATM Working Group. Wireless ATM Capability Set 1. Specification Draft, BTD-WATM-01.10, Dec. 1998.

Biographies



Thor Gunnar Eskedal
Telenor

Thor Gunnar Eskedal studied physics/computer science at the University of Oslo and received his M.Sc. degree in 1990. After a year working as a research assistant at the Norwegian University of Science and Technology, he joined the broadband group of Telenor Research and Development (R&D) in 1991, where he has been working with ATM switching systems, protocols, management and network design.



Arne Folkestad
Telenor

Arne Folkestad received the M.Sc. and Ph.D. degrees in Telematics from the Norwegian University of Science and Technology in 1994 and 1998, respectively. He joined Telenor R&D in 1998.



Bjørn Harald Pedersen
Telenor

Bjørn Harald Pedersen received the M.Sc. degree in Electrical Engineering from the Norwegian University of Science and Technology in 1992. He has been working at Telenor R&D since 1993 with different aspects of next-generation mobile communication systems.



Astrid Solem
Telenor

Astrid Solem received the M.Sc. degree in Electrical Engineering in 1988 from the Norwegian University of Science and Technology. She joined Telenor in 1988 and since 1993 has been working at the R&D department.



Stein Svaet
Telenor

Stein Svaet received the Honours degree in Electrical Engineering in 1986 from Heriot-Watt University in Edinburgh. He then joined Telenor as a radio planner. Since 1988 he has been with the R&D department, working with various fields of mobile communications.

Dietmar Garthe, Clemens Epple and Bernhard Edmaier

Dense Wavelength-Division Multiplexing and Optical Networking

VIAG Interkom introduced dense wavelength-division multiplexing (DWDM) in 1998 for reducing the network cost. This prompted the development of new architectures such as super highway and optical through-channels, whose introduction is currently under study. The demand for flexibility and protection in the optical layer in an increasingly data-driven world paves the way towards optical networking. Optical switch engines may soon be available, but many more technological aspects need to be addressed, particularly in the standardisation bodies.

Introduction

VIAG Interkom installed its first DWDM system as early as the middle of 1997: a four-channel system from Bosch Telecom was trialled and taken into service. This link is still in use with live traffic today. When in 1998 VIAG Interkom built its nationwide backbone, one half of the links were equipped with Lucent's WaveStar 80G equipment with a maximum of sixteen STM-16 channels. This article gives some insight into the reasoning behind this decision and how the system was integrated into the transport network. The article then goes on to focus on optical networking, seen as the next evolutionary step in the area of transport networks, and finishes by discussing technology trends.

DWDM in the Fibre Network of VIAG Interkom

What is DWDM?

DWDM is a technique whereby multiple optical signals can be

transmitted over a single optical fibre by means of colour coding (Figure 1). Each individual signal undergoes wavelength conversion before being multiplexed onto a single fibre via a combiner. At the receiving end, the signals are again separated by means of a demultiplexer (for example, a set of optical bandpass filters), and finally the original optical signal is regenerated.

Apart from better utilisation of the transmission fibre, savings can be expected from the use of erbium-doped fibre amplifiers (EDFAs) instead of synchronous digital hierarchy/synchronous optical network (SDH/SONET) repeaters, which can amplify the compound signal without the need to split into individual wavelengths.

Cost advantage of DWDM and DWDM deployment in VIAG Interkom's network

DWDM nearly always saves money when the number of optical fibres available for lease is smaller than the

Figure 1 – DWDM principle

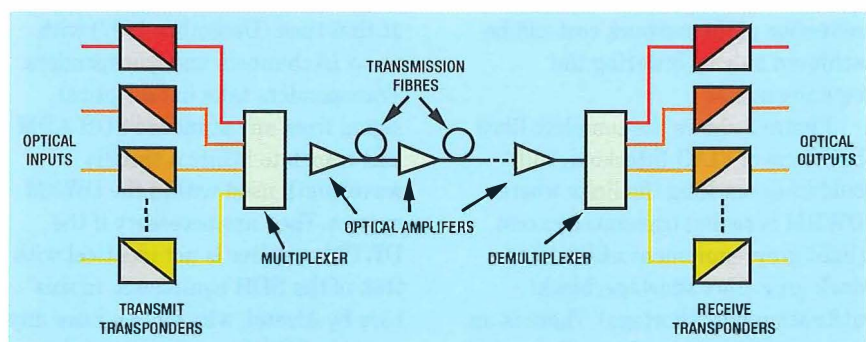


Figure 2 – Comparison of link cost with and without DWDM

number of STM-16 systems to be transmitted. The alternative would be to deploy a new fibre cable, which is expensive (80–200 DM per metre) and is only justifiable for very short links. But even if there is no need to deploy new fibre cables, DWDM can pay for itself by reducing the number of fibres to be leased and the amount of repeater equipment required.

This is illustrated in Figure 2, where the link costs (fibre lease plus repeaters and WDM equipment where applicable, but without terminating add-drop multiplexers (ADMs)) are shown for various system lengths with DWDM (continuous line) and with SDH alone (broken line). The curves are normalised to the transmitted bandwidth, clearly demonstrating the cost effectiveness of DWDM with an increasing number of STM-16 systems. By contrast, the bandwidth cost without DWDM stays constant, as every new STM-16 system requires the lease of an additional fibre pair and the installation of further SDH repeaters.

The curves in Figure 2 are based on real data applicable for VIAG Interkom. The fibre lease is based on an annual fee of 1 DM/m plus investment return (17%), taking into account a lease period of 18 years. Break even is reached between two and three systems for short links and one to two systems for longer links. With increasing system counts, the DWDM bandwidth costs become more and more independent of the system length and system count; that is, the price of the common DWDM equipment becomes the dominant factor. This means that a further reduction of the network cost can be achieved only by lowering the equipment cost.

Figure 3 shows the complete fibre backbone of VIAG Interkom, with bold lines marking the links where DWDM is saving transmission cost (light grey: equipment utilisation, dark grey: fibre shortage, black: utilisation plus shortage). There is an average of six channels per link, with

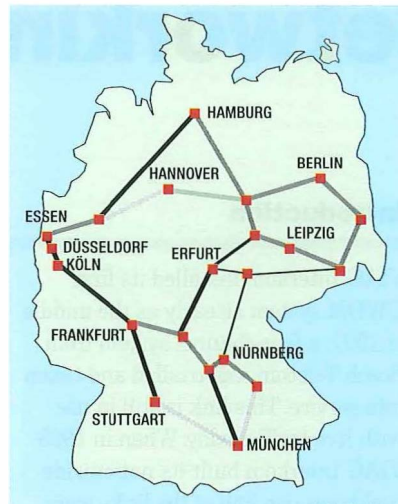



Figure 3 – DWDM deployment in the VIAG Interkom backbone

extremes of nine or even fourteen channels. STM-64 systems were not considered an alternative to STM-16 DWDM because of the limited availability of such systems in early 1998, and some concerns with respect to dispersion and non-linear effects.

To conclude this section, it should be mentioned that DWDM can also help to reduce the cost of network upgrades, as adding new channels requires changes only at the core node sides. No work force needs to be sent out to the repeater sites.

The installed system

The chosen DWDM system was Lucent WaveStar 80G, which was one of the two STM-16 systems available at that time (December 1997) with up to 16 channels and transponders. Transponders take in the optical signal from any standard SDH ADM and translate it into a specific wavelength used within the DWDM system. They are necessary if the DWDM supplier is not identical with that of the SDH equipment, in this case by Alcatel, who did not have any suitable DWDM equipment at the

time. With some combinations of SDH and DWDM systems, transponders are also required at the receive side, for translating the DWDM output into a signal suitable for the input of the SDH ADM. This was not the case in the present project.

Other benefits of the Lucent system were its small size, the availability of high-gain EDFAs (33 dB) and the automatic gain control, which allowed the removal and insertion of wavelengths without alteration of the system. During the implementation, monitor points in the form of fibre taps were inserted for monitoring the signal (wavelength, power level, signal-to-noise ratio, bit error rate) between the SDH and DWDM equipment. Although offered by Lucent, optical protection was not built into the system, as the SDH layer was designed to provide full network protection already and any additional protection would have been a waste of money.

The DWDM is managed via a stand-alone network management platform, which allows the initial set-up of the system, alarm management and performance measurements. The management system is collocated with the SDH manager, and in case of a DWDM failure, alarms will show up simultaneously on both systems, which then enables the operational staff to locate the fault. Integration into a single network management platform has no major benefits and is not foreseen.

New network architectures based on DWDM

Initially, DWDM systems were installed for creating virtual fibres; that is, for using the available transport infrastructure more

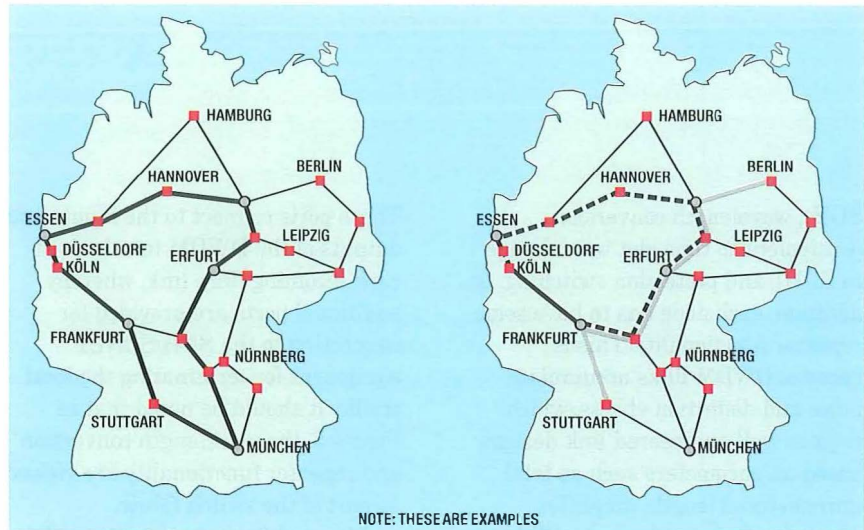
Figure 4—Superring concept (left) and examples of optical through channels (right)

efficiently. Since then, DWDM prompted the development of new network architectures, which are described here briefly.

The *superhighway* forms an additional network layer for the nationwide traffic. A subset of the existing core nodes is chosen to be superhighway nodes and connected via optical channels onto the existing DWDM infrastructure to form a closed ring (Figure 4, left). Nodes without superhighway connection are bypassed by directly connecting the channels from one DWDM system to the other, irrespective of whether these DWDM systems belong to the same SDH ring or not. In this way, the nationwide traffic is concentrated onto a small number of links and nodes, which enables savings in transport and SDH costs to be made. The existing SDH rings then just have to provide regional access to the superhighway, and grow much slower than without the superhighway.

With a DWDM infrastructure already in place, the installation of a superhighway is very economic. The superhighway links just need to be upgraded by one or two extra channels, whereby the output from one DWDM system can be directly connected to the input of another DWDM system (back-to-back connection). Only occasionally is a transponder needed, be it for providing signal regeneration capabilities or changing the wavelength. At present, VIAG Interkom is working on the optimum implementation of the superhighway concept. In the USA, the superhighway concept is adopted widely by long-distance carriers.

Optical through-channels are a product opportunity addressing large business customers or other carriers; for instance, for building virtual private transport networks. Their implementation is very much the same as that of the superhighway links and is as simple. The map on the right of Figure 4 shows examples of a protected connection from Frankfurt to Düsseldorf and an unprotected connection from Stuttgart to Berlin.



Remote access is the connection of a site that is not identical with the DWDM site. With the Lucent system the remote site can be 50 km away from the DWDM system without using any additional SDH technology. In this way the DWDM system itself acts as a traffic collector and distributor in big cities and conurbations.

Optical Networking

Drivers

New developments are usually driven by a desire to create new revenue or to reduce costs. In the

reconfiguration may be done on the SDH layer, but only if all traffic goes through that layer, which is not the case with superhighway traffic and optical through-channels.

More flexibility is certainly the key driver towards optical networking. Another one is network protection. Protection mechanisms should operate as close as possible to the location of the main origin of faults. In the case of fibre networks, this is the fibre, mainly due to building hazards. Protection therefore is required on the optical layer. A single switch action can then provide an alternative path for a huge number of higher-level connections. This

With a DWDM infrastructure already in place, the installation of a superhighway is very economic

case of transmission networks, introducing superhighways and optical through-channels is the first step in the right direction. But what if traffic patterns are changing, and the chosen architecture proves inefficient? Or, even more likely, when new bandwidth customers are acquired, while others discontinue their lease contract?

In all cases, the network architecture needs changing, which is expensive or even impossible under normal operation, if done manually by local service personnel. Relief comes from the introduction of so-called *optical cross-connects*, whose connections may be set remotely from the operations centre and can be switched in milliseconds. Of course,

makes protection switching simple and speedy, and saves costs by avoiding the duplication of 'safe' equipment that never fails. Optical protection becomes a must if the client layer does not provide any protection mechanism at all. A classic example is Internet protocol (IP) directly over DWDM.

Functional and technical requirements

The *functionalities* that need implementing in the optical domain are related to those already obtainable from the SDH and SONET layer. These are wavelength add-drop multiplexing and cross-connection (equivalent to VC-4 add-drop multiplexing and cross-connection in

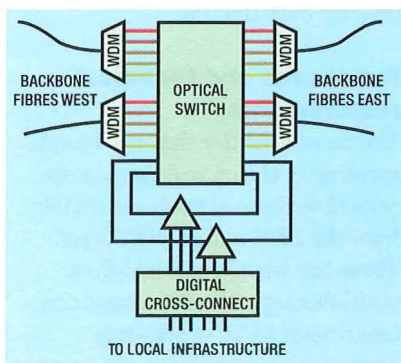
SDH), wavelength conversion (equivalent to time-slot interchange in SDH), and protection switching. In addition, each node has to have some repeater functionality. This is because DWDM links accumulate noise and distortion effects, which require well-engineered link designs based on parameters such as total unrepeated length, amplifier spacing and channel counts. Within optically-switched networks, the total route length may change during operation between some 100 km to some 1000 km as a result of optical switch actions. By integrating a repeater functionality into each node, the unrepeated length becomes independent of the total route length, and each fibre link can be designed irrespective of this parameter.

The main *technical requirements* are a fast switching time (some 5 ms) to keep the total switching time, including the time for processing, under 50 ms, a compact and economical design (space and money are always short), and transparency with respect to protocol and bit rate in order to handle future transport requirements.

Generic node architecture

Figure 5 shows a generic design for an optical network node that can handle wavelength add and drop, wavelength cross-connect, and optical channel protection. The core forms a switch fabric which can provide connections from every port to any other port (full cross-connect matrix).

Figure 5—Generic optical network node architecture



These ports connect to the inputs and outputs of the DWDM terminals of each incoming fibre link, whereby additional ports are provided for connecting to the SDH/SONET equipment for terminating the local traffic. It should be noted that in Figure 5, the wavelength conversion and repeater functionality are viewed as part of the switch fabric.

This architecture operates such that through-traffic is kept away from the SDH equipment by making direct connections in the switch fabric from one DWDM system to the other. Add and drop is performed by connecting the appropriate DWDM port to one of the ports leading to the SDH/SONET equipment. Ideally, such SDH/SONET equipment is just used for terminating the local traffic, and VC-n grooming of through-traffic is kept to a minimum. Otherwise, savings due to the reduction in SDH or SONET equipment may be too small for bearing the additional cost of the optical switch.

Switch technologies

The industry has two principally different answers to the question of implementing the optical switch.

The first is the use of a purely *optical switch fabric* based on mechanical displacement (for example, micro-machined mirrors) or electro-optic effects in semiconductors or polymers (Figure 6). Electro-optic effects enable or disable energy transfer between neighbouring waveguides by means of minute changes of the refractive index. The purely optical switch fabric has the big advantage of being fully transparent to protocol and bit rate. It also has the potential of being an economic solution. Questionable, however, are reliability (new technology), compactness and speed. Moreover, wavelength conversion and 2R regeneration, which would be required at each port, are a real technological challenge.

The alternative is an *electrical switch fabric* embedded between optical-to-electrical (O/E) and electrical-to-optical (E/O) converters (Figure 7). These O/E and E/O converters, which are nothing more than the standard DWDM transmitter and receiver modules, automatically provide the wavelength conversion functionality (in the transmitter) and the regenerator

Figure 6—Network node with optical switch fabric


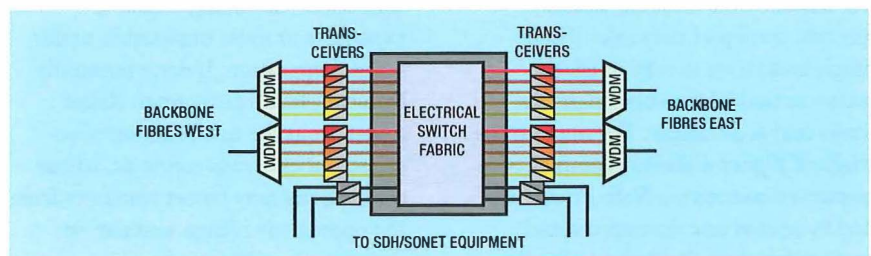


Figure 7—Network node with electrical switch fabric



In recent years it has become obvious that the anticipated increase in data traffic will drive the technology towards optical networking

functionality (in the receiver). Undoubtedly, this so-called *opaque* cross-connect can be compact and fast, but one may be concerned as to whether or not this technique will deliver a transparent switch solution. Another concern is the cost of high-speed electronics.

Network protection

As said before, the architecture in Figure 5 can be used for providing optical channel protection. This would be achieved by foreseeing extra ports at the switch fabric as well as the DWDM systems being served. In the event of an optical channel being disturbed, the switch fabrics on either end of a DWDM link would receive a signal to switch onto one of the spare channels on the same link. There may be fewer protection channels than working channels, resulting in reduced protection cost when compared to implementing optical channel protection via SDH/SONET protection mechanisms. This fact is one of the drivers for the development of optical switches in North America. The architecture in Figure 5 also lends itself to the implementation of *optical switched circuit network protection (SCNP)* in a similar fashion as optical channel protection.

In contrast to optical channel protection and optical SNCP, *optical link protection* is already available with some DWDM systems. This is because such protection does not need an optical switch at all. The compound DWDM signal is forked at the transmit end, and at the receive end, the better signal is selected, for instance, by turning on and off the optical amplifiers located there. At present, mainly linear protection has been available. Only one DWDM supplier has a ring protection concept commercially available (optical equivalent of unidirectional self-healing rings (USHRs)).

Open issues

The main open issue is the *lack of standards*. In particular, there are no

standards for protection schemes. Some manufacturers have optical supervisory channels, others start to introduce optical section overheads akin to the multiplex section overhead (MSOH) of SDH and SONET, but all these attempts are proprietary solutions. The Optical Internet Forum (OIF), which was founded for speeding up the standardisation process, is still in its setting-up phase and has not produced results yet. This forum will also have to discuss issues related to interworking with client layers, be it for deriving a performance measure for the transported signal or for preventing multiple protection switch actions in case of faults.

A technological issue is the derivation of a *measure for the optical signal quality*. This is necessary for triggering protection switches as well as for deriving system performance data. The best measure is currently the B1 bytes of the transported SDH or SONET frames. However, one may ask whether this is a future-proof solution, as no-one knows for how long data signals will have SDH/SONET framing. Alternatives are monitoring some operational parameters such as laser currents and temperatures, or better, the continuous measurement of the optical SNR in each wavelength via a test signal laid on top of the actual data signal. Both techniques are indirect though and may fail under certain circumstances.

Another big issue is *network management* in a multi-vendor environment. If, for whatever reason, the DWDM vendor differs from the vendor of the SDH client, the management systems will usually be incompatible. Setting up an end-to-end connection requires configuration of both the SDH layer and the optical layer (optical switches), and the co-existence of two separate management systems proves cumbersome and expensive. A single connection requires setting up three trails: one on the SDH layer to the DWDM node, one on the DWDM layer between the DWDM nodes, and one on the SDH

layer from the DWDM node. The only choices network operators have at present are either buying SDH and DWDM equipment from the same vendor, or implementing an umbrella management system on top of the existing management systems. There are signs of co-operation between different system manufacturers on the network management issue, but these are rare.

Outlook

In recent years it has become obvious that the anticipated increase in data traffic will drive the technology towards *optical networking*. The use of point-to-point DWDM is state-of-the-art for many new network operators. Soon this optical layer will need switching capabilities for handling the amount of traffic efficiently. System manufacturers have started to recognise this development and will come up with optically switched DWDM solutions in the near future, whereby the opaque switch technology tends to win the race over purely optical switches due to relying on well-established technologies. SDH and SONET will more and more lose their switching functionality, but will still be there for some time to be, mainly for framing purposes.

There is a split in the vendor community with regard to *the integration of PSTN traffic and data traffic (IP)*. Some vendors propose a super node with optical cross-connect, SDH cross-connect, asynchronous transfer mode (ATM) switch and IP router/switch, which would provide efficient management of the available bandwidth for any traffic mix, but may be expensive due to the large amount of equipment needed. Other vendors propose just to connect big IP routers/switches directly with the optical layer via optical ADMs or again cross-connects, thus leaving out SDH and ATM switching entirely. This may be less expensive, but requires all traffic to be converted into IP for efficient

bandwidth management. In all cases, the need for optical protection has been noted, and the industry is trying to overcome the lack of standards by setting up its own standardisation body (OIF).

The great potential of using *DWDM technologies in city backbones* has also been noticed. Key there is the high cost of fibre deployment or lease, and the transparency of DWDM with respect to bit rate and protocol, thus reducing multiplexing costs considerably. However, the so-called *metro DWDM* systems currently available from some of the traditional system vendors are far too expensive for building a business case for city DWDMs. Hope comes from newcomers in the market, who started off in the optical LAN business and now plan systems that produce valuable competition to the traditional system vendors. The greatest technological challenge is the provisioning of reconfigurable wavelength add-drops; even static ones exhibit insertion losses that are still excessive. For protection, the same concerns as for long-haul systems apply.

Conclusion

Today DWDM can be considered as an established technology for reducing the network cost in the transport layer. This is true not only for long-distance traffic but also for city backbones. The drive towards increasing data rates will result in a shift of traditional SDH functionalities into the optical domain, thus requiring optical switching techniques and protection mechanisms. This transition will happen within the next couple of years. SDH will still maintain its role as a framing standard for some time to come.

Biographies



Dietmar Garthe
VIAG Interkom

Dr. Dietmar Garthe read Electrical Communications at the Technical University of Darmstadt. After finishing his doctoral degree in 1992 with a thesis on Fibre-optic Microphones for Airborne Sound, he joined Northern Telecom at its laboratories in Harlow, England, where he became an internationally recognised expert for high-speed optical transmission. He gained world leadership with the first long-haul 320 Gbit/s transmission over standard telecommunication fibre and the first dispersion compensator based on fibre Bragg gratings. Dietmar left Nortel for VIAG Interkom in February 1998, where he is now responsible for Transport Technologies within the Networks and Infrastructure unit. His activities are focused on strategic design and vendor management for the transport layer, including technologies such as radio transmission, SDH equipment, DWDM, optical networks, and broadband transport (IP/ATM).

Dietmar Garthe can be contacted at dietmar.garthe@viaginterkom.de



Clemens Epple
VIAG Interkom

Clemens Epple studied Electrical Communication at Fachhochschule München, where his final project was on multimedia systems for Flughafen München GmbH. He first worked as

a patent engineer for a lawyer affiliation before joining Bayernwerk Netkom in 1997. There his main task was the introduction of DWDM systems into the nationwide fibre backbone of VIAG Interkom, ranging from the tendering process over link planning to system acceptance testing. Clemens is now with VIAG Interkom, where he is responsible for DWDM implementation. He also contributes to the work on new network concepts involving DWDM.



Bernhard Edmaier
VIAG Interkom

Dr. Bernhard Edmaier read Electrical Communications at the Technical University of Munich. He finished his doctoral degree in 1996 (thesis: 'Path Protection Mechanisms with Distributed Control in ATM Networks'). He joined Bayernwerk Netkom in November 1996 and designed an integrated synchronisation concept for the SDH transport networks of Bayernwerk Netkom and VIAG Interkom. Since February 1997 he has worked for VIAG Interkom as a consultant, responsible for the concept of the Germany-wide SDH transport network architecture as well as for vendor evaluation. In October 1998, he joined VIAG Interkom and is now responsible for the Transmission Competence Centre. His activities involve strategic design and cost evaluation of the ATM/SDH/DWDM broadband transport network architecture, resilience concepts, transport product portfolio definition, functional specification of network systems and vendor negotiation.

Frode Beckmann Nilsen and Leif Aarthun Ims

Interworking in the Future Broadband Access Network— A Question about IP and ATM

This article discusses interworking in the future broadband access network from a service-integration perspective. This is a protocol issue and the key question is the relationship between IP and ATM as they are the candidate technologies to accomplish service integration. Both technology interworking and network interworking are addressed. The main conclusion is that ATM is beneficial for technology interworking. However, it is foreseen that the access network will be dominated by IP and it is argued that network interworking is best performed at the higher IP level. The combination of IP and ATM becomes the preferred choice for the future broadband access network, creating new opportunities that affect the role of the access network.

Introduction

The future broadband access network will be different from the existing narrowband copper-based network in three respects:

- technological variety,
- open provisioning, and
- service integration.

Hence, we are facing paradigm shifts along three dimensions that change the architectural requirements of the access network. Technological variety means that fibre solutions, satellite systems, cable-TV networks, radio systems and digital subscriber line (xDSL) solutions over the existing copper network can be part of a unified architecture. An open architecture is required for operating in a competitive environment with different core network operators. Service integration refers to an access network in which all kinds of traffic are mixed by means of statistical multiplexing. This is in contrast to static multiplexing and dedicated connections currently being used.

Service integration

This article focuses on service integration. The objective of service integration is to provide a common access solution replacing the set of specific solutions being used today for platforms like integrated services digital network (ISDN), frame relay (FR), Internet protocol (IP) and

asynchronous transfer mode (ATM). The rationale is flexible, efficient and cost-effective provisioning of broadband access. The challenge is to maintain quality of service (QoS) in the mixed traffic stream.

The vision of a worldwide service integrated network was originally associated with the work on B-ISDN and ATM¹⁰ as the enabling technology. However, the mass market has failed to request any large-scale service except ordinary telephony up to now. For this reason there has been no driving force towards service integration. The recent growth of the Internet as the killer application in addition to telephony has led to a new situation. The challenge today is to develop an overall network architecture which can efficiently provide both a real-time service (telephony) and a best-effort data service (Internet) all the way to the customer premises.

Interworking

In the context of service integration, this article further focuses on interworking. Specifically, the following two issues are addressed:

- How is interworking of the different access technologies handled inside the access network?
- How is interworking handled at the border between the access network operator and various network and service providers?

In both cases, the crucial point is how to perform statistical multiplexing

with QoS guarantees. This is a protocol issue and the key question is whether IP or ATM is the most appropriate technology. In the latter case, it is also a question of the switching and management capabilities of the two networking technologies.

Organisation

The rest of this paper is organised as follows. The next section briefly outlines an architectural model of the future broadband access network. This is followed by a general discussion in on IP and ATM as competing technologies for providing end-to-end service integration. The main body of the article then follows focusing on interworking and the question about IP and/or ATM in the access network. The role of the future access network is then discussed. Finally a conclusion is presented.

Access Network Architecture

The paradigm shifts listed in the previous section change the architectural requirements of the access network. Figure 1 is a generic illustration of how the future broadband architecture is likely to be. The model shown and its terminology is a result of a study performed by Telenor¹. However, the model is perfectly in line with similar work carried out by other operators and equipment vendors²⁻⁸.

The domain of the access operator is marked with shaded areas. The trend towards technological variety is illustrated by three different cases to the right in the figure. The topmost case reflects an ADSL⁹ solution that exploits the existing narrowband structure with copper lines terminating in a remote subscriber unit (RSU). The vertical lines branching off at this point indicate how access to various network platforms is currently realised. The existing solution is characterised by having dedicated connections and no concentration of traffic in the access network. The middle case is representative for a fibre network in

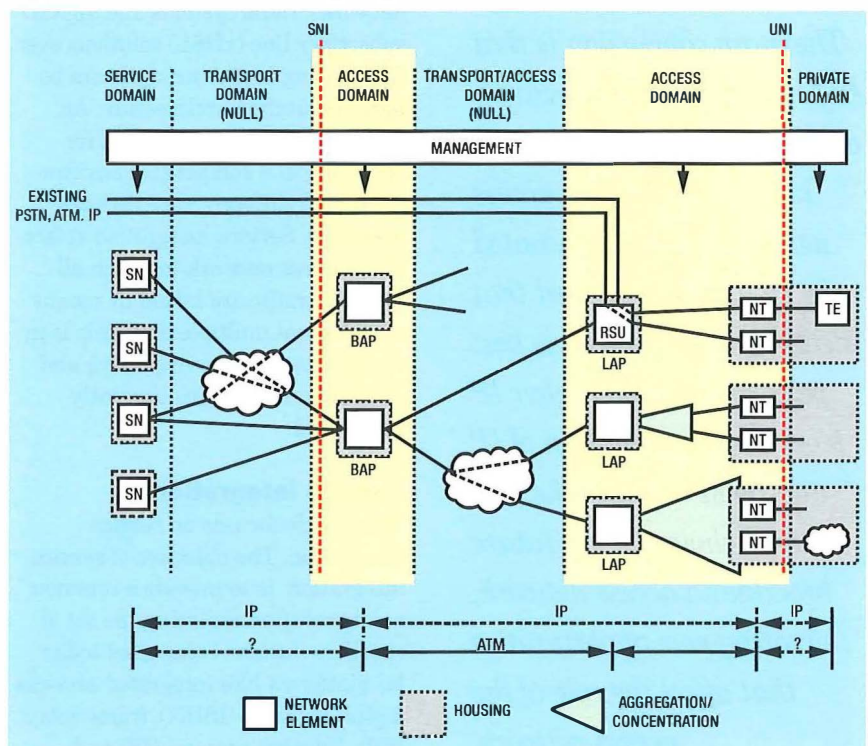
combination with VDSL⁹ over a final copper drop with reduced length. The bottommost case corresponds to a radio solution like local multipoint distribution service (LMDS) where traffic is aggregated in the air all the way to the customer premises.

In the future broadband architecture, access will take place in terms of service nodes (SN) as shown to the left in the figure. The significant difference from the existing solution is that the depth of the access network increases in the sense that traffic will be concentrated in several steps between the end user and the service node. The term *local access point* (LAP) is used for the natural aggregation point covering the customers within a geographical area of moderate size. The point denoted 'broadband access point' (BAP) links a number of local areas on one side to one or more service nodes on the other side. A LAP is linked to only one BAP, whereas an SN can be connected to several BAPs. Hence, a BAP defines the termination of a local access network with the corresponding interface to the service nodes denoted SNI.

End-to-End Service Integration

In order to discuss the role of IP and ATM in the future broadband access network, we need to have a general opinion on whether IP or ATM will be the preferred choice to accomplish end-to-end service integration. This is the most controversial issue in the combined telco and datacommunications industry today. It is important to keep in mind though, that the two industries approach this issue in different ways. Historically, the data-communications industry has been data-oriented without any regard to real-time requirement. It is only recently that this industry has gained interest in supporting real-time applications. The telco industry has traditionally been oriented towards telephony and real-time applications. Even if they have also developed efficient solutions for data communications, it is only recently that this has become equally important in terms of traffic volume and revenue basis. The growth of the Internet is a main contribution in this respect.

Figure 1: Architectural model of the future broadband access network.



The next two sub-sections of this article describe the views on IP and ATM that these two industries represent. The telco industry is ATM-oriented whereas the data-communications industry is IP-oriented. The telco view considers ATM as the universal protocol supporting end-to-end user communication. In the datacom view, the role of ATM is reduced to an underlying transmission technology with local significance. An overlaid IP layer is considered to be the proper level for universal end-to-end user communication. Note, however, that the application of ATM is still beneficial in the latter case since the QoS capabilities can be exploited by the IP level.

We then argue, in the next sub-section, that the IP-oriented datacommunications view is most likely to win the ongoing battle. For comparison, Figure 2 shows the existing situation with parallel technologies at the network platform level. For simplicity, a three-layered view is used in which the intermediate platform level is embraced by an underlying transmission level and an overlying service level.

The ATM-oriented telco view

The key issue regarding the question about IP or ATM is how the Internet and the associated IP are considered. The traditional view of the telco industry is shown in Figure 3 and considers Internet as an important service, but still as a service along with telephony, video on demand (VoD), etc. The Internet is regarded as the killer application that has been missing in order to stimulate the development of a service integrated platform based on ATM. Note also that ATM defines basic transmission as indicated by the lower level in the figure. A number of standards exist that define cell transmission for point-to-point and broadcast media.

The main point of the telco view is that ATM is designed for QoS, thus enabling service integration. In contrast, the existing IP technology

does not have any support for real-time traffic.

The datacommunications view is IP-oriented

The prevailing view of the datacommunications industry is that the IP technology will be further developed to support QoS. This means that resources can be pre-allocated in the network rather than relying on the end-systems to utilise the available resources in the best possible way. The evolution towards IP QoS is a controversial issue, though, and the research efforts go in different directions. There is also a heated debate among the equipment vendors and operators about this issue. Different solutions are being pursued but lack of consensus and standardisation complicate

interoperability across administrative domains. However, the market demand for IP QoS is very strong and it is reasonable to assume that a common agreement will be made.

The situation with a service-integrated IP platform is shown in Figure 4. As opposed to ATM, it is important to note that IP is not a transmission technology. It is rather assumed that IP packets can be encapsulated by a transmission unit offered by an underlying subnetwork. Hence, IP can coexist with other network layer protocols. There are a number of standards for encapsulation of IP packets including both point-to-point links, broadcast networks and switched networks. ATM is an example of the latter case and is increasingly being used for transport of IP traffic.

Figure 2: The existing situation with IP and ATM operating as parallel platforms along with PSTN




Figure 3: The ATM-oriented telco view on service integration





Figure 4: The IP-oriented datacom view on service integration



IP is likely to win

As illustrated in Figures 3 and 4, 'IP over ATM' is a key element regardless of the view taken. The difference is that the telco industry has traditionally considered ATM as the level of service integration. In contrast, the datacommunications industry claims that this should take place at the IP level. We argue that the most likely outcome in the long term is that IP will be used for building a worldwide service-integrated platform with end-to-end guarantees. There are four arguments in favour of this view¹²:

- It is easier to build a worldwide platform based on IP since a large set of underlying transmission technologies can be used. An obvious disadvantage of the telco view is that a homogenous infrastructure based on ATM is required.
- There are implementations of IP for all kinds of end-systems and operating systems, and this software is normally free. Further, TCP/IP is the dominant application development environment.
- IP is designed for distributed routing protocols so that every packet is forwarded on an individual basis. This gives fault-tolerance and simplified administration. In particular, it is easy to build networks crossing administrative domains.
- IP is a simple protocol assuming a minimum of functionality from the network elements. The main responsibility for implementing reliable communication is left to the end-systems.

In sum, these features make IP a flexible networking technology that scales well. A key property compared to ATM is the connectionless nature. Since every packet is forwarded on an individual basis there is no state information associated with the network elements; that is, the

network does not have any notion of which pair of stations are communicating at any specific time. The stateless property ensures scalability, flexibility and fault-tolerance. At the same time, it represents a challenge with respect to QoS. The most important argument in favour of ATM is that resource reservation is simplified when it can be associated with established connections between the communicating parties.

Interworking

Equipped with the earlier discussion about IP and ATM, we are ready to address the two interworking issues raised in the Introduction, and more specifically how IP and ATM can be applied in the future broadband access network. Technology interworking is discussed first followed by a discussion on network interworking.

Technology interworking by ATM

Vendors of broadband access network equipment normally have two types of systems in their product portfolio today. Letting X denote an arbitrary technology like PON, xDSL or LMDS, we distinguish between systems using either ATM or IP as the protocol layer above basic transmission:

- ATM/ X , and
- IP/frame/ X .

In either case, the overlying layer, be it IP or ATM, will ensure interworking of different access technologies. In the first case, ATM cell transmission is defined for the actual medium by a transmission convergence (TC) layer above the physical medium dependant (PMD) layer. For IP-based systems, a variable length frame format needs to be defined first. The frames are in turn used to carry encapsulated IP packets. The two parallel development tracks reflect the general battle

between IP and ATM that is characteristic for the combined telco and datacommunications industry. The interest in IP-based access systems is mainly driven by the need for broadband access to the existing best-effort Internet. The advantages of ATM-based systems are QoS support and increased flexibility. The latter refers to the fact that it is straightforward to use ATM for access to overlaid platforms other than IP.

Despite our general argument in favour of IP we recommend that ATM should be used for statistical multiplexing and service integration in the access network. There are two main reasons for using ATM for technology interworking.

- ATM has better multiplexing properties, in terms of resolution and real-time behaviour, over low-speed links. This is due to the short fixed-length cells used by ATM. The generally longer and variable-length packets used by IP contributes to both latency and jitter resulting in reduced real-time behaviour. This is in particular true for low-speed links as is the case in the access network.
- It is still uncertain when the IP-oriented datacom view will eventually prevail. A service integrated access network should also be able to support other platforms than IP in the short-term and mid-term migration period. Hence, we argue that it is a significant risk to base the access network deployment too strongly on the assumption that a universal end-to-end IP platform will be developed and standardised.

Note that the significance of the first argument is weakened as the line speed increases. Hence, in high-speed backbone networks it is likely that IP will be an appropriate statistical multiplexing technology which can also support traffic with real-time

Figure 5: Two different options for an ATM-based network to operate in an IP-oriented context.

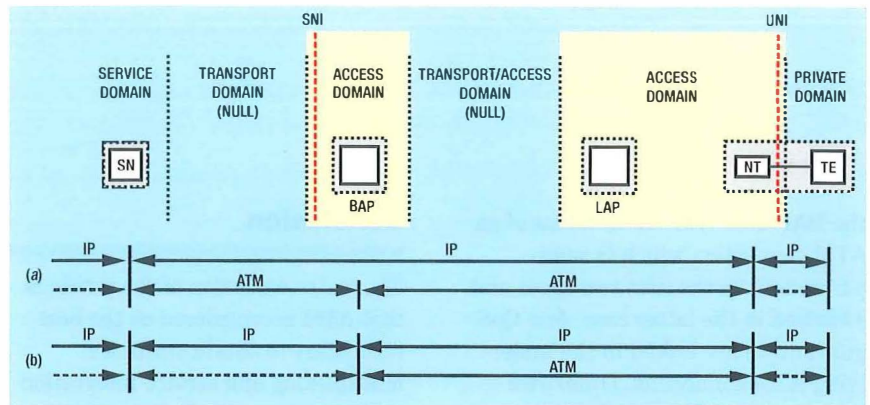
requirements. The key point about the second argument is that the ATM-oriented telco view is not excluded. At the same time, the protocol assumption is compatible with the IP-oriented datacommunications view.

The view that ATM is the preferred protocol for statistical multiplexing in the access network is illustrated by the bottommost double arrow in Figure 1. The small vertical crosslines signify points at which ATM cells are being processed; that is, the figure tells that both BAP and NT will participate in interpretation of the ATM cell stream. Depending on the actual transmission technology it might be additional equipment between these points that also process ATM cells. The typical example is equipment located in a LAP. To simplify the figure, any such intermediate points are suppressed.

Network interworking by IP

As argued earlier, it is most likely that the datacommunications view will prevail in the future so that IP will be the preferred universal protocol for end-to-end user communication. This is indicated by the overlaid level in Figure 1. In sum the figure suggests how an ATM-based access network can fit in a wider IP-oriented context. In accordance with the layered view from Figure 4 we distinguish between a basic transmission level and an overlaid platform level. The dashed line style indicates open protocol issues and the key question is whether IP should be processed at the BAP point or not. An elaborated view of the two different options is explicitly shown in Figure 5 and denoted (a) and (b). Note that the dashed line style is replaced with solid lines in the latter figure to indicate more specific protocol assumptions.

The significant difference between case (a) and (b) is whether interworking with core network operators and service providers should take place at the ATM or IP level. We argue that alternative (b) is



the preferred choice. An evidence of the appropriateness of this view is that all the equipment vendors are now starting to supply 'IP over ATM' nodes that are suitable to be put at the BAP. In case (b) the IP protocol is processed at the BAP, whereas the BAP is transparent for IP in the former case. Hence, ATM transmission must be used also between the BAP and the SN in case (a). Letting the BAP interpret IP as in the latter case gives improved flexibility concerning the transmission technology being used between the BAP and the SN. But, more importantly, it is possible to let the BAP be equipped with a richer set of functionalities as discussed in the next section. The point is that the IP-based approach gives the access provider better management and control capabilities at the BAP point.

Role of Access Network

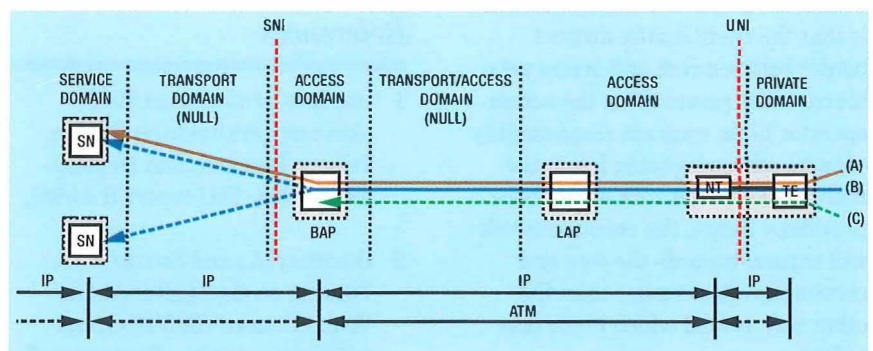
Anticipating that IP will have a dominant role for end-to-end communication in the future, the ATM-based access network will act as a carrier of IP traffic. Running IP over ATM can be accomplished in a number of ways depending on what is considered to be the role of the access network. Since both IP and

ATM are equipped with rich functionality concerning control and switching, new opportunities arise. Hence, we ask if the provision-oriented role of the access network as traditionally carried forward by the International Telecommunications Union (ITU) is obsolete. Different industrial groups like the ATM Forum⁵ and ADSL Forum³ propose novel solutions implying a new understanding of the access network.

Figure 6 shows three interpretations of the role of the future access network. The cases denoted (A)-(C) are distinct in the way customers are associated with SNs. A solid line style is used to indicate fixed associations whereas a dashed line style is used to signify an association that is established on-demand or automatically. Case (A) corresponds to the traditional provision-oriented view with a fixed association. Case (B) represents an evolutionary step in terms of a session-oriented view with free selection of service provider. Case (C) is the most radical scenario. By letting the access network be equipped with routing capabilities it can interact closely with core networks.

A common feature of all interpretations is the existence of an association between the customer and

Figure 6: Different interpretations of the role of the access network.



the BAP. It is realised in terms of an ATM connection which is semi-permanent in the first two cases and switched in the latter case. Any QoS guarantees are linked to the underlying ATM connection. Otherwise there is a significant difference between cases (A) and (B) on one side and case (C) on the other side. In the latter case, there is no association between the BAP and the SN.

The key issue regarding the role of the access network is what kind of functionality the BAP is equipped with at the IP level. In case (A) no processing is performed except reassembly and fragmentation. In case (B) the PPP protocol¹¹ is typically used to perform provider selection and accounting, authentication and authorisation for the access operator. Hence, the BAP takes the role as a generic access server. For each established session, the IP forwarding table in the BAP is set accordingly. In case (C) the BAP is equipped with ordinary routing capabilities allowing the access network to interact closely with core networks. This means that IP packets are routed individually without any association to an SN. The effect is that the access network will comprise an underlying subnet in one or more IP networks. In this context, techniques like classical IP over ATM (CLIP)¹³, next-hop resolution protocol (NHRP)¹⁴, multi-protocol over ATM (MPOA)¹⁵ and multi-protocol label switching (MLPS)¹⁶⁻¹⁹ for providing IP over ATM become relevant.

A consequence of revising the understanding of the access network is that the traditionally distinct border between core and access gets blurred. The point is that the access operator takes common responsibility of tasks otherwise being left to the individual core network and service providers. Hence, the access network will expand towards the core and service providers rather than the other way around which is the case today.

Conclusion

The main conclusion of the article is that ATM is considered as the best technology to obtain statistical multiplexing and service integration in the access network, thus also technology interworking. This is partly due to an argument that ATM gives better QoS performance over low-speed links. But it is also due to the fact that an ATM-based access network is applicable regardless of whether IP or ATM will constitute the universal platform for end-to-end user communication. It involves less risk to start deploying an ATM-based access network since it can also be used in an IP-oriented context.


We argue that the IP-oriented view carried forward by the data-communications industry is the most likely solution for end-to-end service integration. Hence, network interworking is best handled at the IP level. The connectionless nature of IP makes it a scalable, flexible and fault-tolerant networking technology. A pertinent question is how an ATM-based access network can fit in a wider IP-oriented context. The point is that the combination of IP and ATM opens new opportunities affecting the role of the access network. A likely evolution is that the broadband access network will initially be provision-oriented in the same way as the existing narrowband network. The next step will be to implement free selection of service provider per session. The ultimate step is to let the access network become a true IP subnet interworking closely with arbitrary subnets in the core domain.

References

- 1 IMS, LEIF, *et al.* Akseess 2005—akseessnettarkitekturer. Kjeller, Telenor Research and Development, 1998 (FoU report R 44/98).
- 2 DONNELLY, A.; and SMYTHE, C. A Tutorial on the Digital Audio-Visual Council (DAVIC) standardisation activity. *Electronics & Communications Engineering Journal*, 1997, pp. 46–56. <http://www.davic.org>.
- 3 ADSL Forum. *ADSL Forum System Reference Model*. Fremont, 1996. (ADSL Forum TR-001.) <http://www.adsl.com>.
- 4 ADSL Forum. *Requirements & Reference Models for ADSL Access Networks: The 'SNAG' Document*. Fremont, 1998. (ADSL Forum TR-010.) <http://www.adsl.com>.
- 5 ATM Forum. *Residential Broadband Architectural Framework*. 1998. (ATM Forum af-rbb-0099.000.) <http://www.atmforum.com>.
- 6 FSAN. *Full Services Access Network Requirements Specification*, draft G. Ipswich, 1997. <http://www.labs.bt.com/profsoc/access/>.
- 7 DVB. *A guideline for the use of DVB specifications and standards Digital Video Broadcasting (DVB)*. (DVB Report tr-101-200.) <http://www.dvb.org>.
- 8 ITU. *Framework recommendation on functional access networks (AN)*. Geneva, 1995. (ITU-T, Recommendation G.902.)
- 9 IMS, LEIF. *Wireline Broadband Access Networks*. *Teletronikk*, No. xx, 1999.
- 10 DE PRYCKER, MARTIN. *Asynchronous Transfer Mode*. Ellis Horwood, 1993.
- 11 IETF. *Point-to-Point Protocol (PPP)*. 1994. (IETF RFC 1661/STD 51.)
- 12 HUITEMA, C. *Routing in the Internet*. Prentice Hall, 1995.
- 13 IETF. *Classical IP and ARP over ATM*. 1998. (IETF RFC 2225.)
- 14 IETF. *NBMA Next Hop Resolution Protocol (NHRP)*. 1998. (IETF RFC 2332.)

- 15 ATM Forum. Multi-Protocol Over ATM. 1997. (v. 1.0. ATM Forum afmpoa-0087.000.)
<http://www.atmforum.com>.
- 16 IETF. A Framework for Multiprotocol Label Switching. 1997. (IETF Internet Draft, draft-ietf-mpls-framework-02.txt.)
- 17 IETF. Multiprotocol Label Switching Architecture. 1998. (IETF Internet Draft, draft-ietf-mpls-arch-02.txt.)
- 18 IETF. Use of Label Switching With ATM. 1998. (IETF Internet Draft, draft-ietf-mpls-atm-00.txt.)
- 19 DAVIE, BRUCE, *et. al.* Switching in IP networks. Morgan Kaufman, Inc. 1998.


Biographies



**Frode Beckmann
Nilson**
Telenor

Frode Beckmann Nilson holds a Masters degree (1993) and a Doctoral degree (1998) in computer engineering and data communication from the University of Oslo. The academic research focused on methods for performance evaluation of high-speed networks. He is currently employed at Telenor Research and Development as a research scientist working on future full-service network architectures.

Frode Beckmann Nilson can be contacted at
frode-beckmann.nilsen@telenor.com



Leif Aarthun Ims
Telenor

Leif Aarthun Ims is senior research scientist in the Strategic Network Development group of Telenor Research and Development. Since 1994 he has been working on access network development in several RACE, ACTS and EURESCOM projects, and recently as project manager for long-term access network development in Telenor. He is now project leader of the EURESCOM project P901. He is editor of the book *Broadband Access Networks* published last year and has published more than 50 papers in international journals and conferences.

Leif Aarthun Ims can be contacted at leif-aarthun.ims@telenor.com

UK Digital Terrestrial Television

The UK broadcasters for cable, satellite, and terrestrial television have all recently launched digital services to deliver new channels and services to the home. Early in 1998, BT Broadcast Services won the contract to supply to the majority of the UK terrestrial broadcasters the network that would distribute the digital television signals from the studios to the transmitters located around the UK. This contract secured BT's position in the UK television distribution market. To win this business, BT Broadcast Services had to commit to deliver the DTTV network to a very tight schedule so that the broadcasters could achieve live programming for the all-important Christmas 1998 market.

Introduction

The continuing convergence of broadcasting, computing and telecommunications means that the future entertainment landscape will continue to undergo a sea of change. For 50 years, television has become the dominant entertainment medium into the home, delivered by analogue techniques through cable, terrestrial and satellite.

Early in the 1990s, work began on video-coding techniques that were set to transform the television industry. A group of experts, known as the *Moving Picture Experts Group* (MPEG) within the International Standards Organisation (ISO), undertook the work. The arrival of MPEG-2 meant that broadcasters could now extend existing studio digital techniques on to the final distribution transmission to the home receiver.

Europe was quick to realise the opportunities presented and formed the Digital Video Broadcasting (DVB) project in 1993. At the outset, there were more than 80 member organisations, representing broadcasters, telecommunication operators, industry and regulators. This number has now grown to over 200 members from 30 countries around the world. The DVB project established technical committees to specify all aspects of digital video broadcasting from the studio to the home. The DVB project has made as much use as possible of existing or emerging standards in drawing up its own recommendations. Its major decision was the adoption of MPEG-2 as the video compression standard. The work of the DVB has resulted in three specifications:

DVB-C for digital cable broadcasting, DVB-S for digital satellite broadcasting, and DVB-T for digital terrestrial broadcasting.

Digital transmission from end-to-end can improve the quality of the final content viewed at home. It also has the opportunity to accelerate the convergence of broadcasting, computing and telecommunications to bring about totally new services to the home. New services could include such applications as interactive television, where viewers could participate to a much greater level in programmes, home shopping, browsing the World Wide Web (WWW), and other information services. Many of the new services will come about by the integration of Internet protocol (IP) data within the broadcast data streams. The DVB project had foreseen this and had included in its specification means by which IP traffic could be carried within a DVB stream.

The UK government announced the first licences for UK terrestrial digital television in 1996, and the first services went live in November 1998. BT was awarded the contract from five of the six broadcasters to deliver the distribution network in March 1998.

This article outlines the standards for digital television and goes on to describe the requirements of the UK broadcasters for digital television. The design of the BT distribution network for digital terrestrial television is described, together with aspects of the project management. Finally, the article describes new IP services that BT is involved in through the experience and knowledge gained from involvement with DVB.

Standards for Digital Broadcast

The adopted digital compression standard for DVB is MPEG. The MPEG coding involves the reduction of the 'raw' video data rate from 180 Mbit/s to around 4 Mbit/s. The processes used to achieve this level of data reduction involve both 'loss-less' and 'lossy' compression techniques. The 'loss-less' technique means that the original data is recoverable from its compressed form. The 'lossy' technique exploits the human psycho-visual system to introduce artefacts that will not be perceptible when data is recovered from its compressed form; temporal difference and motion estimation techniques form significant parts of this process. The audio data compression is also included in the MPEG specifications, reducing the 'raw' audio data from around 2 Mbit/s down to around 300 kbit/s, while maintaining CD-like audio quality.

The essential feature of the MPEG compression is that the costs of the encoder and decoder are highly asymmetric, with the majority of cost in the encoder, enabling cost-effective home decoders. Consequently, ongoing improvements in algorithms only need to be applied to the encoder, allowing improvements to be made in quality without affecting the consumer.

The evolving Multimedia Hypermedia Experts Group (MHEG) standard is used for the enhanced teletext services.

The DVB asynchronous serial interface (ASI) was developed as one of several options for the interconnection of DVB transport streams. Over the last two years, support for the ASI interface has been growing throughout the industry due to the simplicity of the connection (single 75 ohm BNC)

and its high data capacity (over 200 Mbit/s of data). BT decided on the ASI interface as its preferred DVB connection in early 1997. As well as the advantage of 'looking like' a conventional telecommunications interconnection, the fact that it is a defined video interface rather than a pure telecommunications data-pipe has made regulatory issues in the UK much more clear cut and easily defined. Due to its infancy and lack of absolute definition in some areas, there have been some, relatively minor, inter-manufacturer compatibility issues.

Service information (SI) is a description of the data associated with the video and audio, and is carried within the DVB data stream. It provides the equivalent of an enhanced teletext capability to allow users to determine current and future programmes across channels.

UK Digital Terrestrial TV

In the UK, it is estimated that the number of homes with a television is 25 million. Of these, approximately 30 per cent receive their programming via cable or satellite. The other 70 per cent, or 17 million, receive their content via terrestrial means. Before the start of digital TV, UK viewers had the choice of five analogue channels from four UK terrestrial broadcasters (Table 1).

The UK viewers are also offered regional variations of programming content throughout the day. These regions supply local news and information along with local interest programming. In the case of ITV, there are 16 regions.

For digital terrestrial TV the requirements are based around what is already offered to the UK viewers today; that is, a number of channels offering national and regional content, accessed by their existing aerial.

Benefits

The bandwidth of an existing analogue channel (8 MHz) can support a digital transmission rate of approximately 24 Mbit/s. This can support more than five channels and provides broadcasters with the opportunity to offer additional content choice. It also gives the broadcasters the opportunity to seek additional revenue streams by offering a range of services that are both free-to-air and subscription-based. As more channels can now be transmitted within the same bandwidth as one analogue channel, the spectrum can be used more efficiently, allowing the government to release spectrum for new uses.

The home viewers also benefit from digital terrestrial TV, in that their choice of content increases from the current five channels to at least 30 channels. The challenge for the broadcasters is to ensure consistent quality of content across the increased number of channels. All of these new channels are accessible through existing aerials, with an appropriate set-top box or digital TV.

Additionally, digital terrestrial TV brings with it potential improvements in sound and video quality. Programmes have CD-quality audio associated with them, similar to NICAM stereo. The video quality is more consistent with less degradation due to ghosting and noise. The nature of video degradation changes, with any artefacts now being dynamic and picture-content dependent, due to the compression system. In marginal reception areas, the video is more likely to break up.

UK multiplexes

The government awarded six multiplex licences in 1996, which went to the incumbent broadcasters as well as new consortia (Table 2). A multiplex takes up the equivalent bandwidth of an analogue channel (8 MHz). The total data rate available in each multiplex is 24.1 Mbit/s, which supports at least five programme channels. Each video programme is typically coded at around 4 Mbit/s, leaving capacity for the audio and other programme data.

Table 1 UK Analogue Terrestrial TV

Broadcaster	Programme Channels
British Broadcasting Corporation	BBC1, BBC2
Independent Television Companies	ITV, Channel 4, Channel 5

Table 2 UK Digital Terrestrial TV

Broadcaster	Multiplex	Programme Channels
British Broadcasting Corporation	Multiplex 1	BBC1, BBC2, BBC News, BBC Choice, BBC Text, BBC Learning (late 1999)
Digital 3 & 4	Multiplex 2	ITV, Channel 4, Film 4*
SDN	Multiplex 3	Channel 5, S4C + TBA*
British Digital Broadcasting	Multiplex 4*	Carlton Cinema, Sky 1, Sky Sports 1, Sky Premier, Cartoon Network, Eurosport
British Digital Broadcasting	Multiplex 5*	UK Gold, Granada Plus, Breeze Men & Motors, Sky Sports 3, Sky Moviemax
British Digital Broadcasting	Multiplex 6*	UK Play, UK Style, Horizon, Carlton Select/FN, Carlton Kids, Shop
* subscription channels		

It is possible to support more than five programme channels by statistically multiplexing the video programmes on an intra-multiplex basis. This can increase the number of programmes carried in a multiplex to seven or eight. To perform statistical multiplexing, the broadcasters need to utilise sophisticated look-forward MPEG encoders that can vary the bit rate according to the entropy of content.

The radiated power of each multiplex is different. Multiplex 1 has the greatest power and hence the largest coverage of the UK. The power of the remaining multiplexes decreases with increasing multiplex number. This means that the free-to-air channels, such as BBC1 and BBC2, are more widely available.

The Independent Television Commission (ITC) has mandated that the current level of analogue regional programming must be maintained in the digital world. All DTTV multiplexes source at least some of their content from a central play-out point. However, some multiplexes differ from the others in that they involve the injection of local content, based on the current analogue regional transmissions.

BT Distribution Network

BT has deployed a UK national DTTV network for five of the six multiplex operators: ONdigital, who run three multiplexes, the BBC and D3&4 (see Figure 1). BT contracted with Castle Transmission International (CTI) for the ONdigital and BBC network delivery, and with D3&4 for their network requirements.

The network is predominately fibre-based and links 30 play-out-centres to ultimately 81 transmitter

sites, which are located throughout the UK. At service launch on the 15 November 1998, all of the studios and 23 nominated transmitter sites were required to be operational.

Design considerations

The network design was influenced by a number of requirements, one of which has already been highlighted earlier: the need to preserve the existing regional programming requirement to fulfil the simultaneous obligation of the existing broadcasters. The network had to provide a number of points at which programme content could be injected. It also meant that the network had to be configurable to route the programmes to the correct transmitters.

Another major factor was the service availability targets set, which meant that the network had to achieve a play-out-centre-to-transmitter availability of 99.99% in a rolling six months window. This equates to approximately a 25 minute outage within any six month period. This level of service could be achieved only by having duplicate signal paths through the whole distribution network. Also,

Figure 2—BT's UK DTTV network


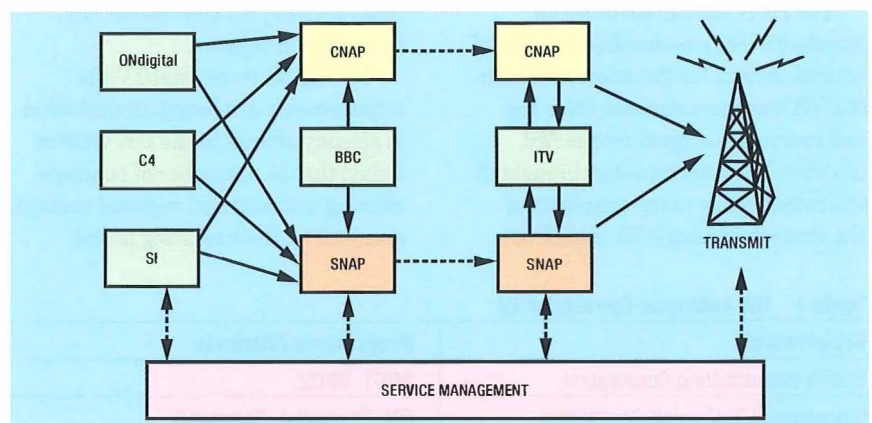


Figure 1—UK DTTV deployment

full network management and monitoring were necessary to determine whether that condition was being met.

Additionally, multiplex operators had their own connectivity requirements that had to be taken into account in the network design.

Network design

In order to achieve the 99.99% availability target, full redundancy in the transmission path was essential. This, together with the need to preserve regional insert points led to a very significant distribution network capacity requirement, which was most easily satisfied by deploying a fibre-based network. The broadcasters also preferred fibre due to a perceived higher reliability than satellite or microwave solutions.

Full redundancy in the network also meant that the redundant feed had to have total separation. Therefore, as well as having physically separate fibre routing, physically separate BT buildings were used as node points for the redundant feed. At the play-out-centres and transmitter sites there was also duplication of equipment.

The network consists of a mesh of access points where access into or out of the network can be effected (Figure 2). These access points are known as *core network access points* (CNAPs) and *separacy network*

Figure 3 – ASI-mux

access points (SNAPs), and are located geographically close to the regional studios of the broadcasters and the transmitter sites.

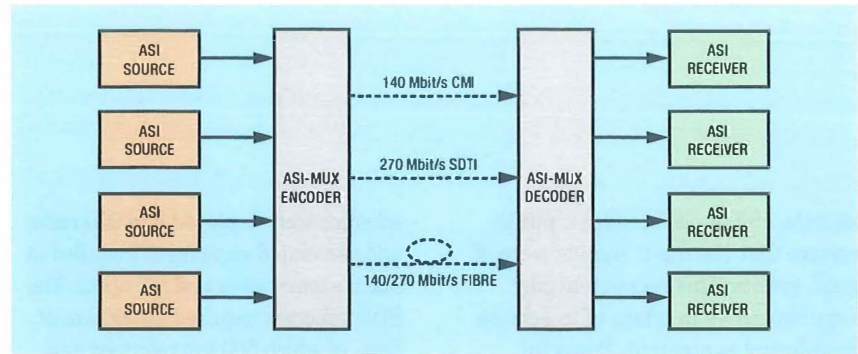
To make the network cost-effective it was necessary to deploy a multiplex at the input, output, and nodal points of the network that could aggregate a large number of MPEG transport streams (TSs). This maximises the utilisation of the synchronous digital hierarchy (SDH) capacity of the network. A suitable multiplex was unavailable, so a device was designed and implemented by BT Laboratories, now known as Adastral Park. The asynchronous serial interface (ASI) was selected as the most robust and easy to handle interface of the available standards. The ASI-mux was developed capable of supporting up to 16 independent ASI inputs. This ability of the ASI-mux to support a large number of TSs, and therefore achieve a high fill factor on the telecommunications link, was critical to achieving the final network design requirement of a cost-effective solution.

ASI multiplex

The DVB distribution requirements for DTTV described earlier led to the search for a network interface that could provide the following:

- cost-effective aggregation of multiple TSs;
- a variety of network interface rates;
- long-range direct fibre connectivity;
- comprehensive stream analysis and service management; and
- the complete isolation from each other of TSs being carried.

The interface chosen was the asynchronous serial interface (ASI). This is a 270 Mbit/s serial data stream common to much of the broadcast equipment. The decision was made to develop an in-house solution at BT Laboratories, because



at that time the above requirements had not been realised in any commercially available equipment. This became known as the *ASI-mux*, see Figure 3. In a very short development period, of just six months, the first ASI-mux had been realised.

The ASI-mux can be interfaced to several sources that comply with the ASI-interface specification and several network connectivity options are offered. Some of the essential features that have been included in the ASI-mux are detailed below:

- a single unit can carry up to four TSs;
- units can be cascaded to provide an expansion capability that aggregates up to 16 TSs;
- 140 (CMI), 270 (SDTI), 34 (HDB3) interfaces;
- ability to monitor TS sync, RS and DVB errors;
- ability to monitor ETR290 TS parameters;
- optional insertion of RS FEC to TS packets;
- full error performance monitoring of network transport using cyclic redundancy checks;
- network and regional ID tagging for all TSs;
- input bit-rate policing;
- comprehensive management interface;
- built-in protection switching capability; and
- output ASI stream data re-timing.

The choice of network interface rates made ensured flexibility of network technology choice and the optimum use of direct fibre routes. The 140 Mbit/s interface provides for direct connection to SDH STM1 systems which form the majority of the network while allowing also for use of plesiochronous digital hierarchy (PDH) links. The 140 Mbit/s CMI data can also easily be driven optically over fibre. The 270 Mbit/s SDTI interface provides the opportunity for a higher capacity of traffic to be carried over fibre while maintaining compatibility and hence routability through SDI equipment.

In any system where data has to be transferred between one party and another, it is essential that the highest level of integrity checking can be performed on the data at that boundary. The analysis carried out by the ASI-mux focuses on three main areas:

- the error performance of the data stream;
- the primary MPEG TS parameters; and
- the data rate of the TS.

The ability to ensure isolation between different transport streams that have to be carried within the same network transport is essential. Several tools have been developed to achieve this aim. The aims of these tools are to ensure that each TS is delivered to the correct destination, and that a TS violating its allocated capacity, does not affect other services being carried.

Service management

The ability to in-service monitor and manage the network was an essential design requirement. Monitoring was clearly required at all transmitters to determine the integrity of the output

signals, and at each studio input to ensure that the input signals were of good quality. This was potentially very expensive in terms of in-service monitoring equipment. Powerful monitoring and alarming capabilities were therefore incorporated within the basic ASI-mux design. Among the many parameters the mux is capable of measuring include:

- input/output TS presence or loss;
- incoming TS bit rate;
- ETR 290 layer one and key layer two parameters;
- telecommunications link bit error rate; and
- errored and severely errored seconds.

To collect the monitored information from the ASI-mux, a BT-designed intelligent site controller (ISC) is collocated with each ASI-mux site. The ISC also has the ability to download the appropriate site-specific operating configuration for the ASI-mux. An overlay primary X.25 and resilient back-up ISDN network provide a low-rate data path from every site back to BT's broadcast-specific 24-hour 365-day service management centre (SMC). The data is filtered and presented to the SMC personnel as a series of user-friendly software tools that have been developed to present prioritised alarms and aid fault location.

Project scale

The contract required BT to deliver the network in three phases. For phase 1, a number of studios had to be connected to 23 transmitter sites around the UK. For phase two the remaining studios and a further 58 transmitter sites had to be connected up. Finally in phase 3 the remaining transmitter sites would be brought on line by December 1999.

From the date when the contract was signed, BT had 6 months in which to complete phase 1. The network for phase 1 required that 250 SDH

schemes were deployed and 250 racks and associated equipment installed at 122 customer sites and BT nodes. The SDH schemes required 20 000 km of fibre, of which 500 km required new ducting, and three fixed radio systems. Approximately 1300 ASI-mux have had to be manufactured, tested and deployed to the customer sites and BT nodes, which only weeks before the contract signing had been at the prototype stage. All of this presented BT with a significant challenge to meet its contractual obligations.

A dedicated project management team was put in place that represented all sections of BT. Extensive use of audio conferencing and the intranet were made to ensure all involved were briefed, had the most up-to-date information and progress monitored. Project planning also had to take into account when site access could be obtained: in a number of cases delivery of the network could not be completed until the customer sites had finished their provision of accommodation.

The network was implemented with the CNAP side first, to allow preliminary overall service testing with no resilience. One month later, the additional SNAP network was deployed to allow for full testing of the service and add full resilience. At the same time, the service management was brought into service.

The introduction of the service management involved sophisticated alarm management. The ASI-mux is capable of producing a considerable number of alarms per unit of which there are approximately 1300 in the network distributed over 122 sites. To avoid a flood of alarms to the service management centre (SMC), which could potentially slow down the speed of the management system, the ASI-mux had to perform alarm masking which needed to be configurable to allow customised profiles.

IP Services Through DVB

The DVB standard has a data broadcast specification that can be applied to protocol encapsulation,

which is the transmission of data communication protocols, such as IP, over DVB networks. The DVB specification is very attractive for carrying data due its low bit-error rate, within the payload, of 1×10^{-11} .

The protocol encapsulation method devised has been optimised for the transmission of both IPv4 and IPv6. The encapsulation supports fragmentation of large IP datagrams for transmission within MPEG-2, private section payloads. The IP datagrams can be transmitted to a single receiver on the DVB network, a group of receivers or all receivers on the DVB network; this is equivalent to unicast, multicast, and broadcast in a conventional IP network. The protocol encapsulation also supports security, with encryption of the packets and MAC address scrambling.

Service opportunities

Recent developments of the Internet have seen the growth of applications based on push technology. Push technology is the process by which content is broadcast from a server, and the client end 'tunes' in to the particular channel to receive the content. The traditional pull approach is where users, through their browsers, search for the content they want and then, when the appropriate web page is located, request the content to be transferred from the server. Push technology is appropriate for use in such applications as news services that may make use of audio and video streaming over the Internet. In many cases, the applications that utilise push technology require sustained data rates across an IP network which, for today's public Internet, is difficult to achieve.

DVB networks, especially those deployed over satellite, offer a significant opportunity to transport IP traffic associated with push applications on the Internet. A hybrid of traditional public Internet connectivity with additional IP traffic via satellite makes for exciting new services to business and the home.

Figure 4—IP satellite delivery via DVB-S

Convergence 1 trial

A trial has been underway in the UK to demonstrate such service opportunities of combining high-speed IP data delivered by satellite to users of the Internet today. Convergence 1, a trial involving Eutelsat, BT and Easynet, delivers TCP/IP data streams to the home via DVB-S. Users of the system make use of a PC, dial-up link to their Internet service provider (ISP), and browser to access the system. This provides connectivity to the World Wide Web (WWW) to send data requests. The delivery of TCP/IP data arrives back from the ISP via a DVB-S stream delivered by satellite (Figure 4). Some modification of the TCP/IP window size is necessary to accommodate the extended round-trip delays associated with satellite links.

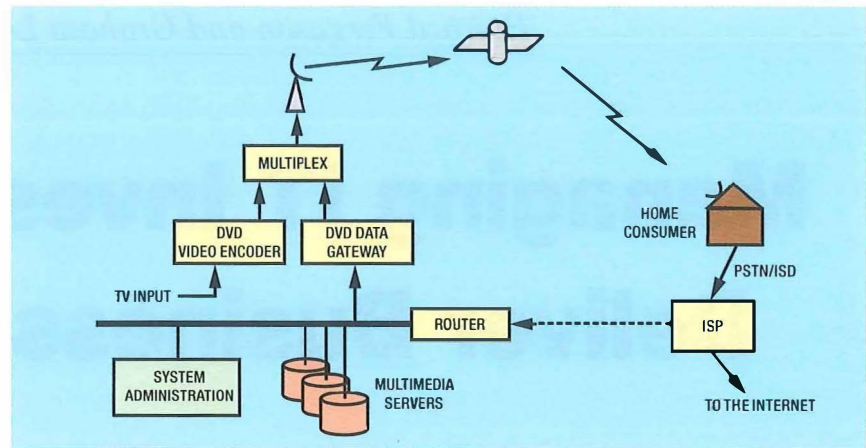
A user of the service would browse the Internet as normal to select the content they wish. When the appropriate web page is reached, the data is delivered back towards the ISP; at the ISP, it is routed to a satellite uplink where there is an IP-to-DVB gateway, and broadcast to the DVB reception equipment connected to the PC at home.

The user would benefit from the high-speed delivery when the ISP has local content of interest, such as audio and video content or large files. The ISP would need to cache content of the WWW to ensure its users are benefiting from the high-speed data delivery.

The sustained IP data rates delivered via DVB could easily be 1–2 Mbit/s. Such a data rate makes it attractive for delivering large files that would otherwise take hours to download using normal modem links to an ISP or for delivering content such as MPEG-1 video clips.

Conclusions

Digital television has just entered service at the end of the 20th century. The stage is set for a new era of services to the business and home that can combine the power of Internet applications with broadcast-quality audio and video.



BT has played a major part in the provision of the UK's national terrestrial network for digital terrestrial television, demonstrating its capability in designing, provisioning, and managing such digital transmission networks.

Acknowledgements

The launch of the BT UK DTTV network was made possible by effective teamwork throughout BT. Major support for the project came from BT Broadcast Services, *networkBT*, *networkbuild* and Systems Engineering. Kevin Rose of the Internet and Multimedia Applications unit in *networkBT* led the overall coordination of the teams.

Biographies

Ian Parke
Networks and Information Services,
BT UK

Ian Parke graduated in 1981 with an honours degree in Applied Physics from Durham University. He received his Masters Degree in Digital Techniques from Heriot-Watt University in 1988. He joined the then BT Laboratories in 1986 to work on the design of a motion estimator chip for the first H.320 videoconferencing codec. He went on to study variable-bit-rate video-coding techniques and represented BT at ITU standards meetings. In 1991, he became a member of the ISO-MPEG group that developed video compression standards, with a specific interest in layered coding approaches for MPEG-2. In 1993, he joined the videotelephony services group involved in the launch of the PC videophone. In 1996, he joined what was then Design

and Build to manage the broadcast and conferencing programme. During 1998, the programme was involved in the delivery of BT's UK DTTV network. He is currently working in BT's Internet and Multimedia Applications sector as the portfolio engineering manager responsible for developments on conferencing and collaboration products and services. He has represented BT on the board of the International Multimedia Teleconferencing Consortium for the past four years and is currently Executive Vice-President.

Ian Parke can be contacted at ian.parke@bt.com.

Andrew Rayner
Networks and Information Services,
BT UK

Andrew Rayner graduated in 1988 with an honours degree in Electronic Engineering from Sheffield University. Since then, he has worked at BT Laboratories, now Adastral Park, in the broadcast television area. He currently leads the team that supports BT Broadcast Services by the provision of technical consultancy, specification, evaluation and bespoke development. He has been involved in all aspects of the evolution of digital television over the last 10 years, focusing recently on the realisation of BT's DVB services which includes delivery for both DVB-T and DVB-S. He is a regular technical presenter at broadcasting conferences and courses.

Andrew Rayner can be contacted at andrew.rayner@bt.com.

Managing IT Investment to Deliver Business Benefit

This article discusses the evolution of IT and IT projects and identifies four guiding principles that should be taken into consideration if IT is to deliver real business benefit. These principles are concerned with the way IT projects are initiated and managed, clarity over measuring benefits at the business level, how the organisation builds IT capability such that new business opportunities can be identified and, finally, how it supports its users of IT.

IT and Business Change

Setting the scene—IT can do anything

In today's competitive world the effective use of information technology (IT) as an embedded element of competitive business strategy is critical. IT has the capability to change whole industries and, with the advent of the Internet, it even has the capability to change social behaviour with more and more products being purchased via the Internet. The Internet is becoming a primary channel for delivering applications, finding information and even for conducting business.

However, the IT industry has an alarming number of projects that fail to deliver and some projects may even be abandoned completely after enormous wasted expenditure. The challenge facing IT managers is to ensure that the projects for which they are responsible are successful, and that this success is recognised and measured as a key and integral part of the overall business strategy. The technology is there, unit costs for computing are falling, and today's software almost guarantees success—or does it? In some cases a project that is a complete technical success runs into problems when deployed, often as a result of failure to ensure that the business embraces the new system. The project then inevitably fails to deliver any significant benefit. Problems may develop in subtle ways; today a common complaint is information overload brought about by the success of e-mail. There is no doubt that e-mail is very effective at moving messages around an organi-

sation, but has it really improved the responsiveness of the organisation? Are people in the organisation better informed than they were before widespread usage of e-mail?

The key consideration is what benefit does IT bring to the business, and related to this is the question of specific importance to the IT managers: is this benefit visible to the senior management of the company? In many companies IT, budgets are under severe pressure and the perception of the business may be that IT is poor value for money or, even worse, the perception might be that IT fails to deliver. The perception may or may not be correct, but the key point is that benefits need to be measured in the business context. This is easy to say and can be very difficult to do because the nature of today's IT projects is to do with business transformation, and the role of IT is to support such a transformation. However, if IT is planned as an integral part of the wider business transformation, the measurement of benefit and the allocation of a component of this to the IT can be very much simplified. The key measurement is the success or otherwise of the transformation, with IT being positioned as one of the enablers of this change. The challenge for the IT manager is then to play a full role in engaging and working with the business sponsors, thus bringing IT into the transformation project and then delivering the IT components on time. These IT components will be highly visible to the organisation, which will then in turn bring IT costs into the context of the overall business benefits of the business transformation.

Figure 1—Evolution of IT projects

ERA	TECHNOLOGY	EXAMPLE	SCOPE
AUTOMATION	CENTRAL MAINFRAME	PAYROLL	ONE FUNCTION
INFORMATION	MAINFRAME/DISTRIBUTED	MRP	CROSS FUNCTIONAL
TRANSFORMATION	DISTRIBUTED/INTERNET	AMAZON.COM	CROSS INDUSTRY

The four guiding principles for managing IT

We have already discussed the need for IT to be seen as an integrated part of managing a business and that it is critical to measure business benefit from projects. This forms the basis for the first two of the four guiding principles for managing IT. In order to achieve the level of integration between business strategy and IT, it is essential that:

- IT is driven by business managers (literate in IT) and not by IT managers.

Having set the first of the guiding principles the second flows naturally:

- The benefit of projects should be measured in business terms with the relationship between the IT deliverables and the business transformation built into the overall project.

The other two guiding principles relate to the way in which the organisation is developed to both foster innovation, generate widespread awareness of IT and what can be achieved with today's IT. In order to raise awareness of IT and the opportunities that it offers, there must be a broad base of IT awareness throughout the business. A good way to achieve this is to:

- embed IT capability into the business, and have IT people reporting into the line on all business change projects. (The third guiding principle.)

Finally, having deployed new IT systems it is essential that users have been well trained and are subsequently well supported, thus the fourth of the guiding principles is:

- Derive maximum benefit from IT by providing users with world-class support services.

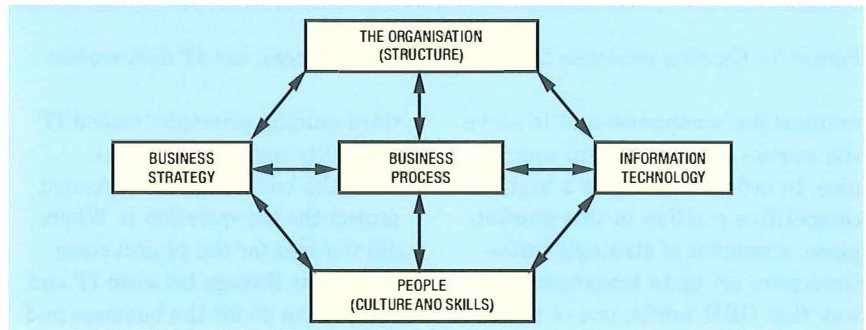


Figure 2—Business model

IT Driven By Business Managers

IT has evolved in a cumulative way such that each era of projects builds on the previous era. The result is that the complexity and interdependency of projects increase. The evolution is summarised in Figure 1. Initial projects automated manual processes such as calculating pay. There were few interdependencies and the systems were designed as copies of manual processes. In the next era the objective of the systems was automation and these operated across a function. Good examples were early airline reservation systems and manufacturing resource planning (MRP) systems in a manufacturing environment. The world has now moved on, and the major IT projects of today are about business transformation. These projects may have an impact across a whole industry and will certainly have a major impact across a company. This is shown by the model in Figure 2, which shows that the dynamics of a company are based on the interplay between the organisation, its processes, people, business strategy and IT systems. A significant change to one has an impact on the rest and in the case of a business transformation all are affected. The change could (theoretically) origi-

nate in any of the components; for example, a change starting in business strategy will have an impact on people, processes and systems, while a major change to systems will have an impact on people and organisation, and may then create opportunities for a new business strategy.

However, a common factor in most examples of highly successful business transformations is that they are driven by the business (business strategy) with an input from other components such as IT. In the case of IT the input is often related to the capability (that is, what is possible using IT)—while it is then the 'business' that recognises the opportunity created and turns it into a business proposition. Wherever the change originates, IT has a critical role, and for this reason IT must be driven by business managers and not by IT itself.

Measure Business Benefit, Not IT Deliverables

The second guiding principle (measure benefits in business terms) is described by looking at a project to deploy knowledge management as a case study.

In BT Global Business Markets (GBM) there is a need to compete with the world's leading telecom-

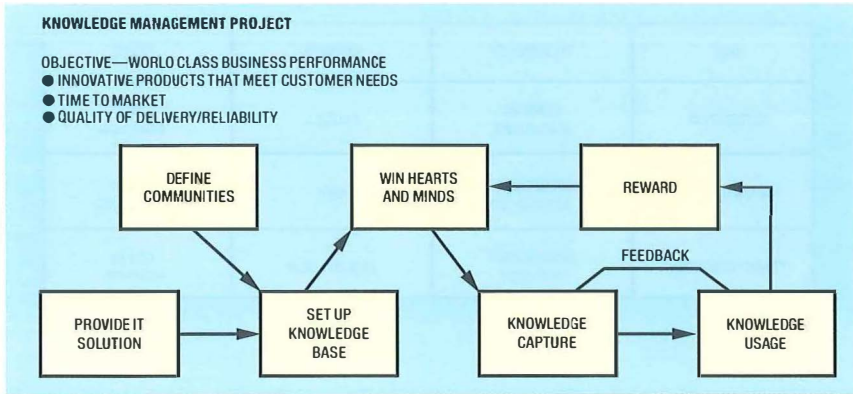


Figure 3—Guiding principle 2: Measure business benefit, not IT deliverables

munications companies and to serve the world's most successful companies. In order to maintain a highly competitive position in this marketplace, a number of strategic initiatives were set up to transform the way that GBM works; one of these strategic initiatives was concerned with how the organisation learns and uses its collective knowledge. The project is using knowledge management and it is one of a group of projects under an overall transformation project called *A Great Place to Work*. It is business led, but it has IT people on the project team. A summary of the project is shown in Figure 3. While the project is dependent on IT for the solution, the objectives of the project are much broader, and success will be measured at the business level in terms of improving the time to market for new products and market share achieved. The initial measures of success will be concerned with the quality of information and, more importantly, the extent to which it is used. This example shows the need for project success to be measured in business terms. The IT component of the project is relatively simple and deploying the IT on time is only a very small part of the project; the hard part will certainly be in getting the behavioural change into the whole organisation, hence level of usage is the key early measure.

Embed IT Capability into the Business

Building on the above example of the knowledge management project, the same project is used to provide the framework for describing the

third guiding principle (embed IT capability into the business).

In the knowledge management project the key question is 'Where did the idea for the project come from?' The linkage between IT and what it can do for the business and business managers seeing new opportunities is critical to success. Opportunities for automation were easy to identify in the first era of IT but it becomes much more difficult to identify business transformation opportunities, even though there may well be many more of them. In order to identify such opportunities a broad knowledge of the business and of the industry in which it operates is required; however, many of the opportunities only arise because of the capability of IT. For that reason it is essential to embed IT people into the business so that they can help to identify opportunities and raise awareness of IT. In the knowledge management project the vision was that the organisation could share its collective knowledge. Turning this into a practical proposition then needed a major contribution from IT people, who also needed to be aware of the overall ambition of the project. The essential ingredient for success was the setting up of a team comprising business visionaries, people who would use the system, and IT people.

The big gains are to be had by exploiting the *technology exponential* shown in Figure 4; new business successes emerge at each stage of the evolution of technology (cost performance and interconnectivity), which has progressed at an exponential rate. In the technology exponential it is not only the power of computing that has increased




Figure 4—The technology exponential

exponentially, but also the ease with which new applications can be developed and the scope of these systems. In order that this is exploited to the full there needs to be rapid sharing of ideas, a common understanding of the business challenges and of the opportunities presented by IT. This needs to be assimilated into an environment of innovation and creativity. Today's computing products put considerable power to develop new systems into the hands of IT-oriented users who, if then closely supported by computing professionals, are able to develop systems of considerable complexity. This process has been used extensively within GBM, and from it have arisen a number of very successful new systems that have subsequently been deployed across the whole of GBM.

Thus real benefit has been derived from embedded IT people. This has included contribution from IT to new business improvement initiatives, creation of new applications and a wider more general increase in the level of IT-based innovation.

Derive Maximum Benefit from IT by Providing Users with World-Class Support

The full potential benefit from IT can be delivered only when everybody is using it correctly, and this requires competence and confidence that can only be built up over time. In order for this to occur and to prevent widespread blocking of IT initiatives it is essential that users are well supported. This is especially difficult to achieve if you are dealing with a mobile workforce, and it is even more difficult with knowledge workers because many of them do not have an opportunity to learn by 'buddying up'.

In the case of GBM the approach has been to use a value-add help-desk service (the Global Front Office (GFO)). However, the way this service is managed is that it follows the same value-add principles outlined in this article. The objectives of the service are set out in Figure 5, and the service has evolved rapidly over its 18 months of operation so that it is now a part of the overall business transformation model of GBM. Fixing faults is a low-value service; the real objective is to provide reliable IT service and to support business change programmes.

The key to continued success for the GFO has been the relationship established with the user community through its service ethos and ability to resolve 80% of user queries on-line. This relationship has now positioned the GFO as a respected source of information and there is ready acceptance that the team can contribute on a wider basis than simply providing help-desk services.

The work of the GFO is now in its second phase of evolution. The first introduced a number of administrative services, many of which are now being automated, and the GFO is becoming, under the second phase evolution, a key support organisation for more general matters associated with the business transformation projects. It is this latter area of work that is high value. It is where the service has migrated as success has been achieved through improving both the reliability of IT and the knowledge of users by on-line training.

Conclusions

The approach described is based on the authors' practical experience in transforming the relationship between Global Business Markets and the IT community. The four guiding principles are simple but effective. In the first stage of the transformation the author had line reporting to the BT Director IS and matrix reporting to Director Global Business Markets. During the first part of this period the author

DERIVE MAXIMUM BENEFIT FROM IT—
SUPPORT FOR IT USERS
HELP-DESK MISSION:
TO FIX USERS' PROBLEMS
OR
TO ADD VALUE TO THE BUSINESS

IMPROVE RELIABILITY OF IT
IMPROVE USER EXPERIENCE
MANAGE IT INVENTORY
SUPPORT BUSINESS CHANGE PROGRAMMES

Figure 5—Guiding principle 4

owned, on behalf of GBM, a strategic initiative to improve IT. While this did not fit with the first of the guiding principles it did establish the relationships around which the overall strategy was then developed. On completion of the IT Strategic Initiative the next step was to embed an IT person on to each of the new initiatives and from this we have seen success grow. One of the key benefits is that we are now seeing a new wave of innovative ideas emerging from the close working between business managers and IT. One area in particular where a large number of new ideas have arisen is in the use of multimedia where GBM has developed extensive practical experience in less than six months.

More recently, and as a result of the success generated through adoption of the principles described, the author now has line reporting to Director Global Business Markets and matrix reporting into Director IS.

The approach has resulted in a culture of partnership between the business and IT people, and from this there has been a very much more responsive approach from IT teams with a much higher level of IT awareness within the business teams. This in turn has led directly to many new business ideas and a major IT contribution to the business transformation projects.

The use of the four guiding principles is simple but effective and, while the order of their full implementation may need to vary slightly depending on the current management framework, their use for managing IT is highly recommended.

Biographies



Howard Ferguson
BT Worldwide

Howard Ferguson is a member of the BT Global Business Markets management team and has responsibility for IT and IT strategy. Prior to taking up this role he managed the internal Desktop support organisation within the BT Computing Service Organisation. Before that he was responsible for design, deployment and operation of the BT internal data-communications network.



Graham Longmate
BT Worldwide

Graham Longmate is an IT management consultant within BT Global Business Markets with particular responsibility for IT systems strategy and process analysis. Prior to taking up this role he managed the billing systems delivery programme within the BT Computing Services Organisation and prior to that was responsible for the delivery of e-mail capability within the BT Office Automation programme.

David Freestone, Jonathan Mitchener, Alan Pengelly and Andy Childerhouse

Extending the Boundaries of E-Business

Transforming Back-Office Operations Between Telecommunications Operators

Most examples of e-business automate standard commercial processes. While these applications undoubtedly bring business benefit, there are many specialist operational processes which are specific to telecommunications and which would also benefit from end-to-end automation. BT has recently conducted a trial for the wholesale carrier business, focused initially on back office operations between operators. This article explains the technical vision and approach towards end-to-end e-business. It describes the trial and its results.

Introduction

E-business is becoming commonplace for retail user interfaces—consumers can buy goods and services through web front ends and companies are ‘web-enabling’ their direct customer processes at a rapid rate—in telecommunications companies as much as other industries (for example, the BT web page for altering discount telephone numbers¹). However *web-enablement* is by no means the whole of e-business. It is only an early, partial example. E-business also includes end-to-end support for industry-specific processes—and networking with suppliers, partners and dealers as well as end customers. It is still the early days of end-to-end e-business and there are still issues to be worked through. However, some industries have already moved ahead.

At BT Adastral Park a programme of work is pursuing the goal of end-to-end e-business for telecommunications businesses—and particularly for supporting business operations when services involve multiple organisations (for example, for geographical or product-related reasons). As part of this work, BT recently conducted a trial for the wholesale carrier business to investigate the benefits of end-to-end e-business for operational processes. The trial deployed a community-of-interest network (CoIN) for data and processing in back-office operations. It focused initially on operations between

operators, covering pre-order forecasts for interconnect, traffic forecasts and customer contact details. It also quickly generated a list of further processes it would be desirable to add.

This trial was the first step in a research and development vision to overhaul the entire approach to business process support and the associated operational support systems. The aims are to reduce costs for operators and customers alike, to improve the service delivered from the processes, and to improve their adaptability. There are a number of different principles in this work. The two most relevant to this early trial were ‘allocation of responsibility’ and ‘dispersal of data and processing’. Both of these rely on distributed systems (especially Internet-related technologies) for their design and implementation.

These principles are radically different from traditional ways of designing and building operational support systems (OSSs), which are usually centralised systems. The telco traditionally guards its control of the processes—and bears total responsibility for the costs. Dispersal goes beyond the current trend for granting direct customer access to systems parameters (for example, telephone numbers for bill discounts) by web-enablement or other technologies (for example, voice recognition) which are still only extra access methods to a system run entirely by the telco. Instead, dispersal moves towards greater partnership in carrying out business

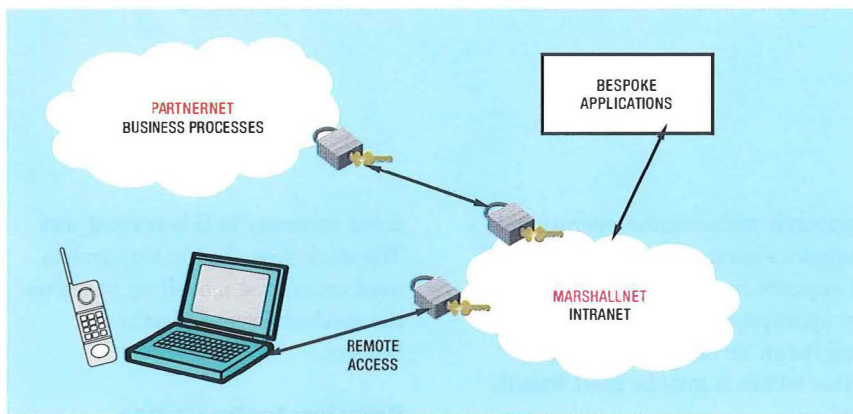


Figure 1—Marshall's e-business strategy

processes—and it is particularly relevant for multi-organisation enterprises.

The remainder of this article outlines the rationale for extending the boundaries of e-business, explains the technical approach for this work, describes the CarrierNet trial and its results, and discusses the implications for other multi-organisation processes in telecommunications operations.

The Case for End-to-End E-Business

End-to-end e-business is an objective for many industries, not just telecommunications. In fact, some of the examples of leadership in this field are from completely different industries, which have used the approach to improve their competitive position by reducing costs and improving the way they handle customer relations. One such example, where information is available in the public domain², is Marshall Industries—one of the four main distributors of electronic components in the United States.

There is intense competition in the American electronic distribution industry—over 1500 companies and rapidly decreasing margins. In 1991, Marshall radically re-thought the way it wanted to do business for the new market conditions—and then completely re-constructed its processes, people skills, systems and technology over the course of six years. As a result it has significantly out-performed other companies in the sector. For example, the company doubled its turnover during the six years (to \$1.2 billion in 1997) at the same time as reducing its headcount—against

the industry trend. This result, with many others, is described in reference 2.

To achieve these results, Marshall innovated customer service activities and invested strongly in state-of-the-art technology to push out the bounds of customer service and to simplify internal processes. The company information, offered customer self-service, and provided 24-hour support via all customer channels.

The technology deployed to achieve all this was Internet protocol (IP) based. Figure 1 illustrates the strategy. Marshall's networked its business internally on an intranet, enhanced with bespoke applications geared specifically to its business, and gave secure remote access to field staff. The final step was to network the business with its partners—suppliers, dealers and joint ventures—over an extranet.

Marshall is just one example of the deployment of end-to-end e-business. There are others also in the public domain; for example, the US Automotive Network Exchange³, which has also demonstrated significant benefit for the organisations involved.

The business pressures on telecommunications operators are similar to those in those two industries. For example, customers (particularly business customers) are increasingly seeking customisation in the way business is done with them, and flexibility to adapt business processes to their requirements. Also, the rate of new product launches is increasing, so operators need more-effective faster ways of building relationships with customers. The comparison is particularly relevant where opera-

tors are working together to serve customers—in a supply chain or in a peer-partnership. There is the same opportunity to innovate and take a lead in the industry, and the technology is becoming ever more cost-effective.

Technical Vision

Responsibility and dispersal

Two key design principles for the BT work were 'allocation of responsibility' and 'dispersal of data and processing'. The work reported in this article has applied these principles to communities of interest in the United Kingdom carrier business. However, the principles apply to broader communities of interest—for example, supply chains—and can include end customers as individuals or organisations within the communities.

Dispersal means decentralisation of responsibility, business processes, information and systems. For example, it could include physically locating parts of the data or processing functionality on another organisation's system (even a customer's system) where that is the more logical or business-efficient place to do so. Organisations doing e-business in this way share responsibility for the end-to-end processes between their organisations, and share selective access to each other's systems to read data from them as needed.

Dispersing responsibilities and allocating them appropriately removes unnecessary process steps (for example, steps to check the changes made by another organisation have been made correctly) and so reduces costs for everyone. At the same time it improves the quality of operational support by eliminating whole classes of problems, and requiring less operational support to be done within the service provider.

Dispersed data and processes are a major enabler for doing cost-effective e-business between organisations over CoINs. It encourages:

- *customer empowerment*: enabling customers (including internal customers in a supply chain) to configure their own service parameters and support options, and giving them *responsibility* for appropriate data and process steps;
- *customisation*: tailoring standard business processes for particular customers or supply organisations, offering service differentiation;
- *end-to-end process automation between organisations*: with human intervention only for complex process steps, exception handling, or for a specific customer;
- *single truth*: only one instance of information shared between organisations, to eliminate costs of replication and replication-related failures; and
- *minimal support systems*: a significant reduction in functionality needed on centralised, back-end systems.

The combined effect of all these effects is to enable closer business relationships with customers and suppliers.

Provisioning, fault reporting/management, charging/billing are all examples of where a dispersed customisable support infrastructure could support better more-flexible operational support processes across communities of interest.

Obviously different degrees of dispersal are possible, and a certain amount of dispersal is already implemented (for example, SIM cards in mobile telephones disperse some responsibility, data and process to the user or customer, but in a limited way). A distributed systems architecture is the optimal platform for fully dispersed operational support, but there are many opportunities to begin a dispersed

approach with current operational support systems. Also, the dispersal of support functions is not necessarily appropriate in all circumstances and it can be targeted at specific cases where it may be most beneficial.

While some customer aspects (such as customisation and flexibility) are made easier by dispersing operational support processes, it might be felt that some corporate management aspects could be made more difficult (for example, aggregation of operational statistics). In fact, this is just a case of design priorities. Just as one would not design a centralised system without non-functional requirements or error handling, one should not design dispersed systems without regard for house-keeping or management information gathering requirements.

Enterprise Modelling

The business model supported by dispersed operational support, together with information sharing and collaborative e-business applications, is a truly radical change compared to the customer-supplier type relationships that are commonplace today. For example, it requires sharing direct access to each other's business information (and, depending on implementation, possibly even to back-end business systems), dynamically and on-demand.

Enterprise modelling provides a formal means to reason about such changes. An enterprise modelling analysis was carried out and underpinned the carrier prototype and trial. The methodology following concentrates on the business responsibilities associated with the roles of the parties involved in providing service. The analysis and exploration of these roles and responsibilities can then lead to a more appropriate design for the systems that support an organisation's business goals. It also provides a stronger basis on which to decide the commercial arrangements for

doing business in this revised way. The work described in this article used enterprise modelling based on the methodology of Martin & Dobson⁴.

Bringing technologies together

Extended e-business with dispersed responsibilities, data and processes can be built because of the availability of Internet and collaborative working technologies.

Internet technology

Within organisations, it is becoming commonplace for organisations to base their internal processes on an intranet. An extranet capitalises on this by allowing secure links to be made between different organisation's intranets (network layer infrastructure). The benefits of the intranet can then be exploited between organisations doing business with each other, forming a CoIN (an application layer infrastructure) over which the organisations share information and business processes.

Extending access to systems by networking techniques is not new. However, a number of desirable aspects of the Internet and World Wide Web's success are preserved by the specific use of extranets:

- The non-proprietary nature of the network makes the interconnection of various organisations easier. The open standards encouraged by the Internet provide an ideal basis for interconnection. The extranet approach also inherently means that existing intranet investments can be sweated for more benefit. This reuse of corporate infrastructure in each organisation is very desirable from a cost perspective.
- The sheer abundance of Internet/intranet/extranet interfaces in products both at the client front-end and the database back-end helps to de-risk extranet-based

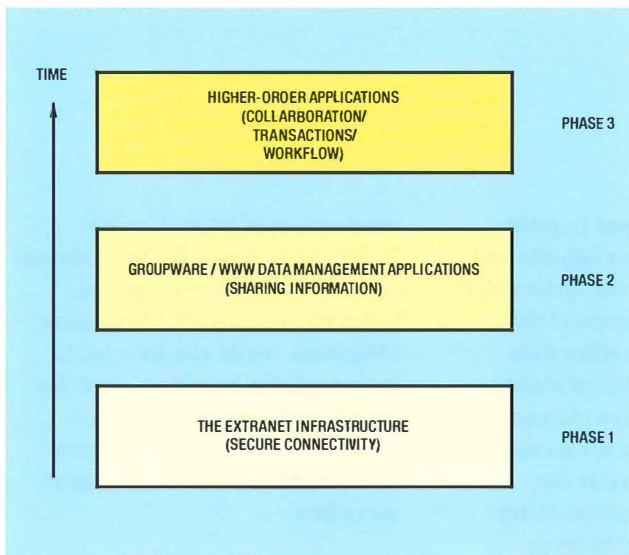


Figure 2—Phases of collaborative working

business support infrastructure against future changes in technology.

- Associated platform-independent technologies such as mobile code (for example, Java) and network computing⁵ make the development of truly distributed business support infrastructure across an extranet an easier proposition.

Collaborative working technology

The same extranet IP-based infrastructure used to build the CoIN, share e-business processes and provide workflow support can also run distributed applications which support managed information sharing and dispersed responsibilities. Applications are also available for conferencing and collaborative working on documents etc, as shown in Figure 2.

CarrierNet Trial

Putting theory into practice

The technical vision presented provides a tantalising glimpse of how things could be in the not-too-distant future. Much of the vision can be built with off-the-shelf technology, so an early technology trial was desirable. The UK carrier business was chosen for this trial because it involves sophisticated back-office operations, because commercial transactions for interconnect necessarily involve coopera-

tive interactions, and because there was a desire by most operators, including BT, to automate processes.

The primary element of the carrier business is *interconnect*, and the key commodity of the carrier business is *capacity*. Operators have to forecast their likely requirements, while the carrier has to ensure that sufficient capacity exists within the carrier network to ensure that the operators' requirements are met.

The general vision, along with the proposal for a technology trial, was presented to an industry-wide forum. It was agreed that a trial should take place so that the UK interconnect industry as a whole could ascertain the benefits. A trial was established, which ran from September 1998 until the end of February 1999, and which implemented selected elements of the technical vision. The primary aim of the trial was to test the general principles, including the concept, business culture and technology of a CoIN. The operators who took part, along with BT, were Cable & Wireless Communications (CWC), NTL, Orange and Scottish Telecom.

System design

The service deployed during the trial was branded as *CarrierNet* and its configuration, based around a hub server for shared information in a *demilitarised zone*, is outlined in Figure 3. The demilitarised zone lay outside the firewalls of all the trialists and was managed by a third-party organisation (in this

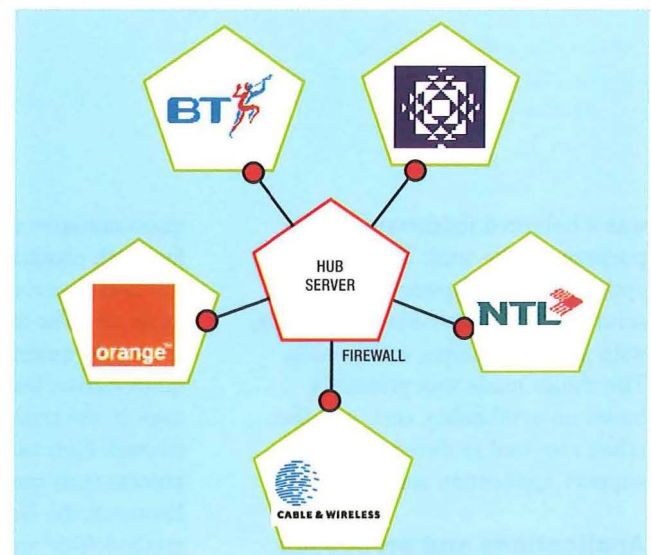


Figure 3—CarrierNet trial configuration

case the Internet and Multimedia business of BT, run entirely separately from the main network business). This design was chosen since a hub is a natural place for truly shared data, particularly in the early stages, when security issues prevent direct access between different operators' back-end systems. It also emphasises that the processes are symmetric—designed to be applicable to any operator supplying any other operator.

There were security controls at both network and application levels. At the network level, there was a firewall around the shared server and secure access links. (The intention had been to use a mixture of dedicated links, dial-up links and encrypted Internet links for different transaction profiles, but, for various reasons, only the dial-up links were actually deployed during the trial.) At the application layer, the data was deliberately partitioned (as part of the design) between organisations, according to responsibilities and access rights. Each partition was stored on the server with password protection and with access authorisation controlled by the partition owners.

The trial used a BT product, Intranet Complete⁶, to provide the extranet on which the CoIN operated. Microsoft Exchange⁷ and Microsoft Outlook⁸ were used to support the end-to-end processes (workflow), using Visual Basic scripts over the Outlook directory structure. The choice of technology

was a balanced judgement for the purposes of the trial. The two main options were groupware-based solutions or browser-based solutions, with supplier choices within each. The choice made was primarily based on availability, cost, and the effort required to develop process support application software.

Applications and processes

Processes deployed in the trial were related to key information about customer requirements: pre-order forecasting, traffic forecasting and contact details. In the trial, the processes involved BT supplying capacity to other triallists. However, they were designed to be symmetric; that is, the same processes would work (in the opposite direction) for BT buying capacity from another operator, or for any two other operators doing interconnects.

Each process was chosen to exercise a particular part of the vision. For example, the document containing key contact information for both BT and the other operator has traditionally been the responsibility of BT. This means any changes have to be made by BT. It would be far more sensible if operators could change their own details. So, in the trial, this was redesigned to be a truly shared document—the operator had read-and-write access to its information but read-only access to BT information, and vice versa.

The traffic forecast (TF), an annual forecast of expected capacity requirements prepared by the operator, and the advanced capacity order (ACO), a quarterly capacity forecast, were examples of supporting and automating shared processes between organisations. These processes involve data submission and validation, negotiation (between organisations and between internal business units), agreement and approval.

Trial results

To evaluate the trial, metrics on usage and system performance were collected automatically. In addition,

questionnaires were used to gather feedback about business benefits and the user experience of the system and approach. The main results of the trial were qualitative rather than quantitative, because the processes used in the trial were low volume (though high value), so the number of process runs was relatively low. However, the Goal-Question-Metric methodology⁹ was used to try to assess the trial and the approach against the objectives. Some of the key results are outlined here.

Results of the CoIN approach

The use of CoINs to support inter-organisation processes generated considerable enthusiasm. This enthusiasm operated both at the level of detailed support processes, and at the level of the changed dynamics of interworking between organisations through a greater degree of partnership and collaboration.

Trial participants frequently came up with ideas for new uses of the infrastructure, to reduce manual intervention or cross-correlate processes to enable enhanced service (for example, correlating advanced capacity orders with actual orders). The trial itself acted as a catalyst for these ideas. The CoIN approach, with its notions of dispersed responsibilities, data and processes, contains some subtleties that only gradually become clear with detailed understanding, and so many good ideas came up only when participants were involved at a detailed level.

Results about technology

Heterogeneity in the desktop equipment of the triallists proved to be much greater than expected, and considerable effort had to be expended resolving individual compatibility issues. Some triallists did not use Microsoft desktop software at all and even those that did had significant variations in the version and patch levels they were using, which led to interworking problems with the hub server. Similar degrees of heterogeneity are likely to be present in other

communities of interest, so it is important to design for that. Interestingly, a browser-based approach, which might be thought to be more ubiquitous, would also have led to incompatibility problems: not every organisation allows browsers on desktops, and there may still have been potential product and version variations.

Benefits and productisation

The volume of transactions in the trial did not allow business benefits to be reliably quantified. However, the qualitative feedback from the triallists was informative, and mainly positive. They were unanimous in their praise for single-truth shared data, to reduce the problems and costs from multiple copies. They recognised definite operational efficiencies even from the trial—however, the efficiency gains from the trial processes alone were limited, and significant improvements would need many more processes running over the CoIN, including high-volume processes. The triallists also supported the cross-correlation of processes and automating specialist processes. On the other hand, the problems with the trial's technical infrastructure had hampered triallists' ability and eagerness to exercise the trial system fully.

The results from the CarrierNet trial are encouraging and they show how the business benefits from extended e-business partnerships can begin to be realised. As a result of the work described here, a product based on CarrierNet has been announced by BT to the UK carrier community. The technology choices for the product have been modified based on the experiences of the trial, and higher volume processes are planned for inclusion.

Towards the Future

There is a significant body of opinion that the organisations which will prosper in future will be those who understand and work most closely

with their customers and suppliers. A natural conclusion from this is that leading companies will change the way they interact with others, to the point where they regard each other as partners. In true partnerships, they will share information and work cooperatively and collaboratively at all levels of operation.

The work described in this article has begun to extend the scope of e-business, applying CoINs to back-office business processes between telecommunications organisations, an area which has not traditionally been designed that way. The CarrierNet trial was an early implementation of the approach. Technically, it showed that e-business can be extended to specialist processes while still using off-the-shelf technology, and can produce a range of significant business benefits.

The work was always intended to be more widely applicable than just the wholesale carrier business. In particular, it is applicable to business operations between operators working together to service customers as a supply chain or as a partnership. The use of CoIN designs can be applied to joint provisioning data and processes, joint fault handling data and processes, direct customer contact data and processes and many other areas of business.

The technology to enable widespread adoption of this approach is becoming available now and is continuing to develop; for example, directory-enablement. This means that the time is right to re-assess the way we do business when operating in supply chains or partnerships. There are still technical issues to resolve; for example, how to support very large-scale operations, and how to disperse management of the CoIN infrastructure itself. Questions like these are being addressed in the continuing programme at BT Adastral Park.

Acknowledgements

The authors would like to thank the many people in BT and the triallists in other operators who have helped during the trial and during the production of this article.

References

- 1 http://www.customer-service.bt.com/friends_family/owa/bestfriend.who
- 2 EL SAWY, OMAR A.; MALHOTRA, ARVIND; GOSAIN, SANJAY; and YOUNG, KERRY M. The relentless pursuit of 'Free. Perfect. Now': IT-enabled value innovation at Marshall Industries. <http://www.simnet.org/public/programs/capital/97paper/index.html>, Sept. 1997.
- 3 www.aiag.org/anx
- 4 DOBSON, JOHN; and MARTIN, MIKE. The Ontology of Enterprises and Information Systems. <http://hsb.baylor.edu/ramsower/ais.ac.97/papers/dobson.htm>, Newcastle University.
- 5 REVETT; KNUL; and STEPHENS. Network Computing. *BT Technol. J.*, Apr. 1997, 15(2).
- 6 <http://www.e-commerce.bt.com/>
- 7 http://www.microsoft.com/products/prodref/599_ov.htm?RLD=21
- 8 http://www.microsoft.com/products/prodref/608_ov.htm?RLD=26
- 9 ROMBACH, H. D.; and BASILI, V. R. The Benefits of Goal Oriented Measurement. *Proceedings 7th CSR Annual Conference on Software Reliability and Metrics*, 1990.

Biographies



David Freestone
Networks and
Information Services
BT UK

David Freestone leads a group in the Business Engineering Laboratory at BT Adastral Park, extending the reach of e-business to include end-to-end business processes, communities of interest, and whole supply chains. The work combines leading-edge technology with radically new kinds of business operations needed for a new-wave telecommunications company like BT. Previous work has included advanced techniques for requirements capture and software production, formal methods for communications protocols and programming support environments. He joined BT Laboratories in 1980 and holds degrees from Cambridge and Oxford Universities.



Jonathan Mitchener
Networks and
Information Services
BT UK

Jonathan Mitchener joined BT Laboratories in 1988, having graduated from Hatfield Polytechnic with a first-class B.Sc. (Hons) in Computer Science. Initially working on architectures for broadband services, he has over 10 years gained a wide range of experience in object-orientation, software reuse research, the Billing 90s programme, and WWW/intranet developments to support Global Sales teams. He has led work on systems and software strategy issues across a variety of projects including currently research into future OSS and total service creation. Recently he has been responsible for applying Community

of Interest Networks to BT business scenarios, including Carrier Services. His current work is looking into the future of e-business for new-wave telcos.

Alan Pengelly

Networks and Information Services BT UK

Alan Pengelly holds a degree in mathematics and a Ph.D. in protocol engineering. He joined BT in 1987 from Ferranti Defence Systems. Working initially on ISDN signal processing, he has worked on a variety of projects including Esprit, Eurescom, 'Inventing the 21st Century Organisation', Billing 90s, and CarrierNet. He recently moved into Solution Design, part of the team responsible for OSS design for BT Joint Ventures.

Andy Childerhouse

Networks and Information Services BT UK

Andy Childerhouse graduated from the University of Hull with an honours degree in Management Systems. He joined BT in 1989 working initially on the NCDB. In 1991 he moved to system security, where he held a number of posts including BT Security Evaluation and Fraud Management. In 1996 he took up a post in Finance Billing Receivables where he was responsible for revenue assurance. This was followed by a term as a team leader within Interconnect Process and System Evolution. He has recently taken on a new role where he is responsible for data wave initiatives for *network*BT.

Mark Fraser

Advanced Global Electronic Messaging

The BT Alliance has one of the most challenging electronic messaging environments in the world.

With members located in every time zone, and with a blend of massive local and global business activities, effective internal communication within the Alliance takes on vital importance. This article describes the business environment, specification, design, and implementation of a sophisticated new global electronic messaging infrastructure for the Alliance. The new service is a Concert commercial product operated by Telenor in Norway. This is the largest global internal IT infrastructure change yet undertaken across the Alliance.

Introduction—Value of E-Mail to the Alliance

Within a few short years electronic messaging has become the lifeblood of the BT Alliance. From its origins as a tool that was used primarily by the engineering and IT communities to communicate with each other, it has become essential to everyone's job—from a receptionist to a chief executive. It is now as valuable as voice communication and has replaced large proportions of postal services and fax. When it works well, it brings enormous benefits of speed, information transfer, and accessibility. But having become so popular and so valuable, if it does not work well the Alliance's core business is immediately and directly affected.

The Alliance's internal e-mail infrastructure sprang largely from the rapid deployment of Microsoft (MS) Mail around the world. Our MS Mail domain became the largest in the world, and stretched the product far beyond its limits. Microsoft recommended that no more than 20 Post Offices were interconnected—the Alliance had approximately 700.

At the beginning of 1999, approximately 750 000 messages *per month* were being sent between Alliance members. This figure excludes messages sent within any of the Alliance members' own domains, or to the public Internet and X.400 domains.

Drivers For Change

The nature and scale of BT and the Alliance's international interests have changed significantly since the

MS Mail infrastructure evolved. During 1998, several compelling reasons to totally reassess global messaging and directory infrastructure emerged, leading to wholesale realignment with the Alliance's international business operations:

- The existing MS Mail infrastructure had become obsolete, no longer meeting user nor administrator requirements. In particular, directory synchronisation globally had become extremely fragile. Most Alliance members wished to move to modern products (particularly MS Exchange Server) but could not do so without risking their connectivity to the rest of the Alliance.
- Business demanded improved communication with customers, suppliers and business partners.
- Alliance members demanded service that differed to that provided to BT UK users, and which took into account their particular local environments.
- Lower-cost better-value messaging services were required.
- Increased local autonomy and control was sought as Alliance members became larger independent companies.
- Increased flexibility to adapt quickly to business change was required.
- Alliance members wanted to make greater use of local simple message transfer protocol

Figure 1—Alliance messaging environment

(SMTP) and X.400 gateways for external messaging.

BT Worldwide formed a consortium with Alliance members and together they established a programme to establish and implement international messaging and directory strategy.

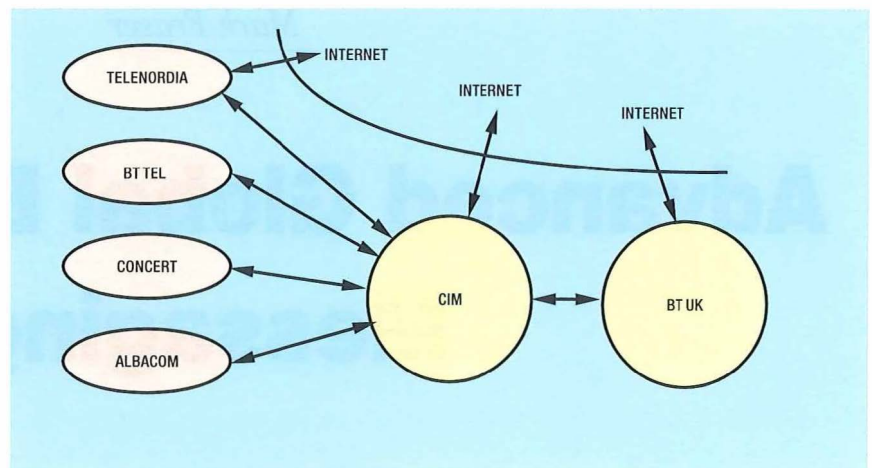
The Market

Corporate messaging products

Only four years ago the market was dominated by simple products: usually 16-bit, DOS-based, and using crude flat-file storage of data. Leading products were MS Mail (which as 'Network Courier' Microsoft had purchased from a Canadian software house) and Lotus cc:mail. These products were designed for use in a limited local area network (LAN) environment and not for use by large enterprises across wide area networks (WANs). Lotus Notes had achieved some notable success but as a groupware platform rather than for e-mail alone.

Three years ago, Microsoft launched Exchange Server with two purposes: as a messaging replacement for MS Mail, and as a groupware competitor to Lotus Notes. Based on a relational database and running on the NT Server operating system, it offered increased flexibility, reliability, functionality and security, while offering superior administration facilities. It also promised organisations potentially lower operating costs. Although earlier versions earned a reputation for non-scalability and poor industry-standard integration, it has now replaced MS Mail as the world's most popular corporate messaging product. However, as intranet technology has developed, so the potential use of Exchange as a groupware/workflow platform has become less attractive.

The massive rise of the Internet for business communication and the commercial world's acceptance of SMTP as 'dependable enough'



(despite quality-of-service issues) has meant that no corporate messaging product can afford to neglect integration with the Internet. Similarly, although X.400 is not directly as popular with users, it is a standard widely used for interconnection that allows vendors to avoid proprietary gateways between rival products.

Message switching and directory services products

From a background of public X.400 switching, a market has emerged to address the interconnection of multiple domains and multiple products. The market segment is dominated by products from companies such as Worldtalk, Control Data Systems, and ISOCOR. Their products are designed to enable a large complex organisation to integrate their various communities usually using industry standards such as X.400 for message transfer and X.500 for directory services. Costs can be very high with such products and they frequently operate as part of an overall service by a specialist supplier that can achieve economies of scale, rather than operated in-house by the user-organisation.

Some of the corporate messaging products (see above) attempt to address this market too, but usually find themselves compelled to modify industry standards to give their own products an 'edge' over their rivals.

The Alliance Environment

Some large organisations operating on a global scale (such as Ford, Unilever, or IBM) are able to ensure a relatively simple, homogenous messaging environment through

exercising central IT design and operations control. However, members of the Alliance have full control of their own IT strategy and naturally wish to deploy messaging infrastructure which is best suited to their local situation. This means that any solution offering effective messaging and directory services across the Alliance must be extremely flexible and must be able to integrate widely differing local environments. It also requires Alliance members' IT people to collaborate on a global scale.

Alliance members quite rightly demanded that control of any replacement for the global MS Mail infrastructure was specified, chosen and deployed in a way that they controlled—not in a centralist fashion that imposed unreasonable working practices or that restricted their ability to progress their own domestic messaging strategy.

Product and Design

Concert Integrated Messaging (CIM)

The heart of the new worldwide messaging service chosen by the Alliance is Concert Integrated Messaging (CIM) v1.5. This is a recent Concert commercial service that is in fact operated by Telenor AS from a mountain installation in Bergen, Norway, using Control Data Systems (CDS) products.

Although not the prime consideration during the product selection process, the Alliance's own internal use of a Concert product is extremely valuable. The cliché of 'use what we sell, sell what we use' is applicable. As an extremely

demanding user in its own right, the Alliance is well placed to influence the product development too, so that it continues to meet the requirements of complex global customers.

Alliance members' corporate messaging systems

Most of the members of the overall Alliance mail domain have moved to MS Exchange Server and have a native Exchange connection to a bridgehead server to CIM. A few are implementing other products (for example, Lotus Notes) and will connect to CIM either as native Notes or as X.400 if appropriate. Many have local SMTP gateways, often using their own-brand domestic Internet services to do so.

Integration with BT UK

BT UK had a large mixed MS Exchange and MS Mail domain using Worldtalk and Hewlett-Packard products to provide an X.400 messaging backbone. It operates an X500 directory fed from its Human Resources systems. The Alliance's CIM-based service exchanges directories and messages (using X.400) with the BT UK service so that all users in all connected countries are able to find each other in their local address books.

Internet and public X.400 messaging

CIM also provides central gateways for SMTP to the Internet and a public X.400 service for those Alliance members who do not use local gateways from their own corporate system. The SMTP addresses provided from the central gateway can be configured with a local format so that recipients are unaware that messages do not originate directly from the Alliance member country.

Internal network

The BT global WAN (GWAN) is used to connect Alliance members to the CIM host in Bergen, and to connect

the CIM host to the BT UK mail domain. The GWAN is part of the PiperNet private integrated data/voice/video network. PiperNet consists mainly of a combination of Concert Managed Bandwidth Service (CMBS), Concert Frame Relay Service (CFRS) and international private circuits where appropriate.

Directories and corporate White Pages

A key requirement of most Alliance members is that they wish to enter and hold user information in their local corporate mail systems only. They do not want to submit user details to a central database and wait for that database to cascade directory information to all connected systems before a new mail userID can be used. Instead, they want automatic bi-directional directory transfer so that once they have entered the userID locally it is synchronised and propagated throughout the Alliance-wide mail domain.

The new service also makes use of the X500 directory to provide corporate White Pages. System administrators can choose to populate fields with useful information (such as Alliance organisation, job title, postal address, job function, telephone numbers) and make this available through CIM. In such a large set of organisations this aids basic communication across the Alliance.

For the first time, the global directory contains universally valid e-mail addresses. This ends the widespread 'Reply All' problem of the previous MS Mail infrastructure where one recipient often cannot mail other recipients due to local systems interpreting non-local addresses wrongly. Similarly, non-Alliance recipients (via SMTP or X.400) are now able to copy to real Alliance addresses rather than the fabricated ones which they are often presented with.

CIM's directory conversion features allows users to see all other

users globally in their own directories as if they were using the same type of mail system. For example, MS Exchange users see all worldwide entries in their address lists in a familiar and meaningful way, even if those users actually use MS Mail or Lotus Notes systems.

Messaging functionality

The adoption of X.400 throughout gives far more predictable and reliable message connectivity. It also allows system administrators to track and log messages whereas the previous global MS Mail infrastructure was impossible to manage in a quality way.

It also allows the use of read- and delivery-receipts which cannot be fully exploited using simple mail transfer protocol (SMTP) or MS Mail. The combination of advanced products such as MS Exchange or Lotus Notes with an X.400 backbone and switching service is very powerful.

CIM preserves attachment icons and filenames between different mail systems. This is very important as it means users do not need to worry about which system their recipients use wherever they are in the world.

Service and Support

Based upon old products that were stretched beyond normal limits, the previous MS Mail-based infrastructure could not achieve very high levels of availability or reliability. It was very difficult even to monitor and measure performance. This led to over-reliance on manual processes and intervention.

The new infrastructure with CIM and Exchange/Notes allows better control. As an Alliance product, faults and queries relating to the service are reported to and handled by the standard customer service centres (CSCs). Local system administrators are able to enter and track faults through web-based tools. The CSCs runs on a 24-hour 7-day basis, as

does the Telenor unit which operates the CDS host in Bergen.

Service level targets are the same as agreed with external customers of CIM; for example, 99.2% service availability. Once messages have left a particular Alliance member's mail system, they typically take less than 10 minutes to deliver to any other Alliance member's system worldwide.

A service level agreement between Concert and BT UK ensures that messaging connectivity and directory issues between the worldwide CIM domain and the BT UK domain are properly handled.

Development and Implementation

At the end of a formal tendering exercise, the Alliance consortium led by BT Worldwide selected the CIM/CDS service above a short-list of alternatives which comprised: a Worldtalk-based X.400/X.500 solution; an MS Exchange global-hubbing solution; and basic SMTP/X.500.

The CIM product as deployed is very similar to the standard commercial product. The principal additional development for this project was the interface of the CIM environment to the BT UK Worldtalk environment. Both CDS and Worldtalk use X.500 directories, but all vendors' implementations are slightly different and X.500 is also highly configurable. Similarly, for the full benefit of X.400 connection between Alliance members it was necessary to ensure the standards were made fully compatible throughout the domains. Close technical collaboration between BT Worldwide, Concert, Telenor and BT UK was necessary for success.

Since many Alliance members were simultaneously implementing MS Exchange, and BT UK was migrating its people to Interchange (its MS Exchange implementation, which was completed in July 1999) there was enormous change occurring in parallel. The size and complexity of the Alliance's messag-

ing domains made it particularly challenging to ensure service throughout the period of change.

Similarly, the test programme was extensive as it had to address the many hundreds of permutations of functionality and mail system connections.

Products like MS Exchange have enormous available functionality and are highly configurable. However a disadvantage of their considerable in-built intelligence is, for example, that they can carry out directory upload/download and replication at an unpredictable times depending upon system utilisation. Such features are largely beneficial but can make large global domains difficult to operate. Considerable work was put in throughout the project to ensure that system administrators retained as much flexibility as possible and would not have to change their own behaviour in favour of their systems.

Future

Groupware and collaboration

At present most groupware and all real-time features of MS Exchange (which is used by virtually all Alliance members connected to CIM) cannot extend beyond each country's own Exchange organisation. However, recently Microsoft has enabled some additional functionality in MS Exchange 5.5 (for example, sharing Public Folders and Schedules) through the release of an updated service pack.

However, the rise of web and intranet products as groupware and collaboration platforms offer alternatives and may be more appropriate for many applications.

Secure messaging and e-commerce

The synchronised X.500 directory of CIM and the corporate systems connected to it make an ideal environment for the introduction of secure messaging and e-commerce

across the Alliance. Apart from secure e-mail within the Alliance itself, potential applications include:

- secure e-mail with business partners and customers,
- virus screening of e-mails,
- reduction of 'spam' (non-business related e-mail), and
- archiving and audit trails.

Directory development and light-weight directory access protocol (LDAP)

It is possible to provide a directory of Alliance users (or an appropriate portion) to other organisations—for example, strategic business partners, retail outlets, or Alliance members who do not participate in the full Alliance CIM domain. They can then use the directory to find Alliance member recipients but use public X.400 services or the Internet to exchange mail with them.

LDAP is a cutdown version of the X.500 directory protocol. LDAP clients can be used to look up X.500 directories such as CIM's, and offer a simple alternative where standard directory synchronisation is not necessary. With appropriate security safeguards, it also offers the prospect of querying the CIM directory from any Internet-connected PC.

Unified messaging

For several years vendors and customers have identified unified messaging as a 'killer application'. The ability to integrate e-mail, voice mail, fax messages, short message service (SMS), and other forms of targeted communication is highly attractive. The products and technology to enable this are now beginning to reach maturity and, again, largely depend upon exploiting a powerful directory such as CIM's X.500 service.

Conclusion

The introduction of the new world-wide messaging infrastructure now gives the Alliance a messaging platform suitable for its level of business and growth. The key concepts to ensure that the business-focus remained at the centre of the solution are worthy of concluding with:

- Alliance members controlling their own domestic mail strategy;
- commercial off-the-shelf products and services where possible;
- use of the Alliance's own products where suitable;
- low-cost high-value service;
- maximum automation and minimum manual process;
- enabling maximum flexibility of connected mail systems;
- industry standards; and
- teamwork across the Alliance.

Biography

Mark Fraser
BT Worldwide

Since 1993, Mark Fraser has held information systems roles in various international BT organisations. He is currently IT Infrastructure Manager for BT Worldwide, leading a team responsible for BT Worldwide's strategy and project implementation in areas such as messaging; intranet; internal global network; LAN; and desktop and remote computing. Working with fellow IT managers from Alliance partners, his objective is to enable the maximum possible level of communication across the Alliance. Prior to this, he worked for VIAG Interkom where he was responsible for the implementation of the company's customer billing system. Previous roles include Technical Operations Manager for BT Europe, and information systems implementation worldwide.

Mark Fraser can be contacted on e-mail at mark.fraser@bte.bt.com

Optimising Investment in the BT UK Access Network

The development of a model to understand and derive the optimum network stock level for the provision of service process within the BT UK access copper network is described. The model explores the relationship for varying network stock levels, between all of the actual costs associated with operating the provision of service process and the required access network investment costs.

Introduction

A study was instigated to establish a model that would define the optimum operating point of meeting orders using existing stock from BT's access copper network. The study also includes the balance of proactive and reactive build costs and how to determine the total cost of moving from today's actual 'met from stock' (MFS) performance to the optimum.

Background

To place the importance of this modelling work in context, BT's access network has been built over the last 80 years, and is worth multi-billions of pounds. This network provides service to approximately 95% of the UK's geography and is now under considerable threat from competition, where customers now have a wider choice of network operator.

In recent years, the underlying investment policy for building the network has been to minimise new build and maximise the use of the existing asset, and reactively deploy pair-gain equipment. This policy was aimed predominantly at improving customer quality of service and optimising the use of capital investment. This has reduced the percentage MFS (% MFS) performance in certain areas and, where this has happened, reactive build costs have grown to unacceptable levels in order to meet the customer quality of service targets. It is in these specific places where the copper network needs replenishing to meet the optimum operating point.

At the same time, increased competition from other licensed operators, has meant that BT is seeing

access network pairs being returned to stock and available for reuse. This balancing effect has enabled BT to maintain a reasonable operating position in some areas only, at a very low cost. The current modelling work summarised in this article incorporates the effect of competition on the overall costs of operating the access network.

What is Met From Stock?

Met from stock (MFS) can be defined as the following:

a provision of service (POS) order that can be routed to a designed distribution point (DP) utilising existing stock (predominantly copper pairs or channels), including copper pairs that are marked as faulty or stopped, or a pair diverted from an adjacent DP using a dropwire.

MFS performance is defined as:

MFS performance =

$\frac{\text{no. of provision orders met from stock}}{\text{total no. of provision orders}}$

× 100%.

An order that is MFS has the cheapest immediate provision cost. However, this must be balanced against the cost of the historic proactive investment that made it possible, which is why it is very important to BT to optimise the stock level in the access copper networks. Recent system enhancements in management information systems (MIS) and the ability to gather actual cost data have provided the means by which MFS can be accurately

modelled for all the key elements of operating and building the network.

The modelling work used actual data for all costs associated with operating and building the access copper network. The optimum operating point was determined by adding all of the proactive build costs to all of the reactive build costs for a range of MFS operating points. The optimum operating point was where the total cost of these two was at the lowest. The potential cost of moving from today's actual operating point to the optimum was calculated from the proactive build costs.

Approach

There were four discrete stages in the design and building of the optimum MFS model:

- identifying all the key cost factors associated with building the access network;
- gathering all the key costs associated with the proactive and reactive building of the access network;
- setting boundary limits to the model; and
- designing the model outputs.

The model was built using Microsoft Excel for flexibility. The use of a more suitable platform that can manipulate the increasing amount of data gathered is currently being investigated. Having used the model to determine the optimum operating point, the cost of moving to that point was calculated.

Identification of Key Cost Factors

The key cost components were identified and split into proactive or reactive build costs.

The key cost components for *proactive build* are:

- cost per pairs added proactively,

- stock depletion,
- planning overheads, and
- network interventions.

The key cost components for *reactive build* are:

- cost per pairs added reactively,
- pair-gain usage,
- CGCS payments,
- planning overheads, and
- network interventions.


These components are sensitive to varying MFS levels in different ways, described in more detail below.

The summation of all the different costs produces an overall cost at a specific % MFS; that is,

$$\begin{aligned} \text{Total operating cost (\% MFS)} = \\ \text{total proactive costs (\% MFS)} \\ + \text{total reactive costs (\% MFS)}. \end{aligned}$$

Thus the % MFS with the lowest total cost would become the optimum. To keep absolute clarity and to establish payback for future investment the categories were finally grouped into current (operating) and capital (investment) costs. These were used to support the net present value analysis as discussed in the commercial value section below.

Figure 1—Percentage MFS annual depletion



The following sections summarise the various cost factors, their effect on overall cost and their variation with % MFS. Initial costs were based on the current performance of BT's access network % MFS.

Cost Factors for Proactive Build

Cost of proactive build for greenfield and organic main side

Proactive build to maintain MFS comprises a mix of greenfield (extending the network to new developments) and organic (increasing the capacity of the existing network) main-side (M-side) build. Greenfield and organic build costs were based on the 1998/99 capital investment programme budgets. These costs are relatively stable due to the steady growth of new developments per annum.

Stock depletion

The model includes an allowance for the annual depletion of MFS capability. This describes the effect whereby, once the network is dimensioned to achieve a level of MFS, allocation of network to customer orders depletes the MFS capability of the network. Further expenditure will then be required to bring MFS performance back to the required level.

A study of the 1997/98 MFS performance allowed a graph to be

plotted based on network depletion over a year at the current build policy. The graph shown in Figure 1 clearly demonstrates that the lower the % MFS the more it will deplete over a year as a given take-up of pairs represents a higher proportion of the spare capacity. Typically, from the graph, this rate of depletion of stock is:

- 1% annually at an 97% level of MFS; and
- 2.5% annually at a 50% level of MFS.

Note: Network stock depletion is inversely linked with the network build policy. The copper network does not deplete linearly with new orders due to the churn of customers within the existing network. The cost of maintaining percentage stock depletion is derived from the same factors as used for restocking the network to a new % MFS level.

Proactive Intervention Costs

Intervention risk has been analysed, and it has been proved that faults rise linearly per network intervention. This intervention risk was based on work previously modelled by BT Laboratories which concluded that there is a consistent risk of introducing an intervention fault when a node is opened (where a node is any point in the network; for example, PCP, DP and joints). Because the intervention fault could be either overhead or underground, the average cost per overall network fault has been used to give the formula:

$$\text{Intervention cost} = \begin{aligned} & \text{nodes opened} \\ & \times \text{risk factor} \\ & \times \text{cost per fault.} \end{aligned}$$

For added copper pairs it was assumed that only 0.1 node would be opened per pair added. This is because pairs will be added in minimum units of 5 or 10 pair cables, but typically 20 pairs upwards.

Planning overhead

Planning overheads for proactive build were set to an assumed 0.25 hours per pair added and then multiplied by the direct labour man-hour rate. This was then universally applied to all of the proactive build and greenfield volumes. The planning overhead is split $\frac{1}{4}$ current (operating costs) and $\frac{3}{4}$ capital investment to allow for accounting apportionment of costs.

Cost Factors for Reactive Build

Reactive network build cost per referral

The capital and current account costs per referral were calculated from actual cost data from the relevant build programmes, along with the pair-gain spend. The costs modelled display a sharp reduction from the current % MFS towards 100 % MFS, and a rise exponentially from the current % MFS towards 50% MFS. Factors that drive these costs include:

- increased use of pair-gain equipment,
- increased use of external contractors as resource demands increase,
- inefficiencies as the volume of POS referral increases to an unmanageable level, and
- increased referrals in greenfield situations where infrastructure must be provided reactively.

Pair-gain usage

Pair-gain usage, based on operational experience, will increase from 13% of referrals at 97% MFS to 50% of referrals at 50% MFS. The pair-gain operating costs refer to the current account expenditure associated with increased appointment time needed to fit the remote equipment. Costs were calculated from the actual provision time increments multiplied

by the direct labour rate to give a cost per pair-gain provided. The capital cost to provide the pair-gain equipment has been included in the reactive network build cost per referral.

Reactive intervention costs

As mentioned earlier under the proactive costs, the intervention risk has been analysed and proved that faults rise linearly with network interventions. Therefore, the changing ratio of pair-gain usage has been taken into consideration. The intervention for pair diverts is calculated at a minimum of two nodes being entered. It is envisaged that the number of nodes opened through referral work will rise to five interventions at 50% MFS from 1.7 interventions at the current % MFS. The average cost per intervention fault was calculated from actual costs. Because the intervention fault could be either overhead or underground the average cost per overall network fault has been used. The formula used is:

$$\text{Reactive intervention cost} = \begin{aligned} & \text{nodes opened} \\ & \times \text{risk factor} \\ & \times \text{cost per fault.} \end{aligned}$$

Intervention risk was based on the same parameters as for the proactive intervention faults.

For pair-gain usage an intervention factor of one per system has been used to allow for the extra activities required for provision. In addition, pair-gain equipment fault rates in the first year have been included at the current rate to calculate the number of additional pair-gain faults in the first year. These were costed at the same value as for intervention faults mentioned above to allow for the ratio between underground and overhead equipment.

Planning overheads

Planning overheads based on actual man-hour data were calculated by dividing the total number of hours by

Figure 2—Total costs of operating the access network

the total number of referrals, and then multiplying this by the direct labour man-hour rate. The planning overhead is split $\frac{1}{4}$ current (operating costs) and $\frac{3}{4}$ capital investment to allow for accounting apportionment of costs.

Customer guarantee compensation scheme (CGCS) payments

Residential CGCS is a standard payment to customers for whom BT fails to achieve the agreed provision-of-service date. The value of the payment is based on the number of days in delay.

For the model, it has been calculated on the actual cost per day. Although overall CGCS costs and average days payable could be ascertained, it was not possible to attribute actual costs to purely network problems. Therefore, based on operational experience, the number of orders for which BT is liable for CGCS payments will exponentially rise from 8% of referrals at 97% MFS to 20% of referrals at 50% MFS.

At the same time, the number of days compensation payable per order will rise, as the volume of work generates an increased backlog, from the current 7.9 days at 97% MFS to 20 days at 50% MFS.


The formula used is:

$$\begin{aligned} \text{CGCS liability} = & \\ & \text{days payable} \\ & \times \% \text{ referrals payable} \\ & \times \text{cost per day} \end{aligned}$$

Note: Although most of the CGCS costing was assumed, varying the CGCS costs does not have a significant influence on the optimum % MFS.

Cost Factors for Restocking to a New % MFS Level

The cost of restocking the network to a new % MFS is a one-off cost and based on the building of the distribution side of the network to support the greenfield and organic M-side build



requirement. The proactive organic build cost per added distribution side (D-side) pair has been set at the current network build cost. This cost is fixed because the cost per added pair would not differ for a varying % MFS.

Targetting this network build is key. Without very specific and accurate forecasts of demand at the distribution point (DP—last node before customer) level, it is acknowledged that, to guarantee a specific MFS performance, more pairs must be added to the network than will be used within the normal provisioning period. The higher the MFS target, the more of an issue this becomes, as targetting DPs for new build will be more difficult. Therefore, the D-side costs have been increased by an organic D-side pair ratio.

Current network build targetting and analysis suggests that this ratio will rise from: 1.6:1 at 50% MFS up to 24:1 at 97% MFS.

This ratio does not currently allow for new initiatives to improve D-side targetting by using market-driven forecast data and new processes for proactively selling stock.

Defining the Model's Boundary Limits

The boundaries of the MFS model were set to show varying levels of MFS at 1% intervals for the range of 50% to 97%. The upper limit was chosen because BT has a universal service obligation for PSTN within the UK, so it would be unrealistic to aim at achieving 100% MFS. The lower limit was set at 50% MFS, because the level of churn would result in too many assumptions, making the model inaccurate.

Design of the Model's Outputs

The model's output was designed as one graph showing the total operat-

ing cost, the total proactive build costs and the total reactive build costs. The lowest point on the total cost line shows the access networks optimum operating point. (See Figure 2.)

In addition to the above graph, a table with all of the total costs was constructed so that discounted cash-flow analysis could be applied to determine the cost and payback for building the network from its current % MFS towards the optimum operating point. This determination is discussed in the commercial value section of the article.

Discussion of model's outputs

Examination of Figure 2 raises a number of useful observations:

- There is a residual level of proactive build at the lowest MFS level considered. This represents the combined cost of greenfield and M-side organic build, which is currently considered inappropriate for reactive provision. Any changes in process or technology which altered this position would need to be reflected within the model.
- These proactive costs only begin to rise at the current MFS operating point. This confirms the very low level of proactive organic D-side build currently undertaken.
- The virtual impossibility of meeting 100% of customer orders from stock would seem to predict that proactive (and overall) costs would rise exponentially as the MFS level approaches 100%. Although the graph shows an upward trend, it is not as steep as might be expected.

This possibly indicates some weaknesses in the model towards the high end of MFS, which is consider-

ably above BT's current operating experience. As the derived optimum operating point is close to this end of the model, it will be important to continually refine the model as experience of operating at higher MFS levels increases.

Practical Application

This model has been predominantly used within BT for setting network investment levels and resource requirements for building and operating the access network. Another use for the model is for strategic development of the access network strategy for build, where flexing of the costs is essential for scenario analysis. In particular, the model clearly identifies those areas where improvements will generate the best payback in terms of 'value for money' investment in terms of reduced capital and current account expenditure. One of the key factors in the model is the organic D-side pair ratio, which compensates for inaccuracies in targetting proactive build. Development work is currently proceeding to use the best available marketing forecasts of overall demand and demand type at a postcode level. By considering the access network in the same geographic way, it is expected that the improved granular forecasts will enable the ratio to be significantly reduced. This, in turn, will reduce the proactive build costs considerably and hence move the optimum operating point further forward towards 100 % MFS, allowing an overall higher economic MFS performance and improved customer quality of service.

Commercial Value

The commercial value to BT is in the increased understanding of the relationships of network operating and investment costs, which can be used to improve customer quality of service.

The investment costs to build the network to meet the optimum were

calculated, and discounted cash flow analysis was applied for each incremental increase in % MFS to yield a true value of the investment. To include the time value of money, cash flows in future years have been discounted using a retail price index (RPI) of 3.5% and a test discount rate (TDR) of 8%. The RPI reflects the general inflation of the industry and the TDR reflects the level of return expected from invested capital over and above the RPI rate. These two figures are then combined to produce an overall discount rate. The net cash flow for each year over a five-year costing period is then converted into present values using the discount rates and the sum of these present values is referred to as the net present value (NPV). The discounted payback period is the number of years it would take for the NPV to move from a negative figure to zero. The model has therefore been used in the production of internal business cases for network investment and the targetting of network stock levels.

Based on BT's current performance and cost data analysed, the optimum operating point is at a MFS level of 95%. This figure is determined from the model based on the assumptions and actual costs analysed where it is the lowest total 'whole life' cost for access, taking into account all of the reactive and proactive build costs. The 95% MFS figure shows that to operate above the optimum would increase costs due to the difficulties in targetting proactive build, and the spreading of added pairs to meet the new orders. Also, if we operate below the optimum there are increasing reactive costs that involve a considerable increase in people for the reactive field force. The development work to increase network targetting of proactive build will reduce proactive investment costs and increase the optimum operating point.

Ignoring the cost to move from today's actual % MFS figure to the optimum and sustain that level for one year would save BT overall

approximately £100M per annum in operating the access network. This figure has to be balanced against the capital investment cost to derive an appropriate build strategy.

Future Developments

The next stage in the development of the model is to test and validate the assumptions made, and identify new data sources to replace these assumptions. The second stage is to develop the levels of sophistication of the model and concentrate on the area around the optimum operating point. At the same time a graphical user interface (GUI) is being developed in line with the investigation for a more suitable platform to operate this model from. The GUI will enable all of the key cost factors to be changed easily, so the model can be used for scenario analysis.

Conclusion

The initial objectives were to establish a model that could determine the optimum operating point, the balance of proactive build versus reactive build and the cost to move to the optimum. These objectives have been achieved as detailed in this article. This work has also shown the value of cost modelling of the various factors involved in managing the access network as a basis for justifying and targetting capital investment. The concept of the optimum operating point is clearly of great use both in strategy and in understanding the effects of changing the way business is done within the existing environment. Any changes to the environment can easily be reflected within the model. The model has proved that it is more cost-effective to proactively target network build than to reactively build the network on the acceptance of a customer order. Reactive network build is only cost-effective in areas which are not suitable for targetted network restock; for example, sporadic churn in rural areas could never be fore-

casted. Therefore, because of BT's universal service obligation, 100% MFS is unrealistic even if it is desirable. Enhancement of BT's network profiling and marketing-driven targetting tools has started to improve the targetting of the D-side network.

BT is currently evaluating the desirability of making the necessary capital investment to move towards the optimum operating point that this model has identified.

Acknowledgements

The authors acknowledge BT colleagues Phil Thompson, Enda McSharry, Lawrence Bickers, Margaret Harrison and David Martin without whom this work and article would have been impossible.

Biographies



Ian Lawrence
Networks and
Information Services,
BT UK

Ian Lawrence is the Strategy and Planning Manager within the Access Infrastructure Unit of BT's *networkaccess* division. He joined BT in 1986 as a trainee technician on network provision and repair. In 1988, he joined Customer Engineering specialising in the repair of small-to-medium business systems and network special fault investigation. In 1990, he moved to BT's Business Systems Technical Support Unit providing dynamic technical support for all customer premises equipment up to 30 lines and special network faults including electromagnetic compatibility and radio interference. During 1992, Ian took on the job of National Pilot Manager for BT's UK Access Division where he was responsible for the pilot trials of many of BT's new access technologies; for example, digital pair gain, HDSL, Mark 2 blown fibre and LA30. He now leads a team of managers responsible for owning, developing and implementing BT's access narrowband infrastructure planning policies and standards, which includes planning policy automation, planning training courses and the sponsorship of development work at BT Adastral Park and BT Development and Training. He is currently studying for an MBA with the Open University Business School.

Ian Lawrence can be contacted at ian.c.lawrence@bt.com



Derek Cottrell
Networks and
Information Services,
BT UK

Derek Cottrell is the Network Performance Manager within the Access Infrastructure Unit of BT's *networkaccess* division. After joining BT in 1972 as an apprentice, he became a Technical Officer on local exchanges. In 1987, he joined the Exchange Conversion Customer Data Management team. In this work area, he introduced many innovative data-comparison techniques to fully exploit the exchange transfer data. As a result of this work, he was promoted to Data Integrity Manager in 1989. His next role was in Network Records where he introduced improvements in the provision-of-service (POS) process. It was here that he developed a management information system (NEMESIS) to measure net from stock. He then became involved in the Fault Reduction Programme before completing a short spell in Finance. In 1996, he joined the centre team where he has been responsible for: introducing meaningful pair-gain statistics, enhancing the now national, NEMESIS and providing BT with a set of concise network performance measures. He is currently exploring the relationship between network performance and productivity. He is also leading on the cost/performance of access network products as well as further developing the optimum MFS model.

Christine Stratton

Investment Decision and Control System for Access Planning

Integrated Business and Network Analysis—A Case Study

This article describes the development and integration of systems to support investment decisions made by BT network planners. The first, designed by the planners themselves, led to the development of a new decision model for network planning. The system provides information on geography, revenue, forecasts, volumes, fault rate, spare network capacity and network capability. It has integrated reports from management information systems and applied intelligence to implement the planning rules and investment decisions at local level.

Introduction

This article describes the development and implementation of systems to support investment decisions made by BT network planners. The first development was an *integrated quick hit tool*, which led on to the development of a new decision model for network planning. Both developments pulled together data on geography, revenue, fault rate, spare network capacity and network capability. Today, these systems provide decision support not only to operational and strategic planners, but also to sales and marketing departments. The system was developed and implemented in all BT offices in less than a year. Success criteria for fast development and a smooth implementation are discussed alongside the barriers that had to be overcome. Finally, the benefits achieved and the enabling management style are noted.

Integration of Network Planning and Business Planning

As BT moved from a monopoly to a regulated competitive environment, it found a need to target capital investment in the access network in a way that had never previously been envisaged. Network planning had been focused on technology and products; it had been sufficient to build service and network manage-

ment systems that were task specific and not integrated.

The new environment demanded network investment decisions that took account of:

- customer requirements (speed of delivery, reliability),
- service profitability (cost of delivery and associated revenue),
- network capability (ability of the network infrastructure to deliver the required service), and
- regulation (accounting separation between network wholesale and retail activities).

Network and business planners needed decision-support systems that would bring together data from a variety of information systems. For BT this meant using its data more effectively.

BT is fortunate in having a wealth of data sources from which to build a better understanding of the customer base and the network. The value was to identify the relevant information from within the mass of data.

The first step was to benchmark BT's access network activities against other leading network operators. The study revealed shortfalls in the provision-of-service areas. A more detailed understanding of the network components and their relevance to costs was essential.

From this work two main objectives were identified:

- to reduce costs by better matching expenditure to revenue, and
- to identify opportunities for better utilisation of existing assets.

To meet these objectives it was necessary to describe the network capability and its capacity at each node in terms of geography, revenue, fault rate, spare capacity and service capability. (A node in the access network would be a local exchange, a primary cross-connection point (PCP or cabinet), a distribution point, or the first 100 pairs from a cabinet—even language could be a challenge!)

Integrated Quick Hit Tools

The first service addressed was PSTN delivered over copper twisted pairs. This is BT's largest service in terms of number of customers and assets used/capital employed. Success with this service would have maximum impact.

The first problem was to know where the primary cross-connection points were. Of course they had not been lost, but never before had there been the need to describe them in terms of eastings and northings. An early decision was made to make the data visible both in table and 'map' format.

The second problem identified the first barrier to systems and data integration. BT's customer locations are identified by postcodes—postcode boundaries do not coincide with local exchange boundaries. An estimate of the correlation had to be made and this proved reasonably successful—subsequently this correlation could be validated by using the mapping tool.

The third problem was revenue. Billing data only provides information about outgoing call revenue. Building an information system which ignores incoming call revenue would understate the value of

businesses such as credit-card validation and mail-order service centres where the number of incoming calls far outweighs the outgoing. Modern data-mining techniques allow the analysis of calls from both directions; BT has this capability and it was used.

To deliver the first 'quick win', the exchange boundary, PCP geographic location, revenue per PCP and business/residential customer-mix information were integrated onto an off-the-shelf mapping package, called *MapInfo*, and coloured the PCP nodes according to revenue quartiles.

The result allowed the network to be described to marketing colleagues in terms of geographic importance according to revenue streams. The simplicity of the presentation was also well received by senior managers who, having seen the strategic importance of this prototype demonstrated, agreed its further development to include other factors.

To be of major use to the company, this system had to be translated into something which would directly influence the investment of capital in the access network. To do this, the way in which access network planners make investment decisions had to be changed.

New Models for Network Planning

The approach was to talk to the planners, observe how they really worked and identify the processes they really used and, most important of all, the systems they had devised to make their job easier.

Some enterprising planners had already begun to use Dbase and Excel to manipulate some of the information they used regularly. We took the best of these developments and migrated them into the more robust Access 2 and later Access 97 media.

The development locked into BT's three-layer system architecture. The first layer comprises the huge databases where the company stores

customer, technology and service data. The second comprises a number of managed data sources where the raw data is processed into meaningful 'chunks' to supply service specific management information. The third level, the application layer, makes full use of the capability of packages such as Visual Basic and Access. It is designed therefore as a dynamic front end with intelligence but no large storage capacity; it is PC-based and consequently can be very user-friendly. It continued to use the expertise of the user and by now a small team of developers from the user community had been gathered.

Requirements for an Automated Planning Decision Tool

Working with the network planners, the development team identified the requirements for an automated planning decision tool:

- common tool in use in all Zones,
- flexible and fast implementation of change,
- decisions which could be specific to the geography,
- local applications,
- two-way data feeds between planners and sales/marketing, and
- accurate data collection from field engineers on completed work (new or modified network assets).

Features required by the central business planners were also identified:

- central master model—for strategic decision making,
- rigorous data change control,
- adequate system security, and
- compliance with the Data Protection Act.

Figure 1—Systems architecture

The Solution

A PC-based system (Layer 3) was provided which interfaced with the managed data layer (Layer 2) through an integrated data platform (Figure 1). This provided a common data layer for the Zones and a central master model, and allowed changes to be simple and, above all, fast.

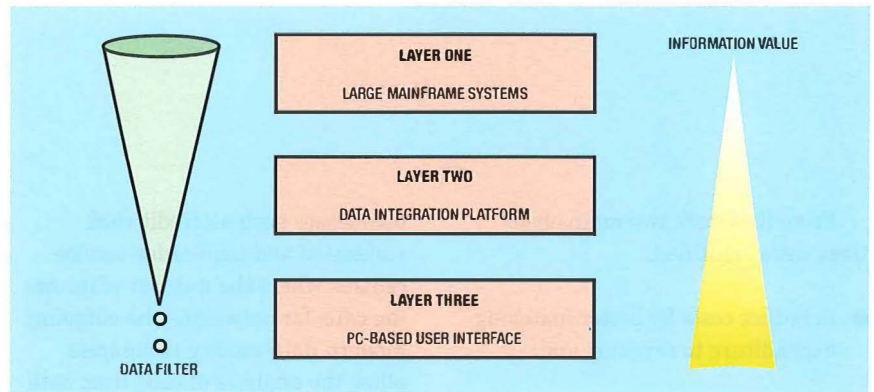
It was critical that the process owners agreed the changes to the application layer. Changes that needed to be made would be verified and validated on the prototype before being released. The advantage of this was that the managed layer maintained the integrity of the data available to other systems, while facilitating changes to the master data to be properly managed.

The managed data layer, the integrated layer, had to be built. This involved establishing interfaces with the source data, the first layer. The reports, which could be integrated on the second layer, were extracted at regular intervals, the frequency of which has had to be increased as the system moved from largely proactive solutions to include service delivery more associated with reactive work.

The management of the security of the system was also easier in this system configuration. The practice of not storing customer and financial data together on one platform ensured compliance with the Data Protection Act. The application layer could be password protected but the integrated layer was maintained under the company's security code.

Recognising that not all data on existing databases is accurate, the system was designed to provide a mechanism for capturing and correcting data errors, thereby improving data integrity. The integration layer provided the opportunity for data comparison and 'cleaning'.

The centre control facilitated cost and price modelling, both strategic and operational, and gave to the planner a system which enabled prioritised planning decisions to be



delivered on a node-by-node basis. Further work linking budget control models with the real-time finance and work control systems will enable the process to obtain real costs for provision, maintenance of service and safety measures. In the prototype, the addition of depreciation brought close a whole-life cost model that was geographically based.

Performance and Evaluation

Once the new information system, investment decision and control system (IDACS), had been established, some measurement of success was needed. Initially, it was felt that the automation of processes would deliver reductions in the number of planners required. However, as the system matured and the applications grew, a number of other benefits emerged and the system became a key enabler for change in organisations and cultures within the network arena.

Asset Utilisation

IDACS presents to the user a number of data-sets representing BT's fixed access network assets together with forecast network growth elements. By using this functionality the system can determine network growth requirements and identify areas of surplus network elements for redistribution (Figure 2).

Figure 2—Dynamic front-end applications



IDACS presents trend information that has been a key to enable the successful cost control of the digital access carrier system (DACs) BT UK budget, and the identification of surplus exchange cards which are then redistributed on a 'lift and shift' basis. This has resulted in significant cost savings to the business, driving up asset utilisation towards optimum levels.

More recently, this functionality has been extended to manage DACs remote units within the external network. Recent trials in the North East Regional Business Area indicate significant capital savings by rationalising single working DACs channels at the same DP. This releases a copper bearer at the DP and enables the recovery of the unit back into stock for future use.

Focused Network Investment

The ability for IDACS to focus the user to areas of network shortfall is a key benefit.

The former North West Zone has adopted the use of IDACS from its conception through to present-day. This enabled the establishment of an *access network capacity unit* to manage the E-side network of some 10 000 PCP nodes.

An early benefit was the ability to map network requirements identified by IDACS across the outstanding planned job stack. This enabled jobs

to be prioritised and, in some cases, cancelled as they had timed out. Significant capital savings were thus made and the outstanding work stack was reduced to more manageable levels.

This was achieved without degradation in terms of met-from-stock (and the number of orders referred due to lack of E-side stock).

Benefits were also achieved in the development environment:

- Testing was carried out initially in the laboratory until the integration layer was stable. Once this was achieved, new developments were moved to prototype very quickly and the delivery into the field for test-bed trials was very much faster.
- Verification testing was not needed to the degree of rigour that developments at the second layer would have required, because the third layer did not change the base data.
- Through the integration of data from separate sources, data inaccuracies became apparent. Correction of these inconsistencies improved overall data quality.

Management Style

The system development described above was not automating what was being done but was being used to change methods and processes.

This delivered operational benefits to the business.

Job management for customer service

Pre-IDACS (See Figure 3)

Proactive network provision, maintenance, and repair requirements are identified by conventional means by reference to a myriad of data storage systems in the planning office environment. This wide and varied storage media delivered a range of data sets for evaluation to determine the network

requirements and record the resulting actions taken. Attempting to source the data upon receiving a customer inquiry presented the planning team with a series of complex analyses in order to answer the pending question. The extraction of management information presented an equally difficult task. This system was not flexible enough to cope with the changes in demand for customer service and reflect the regular changing needs of the customer-service environment through organisational change.

Post-IDACS

Proactive network provision, maintenance, and repair requirements are identified and recorded in IDACS. The myriad of data storage systems in the planning office environment are built into the data-sets contained within IDACS. Data is presented through a front-office module within IDACS upon receiving a customer inquiry to simplify complex analysis in order to answer the pending question more proficiently. The extraction of management information is equally simple. The IDACS system has total flexibility to cope with the changes in demand for customer service and reflect the regular changing needs of the customer-service environment through organisational change.

Success Factors

The success of the system can be attributed to the fact that the team was not made up from traditional system designers but rather planners who had some systems experience and who had a problem-solving approach and no preconceived ideas about the solutions. This probably meant that the team did not avoid some of the pitfalls that may have been avoided had the team comprised fully-fledged systems analysts. The team took a 'can do' approach.

Senior management commitment to the development was essential and ensured several barriers to progress were dismantled.

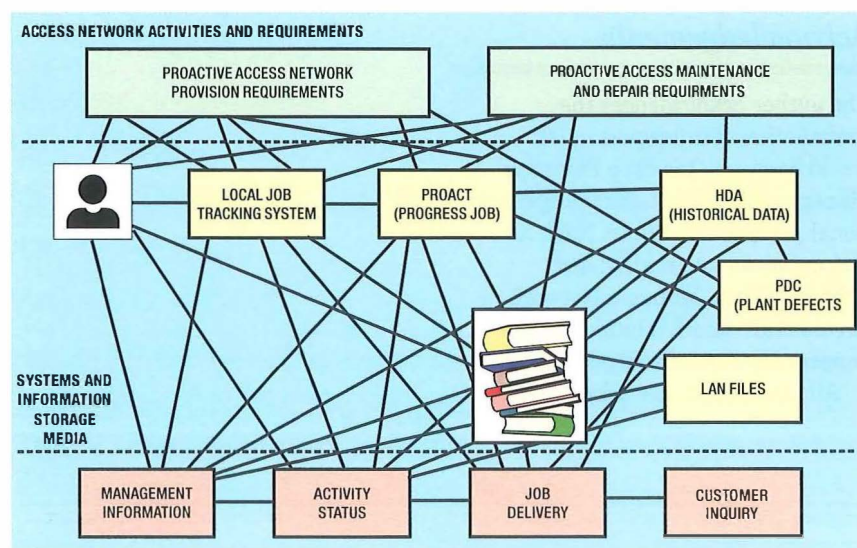
The importance of communication was learnt and in particular the impact of demonstrations both to senior members of the company and to the users. It was critical to the success of the system that the team never lost contact with the user.

Finally, the dedication of the team to the project and their determination to succeed were key ingredients to rapid progress.

Conclusion

The initial application of capital targeting is a reality. Information about fault history, product volumes,

Figure 3—Job management for customer service, pre-IDACS



churn and capacity are applied to revenue information by geography. Application of the decision rules results in a prioritised list of costed jobs for BT *networkbuild* and facilitates the five-year budget build.

The system was built by people who had come from the user environment. Often the local tool designer was matrix managed as part of the design team until the tool was fully integrated into the system. This methodology led to the users saying 'when can we have it?' rather than 'can't see why we need it'. We also had a user forum that provided invaluable guidance.

The system has enabled a more radical approach to network planning. Organisational barriers have been broken down and a culture of customer service solutions implemented.

The key benefit to BT is that for the first time consistent decision making is effected nationally using the latest cost information, and the results can be measured. This facilitates a consistent national, regional and local network health score to be made available throughout the company.

For the future, a total-cost model will be developed based on the PCP which will allow line rental and connection charges within the company to be flexed in line with real network costs.

Acknowledgements

The author acknowledges the contribution of colleagues in BT: David Spencer, One-Step Planning Manager North West, for the operational perspective; Stuart Nash for making ideas into reality; and Anthony Dunn, Manager Network Architecture and Evolution, for his support during the analysis phase.

All trademarks are acknowledged.

Biography



Christine Stratton
Networks and
Information Services,
BT UK

Christine Stratton is currently responsible for developing and implementing business and network decision support systems within the Change Management Unit in *networkBT*. She moved into telecommunications 10 years ago after completing a second degree, in technology, her first being in sociology. Her telecommunications experience ranges from management training where she was responsible for setting up the first industry-based M.Sc., to systems engineering, network evolution planning and product launch. She has managed teams in the UK and America. Recently, she conducted studies analysing the root causes of faults in the digital network and carried out a national review of planning in the BT access network.

Christine Stratton can be contacted at chris.stratton@bt.com.

David Duxbury, Rob Backhouse, Mike Head, Graham Lloyd and John Pilkington

Call Centres in BT UK Customer Service

More and more business is being done via call centres, both within the UK and worldwide. BT has been involved in call centres for many years. This article reviews the current situation in BT UK Customer Service, and considers the next steps in this increasingly competitive and challenging environment.

Introduction

Call centres are a growing industry, in a recent Ovum Report it was estimated that by 2000 5% of the working population in Europe would be employed in call centres. It is also predicted that call-centre revenues globally will be \$6 billion. Call centres have become the core of the service economy in the UK. Counting part-timers, UK call centres currently employ 1.7% of the working population, or nearly 400 000 people. And the numbers are growing. Datamonitor predicts that call-centre positions will double by 2002 before beginning to level off. Since 1994, customer calls to large organisations have roughly doubled, with call centres being largely responsible.

BT UK is a leader in call-centre technology and expertise. It operates its own call centres and manages them for other organisations. BT's own call-centre operations fall into the two broad categories of *outbound* and *inbound*. Outbound operations involve companies' staff calling out to customers, usually to offer new services, while inbound call centres handle enquiries from customers. This article looks at how the inbound call centres are presently used in the following environments, which between them are responsible for handling approximately 1 billion calls per year:

- answering Operator Assistance enquiries (100, 155 and 999 emergency services),
- giving accurate number information (192 and 153 services),
- receiving sales and billing enquiries and taking orders (150), and

- taking fault reports from residential and business customers (151 and 154).

It does not consider some of the more specialist centres (for example, Welsh language, Directory Enquiries for the blind and disabled) or inbound services and business after-sales units, although these all form an important part of the BT Customer Service call-centre management operation.

The article explores 'where we are now' and in particular focuses on systems, switching and queuing capabilities. Finally, it looks at the future of call centres and gives some indication of the way things will change.

Evolution

Although the title *call centre* would not have been used, BT has been taking calls from its customers ever since telephony was introduced. In the early days, these were large operator centres serving the geography around the exchange. Over time, separate centres for activities such as sales and fault reporting were introduced, but these again were geographically based, generally small and supported only by simple technology.

The evolution from Telephone Areas to Districts and then to a more functional organisation significantly reduced the number of centres, but there are still over 100 units and by modern standards many of these would be regarded as small. However, the introduction of both switching and systems technology has enabled BT to maximise the efficiency of these units, although more can and

Figure 1 – Operator Assistance information systems

will be done (see paragraph on call centre futures).

Operator Assistance

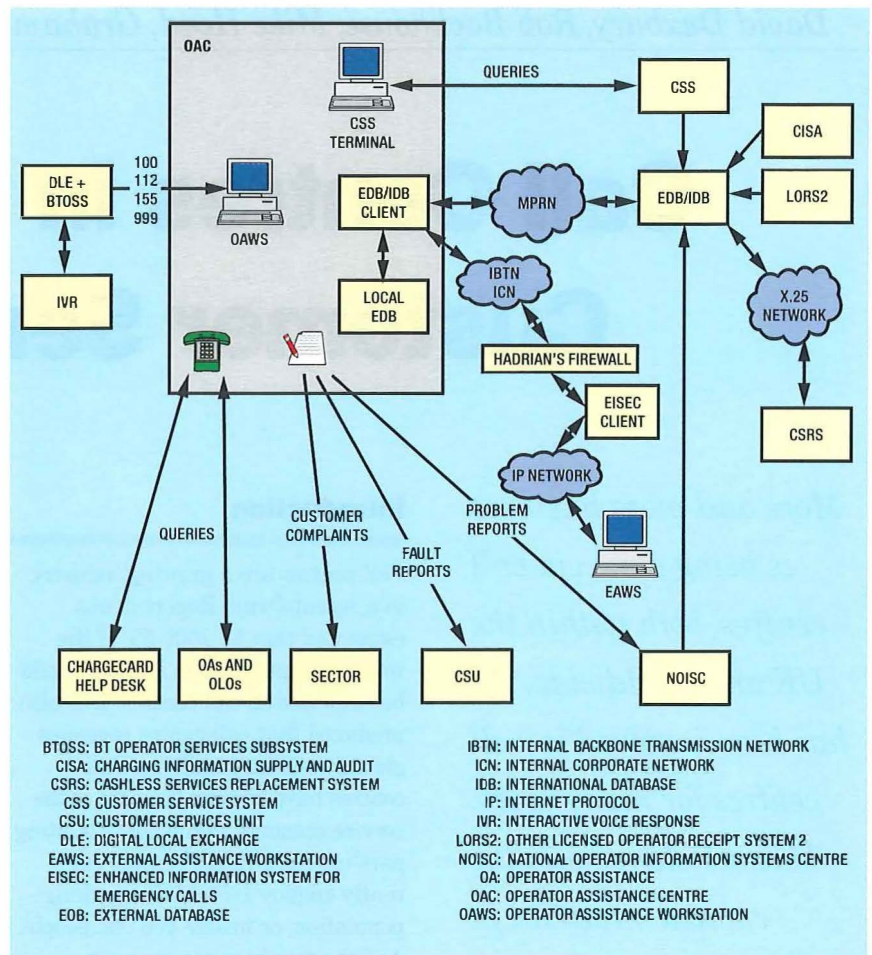
Originally the operator was the only 'connect through' service for all calls. With progressive network automation, this capability is now only used when a customer requires assistance, when a customer encounters difficulty (for example, permanently engaged), or when the customer wishes to use special facilities such as alternative payment options. BT's Operator Assistance covers assistance (100, 155) and emergency (999) services and currently handles 500 000 calls per day at 15 call centres.

Network and systems

Figure 1 illustrates the information systems involved in Operator Assistance.

Common call-centre technology supports all types of call for Operator Assistance (100), emergency services (112, 999) and International Operator Assistance (155). System X BT operator services subsystem (BTOSS) switches route calls to a free operator together with associated network data—such as calling line identity (CLI) and class of service. The operator assistance workstation (OAWS), developed by BT, provides integration of the telephony and information systems, to present operators with a highly integrated and user-friendly interface to the underlying systems. This allows it to:

- obtain information from the network via BTOSS,
- obtain information from the external database (EDB) (state of line, name and address etc.),
- obtain information from the international database (IDB) (time zones call tariffs, routing, Inmarsat), and
- vet BT Chargecards using CSRS (via EDB as a gateway).



To maintain the accuracy of EDB, daily data feeds to EDB are received from both the customer service system (CSS) and other licensed operator receipt system 2 (LORS2).

The design employs object-oriented software and a script engine to achieve a very efficient and flexible user interface to capabilities which embody the required operating procedures. The workstation also incorporates voice technology, which automatically plays a greeting announcement in the operator's own voice on call arrival.

The handling of emergency (999) calls demands extremely high systems resilience and the architecture is designed to achieve 99.9999% systems availability. Information required to handle 999 calls is held locally and is available from remote servers in the event of failure. Further fall-back facilities are also provided to protect against multiple failures. Similarly, multiple levels of network fall-back are used to ensure network capacity is always available for these calls and to automatically route around network and call-centre failures. A recent enhancement has been the introduction of the en-

hanced information system for emergency calls (EISEC). This provides the caller name and address information for a 999 call to the emergency authority through the BT firewall, saving critical seconds in the response to an emergency.

Limited access, via the BT intranet, is also provided to the service management system. This enables operators to check for and report faults when required and transfer callers to an appropriate service department.

Directory Assistance (DA)

BT has a broad range of number information products including the well known Phonebooks and more recent innovations such as PhoneNet—the Phonebooks on the Internet (www.bt.com/phonenetuk). Directory Enquiries is a key part of the portfolio handling some 650 million calls a year at 47 DA call centres. Customer care, speed and accuracy are high priorities in handling Directory Enquiry calls. BT has recently completed the ATLAS programme, deploying state-of-the-art technology into this service.

Figure 2 – Call centre and information systems for BT Directory Enquiries

Network and systems

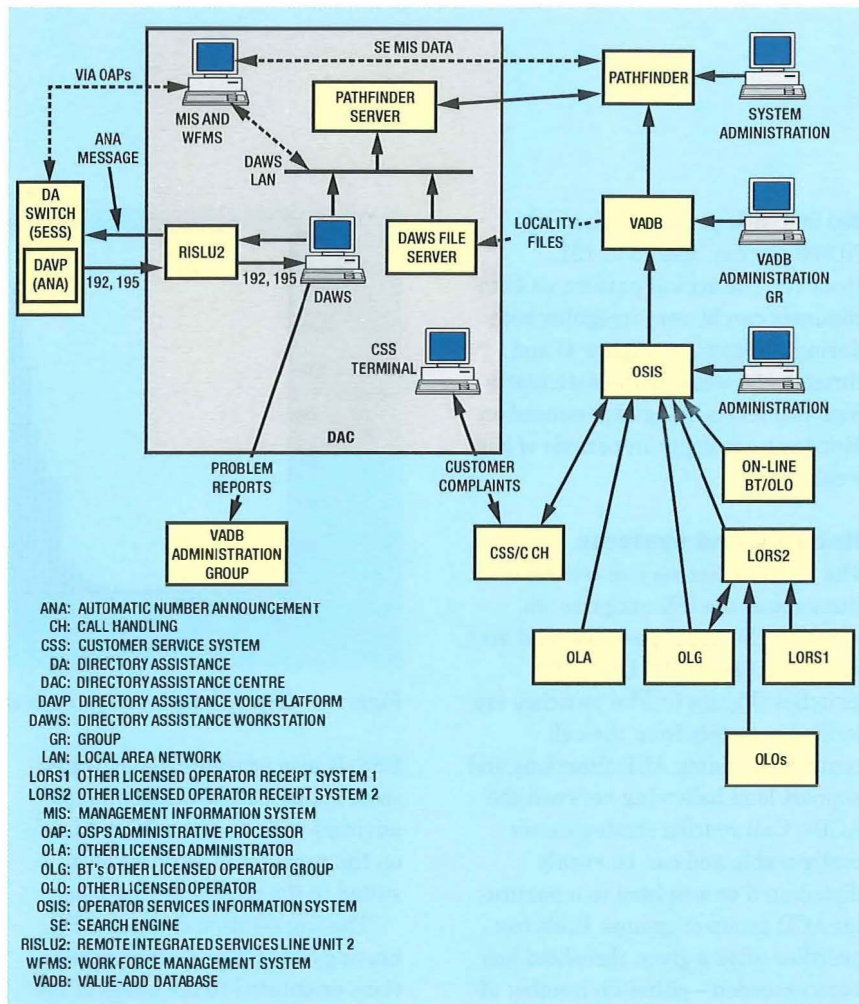
Figure 2 illustrates the call centre and information systems for BT Directory Enquiries.

The DA service uses Lucent Technologies 5ESS automatic call distribution (ACD) switches. Put simply, the purpose of the ACD is to place each incoming call into a single large queue so that the first free operator can answer it, irrespective of where the DA centre is located. However, for the service to operate efficiently and effectively, calls need to be handled by operators who have the best match of skills to the call. This is achieved as follows:

- Incoming calls are allocated to queues based on the locality of the caller (derived from the CLI).
- Operators are grouped into serving teams (skill groups) identified by operator identity or groups of workstations.
- Through configuration data, serving teams are allocated to each queue as the first choice to handle calls from a queue: the *primary team*. The configuration tables (Figure 3) also define the sequence of other serving teams to which a queued call is offered if a free operator in the primary team is not available within the desired answer time (the *secondary team*). The configuration tables can be changed dynamically and set-up to change automatically; for example, when a call centre opens or closes.

Figure 3 – Queue configuration table

Queue:	1	2	3	4	26
Primary Team:	A	B	C	D	Z
Secondary Team 1:	B	C	D	E	A
Secondary Team 2:	C	D	E	F	B
Secondary Team 3:	D	E	F	G	C
Secondary Team 4:	E	F	G	H	D
Secondary Team ...:
Secondary Team 25:	Z	A	B	C	Y



The ACDs are also integrated with management information and workforce management systems which are used to manage service quality and achieve an optimum match of operator resources to incoming call volumes.

Operator workstations

Directory Enquiry operators use the Directory Assistance workstation (DAWS). As with the Operator Assistance workstation (OAWS), this provides integration of the ACD and information systems, and employs an intelligent application. This presents operators with a highly integrated and user-friendly interface to the underlying systems and which supports the business processes. When a call arrives at the workstation, an automatic greeting in the operator's voice plays. The main function of the workstation is to provide operators with access to search the national directory information database for numbers and code information. When the correct number has been located on the database, the operator then transfers

the call to the speech applications server (SAS), which automatically announces the number with the familiar phrase: 'the number you require is ...'.

Pathfinder database

The heart of the system platform is the national Directory Enquiry database (NDIS), supplied by Hewlett Packard and PCPlus. Dimensioned to hold over 150 million listings, it will perform a search and return information back to the DAWS screen in less than 0.5 seconds. The database receives 40 000 updates per night.

Customer Service and Sales, Account Services, Sales and Billing and Residential Repair

With the evolution of call centres described earlier, similar technology has also been deployed into customer service centres (CSCs) supporting 150 (account services, sales and billing) and 151 (residential repair) calls. 150 accounts for approximately

300 000 calls per day with about 70 000 per day offered to 151. However, the arrival pattern on both channels can be very irregular both during the day (see Figure 4) and through the week, with particularly high call levels being experienced on Monday morning or in periods of bad weather.

Network and systems

The 55 customer service centres throughout the UK comprise, on average, 120 agent positions and are served by two Nortel DMS-100 switches (Figure 5). The switches are located remotely from the call centres, providing ACD functions and support load balancing between the ACDs. Call routing strategies are configurable and can be evenly distributed or weighted to a particular ACD group or groups. Calls can overflow after a given threshold has been exceeded—either on number of calls in queue or length of time in queue.

SMART

As with the previous call centres an intelligent agent workstation is employed—SMART. SMART is a three-tier distributed system, which provides a graphical user interface to the billing and order entry capabilities of CSS and to the data of the products and services database (PSD). It supports a set of configurable business rules, which reduces errors and provides a simpler more efficient set of procedures compared with native use of

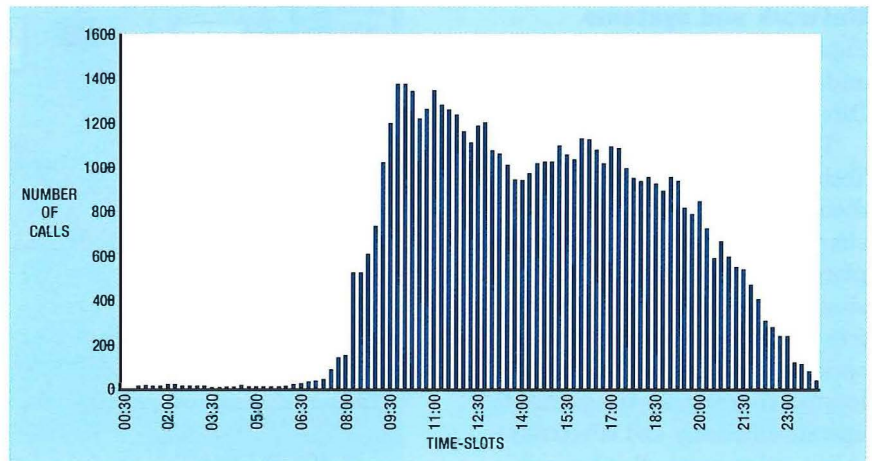


Figure 4—Typical day's call pattern into 151 repair

CSS. It also provides sophisticated sales support tools which prompt advisors with on-line selling 'hints' on the products or services best suited to the particular customer.

The workstation client is a PC, hosting SMART and other applications orientated to the needs of the inbound channels. A mid-tier supports local data, particularly business rules data, a communications gateway to the CSS systems and a set of 'operability' features designed to facilitate remote management of the SMART system across all three tiers.

Management information systems (MIS)

Good customer service depends on having the correct numbers of agents available to deal with the calls arriving at that particular time of the day. The resource and administration teams use management information displays to

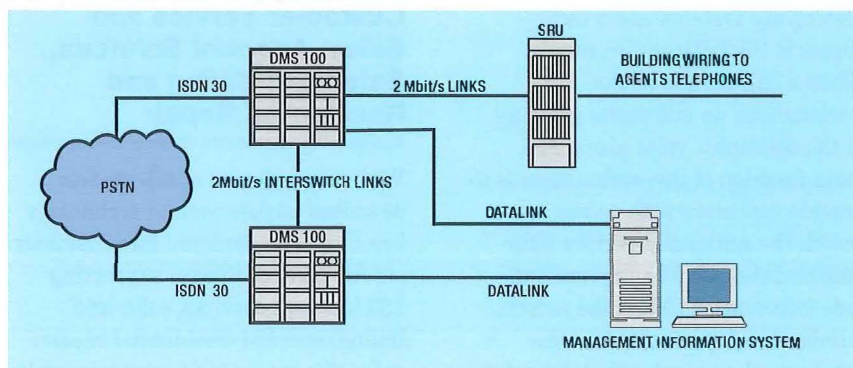
monitor service quality throughout the day. Most widely used in Customer Service is the Nortel RT1000 which displays real-time information on number of calls, number of available agents, number of calls queuing, length of time in queue, etc. Additional advisors are brought on-line when needed; however, a sophisticated forecasting and scheduling package ensures that accurate staffing is achieved based on projections derived from historical data gathered over the previous three years. RT1000 has a real-time adherence monitor, which alarms if the ACD queues are over- or under-staffed.

RT1000 also provides historical reports and is able to change the call flows in real time via load management commands to the DMS ACDs.

Call steering

In order to deal efficiently with the variety of call types being offered, and to route them to advisors with the required specialist skills, automated call steering is used. Syntellect interactive voice response (IVR) platforms are used in conjunction with Genesys computer telephony integration (CTI) servers and the Nortel DMS ACDs to capture details of the calling customer and the service request and then to route the call to the most appropriate team of service advisors. CTI is also used to present the customer account details to the advisors on call

Figure 5—DMS 100 configuration



arrival—so called *SMART screen-pops*. IVR-based call steering (Figure 6) permits a straightforward customer contact strategy while ensuring customers are routed to the correct ACD queue via simple menu selection from their telephone keypad.

Automation

Customer Service has also implemented a number of interactive customer 'self-service' applications, which enable customers to conduct simpler service transactions themselves using IVR systems. Examples include: changing Friends and Family calling circle numbers, simple billing enquiries, initial reporting and progress enquiries for faults. The fault reporting application performs a line test on the suspected faulty line and either reports the fault automatically or connects the caller to a service advisor. If the later option is chosen, the results of the line test are presented to the advisor's screen at the same time as the voice call using the CTI facilities as above. It is estimated that the automated work volume is equivalent to three call centres.

These automation services are also available to customers via the Web (bt.com) and, via the BT intranet, to advisors in other types of call centres; for example, Operator Assistance or Outbound Sales.

Business Customer Service

BT Business Customer Service has five main components comprising 154 Fault Management, 0800 666777 After Sales, Customer Relations, Simplex Management and the BT Own Use Service Centre in Gloucester. The combined units have in total around 1200 service advisors distributed across 27 centres throughout the country. The largest proportion of these people are within the 154 fault management centres (FMCs), which comprise around 600 people in 12 different locations, followed by

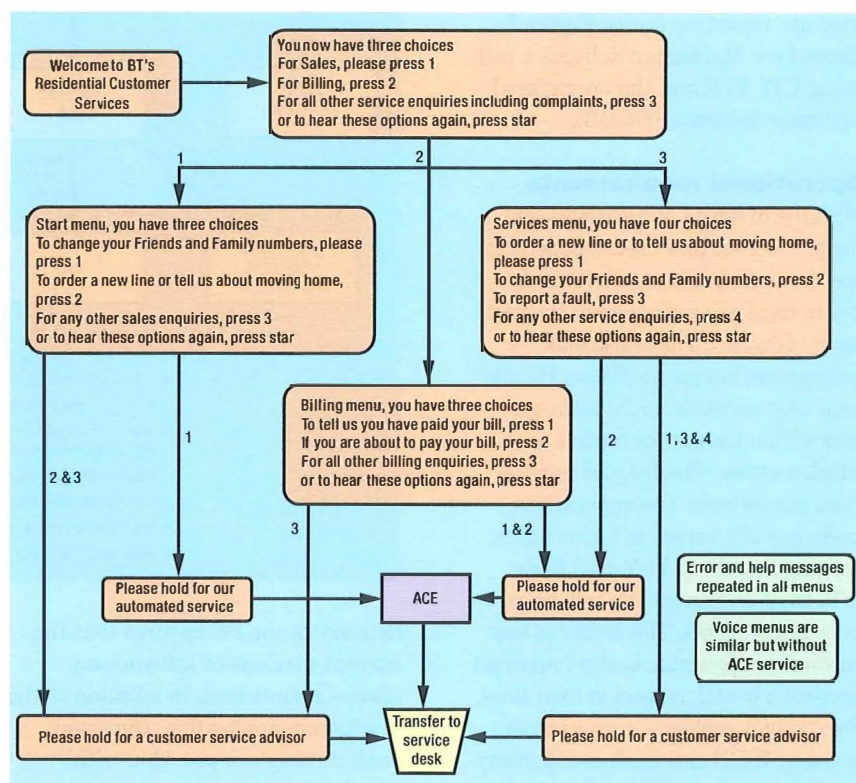


Figure 6—150 IVR call steering schematic summary

After Sales with around 250 people in nine locations.

The customer segments supported by these centres range from single-line business customers through small/medium enterprises up to major customers. The customer base is some 1.2 million customers with some 4 million accounts. The Fault Management and After Sales units are predominantly inbound with Customer Relations and Simplex Management involving outbound calling.

Network and systems

Compared with the residential area, the volume of calls handled are far lower but calls are much more complex. This is being fuelled by the growth in data products and services. As a result, the BT network now carries more data traffic than voice calls.

Calls are presented to advisors by customers calling 154 or via 0800 or dedicated numbers, all of which are all delivered to a network of Nortel Meridian switches: the CSC network.

The CSC network then uses ACD functionality to deliver the oldest call in the network to the longest waiting agent. The CSC network also seeks to deliver calls from customers to their 'home' fault management centre; that is, if a call is identified as being generated from East Anglia, the network attempts to deliver the call to Ilford FMC first.

With the introduction of Marksman to Business Customer Service, calls are still delivered by the Meridian CSC network, but then Genesys CTI takes control of the call and the routing strategy for it. If a customer has an outstanding order, fault or issue with BT, CTI firstly tries to route the call to the advisor who originally dealt with the call. If that agent is not available, then it tries the advisor's team, then the remainder of the centre before looking for the next free advisor in the country. Planned enhancements include the introduction of intelligent call routing (ICR) enabling sophisticated call plans and the use of IVR to identify customers and the numbers

Figure 7—Marksman overview

they are reporting faulty. Figure 7 shows how Marksman delivers a call using CTI, IVR and the operational customer database (OCDB).

Operational requirements

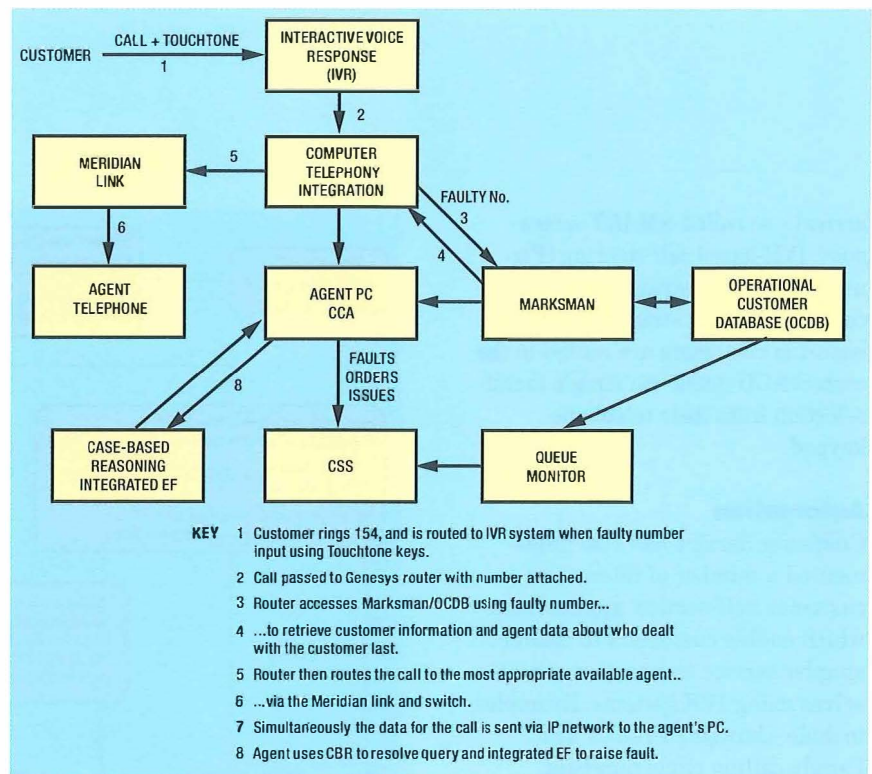
With the diversity of customers and the product set that needs supporting, speed of information retrieval is the key to resolving a problem at the first point of contact. The Marksman programme has put in place a PC and local area network (LAN) infrastructure within the service centres upon which a series of high-speed applications can be built. The applications make use of a variety of technologies, ranging from bespoke Visual Basic (VB) front-end screens to web-based information pages. The ability to test customers' apparatus is also important particularly with respect to their lines. For complex problems such as those posed by ISDN and Business Highway, Marksman uses case-based reasoning (CBR) to improve the probability of a correctly diagnosed fault resolution. Integrated with the CBR is an ISDN line test system (IFETS) which checks the customer's line end-to-end. PSTN lines can also be tested via the Marksman application. In future, line tests will be performed automatically using IVR collected data while the advisor is confirming the customer details. Further benefits this delivers are:

- More problems resolved at the first point of contact without increasing call handling times.
- The application is flexible enabling the introduction of support for new and bundled product sets
- The management information statistics (MIS) allow management of service levels and advisor performance on a real-time basis.

Figure 8 shows the components of the Marksman solution.

Marksman technical overview

The agent workstations are standard desktop PCs. Automated software



delivery to the PC ensures that the current versions of software are always maintained. In addition to the Marksman application, the agent workstations are provided with standard office automation packages and e-mail. The function of the back-end systems is detailed below:

The customer handling intermediate server (CHIS) and common handling of event service system (CHESS) are mid-tier servers providing an interface between the OCDB and CSS. The mid-tier services use event notification to ensure that fault, order and issue data on CSS is updated on OCDB.

Genesys T server is made up of several different components, namely router, stat server, call concentrator, and database server. Together these components provide the heart of the CTI call delivery and drive the intranet-based statistical information on advisor performance.

The five case-based reasoning (CBR) servers provide the interface between advisors and the main CBR case base server. The five interface machines provide an equal load sharing and back-up capability for advisors logging onto case bases for resolving customer problems.

The wall-board server drives the tri-colour wall boards in the centres which are used to inform advisors of the current status regarding custom-

ers waiting or advisors waiting for a call. They also inform management of advisor status; that is, number of positions manned, number of advisors either on a call, not ready or waiting.

The script repository is used to hold scripts which automate repetitive tasks that would otherwise be time-consuming and tedious.

Call Centre Futures

One thing that can be said with certainty is that the information revolution taking place in society today will have a dramatic impact upon call centres.

BT's call centres have evolved in an age where the telephone is king, it forming the principal real time access method into the business. This philosophy is extended into both front- and back-office support systems, which today are organised around customers' telephone numbers.

The deployment and integration of desktop computing and CTI into call centres has brought about significant improvement in efficiency and the customer experience. Simultaneous voice and customer data delivery at the time of call arrival and the ability to route calls depending upon customer service advisor skills levels, are examples in use today.

This incremental approach to improving call centre infrastructure

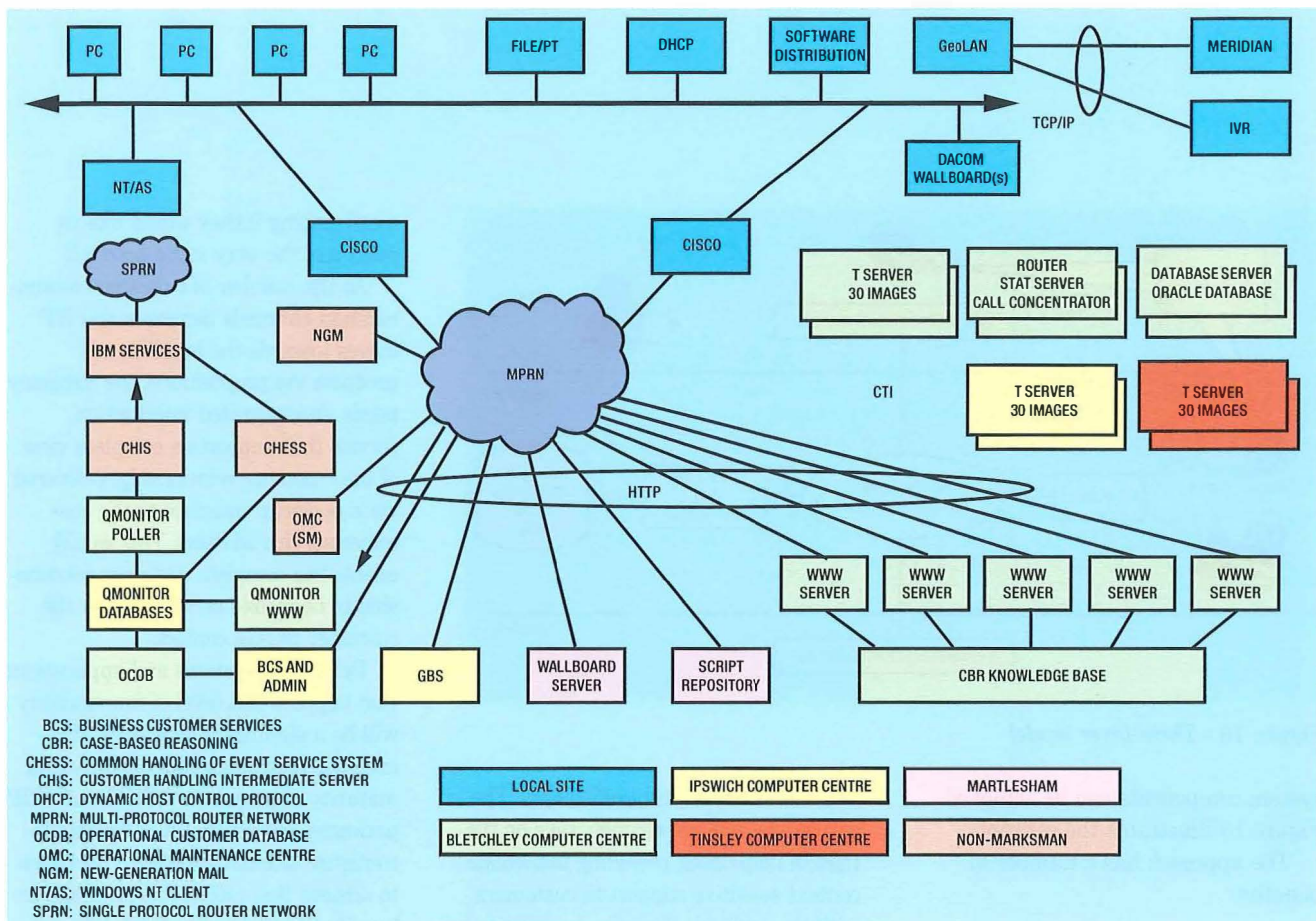


Figure 8 – Components of the Marksman solution

has kept BT in a competitive position, but is it sufficiently radical to leapfrog the competition as BT confronts the information revolution, and will it deliver the efficiencies and capabilities needed to run a world-class operation?

Site rationalisation

After extensive research, both in the UK and abroad, plans are being developed to consolidate the existing call centres. Modular centres, which are sized for the optimum number of workstations will be created to make most-effective use of accommodation and management skills. This will lead to fewer, but larger centres. These will be provided with a common systems hardware infrastructure, offering high resilience and flexibility.

Customer touch points

Today we are witnessing a massive growth in the communication channels that are being made available to customers; the Internet, mobile communications and Internet-enabled TV are all areas experiencing fast customer take up. Figure 9 illustrates the communication channels which will be available to customers.

This spectrum of communication channels will grow, as we come to terms with the transformation taking place in communications, moving from voice to a data centric environment.

How will this change impact on BT's call centres?

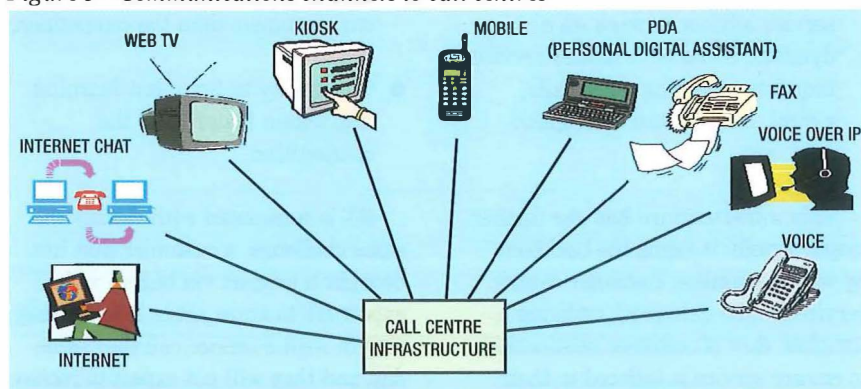
Customers increasingly expect to be able to do business via the communication channel of their choice, with the 'experience' being consistent and content rich.

To satisfy this customer expectation requires a new approach to the way call centre systems are designed and delivered. Historically, these have

been developed largely in isolation to satisfy the specific needs of each channel; for example, 150, 192. This has resulted in a very close fit with business requirement and world-class performance but at the expense of wide flexibility. A key constraint arising from this 'vertical' approach is to prevent the total call centre resource being regarded as a single resource 'pool' at times of peak customer demand and makes organisational change difficult.

By adopting a 'horizontal' approach to the systems architecture of call centres, both new channel integration and reuse of common

Figure 9 – Communications channels to call centres



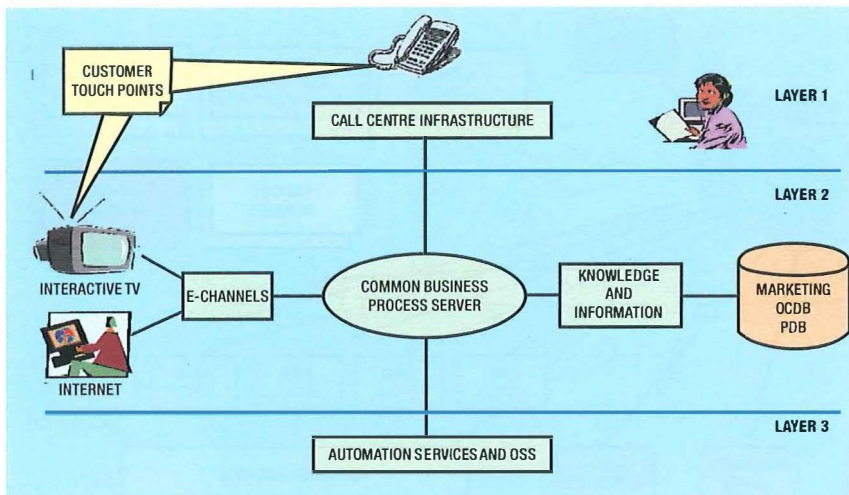


Figure 10—Three-layer model

system components can be achieved. Figure 10 illustrates the concept.

The approach has a number of benefits:

- For the construction of a new service channel, for example, Internet-enabled TV, the re-use of common business services in the layer 2 mid-tier enables both the channel and the existing call centre infrastructure to support this with minimal new development. Hence lower costs and faster time to market result when compared with the traditional approach.
- Dynamic integration at the desktop, of the applications and information required for an individual advisor's skill set increases the flexibility of the call centres and the advisor resource pool.
- Provision of integrated workflow capabilities enables a customer service advisor to work on a dynamic blend of customer service requests including voice calls, e-mail, fax, written correspondence, etc.

This infrastructure has one further major benefit: it forms the backbone by which seamless customer-centric service can be delivered, utilising a complete view of customer relationships to ensure service is tailored to their

individual needs and preferences. The call centres will increasingly take on the roles of help desks, providing immediate context sensitive support to customers, with the majority of service transactions handled in a 'zero-touch' manner. In other words, service will be delivered to the customer without any person being involved except perhaps a customer engineer. In this environment, if the call centre is contacted, it will be increasingly in a secondary 'assistance' role rather than a primary 'transaction' role.

Transactions to relationships

In a business environment, where BT's core network is no longer a competitive advantage and where it is possible to build a new telco from scratch in under five months, how does BT compete?

Jack Welch, CEO of General Electric, made the following statement, 'We have only two sources of competitive advantage:

- The ability to learn more about our customers than the competition,
- The ability to turn that learning into action faster than the competition.'

BT is confronted with exactly the same challenge, a customer that has brought a product via bt.com will expect BT to know about it when they ring in with a service call that same day, and they will not expect to receive

a call asking if they would like to purchase the very same product!

As the number of customer communication channels increases and BT moves towards the bundling of products via propositions, the company needs an integrated information service that supports a complete view of the customer relationship, delivered via a powerful interface to the customer service advisors. This would enable the complete customer relationship to be visible no matter how the customer makes contact.

Delivering systems and applications that support this level of functionality will be a significant challenge. However, the increasing functionality and maturity of proprietary CRM and ERP packages, means that it is feasible to configure standard software packages to achieve this rather than embark on bespoke development. Availability of 'out-of-the-box' solutions will be essential to support the evident rate of change in the marketplace.

Several initiatives are underway looking at how BT can satisfy this requirement, with vendors of both front- and back-office applications being assessed.

In summary, BT's Call Centres will change dramatically over the next few years as customer service, BT's true competitive advantage, is optimised.

Biographies



David Duxbury
Customer Service,
BT UK

David Duxbury joined the British Post Office in 1961 as an apprentice in the Leeds Telephone Area. He was later involved with exchange design in the North East Region before moving into Engineering Management Services within headquarters. He later became involved in the setting up of Account Management and Technical Support before moving to Newcastle where, with the formation of Districts, he

became Deputy District General Manager for the North East. With Project Sovereign he moved south and joined the Personal Communications Marketing Division and went on to become the Director of Retail, Consumer Products, and Cable Television Services in the old Consumer Division. On 1 October 1998 he took on the role of General Manager Call Centre Management in Customer Service. David is a Chartered Engineer and a member of the IEE.



Rob Backhouse
Customer Service
BT UK

Rob Backhouse joined BT in 1972 as an apprentice. He progressed

through roles as linesman, PBX and telephone exchange engineer, PBX and network installer, project coordinator, and project manager. He was promoted to manager in 1988 in Colchester, responsible for East Anglia Districts internal PBX and ACD systems. He then joined a national group performing a similar role. He led a team which implemented Meridian 1 ACD systems in the customer service call centres throughout the Home Counties. In 1993, he moved his job to Personal Communications division in London and was soon running the project to replace the Meridian ACDs with DMS-100. He has represented BT at an international conference in Florida, USA, which involved presenting a paper describing the use of the DMS-100 facilities most efficiently in supporting networked ACD solutions. He has also assessed the suitability and specified call centre solutions for BT joint ventures in Italy and Spain.

Mike Head
Customer Service,
BT UK

Mike Head joined the General Post Office in 1966 as an apprentice in London working on repair and

calibration of equipment used in repeater stations and transmission centres. Over the subsequent years, he worked on a variety of products, projects and programmes varying from loudspeaking telephone and noise-cancelling microphone design through to systems implementation. He spent four years working with Government National Accounts as a presales engineer selling computing solutions to the army and navy before moving to BT Laboratories in Martlesham to run a computer support and in-life software development group. He currently works in BT UK Customer Service as the Marksman programme manager delivering a culture change programme to Business Fault Management and After Sales.



Graham Lloyd
Customer Service,
BT UK

Graham Lloyd joined the then British Post Office in 1972 and progressed through the three year engineering apprenticeship scheme, taking up a Technical Officer post in business systems planning at Cambridge. He completed a number of major projects including the design of a replacement ISDX switch for Addenbrookes Hospital. After promotion to a managerial grade in 1988, he was given responsibility for the voice communications needs for East Anglia's 8000 workforce. Key achievements included upgrading the District voice network and designing a portfolio handbook with product and service descriptions for 800 managers. In 1991, after further promotion, he played an active role within the London Front Office project, being responsible for the design of voice systems technology utilised by Customer Service. He acted as a single interface with Nortel and designed the DMS100 network to support a single virtual queue of 950

advisors. The configuration was the first application of its kind in Europe, and the largest single node ACD in the world. In 1996, he was selected to set up a voice systems team to deliver infrastructure for the telemarketing expansion project, which subsequently achieved the fastest ramp up of call centre capacity in the history of BT—3300 advisor positions in 14 months. He is currently working in the Business Development and Operations team working upon solutions design approval process and influencing the technical direction of BT Customer Service's call centres.



John Pilkington
Customer Service,
BT UK

John Pilkington graduated from Leicester Polytechnic with a joint honours degree in Physics and Computing. He later went on to gain a Diploma in Management Studies from Leeds University. He joined BT Laboratories in 1987 and spent several years working on international standards where he led various CCITT sub-committees developing the X.400 messaging and X.500 directories protocols. In 1992, he moved to BT's Operator Services Directorate, whose headquarters are in Leeds. In 1997, he was asked to create a new team responsible for the requirements capture, design and integration of a platform that would replace the DA service before the millennium. With the DA programme successfully deployed on time, John's role was recently expanded. He currently leads a multi-discipline team responsible for the end-to-end management of all hardware and software releases onto the DA, Operator Assistance and 999 platforms.

Enquiries about this article should be directed to Lyn Windmill at lyn.windmill@bt.com.

Jerry Bowskill, Craig Wisneski, Barry Crabtree, Peter Platt, and Stuart Soltysiak

Telempathy

This article introduces telempathy, a complementary form of telecommunications in which the presence and activity of remote people is presented using ambient displays. This paper focuses on the 'awareness' of people's activities as a precursor to more conventional forms of telecommunications. Initial research is described in which a framework for software agents is used to gather real-time availability information, displayed using visual and auditory interfaces on networked wearable computers. The potential market for telempathy services is discussed. It is believed that such services will promote telephony usage in a variety of commercial and residential applications.

Introduction

Our senses are constantly giving us peripheral information that we use to focus our attention or monitor the 'ambience' of our environment. For example, at work we use cues like the sounds of doors opening, typing on keyboards or lights in the office to help subconsciously understand the activities of other people. We also use ambient cues, such as the posture we adopt in our chair, to communicate our own activity and availability. Why then does current technology not exploit the human ability to process subtle changes in light, sound, even smell and temperature?

Just consider modern telecommunications. Today people have a variety of 'explicit' communication channels available. While this enables us to become available more of the time, a disadvantage is that it also increases the likelihood that we telephone 'at the wrong time' or have to use a succession of communications types to contact people. E-mail, for example, has become the default method of communications for many even though a 'synchronous' conversation via a telephone call conveys more in a shorter time, if of course the person is there to pick up the telephone. It is perhaps more likely that the user will be connected to voice mail. Keeping people 'in touch' demands more than current telecommunications provides.

The authors are exploring alternative communications interfaces, in particular those that promote *telempathy* within distributed communities. As reported by Preece¹ the term *empathy* has at least three different meanings. It can mean *knowing* what another person is feeling, *feeling* what another person is feeling, or *responding* compassionately to another person's distress. The term 'telempathy' is used here to refer to a system that allows a person

to *know* what a *remote* person is feeling, their activity, availability or other 'contextual' cues via a peripheral or ambient interface. Awareness is essential in collaboration to coordinate the activities in a shared workspace. Technological devices such as remote cursors, multiple scroll bars, audio cues, and low-frame-rate video have been proposed to support the awareness of remote collaborators' activities. Dourish and Bly's Portholes project² is an example of an awareness support system using low-resolution low-frame-rate video.

There are a number of awareness cues that allow telecommunications, particularly within distributed teams, to be made more efficient. For example, who is present or on-line? Where are they and how can I reach them? Are they available for synchronous or asynchronous communications? Thus telempathy seeks to provide a subtle, or ambient, sense of presence rather than a direct communication medium such as synchronous telephony or asynchronous messaging. In a physical environment, people use their ambient awareness of others as a precursor to formal communications. The majority of people quickly become astute at personal communications based on an interpretation of ambient, as well as direct, visual and auditory cues. Current telecommunications technologies offer little support for astute personal telecommunications. Methods of telempathy are required.

Tangible Bits

Before describing our research we shall consider the previous work which stems from the MIT Media Laboratory. Ambient displays are part of the broader 'Tangible Bits' vision of Wisneski *et al.*³. Notions of ambient media were, in part, inspired by the Fields and Thresholds work of

Current BT research into telempathy is focused on supporting collaborative working within distributed project teams.

Dunne and Raby⁴ and the Live Wire of Jeremijenko⁵. These projects explored a theme of peripheral awareness of external activity, especially of activity attributable to people. The AROMA project⁶ further investigates ideas of peripheral awareness.

Several experimental tangible interfaces have been demonstrated at MIT that blur the boundary between physical and digital worlds. The focus for much of this research has been the creation of tangible interfaces onto intangible data. However, there are examples of ambient displays being used to convey remote presence albeit within the same physical building. The ambientROOM is based on the Personal Harbor™ product from Steelcase Corp., a 6 ft × 8 ft enclosed mini-office (Figure 1). The ambientROOM surrounds the user within an augmented environment—‘putting the user inside the computer’—by providing subtle augmentations to activities conducted within the room. People are made to feel connected to others through a number of ambient displays:

- **Active wallpaper** Electric field sensors are used to measure the level of human activity in a work area. They are set up to sense the number of people in an area, and how often people generally come and go. This activity is represented by a pattern of illuminated patches projected onto an inner wall of the ambientROOM. When activity is low, the movement of these spots is minimal, but as activity levels increase, so does the motion of the spots, providing a visual display of remote activity.
- **Ambient sound** A ‘soundscape’ arises out of activity on the digital whiteboard in a remote workspace. When the board is in use, the sounds of the dry-erase pens rubbing against the board are transmitted into the ROOM in a low-volume subtle way. This gives the inhabitant of the ROOM some awareness of activity in the central workspace.
- **Water ripples** One display in the ambientROOM allows the user to have some awareness of remote


activity, which is transferred to the motion of a solenoid in a shallow water tank. A lamp reflecting off of the water then produces rippling shadows on the ceiling.

On first impressions it may be difficult to imagine working within an ambientROOM, and long-term studies are required before effective displays are revealed, but consider the potential impact on telecommunications. The whole room becomes a telepresence interface onto telephone and computer services, without the intrusion of permanent video connections or the isolation of current-day ‘click to talk’ teleworking solutions.

Ambient Telepresence Through Wearables

Current BT research into telempathy is focused on supporting collaborative working within distributed project teams. We believe that the most appropriate platform for ambient awareness interfaces are personal devices such as ‘wearables’ and mobile telephones. While some potential forms of telempathy data are public, temperature for example, many forms of data relating to collaborative working are personal. For example, the real-time availability of a person is privileged information which should be released only to trusted parties, and then only within limits. Thus interfaces need to be personal and in some cases private. An installation such as ambientROOM may be appropriate for personal work cells; however, a personal device has the advantage of being able to deliver personal information within a large shared office and tailored to individual preferences. It is easy to imagine that if everyone in an office had ambient audio interfaces these would soon distract and annoy. Personal devices also allow access to context specific information when mobile; for example, location-based availability information as discussed in the

Figure 1—Ambient media displays and controls






Figure 2—Wearable computer

following section. Lars Holmquist has similarly embraced the use of *interpersonal awareness devices* (IPAD) as communications interfaces⁷. His definition of IPADs is that they are 'small devices, designed to be carried at all times...that give an indication—aural, visual, tactile, etc.—when another IPAD is in the vicinity'. The *LoveGetty* is an example of a successful commercial IPAD, a device the size of a key fob that can be programmed with user interests. When another *LoveGetty* comes in range on the device's infrared link the interests are compared and the device flashes when the people are a good match. The *Hummingbird* has been developed by Holmquist to explore the issues of personal awareness devices and can, for example, be used to give aural and visual indications of when other group members are in the vicinity. Whenever two or more *Hummingbirds* are in the vicinity of each other, they chirp, and the identity of other *Hummingbirds* in the vicinity is shown on a display. Thus a *Hummingbird* gives its user continuous information of which other *Hummingbird*-using group members are close, whether in the place they usually meet, or somewhere else. 'Close' in this case is approximately when users are in the same building when indoors, or within 'shouting distance' when outdoors.

The first commercial personal-availability tools are likely to be variants of buddy lists delivered to smart telephones on the current global system for mobile communications (GSM) network. Buddy lists such as ICQ⁸ and PAL⁹ allow groups of computer users to stay in touch. Users register their details with the service provider who maintains a list of all connected clients. A user is able

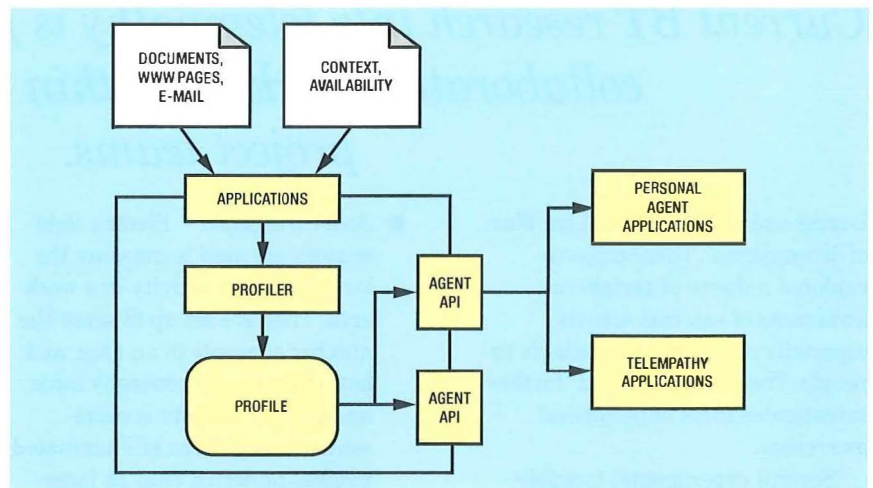


Figure 3—The personal agent framework and its role in telempathy

to select the names of other registered users whose availability they wish to be aware of. Buddy-list tools also support direct communications by providing a wrapper to e-mail, text chat and even Internet protocol (IP) telephony applications. Ericsson and other companies have launched smart telephones and are developing cellular JavaPhones. The handset has a liquid crystal display (LCD) screen and will be able to download Java-based applications connected to services in the network. Munch-Ellingsen *et al.*¹⁰ from Telenor are developing a buddy list or *people's browser* to support mobile virtual co-location in groups.

Within BT, the authors have been using wearable computers as a platform for telempathy research. These have the specific advantages of supporting both visual and auditory displays, significant local processing power and a range of network interfaces. Our current wearable system is shown in Figure 2 and consists of a body-worn Toshiba Libretto computer, wireless local area network (LAN) and GSM network interface, spectacle-mounted display, ear piece and corded 'twiddler' keyboard.

A particular advantage of wearable computers is their ability to collect local contextual information. The wearable could be collecting data such as local light, sound, temperature and movement levels and convert that into useful higher-level information to determine the user's context: whether they are moving or still, talking or in discussion, inside or outside. Making some or all of

these cues available to remote colleagues would allow more natural communications. In the office, we know if we can talk to or interrupt a colleague just by looking over. In the remote environment we do not have the cues, and interrupt at inappropriate times. The current research of the authors is focused on providing experimental audio and visual interfaces onto availability, or other forms of context data, to support collaborative working in distributed teams. Raw data is being collected via the BT personal agent framework, described below.

Personal agent framework

The personal agent framework (PAF)¹¹, developed at BT Adastral Park, is a unified environment in which many personal agents can be integrated. The core part of the personal agent framework proposes a secure and private registered user environment, which holds dynamic user-profile information that can be shared among agents in the framework. A fundamental principle of the PAF is the provision of personalised services with minimal burden on the user. Within the PAF, a user has a personal profile, which describes his/her interests and preferences. This profile is then available for a number of agent applications to find and present information to the user, as shown in Figure 3.

Some of the personal agent applications we currently operate are concerned with locating and providing the right information to the right people at the right time. Thus having a user profile that is rich enough to




Figure 4—Availability information displays—the radar scope and the graphical display

contain not only the interest areas, but also contextual information such as when, where and in what situation the various interests are relevant can greatly enhance the effectiveness of this type of system. Personal profiles can also be available for use by telepathy applications—to bring to an individual's attention those topics of interest or importance to another person according to the presence of that physically remote person.

Within the PAF, the availability agent (based on the APIA work of Duke¹²) enables an individual's profile to be updated to reflect the current status of his/her availability. Availability takes many forms, from a simple presence to more complex details of the tasks and loading of the individual. The availability agent attempts to provide information on a number of such features. To determine the simple presence of an individual, the availability agent monitors the keyboard and mouse activity, to give a crude availability 'score'. To provide more detailed information, the availability agent looks at the applications running on the machine.

Visual availability awareness

We have discussed the advantages of making a remote user's context available to his/her colleagues so that interruptions/discussions could be timed to mutually beneficial times.

Even the simple availability awareness of when the user last worked on his/her machine is beneficial to remote interaction. The availability information gathered within the personal agent framework can be provided explicitly or implicitly. Explicit provision can be seen in Figure 4, where the availability score is mapped in real-time onto the desktop. The two applets in the figure show two aspects of the availability score—the radar scope is a snapshot of the individuals' availability, whereas the graph shows an 'historical' view of availability over a period of time. The radar scope allows for an individual to have associated with them a personalised image for display (if none is available then a blank profile is used), as well as a personalised audio signature, to enable the person monitoring to be aware of who is available. This feature merges the explicit provision of availability information into implicit provision.

Implicit provision of the availability takes the idea of personalised visual and audio cues a step further into integration with the environment. Thus presence and activity can be conveyed using ambient visual and auditory stimuli, to enable telepathy to be based on real-time, accurate information. Even using the radar-scope and availability graph for a short

time influenced the way we interacted with colleagues. When someone suddenly comes on-line we have the best chance of contact, and especially when working with remote colleagues (across the Atlantic) we have a much better feel for their availability.

Audio availability awareness

Our efforts have not just concentrated on visual displays of people's availability, we have also been investigating a wider use of auditory display in such systems. The use of audio is a natural approach to delivering ambient cues as much of our peripheral awareness is based on our auditory environment. An experienced car driver knows when it is time to change gear based on the auditory feedback from the engine without having to consciously attend to it. This has the benefit of leaving the visual field free to concentrate on the current task. The human ability to monitor the auditory environment as a 'background' task is well documented; the 'cocktail party effect'¹³ and Gibson¹⁴ show that people can be aware of auditory stimuli without it severely affecting their conscious attention.

Organising auditory streams

To effectively convey information through a graphical display, a lot of attention is paid to the layout and appearance of that information. In the

same way, the elements of an auditory display that may consist of a number of different, simultaneous sounds need to be organised in some way to prevent a meaningless cacophony. We have been considering two ways of doing this: the use of auditory scenes, and the use of musical rules.

Auditory scenes are created by associating the availability parameters with particular sounds that can be mixed together as part of an auditory virtual environment. For example, Audio Aura¹⁵ uses a beach metaphor with the number of e-mails waiting to be read associated with the sound of seagulls, and general group activity associated with the frequency of surf hitting the beach. We have used a similar metaphor where the availability of people is associated with a herd of cows in an auditory scene with birds singing and other natural noises associated with other states. As more people become available the herd of mooing cows becomes larger. In this instance it is not possible to detect exactly who is available at any time, the aim is to provide a general awareness of the number of people currently available.

Our other approach to the problem of an organising structure has been to use the rules of harmony, rhythm and melody to bring together sounds in a coherent way. By associating the data we want to convey with musical sounds and passages that fit together and make musical sense, we have created musical soundscapes in which changes in individual elements can be detected. To prevent the presentation of different passages from becoming repetitive, we are using a generative music composition tool called *Koon*¹⁶ to generate the music in real time. The system works by associating a person's availability rating with a motif played on a particular instrument. As the person becomes more or less available, the rhythm of the motif is varied. When they have been unavailable for some time, their motif is removed completely. In this way it is possible to determine who is

currently available and to what degree from the collection of motifs making up the whole soundscape. This is similar to Wagner's use of leitmotif and Prokofiev's identification of animals in 'Peter and the Wolf'. We are also using music to 'build up' to scheduled events by raising the user's expectation that an event is about to occur, such as a meeting or appointment.

Musical displays

The use of music has some advantages over auditory scenes. Due to its more abstract nature, the musical approach is more scaleable as it does not require the designer to find sounds that fit the metaphor in a meaningful way. There is also less danger of a musical environment interfering with the real environment, be it visual or auditory; one can imagine real-world seagulls being confused with waiting e-mail.

It is also possible to imagine systems where the use of music is taken further than direct associations, and is manipulated to convey emotions, moods and other states that are hard to display graphically. This is similar to the use of scores in feature films and advertising where events in the past, present and future are brought together by music to enrich the understanding of the scene¹⁷. Also, we must not forget the power of music to affect the way we perform tasks. Research into the use of muzak in shopping centres and restaurants suggests we can be strongly influenced by the style and pace of the ambient music we hear^{18,19}.

Business Potential

So why are we researching telepathy and how could it impact telecommunications? There are two compelling examples of how people's lives could be positively influenced. Firstly, consider business. Telecommunications has revolutionised business, and in various sectors has enabled distributed team working and even virtual organisations to flourish.

Within many companies teleworking is actively promoted, and few can resist working some, if not most, of their time outside the office. However, there are problems. In particular, teleworkers can feel isolated from other team members. While communications via a variety of methods are available they tend not to support the informal repartee and exchange of ideas that occur when people share a physical workspace. Now imagine that each member of a distributed team had an awareness of each other; for example, they could sense the availability of others using the same background cues as if they shared their physical space. This would actively promote communications within the group with all the positive effects on team cohesiveness that this implies. Perhaps the biggest potential benefit is to encourage people to speak rather than using e-mail excessively. Because telephoning is 'hit and miss', e-mail tends to be used as the first option even when communicating with others in the same office. In most cases, e-mail can be avoided if you know when to telephone, and telephone conversations are certainly richer in information.

Secondly, consider the potential impact on residential services. Routine social interaction is important to the maintenance of relationships. This becomes increasingly difficult to maintain as friends and family move away from each other. People have a need to feel connected to others, especially loved ones, and we believe ambient displays can aid in this connection. A recent pilot survey of BT customers, reported by Anderson *et al.*²⁰, reveals much about how we communicate at home, and proposes taxonomy of call types. Findings also indicate several opportunities for telepathy. People are creatures of habit, and their calling patterns are no exception. For example, all survey participants reported making *duty calls* to family members. These 'obligation' calls are often made on a weekly basis and tend to be made at regular times,

giving people an ambient awareness of other remote people or events will complement existing telecommunications

usually during cheap rate periods and especially on Saturday mornings. Extra duty calls are also made on birthdays and anniversaries. It is easy to imagine that if members of a family could receive a constant stream of empathic bits this would stimulate more frequent and timely communications. An awareness of others may even stimulate men to make *maintenance calls*—that maintain a friendship—as they frequently wait for a reason to telephone, rather than women who are more likely to phone ‘for a chat’ (in fact women are responsible for twice as many calls as men). Other call patterns may also benefit from availability cues. Several participants mentioned *shared calls* in which they pass the telephone to another family member during the course of a call. This occurred most often when the call was from another member of the family, particularly in the case of grandparents talking to grandchildren, and more rarely in three-way conversations with another member of the family ‘listening in’ and making comments on another telephone. If members of a family, or group of friends, were aware of calls between members then this could promote shared calls, with other people joining conference calls in the same way that members of a group tend to be drawn in to a conversation in a shared physical space. *Pseudo maintenance calls* were also reported and are made to friends in order to keep in contact and maintain the friendship. However, they are made deliberately when the caller knows that the call target is very unlikely to be there and the object of the call is simply to leave a message. Thus remote-availability information could even be used to prompt people to call when they know the recipient is unavailable. Conversely, availability awareness could also reduce communications. Many participants reported that if they make a call to a target who is unavailable they often make a call to a second, third or even fourth friend until they find someone

to whom they can talk (a so-called *domino call*). This behaviour was much more common among female participants than males. However, we believe that an awareness of others will, by and large, promote communications not restrict them. This hypothesis is to be quantified during future research; however, it has already been validated on an extended internal trial. Four members of a project team distributed on a single site have used the web-based availability radar and auditory interface for three months. During this time participants made increasing use of the telephone (by as much as 300%) and physical ‘chance encounters’ based on the sense of presence of others. E-mail use among the group decreased. While this evidence is inconclusive it does give grounds to perform a larger user study, in particular to assess the impact of different interface devices.

Discussion

There are many examples of telecommunications being invasive or inappropriate. We believe that giving people an ambient awareness of other remote people or events will complement existing telecommunications. For example, a recent trial demonstrated the possibility of gathering life-style data on elderly people living in their own homes. The system is able to alert care staff when uncharacteristic behaviour is recognised, perhaps a lack of activity in the kitchen first thing in the morning. While the benefits of such a ‘reactive’ system are undeniable, consider for a moment if friends and family were given a continuous ambient awareness of their relative’s well-being rather than an ‘alert on emergency’. In an ideal case, such ‘telepresence’, via a trickle of empathy bits, would perhaps prompt social telecommunications and physical calling. Of course it is equally possible to predict many of the problems with such interfaces.

When we look at the idea of telepathy and the need to project

remotely the ambience of a group or availability of a person, a major consideration is that of personal or group privacy. Because we are using personal information to project awareness we need to ensure that this is used in a controlled and acceptable manner. In the personal agent framework all the information is protected and it is up to each individual to decide who else should be able to see what parts of a profile. For availability information we may allow users to display different granularities of information depending on who is asking. For example, I may be willing to allow my close colleagues to know what my availability is down to the nearest minute, whereas for a total stranger I may only wish them to know that I am/am not in the office that day. Addressing privacy and security issues early on will ensure that any solution generated will be more widely acceptable.

In addition to privacy and security, the style and modality of interface is a critical factor in the success of ambient displays of any kind. The choice of modality for the background media should be considered with the person’s foreground task in mind. For example, if a person is in his/her office performing a visually intensive foreground task—say, writing software—a visual ambient display might not be as effective as an auditory display. Similarly, when a person is performing an intensive auditory task, like talking on the telephone, ambient information might be better presented through non-auditory ways. However, if the modality and spatial configuration of a particular ambient display is substantially altered during the course of use, people may become confused about the display mapping being employed. Mynatt *et al.*¹⁵ observed that the type of auditory display also impacts the cognitive load of the user. For example, vocal feedback is familiar but usually demands more attention than a non-speech sound and musical ‘earcons’ demand less attention providing their meaning has been learnt, which is not

always easy. Familiarity with the interface can also prove distracting. For instance, Wisneski *et al.*³ report that the sampled looped sounds of water in the ambientROOM became annoying after months of repetition. In particular, sounds that are looped in some kind of discernible way tend to become annoying very quickly. Also, many media elements going on at once can cause an 'information overload'. Display style and modality can also impact user privacy. For example, researchers at MIT experimented with placing an infrared camera in a physical space and projecting its out-of-focus image on the wall of the ambientROOM. It is reported that many people thought the camera-based display felt intrusive, as the mapping was too literal and privacy had not been respected.

Conclusions

The authors have presented telepathy as a complementary form of telecommunications, which we believe will be the basis of future services. While the commercial implications of telepathy are largely unknown, Internet buddy lists illustrate how computer users are starting to keep in touch. We predict such services will spread rapidly to mobile devices, initially delivered to conventional audiovisual displays. The authors have developed a number of visual and auditory displays that are being used on both wearable and desktop computers to inform team members of their presence and availability. Ambient displays are particularly relevant to viewing availability information where privacy is a major concern. For example, using our Koan-based audio-availability awareness display it is possible to hear a representation of real-time availability information, giving a sense of presence of a distributed community of people without revealing individual identities. The evaluation of such displays represents a significant task. Mynatt *et al.*¹⁵ performed a small-scale

evaluation of audio aura and concluded that, in general, sound was a good choice of display, although some found the meaning of sounds difficult to remember. We have yet to formally evaluate our current displays, although the visual displays are in continuous use within a distributed project team and regularly initiate telephone calls, physical meetings and even e-mails. This is promising, and we intend a formal evaluation by comparing the communications habits of people with and without access to a variety of availability displays. The BT research is focused on developing less intrusive display techniques for multi-modal wearable or mobile devices. In the future we would like to be able to display availability information in a modality and style appropriate to the user's task, location, and environmental constraints. This demands a high degree of 'contextual' reasoning and longer-term research.

Acknowledgements

The authors acknowledge all in the Tangible Media Group at the MIT Media Lab and members of the Emerging Technologies Laboratories at BT Adastral Park.

References

- 1 PREECE, JENNY. Empathic Communities: Reaching Out Across the Web. *Interactions*, March 1998, pp. 32–43.
- 2 DOURISH, P.; and BLY, S. Portholes: Supporting Awareness in a Distributed Work Group. *Proceedings of CHI '92*, pp. 541–547. ACM Press, 1992.
- 3 WISNESKI, CRAIG; ISHII, HIROSHI; DAHLEY, ANDREW; GORBET, MATT; BRAVE, SCOTT; ULLMER, BRYGG; and YARIN, PAUL. Ambient Displays: Turning Architectural Space into an Interface between People and Digital Information. Published in

Lecture Notes in Computer Science, First International Workshop on Cooperative Buildings, Vol. 1370, Springer, 1998.

- 4 DUNNE, A.; and RABY F. Fields and Thresholds. Presentation at Doors of Perception 2. 1992. <http://www.mediamatic.nl/Doors/Doors2/DunRab/DunRab-Doors2-E.html>
- 5 WEISER, M.; and BROWN, J. S. Designing Calm Technology. 1995. <http://www.fxpal.xerox.com/chi97/white-papers/Mark%20Weiser.html>
- 6 PEDERSEN, E.; and SOKOLER, T. AROMA: Abstract Representation Of Presence Supporting Mutual Awareness. *Proceedings of CHI '97*, pp: 51–58. ACM Press, March 1997.
- 7 HOLMQUIST, LARS ERIK. Supporting Group Collaboration with IPAD:s—Inter-Personal Awareness Devices. Handheld CSCW workshop, CSCW '98, 14 Nov., Seattle, USA. <http://www.teco.uni-karlsruhe.de/hcsw/sub/101.Holmquist/101.holmquist.html>
- 8 <http://www.icq.com>
- 9 <http://talk.excite.com/communities/excite/pal/home/>
- 10 MUNCH-ELLINGSEN, ARNE; MALM, PÅL; GRAV, JAN; and STENVOLD, LILLY ANN. Position paper for Handheld CSCW Workshop, CSCW '98, 14 Nov., Seattle, USA. <http://www.teco.uni-karlsruhe.de/hcsw/sub/118.Munch/att00003.html>
- 11 CRABTREE, I. B.; SOLTYSIAK, S. J.; and THINT, M. Adaptive Personal Agents. *Personal Technologies Journal* special issue on Personal Information Agents, Springer-Verlag, 1998.
- 12 DUKE, ALISTAIR. A People and Information Finder for Virtual Construction Project Teams. *Br. Telecommun Eng.*, Oct. 1998, **17**, p. 198.

- 13 CHERRY, E. C. Some Experiments on the Recognition of Speech with One and Two Ears. *Jnl. of the Acoustical Soc. Of America*, Issue 25, 1953, p. 975-979.
- 14 GIBSON, J. J. The Ecological Approach to Perception. Lawrence Erlbaum Associates, 1986.
- 15 MYNATT, ELIZABETH; BACK, MARIBETH; WANT, ROY; BEAR, MICHAEL; and ELLIS, JASON. Designing Audio Aura. CHI '98, pp. 566-573.
- 16 <http://www.sseyo.com>
- 17 BURT, G. The Art of Film Music. Northeastern University Press, 1994.
- 18 NORTH, A. C.; and HARGREAVES, D. J. The Musical Milieu: Studies of Listening in Everyday Life. *The Psychologist*, July 1997, p. 309-312.
- 19 ARENI, C. S.; and KIM, D. The Influence of Background Music on Shopping Behaviour. *Advances in Consumer Research*, 20, p. 336-340, 1993.
- 20 ANDERSON, BEN; MCWILLIAM, ANNABEL; LACOEHEE, HAZEL; CLUCAS, EILEEN; and GERSHUNY, JAY. Family Life in the Digital Home: Domestic Telecommunication at the end of the 20th Century. *BT Technol. J.*, 1999.


Biographies



Jerry Bowskill
Networks and
Information Services,
BT UK

Jerry Bowskill joined the Advanced Perception Unit within the Applied Research and Technology department at BT Laboratories, now known as BT Adastral Park, in 1996. Before joining BT he worked as a Research Officer within the Depart-

ment of Electrical Engineering at the University of Brighton. In 1990, he was awarded an honours degree in Microelectronics and Information Processing. After graduating, he started his research career at the University of Brighton as a research assistant working on machine vision systems. He was awarded a Ph.D. in 1994. During the past six years he has been researching augmented reality and has managed BT's 'wearables' project for the last two years.



Craig Wisneski
Presto Technologies

Craig Wisneski recently received a Masters degree from the MIT Media Laboratory under the guidance of Professor Hiroshi Ishii. While at the Media Lab, he worked in the Tangible Media Group on projects such as the ambientROOM, Ambient Displays and PingPongPlus. Craig received an undergraduate degree from MIT in Brain and Cognitive Sciences, and now works for Presto Technologies, a Media Lab spin-off, working on radio tagging systems.



Barry Crabtree
Networks and
Information Services,
BT UK

Barry Crabtree joined BT in 1980 after completing a Physics degree at Bath University. After a number of years in software development he moved into the field of artificial intelligence (AI) and helped begin BT's systems in automated diagnosis and workforce management. He moved on to work in distributed AI, optimisation problems and adaptive systems. He is currently a Technical Advisor in Advanced Communications Research (ACR), and has particular interest in intelligent agents and wearable computers.



Peter Platt
Networks and
Information Services,
BT UK

Peter Platt is a research designer in the Internet and Multimedia Applications Unit at BT Adastral Park, Martlesham. His two main areas of interest are on-line collaborative environments and wearable technology. He is particularly concerned with how to bridge the gap between being on-line and a full existence in the 'real' world. To this end, he is interested in bridging the gap between wearable technology and on-line presence. He also worked on the Mirror project and is a key member of the Forum team.



Stuart Soltysiak
Networks and
Information Services,
BT UK

Stuart Soltysiak obtained a B.Sc. in Command and Control, Communications and Information Systems from Cranfield Institute of Technology (1992), followed by an M.Sc. in Cognitive Science (1993) and Ph.D. in Computer Science from Essex University (1997). Since joining BT in 1996, Stuart has been working on personal software agents, including machine learning of user interests, and security and privacy issues related to personal agents. He is now leading an international collaborative project for agent-based information management systems. Alongside his current work with software agents, and artificial intelligence in general, Stuart retains research interests in cognition, particularly cognitive models of language processing.

Alwyn Lewis and Graham Cosier

Whither Video?—Between Sentiments and Signals

This is the fourth of five articles about the future role of video pictures in telecommunications. Previous articles¹⁻³ have outlined the history of television and videotelephony, reviewed human behaviour in communications and explored engineering and artistic conventions in the use of pictures. This article discusses the implications for new videotelephone, telepresence and 'virtual-world' communication systems.

A New Era

In the late 1960s, AT&T planned a new era for telecommunications, where telephones would be replaced by videotelephones. Those pioneering plans were wildly optimistic and videotelephony has proved paradoxical—successful in niche applications yet somehow barred from universal appeal.

With the benefit of hindsight, the authors are not surprised that videotelephony was not as successful as the introduction of the telephone. Technical and commercial limitations played a part, but more fundamental inhibitions lay elsewhere. This series of articles attempts to understand this paradox and what it suggests about future developments in human telecommunications.

Reprise

The first article¹ described the prodigious creativity of early television technology, that helped make television a plausible model for early videotelephony. The second article² posed the question 'Do minds meet more easily when voices are disembodied?' and described human perceptual abilities in sight and sound. It contrasted a human-centred view of videotelephony with the more traditional technology-centred view. The third article³ discussed pictorial culture in terms of the multi-modal interactions between physical optics and human psychology. It also explained why the pictorial culture of television, widely successful as entertainment, is a bad model for videotelephony.

Paradox in Perspective

Human culture has exploited the subtle power of pictures for tens of thousands of years, since the days of the cave painters (Figure 1). In the last hundred years or so, we have been able to make our pictures move⁴. Real-time moving picture technology has been developed within living memory. But all these different kinds of pictures, live and still, share one common characteristic. Someone in the role of artist, photographer, designer or director deliberately selected, created or composed each picture. Even a surveillance camera shows a consciously chosen viewpoint.

What's Novel?

Artists find self-portraiture difficult. Photographers take pictures of other people. Live self-photography is a new idea. We have little or no cultural adaptation to self-portraits we take to represent ourselves remotely, or to live pictures taken without our control by an autonomous machine. We should not be surprised that videotelephony still evokes mixed reactions from its users—it differs fundamentally from the conventional use of pictures and needs a new approach.

Figure 1—26 000-year-old cave-painting at Lascaux, France



Figure 2—Process hierarchy for telephony

Videotelephones have often been classed as a novel application (personal communication) for an existing technology (television). The authors think this classification is misleading. Videotelephony is better seen as an existing application (telephony) for a novel technology (live self-photography). The area with less knowledge is likely to have more novelty.

The Way Ahead

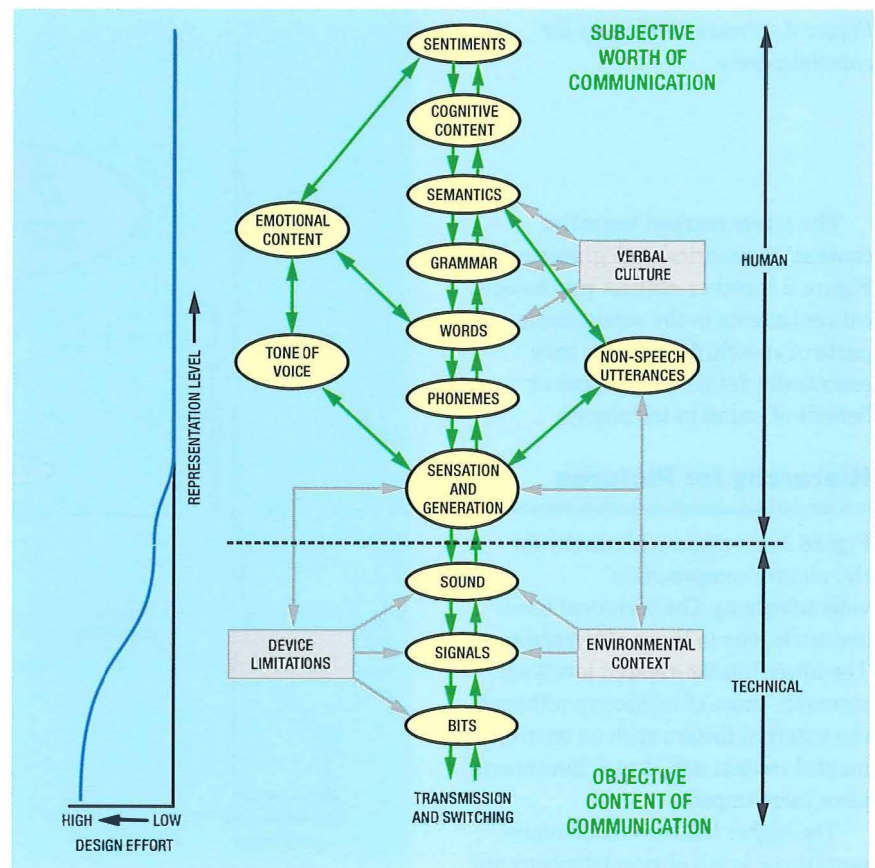
There is now good evidence⁵⁻⁷ to suggest that video pictures can enhance the effectiveness of human telecommunications, if used appropriately. The design and use of videotelephones should evolve so that users feel more confident about how others see their self-photography image.

Videotelephones can be more effective than simple telephones, but they are nowhere near as simple to get right. The significant problems are no longer technical but cognitive, in understanding the ergonomic aspects, social norms, cultural adaptations and psychological requirements of videotelephony. Anderson⁸ proposes an outlook for telecommunications that focuses on task not technology, fully understanding why something is useful before defining how it should be achieved.

The future growth of interactive telepresence and multimedia communications requires greater understanding of the roles and benefits of pictures in real-time, bi-directional systems. It also requires a fresh start, free from preconceptions. Why has the telephone been so overwhelmingly successful? What are the factors that connect psychology with engineering in telephony and videotelephony? What lies between human sentiments and electrical signals?

Hierarchy of Representation

Figure 2 shows a process hierarchy for telephony, composed of various



levels of representation between sentiments and signals. Note that, in this context, the word 'sentiment' has no association with sentimentality. It is used as an alternative to the word 'motive', because human communication occasionally occurs without precise motives.

The technical signals are at the bottom levels of Figure 2, with the human-related processes of telephony above. Each level is in two-way contact with its immediate neighbours, forming a serial chain running from bits through sounds and phonemes to semantics and sentiments; and vice versa. External factors, such as environmental context or verbal culture, can link with more than one level.

Side Channels

There are some processes which 'jump over' this serial arrangement of levels in telephony, such as the interpretation of tone of voice. This can result in semantic or cognitive content linking directly to a representation in sound. Such side channels can carry significant information, but the most important channel in telephony is the serial chain of representations that form speech. This is probably no

accident, since sound is a single stream of successive events with only one chance to observe each event as it happens. The nature of sound seems, in part, to have dictated the nature of its cognitive processing. The process hierarchy for telephony is therefore largely single-channel and time-invariant.

Benefit of Sound

The scale at the left of Figure 2 indicates the relative amount of design effort that has been involved in telephony. Only the lowest human level, marked 'Sensation and generation', has required any expenditure of effort, to match the telephone to human physical and psychoacoustic characteristics. The largest design effort in telephony has been devoted to the technical levels of transmission and switching. Almost all the human-related levels of representation in telephony have been inherited, without modification, from face-to-face conversation. No engineering design effort has been needed to shape these representations for telephony. That is why the '1 metre air-path' has been such a successful performance metric in telephony⁹.

Figure 3—Process hierarchy for videotelephony

The levels marked 'cognitive content', 'semantics' and 'grammar' in Figure 2 together connect psychological sentiments to the mechanistic parts of speech. As a group, they essentially define the purpose or benefit of sound in telephony.

Hierarchy for Pictures

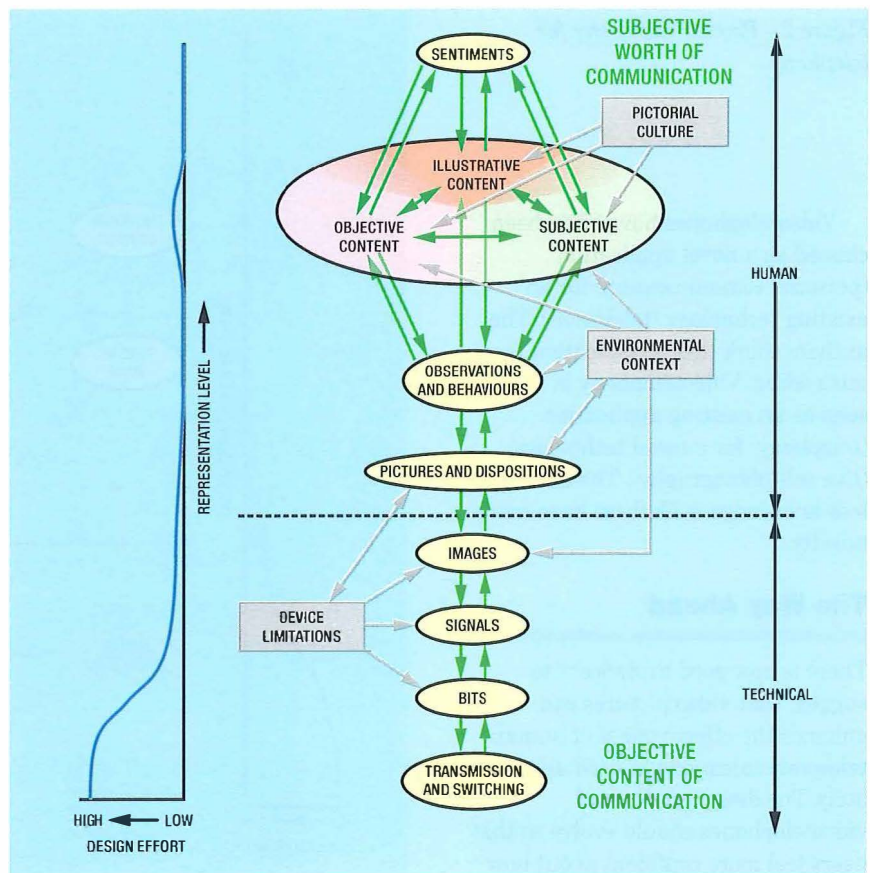
Figure 3 is a process hierarchy for the picture component of videotelephony. The technical levels are analogous to those of telephony. The lower human-related levels also resemble those of telephony, although the external factors such as environmental context and device limitations have more impact.

The higher human-related representational levels of videotelephony are unlike those of telephony. Pictures are unlike sounds, being more than one thing at a time and more than one time at a thing. As with telephony, the nature of pictures seems to have dictated many features of our cognitive processing. In the multi-stream, parallel world of pictures there is a complicated relationship between the 'sentiments' and the 'observations and behaviours' levels. A parallel arrangement of channels thereby emerges, with no dominant channel and no definite side channels. Instead, the importance of a given channel changes with time, depending on the tasks.

Environmental context mainly alters the human-related levels rather than the technical ones in videotelephony. Environmental factors can directly modify the relevance of observations, the interpretation of behaviours and some aspects of the usefulness of pictures.

Benefit of Pictures

Instead of a serially-connected group of three functionally-distinct layers, that define the benefit of sound, Figure 3 represents the benefit of pictures as a single layer. This layer has multiple overlapping and interdependent functions and connections, whose relevance changes with time. The links



between these functions, within the layer, can be as strong or stronger than the links to neighbouring layers, depending on the task. At the 'benefit' level, where the use and purpose of the picture is represented, videotelephony has a process hierarchy that is multi-channel and time-variant.

Three kinds of functions are shown within the 'benefit' layer in Figure 3, grouped together as objective content, illustrative content and emotional content. These groups, shown in more detail in Figures 4–6, are not meant to imply hard-edged or exclusive boundaries but to illustrate the range of functions in a complex continuum. For example, objective content can have emotional aspects, and illustrative content often has objective aspects.

The relative design-effort scale, at the left of Figure 3, shows that some effort has been devoted to human-related representation levels in videotelephony, especially for objective content. However, less effort (compared with telephony) seems to have been expended on the 'image' and 'picture' levels, either side of the human-technical interface.

Early trials of videotelephony tried to adopt the 'benefit' structure of telephony, with some additions from television culture. Except in

special cases, this approach did not prove successful and the complexity of the videotelephony 'benefit' layer may be the reason. Its multi-channel, time-varying nature may also explain why no satisfactory equivalent of the '1 metre air-path' has been adopted for videotelephony.

Objective Content

Figure 4 shows some aspects of the objective content of a videotelephony picture, classified by information, interaction, transaction and ergonomics. Many of these aspects can be represented in a simple telephone call or audio conference. Videoconference equipment is usually designed for a wide range of objective content but videotelephones often focus on social interaction through facial expression, with little provision for other kinds of content. Transactional content is expressible in speech, but is much more intuitively expressed through body language.

Illustrative Content

Aspects of illustrative content in videotelephony are shown in Figure 5, classified by demonstration, behaviour and circumstance. These

Figure 4—Objective content of a picture, grouped by type

aspects are very difficult to represent by speech or sound. This kind of content is typically the subject of carefully-directed video productions for training or education, but videotelephony can offer similar content in a spontaneous or informal way. This requires considerable flexibility of camera positioning and control, by the local or distant user.

Emotional Content

Lewis and Nightingale discuss emotional content in videotelephony using the terms *expression* and *projection*¹⁰. Expression is used to describe emotional properties of the picture that are obvious to an independent observer, while projection describes properties that are truly interpreted or experienced only by the participants.

Figure 6 shows an alternative view of emotional content in videotelephony, categorised by external situation or internal state of mind. Some of this information can be conveyed by telephony, in tone of voice or tempo of conversation. But most of the aspects in Figure 6, such as misgivings or distractions, are more easily and accurately portrayed by pictures. Negative aspects, about what is not said or done, cannot usually be conveyed without certain kinds of pictures. Body language is usually a good indicator of state of mind, even to the extent of contradicting the semantic content of speech.

Surrogacy

Emotional content in videotelephony is strongly related to a sense of surrogate presence, or presence by proxy, at a remote location. This surrogacy might involve psychological empathy or a shared state of mind. Alternatively, surrogacy can mean a shared awareness of situation without any marked emotional involvement. The two categories of Figure 6 reflect these two kinds of surrogacy, which often have contradictory implications for picture taking. Creating empathy

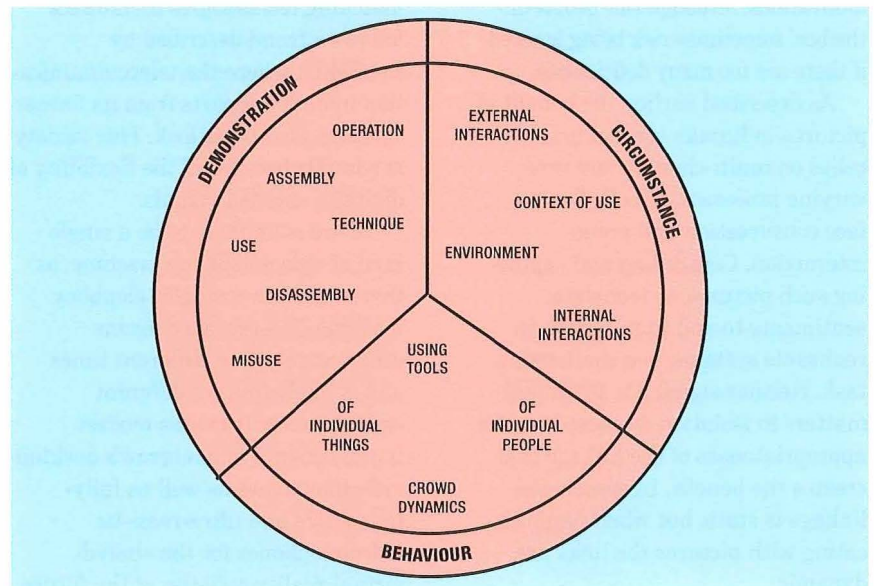
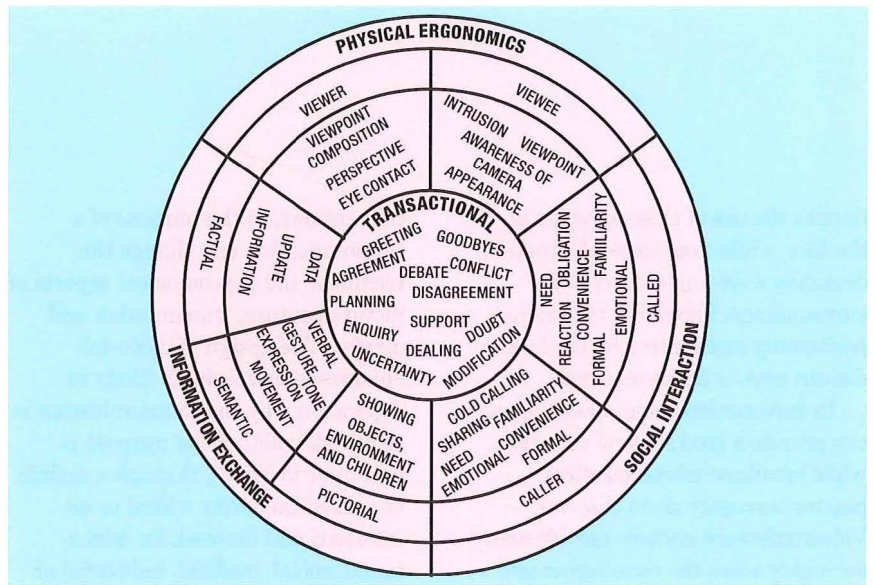
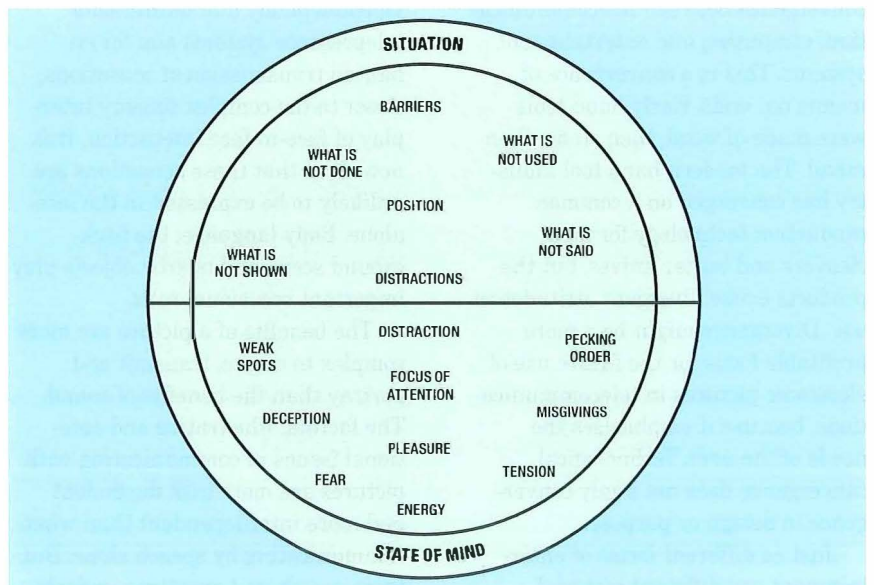


Figure 5—Illustrative content of a picture, grouped by type

Figure 6—Emotional content of a picture, by external or internal category



implies the use of close-up views of the face, while awareness of situation demands wide-angle views of the surroundings. Flexibility of camera positioning and control, by the local or distant user, is again required.

In audioconferencing, speech alone can provide a good sense of surrogacy, while broadcast television offers passive surrogacy as an observer. Videoconference systems provide active surrogacy when the room layout and equipment performance satisfy user motivations, although the 'people on the box' sometimes risk being ignored if there are too many deficiencies.

As described earlier, the benefit of pictures in human communication relies on multi-channel and time-varying processes, in both face-to-face conversation and group interaction. Controlling and capturing such pictures, to translate sentiments to and from signals in realisable systems, is a challenging task. Neither signal nor sentiment matters in isolation, because it is the appropriateness of the linkage that creates the benefit. In speech, the linkage is static but when communicating with pictures the links are dynamic.

Converge or Diverge?

The improvement of videotelephony is one element in a greater trend of convergence between telecommunication, computing and entertainment systems. This is a convergence of means not ends. Early hand tools were made of wood, then stone, then metal. The modern hand tool industry has converged on a common production technology for meat cleavers and butter knives, but the products evoke divergent attitudes in use. Divergence might be a more profitable focus for the future use of electronic pictures in telecommunications, because it emphasises the needs of the user. Technological convergence does not imply convergence in design or purpose.

Just as different forms of entertainment use different pictorial

conventions, so the purpose of a communication can change the technical and psychological aspects of picture capture, transmission and display. The design of video-telephones is increasingly likely to depend on why the communication is required, whether the purpose is formal or informal, through a mobile or a fixed link, over a local or an international network for educational, social, medical, industrial or business needs. Such a variety in matching technologies to markets follows a trend described by Bonfield¹¹, where the telecommunication industry departs from its former 'one size fits all' outlook. This variety is possible because of the flexibility of digitally-encoded signals.

There is unlikely to be a single kind of videotelephony machine, as there has been a single telephony machine. Telepresence means different things at different times and is likely to have different embodiments. It means modest improvements to tomorrow's desktop videotelephones as well as fully-immersive and ultra-realistic videotelephones for the shared-virtual-reality societies of the future.

Common Aims and Issues

Yet such widely different implementations will share common aims. All videotelephony and multimedia telepresence systems aim for enhanced transmission of sensations, closer to the complex sensory interplay of face-to-face interaction. It is now clear that these sensations are unlikely to be expressed in the face alone. Body language, the background scene and nearby objects play important occasional roles.

The benefits of a picture are more complex to create, transmit and portray than the benefits of sound. The factual, illustrative and emotional issues in communicating with pictures are more task dependent and more interdependent than when communicating by speech alone. But there are shared questions underlying

ing these design issues as Pickard and Cosier discuss¹².

When exploring the features of the signal-to-sentiment linkage in the evolution of improved videotelephony, the authors suggest four generic questions could provide illumination. They resemble the 'four whos test' used by economists:

- Who benefits from a feature, when and why?
- Who loses by that feature, when and why?
- Who provides the control decisions for the feature, what, when and how?
- Who pays for a feature, in what way and how much?

There are also four general issues about the interworking compatibility of bi-directional videotelephony systems with existing communication devices:

- *heterogeneity*—the inter-operability of different technologies;
- *representation*—the expression of information at a selected level of abstraction;
- *symbiosis*—a balanced combination of presentation, design and meaning; and
- *immediacy*—a low latency in information delivery.

A fresh approach to the interface between signals and sensations in videotelephony is likely to be market-specific and should make a fundamental re-examination of the purpose of the picture. What should it be a picture of, and why?

Conclusions

The authors wish it were otherwise, but they know of no magic potion to increase the success and popularity

of videotelephony. They see a hard road ahead of greater psycho-ergonomic understanding and more subtle engineering development, exploring the complexity within the signal-to-sentiment linkage of pictures in telecommunications. Their nearest approach to a magic potion is an attitude of mind for the journey.

The authors believe that too little is yet known for detailed engineering design or even for academic study of many issues in the signal-to-sentiment linkage. A playful outlook with an undamped, broad-focus curiosity and frank user feedback should prove rewarding in this poorly-known yet fertile territory between communication technology and its customers.

Coming Next

The next article will conclude the series with a collection of pictures, illustrating some of the issues raised and suggesting directions for exploration.

References

- 1 LEWIS, ALWYN; and COSIER, GRAHAM. Whither Video?—Television and Telepresence. *Br. Telecommun. Eng.*, Oct. 1998, **17**, p. 158.
- 2 LEWIS, ALWYN; COSIER, GRAHAM; and NIGHTINGALE, CHARLES. Whither Video?—Videotelephony in Perspective. *Br. Telecommun. Eng.*, Jan. 1999, **17**, p. 245.
- 3 LEWIS, ALWYN; and COSIER, GRAHAM. Whither Video?—Pictorial Culture and Telepresence. *Br. Telecommun. Eng.*, Apr. 1999, **18**, p. 30.
- 4 CHRISTIE, I. The Last Machine—Early Cinema and the Birth of the Modern World. British Film Institute and BBC Educational Developments, 1994.
- 5 SHANKAR, B. Breaking the Ice in Greenland. *Telecommunications*, July 1997, pp. 57–59.

- 6 <http://www2.state.ga.us/Departments/DOAS/GSAMS/> (distance learning).
- 7 <http://www.mayo.edu/telmed/history.html> (telemedicine).
- 8 ANDERSON, R. Value Extracting Services—A Better Deal for Customers. Supplement to *Competitive Edge* (BT internal publication), July/August 1997.
- 9 RICHARDS, D. L. Telecommunication by Speech. Butterworths, 1973.
- 10 LEWIS, ALWYN; and NIGHTINGALE, CHARLES. The Paradox of Videotelephony—Unconscious Assumptions and Undervalued Skills. *Br. Telecom. Technol. J.*, Jan. 1998, **17**, pp. 47–58.
- 11 BONFIELD, P. From Marconi to Multimedia. *Executive Voice*, Nov. 1996.
- 12 PICARD, ROSALIND; and COSIER, GRAHAM. Affective Intelligence—The Missing Link. *Br. Telecom. Technol. J.*, Oct. 1997, **15**, pp. 151–162.

Biographies



Alwyn Lewis
Networks and
Information Services,
BT UK

Alwyn Lewis is an advisor in Advanced Applications and Technologies, in Networks and Information Services, BT UK. He gained a B.A. from Cambridge University in 1971, and an M.Sc. from Essex University in 1972. After a spell with Plessey working on format-tracker vocoders, he joined BT Laboratories to work on acoustics, telephotonometry and telephone design. In 1984, he became head of a team developing a duplex hands-free telephone with a custom VLSI chip-

set and a DSP microprocessor. He then led the speech coding team, helping develop and test DSP software for CallMinder and for Skyphone facsimile and data services. His interests include beam steering microphones, speech coding and enhancement, adaptive audio signal processing and the social impact of electro-technology. He is a Chartered Engineer and a Member of the IEE and IEEE.

Alwyn Lewis can be contacted at alwyn.lewis@bt.com



Graham Cosier
Networks and
Information Services,
BT UK

Prior to 1984, Graham Cosier worked in the area of electro-acoustics and published several papers on auditory human perception which contributed to the British Standard for hands-free telephony. In 1986, he became involved with the emerging digital mobile radio system in Europe and was elected chairman of one of the technical groups that standardised GSM. In 1988, he became head of the group concerned with the assessment of coding methods, channel optimising techniques and echo control systems for increased utilisation of services using telephone networks. In 1994, he was appointed UK coordinator for ITU SG12, where he led the formation of the international standard for digital telephony. In 1995, he was appointed as the BT technical manager for the MIT Media Lab and was involved in establishing the UK Creative Art and Technology Centre. He is also the business support manager for BT Global Marketing, championing a vision of 'through the screen telepresence'—the ultimate collaboration space, that will provide a 'better than real life' experience between remote locations with the vision that 'you don't have to be there, to be there'.

Ian Marshall

Active Networks—Making the Next-Generation Internet Flexible

Active networks is a popular area of research that aims to enable the addition of new services to the Internet by the users, using simple components that anyone can build. The techniques involved also allow the network infrastructure to automatically adapt to any changes in traffic patterns caused by the new services. This will increase network usage and revenues, while radically reducing the cost and complexity of service and network management.

Introduction

Network users always want more services, better services or different services. To attempt to keep their customers happy, network operators have to support a bewildering array of services and features, leaving managers in the same position as a magician keeping a large number of plates spinning on top of long poles. Clearly there is a limit to how many plates can be kept up without any embarrassing failures. If the audience demand more plates, extra magician's assistants must be trained, and the audience must wait. The same is true in network management—new services cannot always be introduced without some delays due to equipment change or staff training. This makes the customers unhappy.

Active networks^{1,2} is a worldwide programme of research that aims to address this problem. The idea is to enable customers to make their own services by sending special instructions to network devices along with the information they are sending across the network. This is a bit like the magician teaching the audience to spin their own plates. Unfortunately it is not so easy, since network users are much more likely to get in each others' way than the magician's audience. It is not obvious that if I ask for fast delivery you will get slow delivery, but that is often the case in practice. One way to reduce the possibility of interference is to insist that the user's instructions run on a protected virtual machine. This is

known as *application layer active networking*³.

Researchers have already demonstrated benefits using application layer active networking for a range of applications including: transcoding web content, smart web caching, multicast bridging, reliable multicast, link-by-link compression, active web pages, information filters, alternate path routing for improved quality and security, security proxies, and personal profiles for mobile users. Nevertheless, as might be expected, researchers must address many issues before this radical new proposal can be used in anger. The most obvious outstanding requirement is to develop a management system that will help network operators and their customers develop, deliver and maintain products and services with the speed, flexibility, quality and economy necessary for future multiservice networks.

This article presents an overview of active networks research, and reviews the benefits of active networking to network operators. As an illustration of current active networks research, the evolution and activity of the Alpine project⁴ are described. The article concludes with some thoughts about the probable introduction schedule for active networks.

Active Networks

Existing networks are very expensive, and the deployment of new communication services is currently restricted

by slow standardisation, the difficulties of integrating systems based on new technology with existing systems, and overall system complexity. The biggest cost is management. The network of the future will need to be kept as simple as possible by using as few elements as possible, removing duplication of management overheads, minimising signalling, and moving towards a hands-off network.

The simplest (and cheapest) current networks are multiservice networks based on powerful asynchronous transfer mode (ATM) or Internet protocol (IP) switches. New transport networks are designed on the basis that nearly all applications will eventually use Internet-like connectionless protocols. The difficulties of adding services such as multicast and quality of service (QoS) to the current Internet demonstrate that even these simpler IP-based networks will require additional mechanisms to enhance service flexibility. Improving information flows so that customers can more easily express their desires, while providing the support that enables the desires to be realised, is the most obvious way to successfully enhance flexibility. Figure 1 shows the information flows between the stakeholders of a network. Currently, if a consumer requests some content, his/her request goes to a provider who

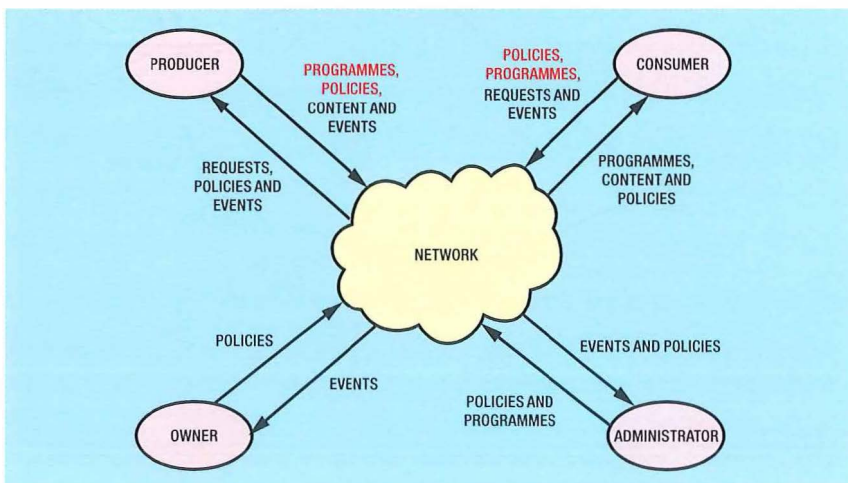
then supplies the content. Both parties have to formulate their request and response in terms laid down in policies specified by the administrator, and both generate usage events which will be seen by the administrator (and the owner if they are very important). Active networking¹ adds extra mechanisms (highlighted in red) to enable customers (consumers and producers) to add programmes and policies of their own, rather than being too tied to those provided by the operator as at present. For example, the customer can supply an information filter of his/her own choice, avoiding the risk of only being sent information that earns revenue for the network, or for a provider (as with the Amazon.com agent), while minimising his/her transport costs by filtering at the optimum site in the network. Those customers requiring global service have a further advantage, as they can supply programmes into all the networks their service crosses, and are not restricted to the lowest common denominator service that is currently available. In other words, active networking avoids the need to standardise services and enables customised services with global reach.

Active networking was originally a proposal, by Tennenhouse at MIT¹, to increase network flexibility by adding programmes to the packet

header. These programmes are intended to run on network devices that the packet encounters. This is referred to as the *capsule* approach. There are a number of problems:

- The maximum transport unit (MTU) size in the Internet is typically 565 bytes, so if there is to be a program embedded in every packet the program must be very small.
- The proposal envisages programmes being supplied by network clients. However, service operators will never permit third-party programmes to run on their equipment without a strong guarantee that the program will not degrade performance for other users.
- The programmes are added to the switch control kernel in the router; however, all the known approaches to making the kernel extensible degrade performance.
- Undesirable interactions between programmes and network features are almost impossible to predict and control.
- Standards for the interface offered by active routers must be developed before any service based on this proposal can be offered. Appropriate standards are not even being discussed at present.

Figure 1 – Network information flows. The flows available in active networks but not traditional networks are highlighted in red



Despite the manifest difficulties inherent in this proposal, it has succeeded in highlighting an important requirement, and stimulating discussion among a previously disparate community of researchers attempting to develop more immediately realisable means to resolve the requirement. The main alternatives are summarised below.

The first alternative is a somewhat different flavour of active networking, in which the packets do not carry programmes but flags indicating the desirability of running a programme².

This approach attempts to resolve the issue of restricted program size, and potentially gives network operators the freedom to choose an appropriate program of their own which has been tested. However, the range of flags cannot be large, as the space available in the transport layer header is tiny.

The second alternative is usually known as *programmable networks*⁵. In this system, clients can download their own software onto network devices before using their application. However, this presumes the client knows which devices will be used. In the Internet the client typically has no idea where his/her packets will go, indeed it is quite possible for packets in a flow to follow a number of routes.

We have proposed a third alternative, known as *application layer active networking* (ALAN)³, which is perhaps the most immediately realisable. In this system, the network is populated with active nodes that can run user-provided programmes on top of the network stack rather than within it. The active nodes are offering a general-purpose processing platform, and they are logically end systems rather than network devices. The approach relies on redirecting selected packets onto the processing platform, where user space programmes can be run to modify the content or the communication mechanisms. Packets can be redirected using a single active packet tag in the transport-layer header, or on the basis of the mime type in the application-layer header. The programmes can be selected from a trusted data source (which may be a local cache) containing only well-tested or specified programmes, and can be run without impacting router performance or requiring operator intervention, since they do not impact the control kernel for normal packets. Programmes can be chosen to match the mime type of the content (in the application-layer header), so again no additional data or standards are required. Alternatively, a more detailed specification can be supplied using XML, the

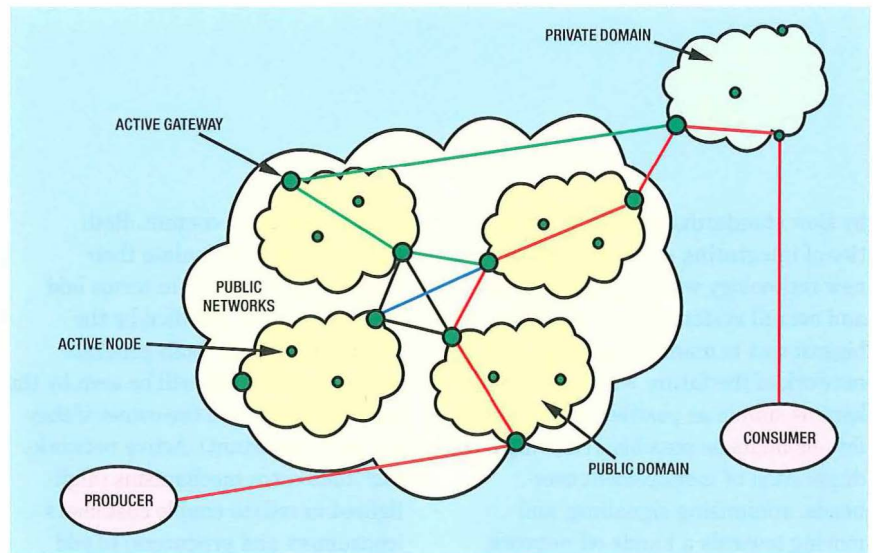


Figure 2—Schematic of an ALAN overlay, showing alternate path routing (green) for the producer's response to a customer's request

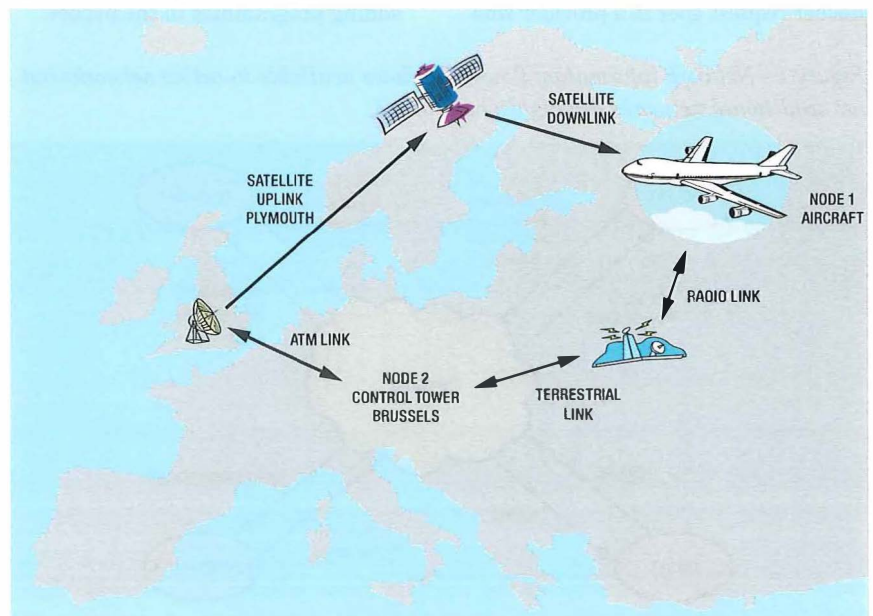
infinitely extensible new language for the World Wide Web⁶. There is a small performance penalty associated with the redirect operation, but this is acceptable for most applications. The major outstanding issue is the interaction between dynamically loaded programmes, and this is a priority for ongoing research.

Figure 2 illustrates an ALAN implemented as an overlay on the existing Internet. The basic idea is that return traffic will follow a different route (shown in green) from the outgoing traffic (shown in red). The re-routing allows the traffic to be sent via a site where enough computing resource exists to perform any extra operations specified by the consumer

or producer. These could include transcoding for low-capacity receivers, link technology specific compression, establishment of trust relationships, or using the nearest cache that already has the requested content.

In general, modern networks must frequently optimise the management of communication among nodes that are interconnected by diverse paths. These paths may be based on heterogeneous technologies with divergent properties. This makes smart path choice an essential feature in order to provide quality services. An interesting application that highlights many of the issues is the aircraft services application illustrated in Figure 3. There are two basic services: commu-

Figure 3—Aircraft services application



nication with the flight deck and e-commerce/WWW access for the passengers. The aircraft (Node 1) has a 64 kbit/s radio-based bi-directional link to the control tower and a very small aperture terminal (VSAT) based downlink (2 Mbit/s). The control tower can use wide-area connectivity to establish a high bandwidth link to the aircraft using ATM and satellite combined. Path selection is performed at the control tower by examining the meta-information of the packets, deciding which is the most appropriate path and adding security if needed. QoS management will also be needed in the plane, to ensure that life-critical information from the flight deck is delivered into the bandwidth-restricted downlink before any traffic originating from passengers. For this application, QoS can be characterised in terms of bandwidth, latency, security, strength of guarantee and uni- or bi-directionality. Current Internet routing protocols, for example, OSPF and RIP, use shortest path routing that is optimised for a single metric, the hop count. The aim here is to enable route decisions to be made on a more application-specific basis. In this design, the application requirements are expressed in XML, a language that is rich enough to express the application requirements and force a correct route choice. Our scheme also allows return traffic to be correctly routed using an appropriate dynamically loaded program, supplied by the application and run in the control tower or some other more appropriate intermediate node near the head-end of the satellite link.

Benefits of Active Networking

The key benefit of active networking is the ability to support the deployment and execution of network programmes as and where required by operators and users. Active network services enable rapid deployment of 'value add' that is apparently part of the network, but can be sourced from third parties.

This will enable market response to be quickly tested, and market presence to be rapidly established with minimal development costs. There will also be real benefits to users compared to 'plain vanilla' Internet service provider services, which will tend to increase network traffic. In addition, there will be many opportunities for novel forms of commercial activity by both operators and their customers. Some key revenue enhancement possibilities are:

- provision of active network services;
- service and code brokerage;
- directories for active programs and their sources;
- service 'operating system' leasing (processing time to be charged);
- active network program distribution and storage (information logistics);
- facilities management of third-party services (advertising, access control, usage logging, billing, revenue collection etc.);
- rapid launch of new services;
- additional flexibility to facilitate fast service enhancements;
- the range of services will increase (This is a direct consequence of the programmability and third-party involvement—when everyone can develop services, rather than just those developers employed by operators. Many more services can be developed in response to user needs.); and
- customised services will increase.

Active networks also promise diminished network operating costs. Central management overheads are reduced and service management is

simplified. Rather than operators being required to manage a huge range of Internet services, users are free to compose and manage their own services. Operators only need to manage the platform. The benefits of active networks are clearly substantial if they can be realised. Many operators have now set up research projects to help develop the technology required and resolve the outstanding issues. To illustrate the work, and the radical changes of thinking it represents, the Alpine project set up by BT is described below.

A History of the Alpine Project

The Internet opened up to commercial users in 1991, offering an anarchic model of service provision. This was largely ignored by established operators, as the Internet (due to problems with security, quality and reliability) was seen as an expensive academic toy, not as a credible option for major business users. In 1993, the first Mosaic browser appeared, providing for the first time a genuinely useful graphical interface to Internet features. This provided a compelling demonstration of what the Internet could offer, and enabled telco-based researchers around the world to convince their senior managers and strategists to take the Internet seriously as a future option.

One of the impacts of this change on the BT research programme was to set up a major university research initiative (the management of multiservice networks (MMN) University Research Initiative (URI)⁷) which combined most of the UK universities with prominent work on internets, management and distributed systems. The URI was first discussed in late-1993 and eventually started in September 1994. The initial work of the URI was largely concerned with QoS architectures and management, and, with the honourable exception of the prime contractor

at UCL (Jon Crowcroft) who was pushing an end-system-based approach (derived from his Internet Architecture Board work), tended to follow a somewhat centralised control model. Moving away from centralised control is hard for scientists and engineers who have been using it successfully for many years, but is essential to make the best use of active networks. The URI succeeded through absorbing new people and new ideas as described below.

In mid-1995 the MMN researchers were joined by Mike Fry from the Sydney University of Technology, who had been working in the ESPRIT project HIPPARCH on the application layer framing (ALF) and integrated layer processing (ILP) ideas proposed by Clark and Tennenhouse at MIT⁸. The discussions leading to this new contract had dramatic effects.

The potential of the ALF/ILP ideas was so enticing that we were able to embrace customer empowerment through adaptive applications and an adaptive IP-based infrastructure in our main research project. We concentrated on enhancing the WWW by adding a distributed processing capability to enable transactions, messaging, streaming and dynamic applications. The URI partners refocused their work. Lancaster moved from QoS architectures with global knowledge to architectures based only on local knowledge, and concentrated on the local resource control necessary to deliver QoS. Imperial moved from centralised configuration management using policies to a more distributed role-based model, and concentrated on dynamic stack building in end systems. Cambridge moved from considering how to securely manage the switches in an operator domain to enabling customers to manage switchlets. Oxford Brookes moved from available bit rate (ABR) to heterogeneous multicast. Loughborough continued their work on Internet control message protocol (ICMP) based monitoring but focused

on monitoring switchlets. UCL remained focused on end-system-based network admission control. UTS commenced by building an audio application that used Java in an adaptive integrated layer protocol stack controlled by ALF. We also joined the HIPPARCH II project, which is also part of the ESPRIT programme, and later the COIAS project in the ACTS programme, in order to develop links with research and development by potential suppliers.

By Christmas 1996, the team was exploring the use of network intermediaries based on http caches. The idea was that caches could serve as the platform for the web extensions we wanted. Discussions among the team led to the first enunciation of the concepts now called *ALAN*⁹, which were based on developing the ALF/ILP ideas for the context of interactive multimedia services in a cache hierarchy. A year later we had the first ALAN demonstration working. We are planning a full-scale test open to real users on the public Internet in three years time. The research that must be done first is the ongoing work of the Alpine project.

Some Issues Alpine is Addressing

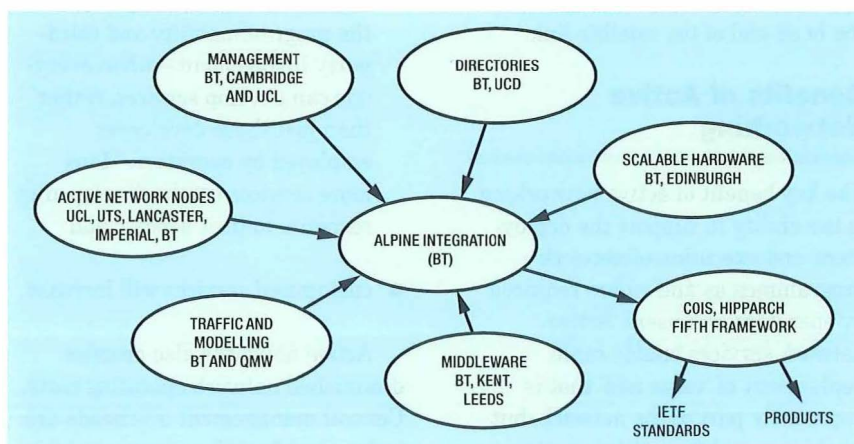
Alpine is coordinating a range of tasks (highlighted in Figure 4) with a

wide variety of partners. Some of the key activities are summarised below.

Preventing programmes or policies supplied by a user from interfering with the operation of policies and programmes supplied by other users, unless the interference is authorised by both users, is the most obvious unsolved problem in active networks. Both direct interference (as in hacking attempts) and indirect interference (through resource grabs, name conflicts etc.) need to be considered. It is also important to recognise that most cases of potential interference are not deliberate, but are the result of programming errors. The research community has so far concentrated on two approaches: 'safe' programming techniques and languages, and protected environments. Our own research uses an enhanced Java virtual machine (JVM) and operating system to provide a strong protected environment. Our reasoning was that safe programming is not good at dealing with platform code errors, and in any case users will not be keen on the languages. Currently, proprietary language-based active networks are losing popularity elsewhere too.

Management of active networks is a second major problem, which has not been widely considered as yet. Although active networks considerably reduce the service management overhead of today's networks, the element (router/switch, pipe and

Figure 4 – Task structure of the Alpine project



processor) management needs are still considerable. Management architectures such as telecommunications management network (TMN), which were designed for relatively static network provision scenarios, do not cope well with a requirement to continuously add management routines for new programmes supplied by users. In addition, they constrain users to design services which conform to a single service model (designed by operators, for operators), which in many cases does not match the users' needs (for example, the intelligent network (IN) call model does not allow changes like upgrade quality to be made during a call, and the management system does not allow users to define their own call models). We have proposed a highly distributed system based on established research in distributed processing. Our proposal uses policies to provide the dynamic programmability, and a combination of message-based push and demand-pull, through a cache hierarchy to distribute the policies. To avoid imposing a model, users are left to manage their own policies and programmes in their own way, while the operator simply polices and controls resource usage on the devices he/she owns.

Performance is a further area that raises unanswered questions for active networks. Scaling the throughput of routers to billions of packets per second has been a long hard struggle, and there is a risk that the packet-handling overhead implicit in active networking could eat away the performance gains that have been achieved. In application layer active networking² we have attempted to sidestep this problem by minimising the change in the forwarding path, and by not allowing the forwarding rules to be dynamic. In ALAN a single header byte can be used to mark the packet as active. The only change in the forwarding path is then to add a forwarding rule which routes active packets to an internal network port attached to an active proxy

server where the active services are implemented. In the future, use of adaptive hardware such as field programmable gate arrays will enable this redirect to be dynamic without reducing performance.

Resource discovery is emerging as a hot topic in all varieties of networking at present. The debates around directory-enabled networks, WWW search engines and IP telephony have all highlighted the need for effective mechanisms for users to identify who and what are connected when and where. In active networking it is also necessary to discover programmes which can be used to build services. Our current approach is based on indexing the contents of caches at the active nodes, and distributing the indexes to neighbouring nodes. Each node will have local policies (based on usage) to decide which entries to keep and which to pass on further, and users can push entries into their local index. The resulting directory can resolve enquiries for popular items with a very localised search, but can also support more exhaustive distributed searches on the few occasions when they are necessary.

Traffic in data networks is highly unpredictable, and designing the network to be efficiently utilised, without incurring unacceptable loss rates on popular routes, is very difficult. Active networking is likely to make it even harder to design a good network. We believe it is vital to obtain an improved understanding of traffic statistics in data networks, and develop improved models to assist network designers. With this in mind we have studied a variety of Internet traffic types in detail, and are using the results to develop a range of models. One interesting (but speculative) approach is to test control rules which enable a stable network to emerge from a process which randomly adds nodes to the network (rather like growth in the current Internet), and avoid the need for an a priori 'design' altogether. Of course, in an active network the nodes need not

have fixed properties; nevertheless, the network can still be stable if the control rules are well chosen.

Currently, Alpine plans to have demonstrated usable solutions to all the above within four years.

Possible Introduction Schedule for Active Networks

It is highly probable that a limited form of active networking will be deployed on the public Internet within two years. This is likely to involve active nodes at the edges of some network domains that run simple programmes in a protected virtual machine. The programmes will not interact with each other, and will primarily be proxies for application operations currently carried out more remotely, such as transcoding, compression, authentication, advert selection and directory search.

About five years from now, when many of the research issues are resolved, a second generation of active network will be available. This will be based on a network-wide overlay of high-performance active nodes that can handle a wide range of programmes. Additional enhanced services will include personal agents and mobility services, reliable messaging and multicast, flexible service level agreements and QoS mechanisms, a direct management interface for users and dynamically upgradeable protocol stacks.

Five years later, in 2009, there will certainly have been further developments. While it is difficult to make meaningful predictions in an area of technology moving as rapidly as the Internet, it is clear from current research that dynamically configurable hardware will be available. Such hardware will increase the flexibility and performance of active networks and will certainly be used in some way. It is also possible that a modified routing architecture, designed to enable continuous modification of the forwarding rules, will be in place.

Any operator wishing to take advantage of these opportunities will need to ensure its network designers and engineers adopt the decentralised approach implicit in active networks. As illustrated in the development of the Alpine project, this is certainly possible, but will not be easy. I hope this article will act as a catalyst for some of its audience, and encourage the development of new approaches to network engineering, including active networks.

Further Reading

- 1 *IEEE Network*, special issue on Active and Programmable Networks, May/June 1998.
- 2 *IEEE Computer*, special issue on Active Networks, April 1999.
- 3 FRY, M.; and GHOSH, A. Application Level Active Networking. To appear in *Computer Networks and ISDN Systems* and <http://dmir.socs.uts.edu.au/projects/alan/>
- 4 <http://www.alpine.drake.bt.co.uk>
- 5 *IEEE Communications*, special issue on Programmable Networks, Oct 1998; and CAMPBELL, A. et. al. Open Signalling for ATM, Internet and Mobile Networks. *Computer Communications Review*, Jan. 1999, **29**(1) and <http://comet.columbia.edu/opensig/activities/opensig98.html>
- 6 BOSAK, J.; and BRAY, T. XML and the Second Generation Web. *Scientific American*, May 1999. <http://www.sciam.com/1999/0599issue/0599bosak.html>
- 7 <http://www.cs.ucl.ac.uk/staff/jon/mmn/mmn.html>
- 8 CLARK, D.; and TENNENHOUSE, D. Architectural Considerations for New Generation Protocols. *Computer Communication Review*, Sept. 1990, **20**(4).

Biography



Ian Marshall
Networks and
Information Services,
BT UK

Ian Marshall is a technical group leader at BT Adastral Park. Since becoming a group leader in 1992 his research has concentrated on distributed systems and the Internet. Currently he is leading a project on active networks, involving BT and a number of partners from universities in the UK, Europe and Australia. He is also involved in the ESPRIT project HIPPARCH and the ACTS project COIAS. Previously he worked in optical networks, broadband networks and network strategy. He is a fellow of the Institute of Physics, a chartered engineer, and a senior member of the IEEE. He serves on several institute committees, and on a number of EPSRC and European research panels. He is also an industrial adviser at several UK universities and a programme committee member at several major conferences. He publishes frequently in international journals and conferences.

Martin Hind

Simulation Modelling in BT Operations

With the advances in technology and the opportunities that follow within the communications world, BT's operational environment is becoming increasingly complex. Traditional process management methods are becoming outdated and BT is updating its process management tools and techniques to ensure its operational environment remains state-of-the-art. Discrete event simulation is one of the tools being introduced into BT's toolkit to provide a scientific approach to operations management.

Introduction

BT processes and systems have evolved over the last few decades to support the rapid increase in volume and complexity of new technologies and services being introduced into its network. Consequently, the operational environment as a whole has become very complex and very difficult to understand and manage.

The Operational Design team, part of BT Advanced Communications Engineering, has been identifying and demonstrating tools and techniques that would help improve process management within the company. Among the tools and techniques evaluated, discrete event simulation was found to be a very powerful and effective tool. This article presents a brief introduction to discrete event simulation and its application to BT operations.

Simulation

A simulation in this context is a computer model that mimics a real-life process and can be used to study the behaviour of the process in a scientific way. The computer model consists of a set of assumptions which describe how the process works in the form of mathematical and logical relationships. If the relationships are simple then it is possible to use mathematical deterministic methods to study the behaviour of the process. However, most real-world processes are too complex for these methods, and this is where simulation can help. A simulation brings together the components of the process that interact to produce the overall characteristics which the process displays. These include key process

activities and their relationships to each other, the activity times, the people in the process, their responsibilities and skills, the orders and their attributes such as product type, priority and agreed completion time. Having built a simulation model, experiments can be performed by changing parameters and predicting the response. Experimentation is normally carried out by asking 'what-if' questions and using the model to predict the likely outcome. However, it is important not to overlook the fact that simulation is a decision support tool and does not seek to directly optimise a process.

Simulation models can be categorised as being either *discrete* or *continuous*. A discrete model is one in which the state variables have step changes in line with certain events at separate points in time, such as orders joining or leaving a queue. In a continuous model the state variables change continuously with respect to time; examples are the speed and direction of a car. As most processes contain aspects of both discrete and continuous behaviour, it is possible to use either type of model. However, the behaviour of the operational processes being considered within BT fall mainly into the discrete category and therefore discrete event simulation is of most interest.

Though discrete event simulation has been used for many years in most business sectors, including both manufacturing and service industries, its use within BT has been limited to a number of small projects up to now. The Operational Design team at Adastral Park, Martlesham, is the first group to build the skill into the team and apply it widely throughout the company.

Why Use Simulation?

Simulation models can be used to facilitate improvements such as reduced cost, reduced process lead-time and improved quality of service. For many applications, simulation is the only tool that is capable of providing scientific analysis to support these types of improvement.

There are many other benefits from using simulation. The procedure of constructing a simulation model itself leads to a better understanding of the overall process. Then any improvements identified can be tried out in a safe environment before making the change to the real process. This gives greater confidence and reduces risk.

With a model, the user has control over the rate the simulation is run, therefore the impact of a seasonal type trend can be evaluated in minutes using a simulation model instead of waiting several months to see how the real process will react. The model is also repeatable, and so the same situation can be tested under slightly different scenarios to identify the most suitable process behaviour.

Other benefits of simulation include the power of its communication and creative environment. Not only is it very good at demonstrating what is currently happening, it is also effective at evaluating and demonstrating new ideas. It fosters creative attitudes and encourages end-to-end solutions.

Application Areas

Simulation can be used for a wide variety of problems and applications. The main application areas within BT currently include the following:

Process re-engineering

This involves constructing a model to help understand the behaviour and characteristics of an existing process or components of an existing process. This model is then used to experiment with scenario changes to help

define and quantify suitable improvements.

Process engineering

When designing a new process, simulation can be used for communication and to ensure that the new process behaves as required. By imitating the operation of the new process, potential bottlenecks can be identified, shortages found and solutions designed.

Operational planning

There are many ways in which simulation can assist in making the best use of available facilities through day-to-day planning and scheduling support. It is good at representing the transient effects of processes and so can be used, for example, to look at existing work stacks and look ahead several days to predict the expected outcome of a particular resource decision.

This article now presents two simulation models that are currently being developed, one applied to process re-engineering and the other applied to operational planning.

SDH Circuit Throughput Model

Process re-engineering is currently one of the main application areas for simulation in BT. Given a specific objective, the first step is to translate the process requirements into information requirements. The information is then captured, usually in the form of an activity process flowchart and associated data such as activity times, percentage failures, number of people, volumes of orders, order profiles etc. This information is built into the simulation tool and validated to ensure that it is a reasonable representation of the actual process. After an iterative cycle of data capture, validation and refinement, the model can be used for analysis and scenario testing. This type of model is very powerful because it consists of real data and

information and can greatly improve the understanding of a process operation and behaviour.

The production of a model for the synchronous digital hierarchy (SDH) core transmission technology is a process re-engineering exercise which is required to understand the existing behaviour of the SDH provision process and to identify the best strategy for quickly increasing the daily average throughput volumes for new customer circuits.

The SDH model is built around the process of the network operations units (NOUs)[†] in Manchester and London, the central operations unit (COU)* in Oswestry and the planning assignment and configuration system (PACS). The process is broken down into the three main process function areas: *routing*, where the customer circuit is designed through the network; *assigning*, where the network databases are updated to register the new circuit; and *configuration*, where the network elements are configured remotely to enable the new circuit. The process flow and functional areas are shown in Figure 1. The process flow follows the PACS system state transitions from placed (PLC) to awaiting route (ART) to routing (RTG) etc. This allows the PACS system operations database log of all historical circuit provisions to be used in providing the timings and success/failure percentages for all system functions. Validation takes place by comparing the model performance to the real-life measures; the model is then refined until correctly calibrated.

The procedure of constructing a model such as this effectively provides a framework in which to gather information. This results in a very

[†] The network operations units are responsible for providing, configuring and maintaining the SDH network.

* The central operations unit is responsible for planning and managing core network capacity, designing, configuring and maintaining the core network.

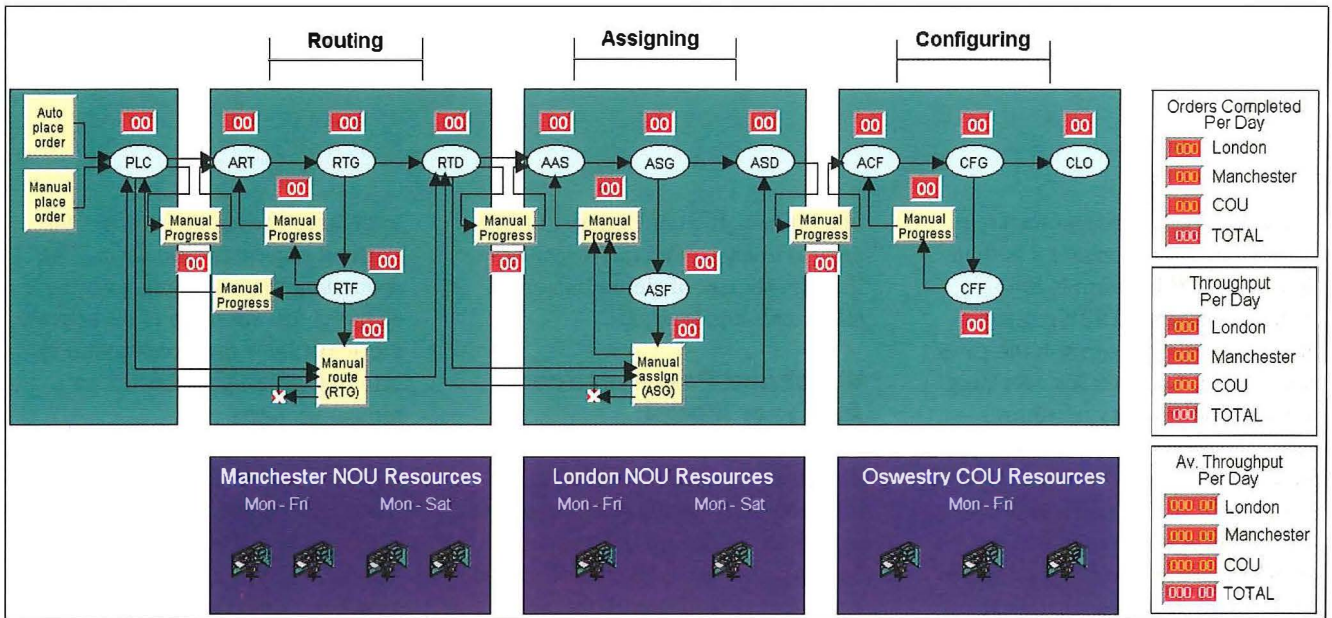


Figure 1—Representation of SDH circuit throughput process in the simulation model

thorough assessment of available information, which in itself provides a very good understanding of the process. In addition, specific analysis can be used to understand the dynamics of the system. Some examples of this analysis are now given.

An example of the initial analysis† of the process sensitivity to volume increases is shown in Figure 2. This shows how average throughput increases as the volumes of orders increase, but only up to a certain level. This equates to the maximum circuit throughput volume. The graph shows that component 1 reaches 100% utilisation first and hence is the bottleneck in the process. Also it shows that increasing order volumes beyond this point can have a detrimental effect on circuit throughput. Action can be taken to increase the capability of component 1 and remove it from being a bottleneck.

The next stage of the analysis is to produce a utilisation profile for component 2 by artificially removing component 1 as a bottleneck. This is achieved by running separate simulations of the model for increasing levels of component 1 capacity, and then recording the maximum average throughput and component 2 utilisation level for each run. Figure 3 shows the result of bringing together the information from these simulation runs.

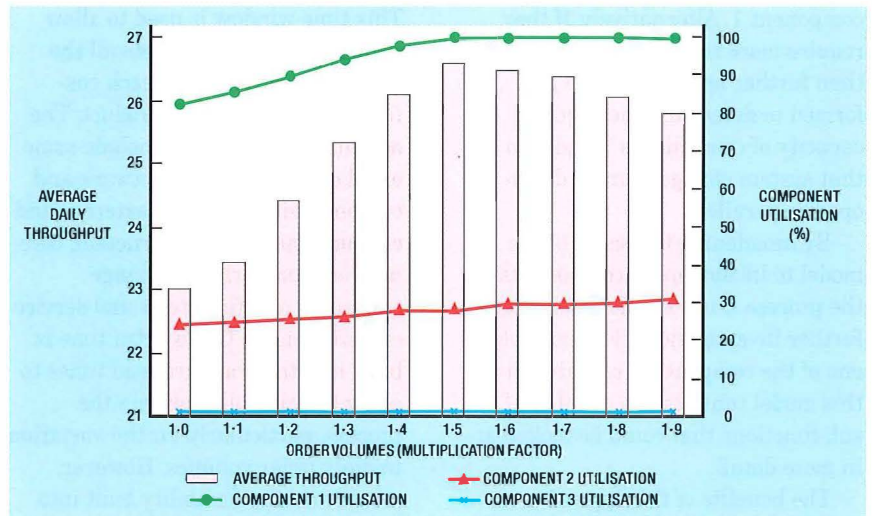



Figure 2—Process sensitivity to volume increases (based on component 1 with capacity of 1)

Figure 3—Theoretical maximum circuit throughput for component 2



† The data used in this example is for illustrative purposes only

The theoretical maximum circuit throughput for component 2 is then identified; in this example the maximum is around 90 circuits per day. Additionally, the graph indicates that when component 2 is operating at maximum capacity it requires component 1 to have a capacity of at least 9.

Hence the user can identify the amount of capacity required for a given number of circuits per day and is now in an informed position to decide if a maximum of 90 circuits per day is acceptable and, if so, how much capacity is required from component 1. Alternatively, if they require more than 90 circuits per day then further analysis can be performed to determine the required capacity of components 1 and 2 so that system changes can be developed in parallel.

By broadening the scope of the model to include more components of the process it is possible to carry out further investigations. For example, one of the components contained in this model consists of a number of sub-functions that could be looked at in more detail.

The benefits of this type of modelling are to greatly improve the understanding of the process behaviour, to highlight bottlenecks within the process and to quantify the volume of orders at which these components become a problem. Often the process component changes required to increase the capacity of these components involve system development changes and people training. Both of these changes incur a significant time delay, so using simulation to understand and quantify requirements initially is a great benefit.

The use of simulation in this project has been greatly assisted by the quality and availability of process audit information obtained from the PACS system log. The accuracy of any simulation model is dependent on the quality of the information available, thus if more operational systems had suitable instrumenta-

tion built-in, it would support future process analysis and improvements through the use of simulation and other tools and methods.

Wideband Performance Analyser Model

The objective of this model is to illustrate how simulation can be used to address resource management barriers and maintain the prescribed level of quality of service for provision of a product to customers.

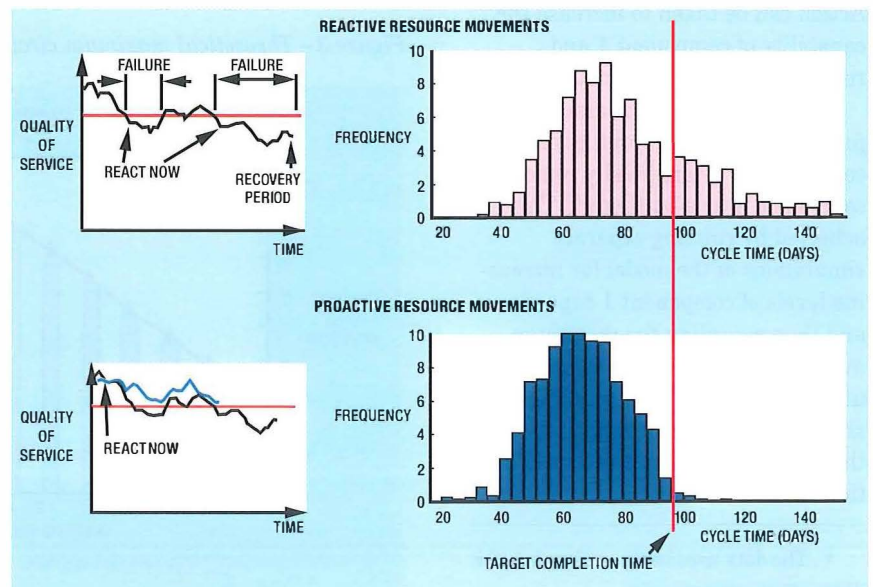
Consider a provisioning process which has a lead time of 30 days. This time window is used to allow sufficient time to complete all the necessary activities for each customer for their chosen product. The activities required can include some or all of the following: network and customer site planning, external and customer network construction, core network jumpering, exchange equipping and final tests and service commissioning. Often extra time is built into the standard lead times to allow for variability within the process, particularly for the variation in daily order volumes. However, sometimes the flexibility built into the process through the standard lead times is not sufficient to cope

with large variations in orders for particular geographical areas. The desired method of coping with these geographical surges is to proactively move available resource capability into that geographical patch on a temporary basis.

If there is a limited amount of information available, combined with complexity of products and a number of different parallel activities, it is very difficult to manage the work and the people in a proactive way. This can result in a time lag between identifying potential failure to meet a deadline, and the resource management teams recognising and reacting to this trend.

This results in a situation where the resource teams do not realise they are about to fail to meet their target delivery times until it is too late to react. Consequently the process has the potential to go into a failure mode for a period before being brought back on track, see Figure 4. The reactive resource movement graphs demonstrate the situation where the resource management teams do not have sufficient information to react until it is too late. This is reflected in the performance distribution graph where the cycle time often exceeds the target. The second set of graphs showing proactive resource move-

Figure 4 – Rationale behind moving from reactive to proactive resource management



ments illustrate how the situation can be improved if more information is available to the resource management teams; they can then react earlier and reduce the amount of process failure.

Simulation can be applied to this type of problem. The simulation uses available information about existing work stacks, people availability and activity assignment rules to predict future resource capacity requirements on a geographical basis. The simulation will also highlight surpluses and deficits in resourcing and thus facilitating proactive resource management of the existing people within the process.

The type of simulation model which could be used to provide decision support to the resource management teams is shown in Figure 5. This shows the graphical interface to a model and

allows the user to view the volumes of activities as they flow through the various resource team queues within the process and to view the utilisation of the people within the individual work teams.

There are potentially a number of significant benefits in using simulation to assist in decision support of a process as described above. These include:

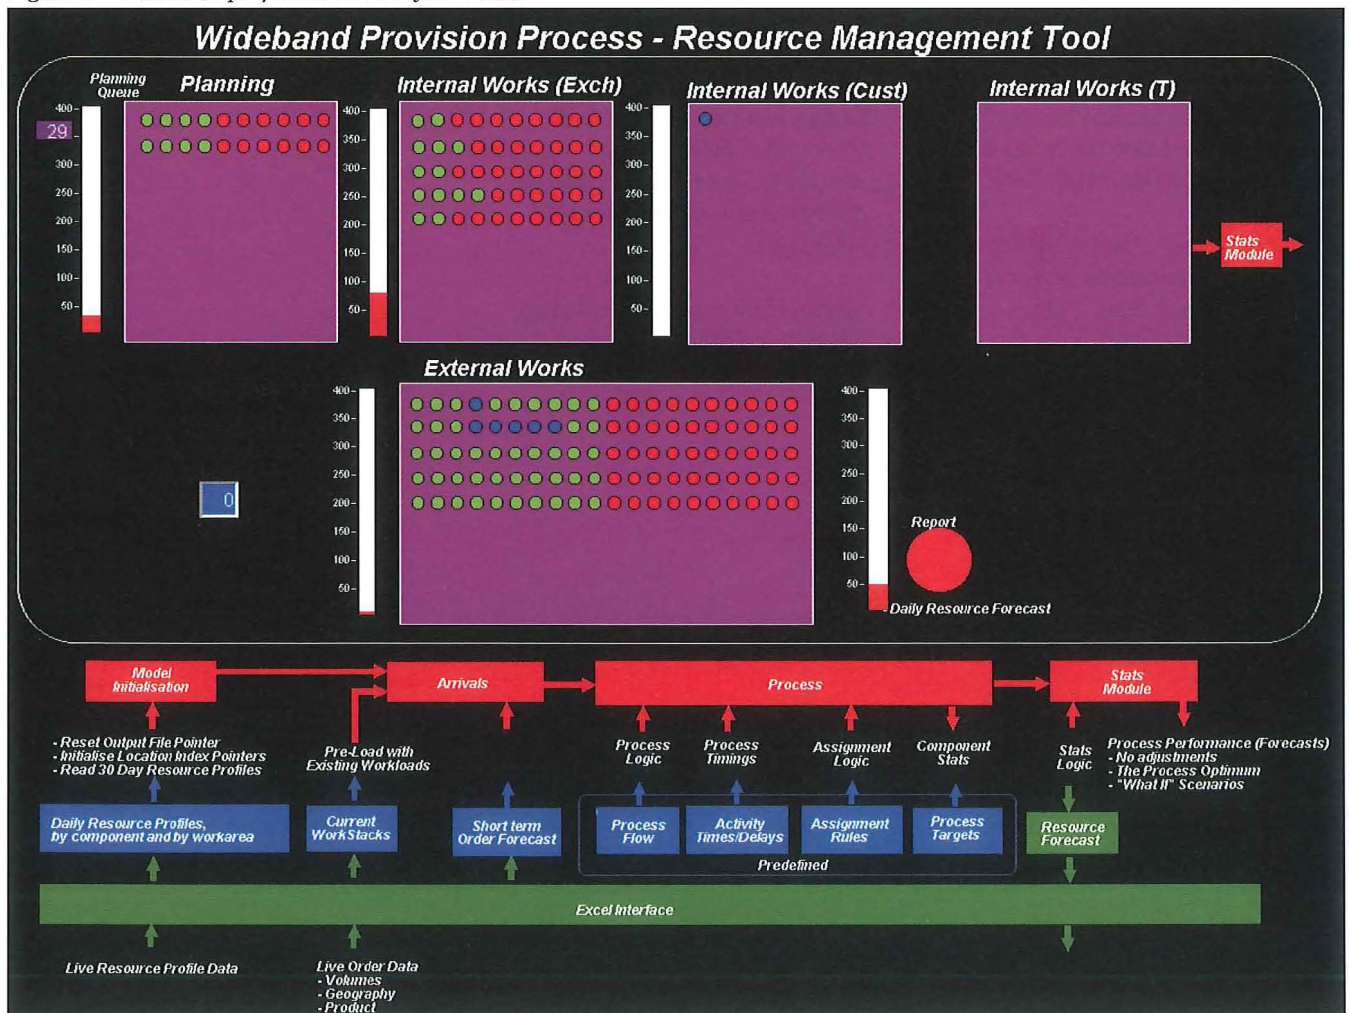
- improved performance in terms of delivering to agreed dates (deadlines are not missed, so compensation is not an issue);
- increased dependability, leading to –more business and thus increased revenue, and –improved customer and employee satisfaction; and

- improved resource utilisation leading to lower unit costs.

Conclusions

Simulation has proved to be a valuable and effective tool in process management. It is a very powerful tool that, when used appropriately, is a great enabler to understanding chosen areas of today's complex operational environment. However, the quality of each model is dependent on the information available and the user's ability to interpret the information in the correct way to produce a valid model. The ability to operate a simulation tool can be picked up quite quickly. However, the application skill required to use the tool for particular problems has a much longer learning curve, and

Figure 5 – Wideband performance analyser model



hence there is a risk that an inexperienced user may produce invalid models, possibly without even realising it. Therefore care should be taken when using these tools. The Operational Design team is working with the rest of BT to ensure that the tools are applied to appropriate applications and in the correct way.

Acknowledgements

The author would like to thank his colleagues for their support and contribution to the simulation projects included in this article and in the simulation strategy work, particularly, Benjamin King, John McIlmurray, Simon Dennis and Monty Gittoes.

References

- 1 ROBINSON, STEWART. *Successful Simulation: a Practical Approach to Simulation*. McGraw-Hill, 1994.
- 2 LAW, AVERILL M.; and KELTON, W. DAVID. *Simulation Modeling & Analysis*. McGraw-Hill, 1991.
- 3 PIDD, MICHAEL. *Computer Simulation in Management Science*. John Wiley & Sons, 1998.

Biography



Martin Hind
Networks and
Information Services
BT UK

Martin Hind joined BT in 1989 after graduating from Sunderland Polytechnic with a B.Eng. (Honours) degree in Digital Systems Engineering. He began his career in BT in the Copper Access Systems unit, designing a payphone fault detection system, and then joined a team developing an advanced diagnostics fault location system. A few years later he joined the Network Transport Operations and Maintenance unit specifying operations and maintenance requirements and process designs for new technologies. In 1996, he joined the Operational Design team in the Advanced Operational Solutions unit, where he now leads a small team applying discrete event simulation to BT operations.

Martin Hind can be contacted at martin.hind@bt.com.

Jonathan Malpass

Statistics—An Indispensible Management Tool for Telecommunications

The current trend for management by statistics can be flawed by a lack of appreciation of the complexities of a subject that appears straightforward. Statistics, as a subject, goes much further than the common interpretation—that of summarising values of data. This article introduces the reader to the subject of statistical analysis, and presents case studies describing how it may be applied by managers, with particular emphasis on management within telecommunications companies.

Introduction

The world in which we live today is full of facts and figures which are used to advertise products, manage businesses, provide proof in court, or demonstrate that one sportsman is better than another. Modern society has become so accustomed to 'statistics' that it seems that everyone can use them to make some point. Using statistics can be a very worthwhile exercise—used judiciously statistics can be employed to tell the irrefutable truth. However, as Disraeli famously said, 'There are three kinds of lies: lies, damned lies and statistics.' Advertisers may tell us that 9 out of 10 people preferred a certain brand, but we don't know who these 9 people are, and we certainly are not told whether it is 9 out of every 10 people. Governments may inform the public that the rate of inflation (RPI) is decreasing, but what they rarely tell us is the way in which the RPI is calculated, and how it can be adjusted every month so as to mask the rise in prices of a particular commodity.

Statistics can be used to inform, but they can, and often are, used to mislead the public. Unfortunately, inadvertent deceptions can often be the result of the good intentions of enthusiastic employees.

Statistical analysis is applied in most science and engineering subjects to study variation in the measurement of experimental parameters. With the more exact sciences, such as physics and chemistry, most of the variation present in experimental results is due to

experimental error. Even the most precise experiments are subject to tiny amounts of variation. This irreducible variation can normally be averaged out over several repeated determinations. In other subjects such as biology, psychology and management science the inherent variation has to be accepted. However, experiments in these sciences are subject to greater variation than that caused purely by experimental error and, consequently, averaging observations can disguise very important specific instances that result in different outcomes. As a result it is very difficult to determine whether any differences are caused by some real phenomenon or purely by chance.

This is when the scientist should turn to statistical methods. It is a common misconception that as variation in data increases, the less useful statistical methods become. In reality, the opposite is true. The greater the variation present, the greater the need for the powerful tools available to the statistician. Many of these tools are relatively straightforward to apply and, with the widespread availability of statistics within spreadsheets and similar packages, are regularly used. Herein lies the problem: to apply many of the statistical techniques little or no mathematical knowledge is necessary but, virtually without exception, these methods require fundamental conditions to be met before they can be used. It is therefore of paramount importance that the user of statistics is aware of such conditions. Further-

more, the user should have the understanding that when a non-standard situation occurs the usual test may be invalidated. Finally, and perhaps most importantly, the user has to appreciate that if the technique is correctly applied then the result is valid, even if it is undesirable.

This article serves two purposes. First, it is to introduce the power of statistics and the possibilities that the use of the method can give rise to, and to illustrate the pitfalls of using statistics! It is common to misuse some techniques, and to the unscrupulous, it is possible to abuse the methods. Second, while this article will aim to illustrate ideas and concepts, it is impossible in the margins of such an article to provide detailed theory. Any reader who is interested in developing a greater understanding of statistics is directed to any of several excellent texts that are referred to throughout the article. A BT intranet web site is also available for BT people at <http://odta.aom.bt.co.uk/acp/projpub.htm>.

What is Meant by Statistics?

Historically, *statistics* developed from *state arithmetic* or *state-istics*¹ which was the collection of population and economic information vital to the state. This kind of data collection has some very famous examples—the census that required Mary and Joseph to journey to Bethlehem in the Biblical story and William the Conqueror's Domesday Book are examples of how population and resources could be determined. This was (and still is) essential to the running of a state. Without knowing the size and composition of the population, it was impossible to know how many men could be called upon at a time of war; or how many taxes could be raised to finance that war. However, the raw information cannot easily be used, and so it requires summarising. As a result these census counts led to summary values called *statistics*.

By definition a *statistic* is any single figure that describes a set of

data², and summary figures such as the average (the measure of central tendency) or the variation (the spread of data) are typical of this definition. The word *average* is derived from the Latin *havarria*, which was the name given to the compensation money paid to the owners of merchandise carried by old sailing vessels that was lost at sea. The total cost of the lost cargo was divided equally among those whose cargo was being carried³.

Some modern-day interpretations of statistics, such as sports averages and workers' performances are essentially summary values of an entire group, or *population* (for example, a cricketer's career batting performances). Other statistics, such as TV viewing figures, are based on a small section, or *sample*, of the population, and values derived from the sample are used to infer characteristics of the population. This *statistical inference*, just one branch of the subject, is widely used as economic or time constraints rarely permit censuses to be carried out.

Statistics these days has come to refer to a much broader subject than summary values. Widely-used tools such as charts and averages, known as *descriptive statistics*, are useful to provide a snapshot of a set of observations. Other methods, such as those used in statistical inference and other analytical tools, are under-pinned by complex mathematics and are reliant upon fundamental assumptions being met for any results to be valid. This mix of descriptive and analytical methods is better described as *statistical analysis*, and it is with this set of techniques that the statistician approaches both the simple and complex problems that occur in everyday life and in the applied sciences.

Basic Concepts

To a manager in a telecommunications company, statistics may mean a series of performance measures with which to monitor the engineers in his/her patch or the agents in a call centre. Statistics

may also be the number of jobs that remain in the work-stack at the end of each day, the number of faults per day per exchange connection, average travel times or the number of jobs that fail to meet the commitment time. There may be literally hundreds of measures, each one a statistic. However, these figures are summary or descriptive statistics. While recognising the importance of such information, this article is more concerned with the tools available that enable the statistician to place such values in context, interpret results and make legitimate and correct decisions based upon the analysis of such data.

It may be pertinent to present a more formal definition of statistics. *Statistical analysis is a method that draws upon a set of tools that can be used for the collection, analysis and presentation of data and results.* The complexity of each of these three areas ranges from the simple to the intricate depending upon the task.

It is important for users to understand the nature of the problem that they face, the tools at their disposal, the limitations of those tools, and the target audience. Before setting out to perform a statistical analysis, the practitioner should ask him/herself four questions:

- What do I want to investigate?
- What kind of data do I need?
- What method of analysis am I going to use?
- How do I best get the point across?

Deciding upon the aim of the analysis prior to collecting the data is essential. Clear definitions of the problem and how it may be solved can save time, effort and money in the data collection phase. Knowledge of the goal of the analysis highlights what data should be collected and which techniques can be used.

Data collection is as important to accurate statistical analysis as the calculations that follow. Moreover,

while it may be possible to collect a very precise set of data, the overall analysis can be rendered useless if the wrong analysis is performed. Likewise, the most accurate calculations on poorly recorded data are useless. Even so, there is more to data collection than accurate recording of experimental results. Many statistical analyses are performed on data that has been recorded without any forethought. Specifically designed experiments may be required to show clearly the true effects of the factors under consideration. Other statistical analyses are performed in hindsight on data that has already been collected, simply because it was possible to measure a factor. A problem prevalent in some sciences is the injudicious collection of data and subsequent 'blind' analysis in the hope that some statistical relationship will appear. The risk here is finding a chance result as opposed to a true relationship.

Issues surrounding sampling further confound the problems of data collection. If a census cannot be performed, then an unbiased sample that reflects the population must be used. Often a random selection of subjects is sufficient, but more complex sampling techniques are sometimes required to counteract bias or to extract as much information from a sample as possible. A well-constructed sample (even if it is only a truly random one) will provide data that can be analysed easily and should yield results that truly reflect the phenomenon being investigated.

Given the data has been collected, which are the best analytical methods for the situation? Most methods require certain assumptions to be true. If the data does not satisfy these assumptions then the method may produce erroneous results. Choosing the right method therefore is paramount and failing to understand the underlying principles is a common problem. While it is not necessary to be an expert to perform statistical analyses, it is essential that all users understand the limitations of the methods at their disposal.

Finally, if the message fails to reach the audience, all previous efforts to achieve an accurate analysis of well-collected data can be rendered useless. For instance, graphs can show much more information and are easier to interpret than a table of numbers. Individual figures are meaningless, and a far more powerful result will be achieved if the result is placed within a proper context. Most importantly, if the analysis has been performed correctly, then the result, whatever it may be, must be presented.

Case Studies

To illustrate the benefits of using an analytical approach, and possible pitfalls, the rest of this article looks at how statistics can be applied by managers in a telecommunications company. Each of the three cases described below is hypothetical, but all are representative of situations that could easily arise. While the data presented is typical, being based on real figures, all cases have been modified to highlight potential problems.

Case study 1—the CST manager

The day-to-day management of field engineers is divided into geographical units called *customer service teams* (CSTs). The CST manager is responsible, among other things, for ensuring that adequate resources (engineers) are available to meet the number of jobs in the daily workstack. There are several factors that influence the size of the workstack, the number of engineers available and the number of jobs that can physically be performed, and these influences such as the type of job, the travel times between jobs and the skill level of the engineers have to be monitored to ensure a smooth operation. One statistic in which the CST manager is interested is productivity† as it reflects the amount of work that is likely to be performed. Thus productivity

influences the number of jobs that can be completed on a given day by a given number of engineers and the number of jobs that are carried over to the next day. If productivity is known and the workforce is managed effectively then the number of jobs that fail to meet their targeted completion time is reduced and so customer satisfaction may improve.

Consider the manager of a CST which covers six patches, with each patch comprising a team of engineers, or work team members (WTMs), themselves supervised by a field manager (FM). When examining the last week's performance figures (Table 1) of the teams under his control, he notices that Patch 5 has the highest average at 3.78, and that Patch 2 has the lowest at only 2.62.

Table 1 One Week's Average Engineer Productivity of the Engineers in Each of Six Patches in the CST

Patch	Average Productivity
Patch 1	3.02
Patch 2	2.62
Patch 3	2.93
Patch 4	3.14
Patch 5	3.78
Patch 6	2.97

Given these performance figures, what conclusions can he draw? He may think that Patch 5 is much better than all the other patches, or that Patch 2 is much worse than the others. In fact, the manager has already concluded all that is reasonable to infer—that last week Patch 5 had the highest productivity of all the patches under his control, and Patch 2 had the lowest.

The manager is aware that the target for engineer productivity is

† Productivity is a function of the number of jobs completed by engineers, the number of hours actually worked by engineers and the number of hours in a standard working day. The figure can be calculated for individuals or teams.

3.08 jobs per day. He is pleased that Patch 5 performed well above target, but can he be sure that this is typical of that patch's performance? Conversely he might be concerned that Patch 2 performed well below target. Should he take action to ensure that Patch 2 is more productive in the future?

It is important that the manager does not react to these performance statistics without knowing the context within which they sit. For instance, just from the performance figures in Table 1, he can see that four of his six patches failed to meet the target last week. Statistically speaking, unless he is able to see a wider picture he cannot possibly say that Patch 5 is better than Patch 2. For this he needs to consider the average performances of the patches.

Ideally, the manager needs to consider several weeks' performances for all the patches under his control to get an idea of the big picture. If he then finds that Patch 2 is still underperforming then he may take action. The CST manager wants to answer two questions:

- 1 Is the mean productivity level of Patch 5 significantly better than that of Patch 2?
- 2 Are the mean productivity levels of Patch 2 and Patch 5 significantly worse or better than the target?

These questions can be answered using statistical tests, but to be able to perform these tests the manager requires more information. He needs to develop an understanding of the average performance level of each patch and the variability of the performances. The manager decides that rather than collect the last few week's performances he would be better served by considering productivity levels over a greater length of time. This in itself presents problems, as he does not really want to spend too much time collecting data. To overcome this, he chooses a random

Table 2 Average Productivity of Patches 2 and 5 for Seven Weeks, Six Chosen at Random from the Last Year

Week	Patch 2	Patch 5
3	2.88	3.48
22	3.25	3.86
26	3.13	3.54
40	2.99	3.88
2	2.30	3.34
32	2.72	3.48
42	2.82	3.19
13	3.04	3.69
10	3.51	3.79
Last Week	2.62	3.68
Mean	2.93	3.59
Standard Deviation	0.36	0.24

sample of nine weeks from the past year (Table 2), and then includes last week's performances. The CST manager can now see that while last week's performances of the two patches are not that unusual, it does appear that the two patches may be very different in their overall performances.


To answer his questions the manager employs two different statistical tests that enable him to work out the chance of getting a specific result by chance. The tests are based on the *normal probability distribution*† (Figure 1), which has a distinctive bell shape and approximates many natural phenomena. By translating the user's questions into mathematical relationships, or *hypotheses*, then, providing the mean

and standard deviation of the distribution are known it is possible to determine the probability of specific values occurring. If this probability is large then the likelihood of the result occurring by chance is low, and so the hypothesis is 'not rejected'; that is, it can be taken as being true based on the data used in the calculations. If, however, the probability is small then the hypothesis is said to be 'unsupported by the data', and so can be discounted.

The CST manager translates his two questions into hypotheses, known as *null hypotheses*. To answer his first question, to investigate whether the two patches are different, the manager uses a test that compares the difference between pairs of values. The manager assumes that the differences between the two patches should be zero, and so, in statistical parlance his *null hypothesis* is that the mean productivity level of Patch 2 equals the mean productivity level of Patch 5. He will test the null hypothesis against an *alternative hypothesis*: that the mean productivity levels of Patch 2 and Patch 5 are different.

† Each normal distribution has a mean, μ , and standard deviation, σ . The properties of the distribution are such that approximately 67% of the distribution lies within $\mu \pm \sigma$, approximately 95% lies within $\mu \pm 2\sigma$, and approximately 99.74% lies within $\mu \pm 3\sigma$. Figure 1 shows the standard normal distribution, with mean $\mu=0$ and standard deviation $\sigma=1$.

Figure 1 – The normal distribution with zero mean ($\mu=0$) and unit standard deviation ($\sigma=1$). 95% of the distribution lies within ± 2 standard deviations of the mean



The Paired t-test^{4, page 450} gives a *test statistic* that can be compared with a value from the normal distribution, called the *critical value*[†], to see whether or not the result is statistically significant. (The tests described below use critical values based on $\alpha = 5\%$.) The null hypothesis is said to be 'not rejected' if the test statistic is less than the critical value. If the test statistic exceeds the critical value, the null hypothesis is 'rejected'. Note that it is said that the hypothesis is 'not rejected' as opposed to 'accepted' as the test is only for a given sample; a different sample may give a different result.

The results given in Table 3 are obtained from using the 't-test: Paired Sample for Means' in Microsoft Excel. The important figure is the t-stat, which is compared with the t-critical two-tail value.

[†] The critical value corresponds to the limit of the normal distribution beyond which the user deems a value to be unacceptable. The user defines a value (α) for a given test which is equivalent to the percentage of times the user is prepared to allow for an incorrect result, and as α tends to zero so the test becomes more and more strict. In one-tail tests the $\alpha\%$ of values that are at one end of the distribution, in two-tail tests those values that lie in the $\alpha/2\%$ of both ends of the distribution. Tables of critical values for the normal (or t-) distribution are available in most statistical test books.

Table 3 Partial Print-Out of Excel's t-Test: Paired Two Sample for Means

	Patch 2	Patch 5
Mean	2.926	3.593
Variance	0.3411	0.2270
Observations	10	10
Hyp. Mean Difference	0	
Df	18	
t-stat	-5.149	
P(T ≤ t) two-tail	0.0001	
t-critical two-tail	2.26	

As the CST manager is interested in simply seeing whether or not the patches are different, he examines the absolute value of test statistic (t-stat) which is 5.15 (he ignores the minus sign as this indicates that Patch 2 is less than Patch 5). As this value is greater than the critical value, t-critical two-tail = 2.26, he concludes that average performance of Patch 5 is statistically different from the average performance of Patch 2. By looking at the value of P(T ≤ t) two-tail he sees that the probability of getting the same result by chance is 0.0001; that is, very small. He can think of it as the chance of Patch 2 and Patch 5 being equal as 1 in 10 000 (0.0001).

Now that the manager has established that the performances of the two patches are different, he wants to know whether it is true to say that they are both achieving the target productivity of 3.08. He uses a t-test^{4, page 387}, performing the test twice, once for each patch. This time his *null hypothesis: that the mean of each patch is 3.08* is tested against the *alternative hypothesis: that the mean is not 3.08*. Table 4 gives the results for both patches.

As the t-test is performed on both patches the CST manager has to make two sets of comparisons. First, he looks at the test statistic for Patch 2 and sees that the absolute value of t-stat is 1.42 (again he ignores the minus sign), which is less

Table 4 T-test Results for the Null Hypothesis: The Means of Patch 2 and Patch 5 are Equal to 3.08

	Patch 2	Patch 5
Mean	2.9270	3.5925
Stand Deviation	0.3411	0.2270
Observations	10	10
Standard Error	0.1137	0.0757
t-stat	-1.4181	7.1396
Critical t	2.26	2.26
P < t	0.1898	0.0001

than 2.26, the critical value. He therefore concludes that the performance of Patch 2 is not significantly different from the target of 3.08. When he looks at the test statistic for Patch 5 (t-stat = 7.14) he sees that is much greater than the critical value (2.26) and so infers that the engineers in Patch 5 are performing significantly better than the target. The probability of seeing this result occurring by chance is 0.0001; that is, 1 in 10 000.

Overall, the manager finds that by using two different statistical tests he has prevented taking action against a patch that appeared to be under-performing. His natural reaction might have been to intervene, but doing so could have had a negative impact. Telling someone who has just completed a very difficult job successfully that they should be working harder is not going to help the morale of the team⁵. The manager should, given the statistical analysis, examine the reasons why Patch 5 is so much more productive than the other patches and look to see if there are any working practices that can be transferred to the other patches to aid productivity.

Case study 2—the call centre manager

Call centres are managed to meet various targets and produce a number of statistics with which to assess how a call centre is performing. Each procedure that is performed within a call centre is assessed in some fashion, and if that process is failing then the call centre manager or agents are put under pressure to improve, often without any changes to the process or assistance with the current procedure. Some targets are set at what appear on paper to be achievable levels but, for practical reasons, are rarely revised if they are never met. These targets are often based on management beliefs, rather than considered estimates derived from current performances.

The manager of a call centre has been asked to make a presentation to other managers describing the statistics he had used to demonstrate how changes to the call centre were necessary, and how those changes resulted in dramatic improvements.

A year earlier the call centre manager had been under pressure as, during one particular month's performance review, it was seen that the average percentage of promptly answered calls† in his call centre was approximately 88%. He was criticised for the poor performance, and told that for promptly answered calls the critical value was 93% and the target value was 96%. Although he was aware of these figures he felt it impossible to meet them, given the constraints he had to work under.

At this point the manager had decided to use statistics to explain his problems at the next performance review. From his background in the manufacturing industry he was aware of the power of *statistical process control* (SPC), and how it could be used to illustrate the capability of processes.

The control chart is a graphical tool used to show when processes are in, and out of, control; that is, operating within the expected limits of the process. To develop such a chart, the manager had collated the performance of all the agents in his call centre for the previous 13 weeks, which was to act as the training data. After taking a random sample of four agents for each week, he calculated the sample means and control limits*6 and plotted them on a chart, along

† The industry standard definition of a promptly answered call is that which is answered within 15 seconds.

* The control limits are based on the range of values used to calculate the average percentage of promptly answered calls at each time period. The upper and lower performance limits are defined as $\bar{P} \pm (0.48 \times \bar{R})$, where \bar{P} is the average of all the performances considered in the construction of the graph and \bar{R} is the average range of values, each week's range is the difference between the largest and smallest values for that time period. For a full description of the calculations see reference 6, page 390.

with the critical and target values. He continued to do this for each week up to the performance review.

The control chart (Figure 2) clearly shows that the call centre is not attaining the standard required. The mean performance levels are all around 88%. However, the process can be said to be 'in-control' as all but Week 2 and Week 10's performances are within the control limits. The manager explained that the graph demonstrated that the call centre had been operating within its *process capability*, and as such the performance is unlikely to exceed the upper performance limit of approximately 89%.

The call centre manager explained that because the process under consideration was 'in control' then to meet the targets imposed upon him he had three alternatives:

- 1 improve the system,
- 2 distort the system, or
- 3 distort the data⁷.


While only the first of these options is desirable, the manager described several instances of control engineers and managers who, under pressure to meet targets, 'fudged' calculations or redefined control limits so that the targets were achieved. In order to meet his targets, the manager went on to explain, required an undertaking in either revising the processes by which the call centre operated, or by employing more staff to cope with the demand. The manager had identified that both a process re-engineering exercise and more staff were necessary, and even then it would take several weeks to meet the targets.

After being given the go-ahead to redesign the processes, but not to recruit, the manager continued to collect the measures of promptly answered calls measures. In this hypothetical example, the new processes went live after three weeks, and the effects were first seen in Week 21 (Figure 3). The performance improved (93.4% in Week 22), but the

Figure 2 – Control chart for the call centre's promptly answered calls



Figure 3 – Control Chart for 43 weeks of percentage of promptly answered calls



re-engineering alone was not enough to reach the target levels required. Because of the vast differences in how the processes operated, the manager had re-calculated the control limits for the period since the re-engineering and reported his findings. These improvements were enough to be able to convince his superiors to allow a one-week trial (Week 32) during which additional staff were employed. (Note that Week 32 is not included in the calculations of the control limits for this phase.)

The evidence of the improved performance in Week 32 was enough to convince his superiors that additional staff were required all the time. After several weeks in which the performance level was maintained at an average of 93.1%, the manager's request for more staff was granted. The new staff, once fully trained, were employed from Week 37, and the performance improved to 96.5%. After a further six weeks, the manager calculated another set of control limits based on an average performance level of 96.2%.

The control chart for the period of change is given in Figure 3. Note that the three phases are clearly observed in the chart, and that the control limits clearly change as, first the process re-design and then the additional staff take effect. While the performance remains within the control limits the process is said to be 'in-control'.

While SPC can be used to demonstrate when processes are operating within their natural capability and when there are changes in the mean level, SPC can also be used to show when the inherent variation in the process changes. There are, however, many rules for understanding whether a process has changed, too numerous to be addressed here. Several of these tests detect whether observations fall outside of the control limits. However, in some cases action must not be taken. In this case study, Week 32 stands out

as being unusual and in SPC this is known as a *special cause*. Action would not be taken here as an explanation can be found. Typical special causes might include a high proportion of untrained staff or low staff numbers due to widespread illness. In normal process conditions if a valid reason for an unusual observation can be found then it should be noted but ignored when calculating control limits.

Case study 3—the distribution warehouse manager

Distribution warehouses operate according to several performance measures. Some of these measures are dependent on the productivity of the staff, which is affected by the number of staff working each day. While the manager can resource the warehouse according to how much work is anticipated, he/she also has to make allowances for how many staff will be on sick leave. The nature of sick leave is such that it cannot be anticipated, but the manager can make allowances for the number of people expected to be on sick leave as a result of accidents that occur within the warehouse.

Until recently a manager has estimated the number of accidents that occur each month. The manager is aware that there are two types of

accidents that need to be considered, non-sick absence causing accidents (NSACs) and sick absence causing accidents (SACs). He estimated that there would be five NSAC accidents and three SAC accidents every month. Using these values he has been able to work out the amount of overtime that would be needed, and manage his staff accordingly.

However, the latest directive aimed to introduce target values of 2.5 NSAC and 1.3 SAC accidents per month. Obviously, these targets are much lower than the levels he currently works to. Unsure of what to do to tackle the problem, he decides to investigate how appropriate his estimates are and how reasonable the targets are.

First, he collects 12 months of accident data, and calculates the monthly average (denoted by the Greek letter 'mu', μ) of both SAC and NSAC accidents. He finds that the average number of SAC accidents is $\mu_{\text{SAC}} = 2.8$ per month, and the average number of NSAC accidents is $\mu_{\text{NSAC}} = 3.5$ per month.

The manager then uses probability theory, specifically the *Poisson probability law*^{4, page 235}. This law enables him to work out the probability of a certain number of accidents occurring each month. First, he has to ensure that the

Table 5 Probability Values for $k = 1, 2, \dots, 10$ Accidents Occurring in a Given Time Period, Based on the Mean Levels of SAC and NSAC Accidents

k	Sick Absence Causing		Non-Sick Absence Causing	
	$\mu_{\text{SAC}} = 2.8$		$\mu_{\text{NSAC}} = 3.5$	
	$p(k)$	Cumulative	$p(k)$	Cumulative
0	0.0608	0.0608	0.0302	0.0302
1	0.1703	0.2311	0.1057	0.1359
2	0.2384	0.4695	0.1850	0.3208
3	0.2225	0.6919	0.2158	0.5366
4	0.1557	0.8477	0.1888	0.7254
5	0.0872	0.9349	0.1322	0.8576
6	0.0407	0.9756	0.0771	0.9347
7	0.0163	0.9919	0.0385	0.9733
8	0.0057	0.9976	0.0169	0.9901
9	0.0018	0.9993	0.0066	0.9967
10	0.0005	0.9998	0.0023	0.9990
> 10	0.0002	1.0000	0.0010	1.0000

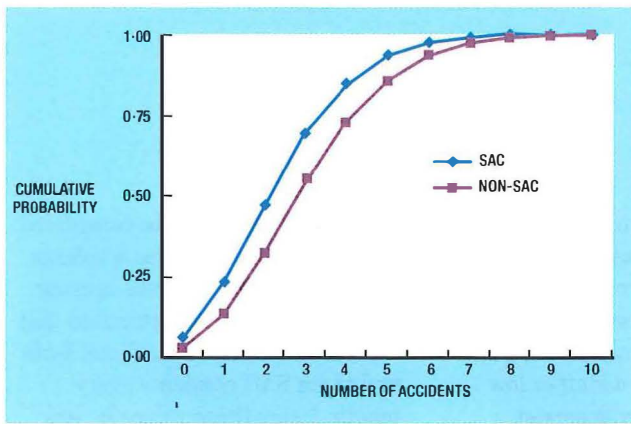


Figure 4—Plot of cumulative probabilities for sick absence causing and non-sick absence causing accidents




Figure 5—Cumulative probability plot for four different values of μ

situation he wishes to model satisfies the assumptions of the Poisson distribution, namely:

- that only one accident can occur at a given time,
- that the chance of an accident occurring, say, in any one hour is 24 times less than an accident occurring during a 24-hour period, and
- that if an accident does occur then it has no effect on the chances of another accident occurring.

The manager decides that, on the whole, the accidents are independent of each other, that they are likely occur at all times of the working day and at any time during the year with the same probability. Hence the three assumptions are valid.

A Poisson probability is calculated by

$$p(k) = \frac{\mu^k}{k!} e^{-\mu}$$

which is the equation that he now uses to work out the probability, denoted $p(k)$, of a given number (k) accidents occurring in a given time period. By adding values together the probability of less than a specific number of accidents occurring can also be determined. Table 5 gives the values for $k = 1, 2, \dots, 10$ and for $k > 10$. It is important to note that these probabilities will change for different values of μ .

The manager can now look up any value in the table to determine the chances of a situation occurring. For instance:

- the probability of exactly two SAC accidents occurring during any

given month is SAC $p(k=2) = 0.2384$; and

- the probability of four NSAC accidents occurring in a month is NSAC $p(k=4) = 0.1888$.

By looking at the cumulative probabilities he can see that, for example:

- the probability of four or less SAC accidents occurring is SAC Cum. $p(k < 5) = 0.8477$; and
- the probability of less than three NSAC accidents is NSAC Cum. $p(k < 3) = 0.3208$.

Using these calculations the manager can decide whether his estimate of the current situation and the targets he has been given are reasonable. He can see that his estimates of three SAC accidents occurring is reasonable compared to the mean, but only has a probability of 0.22; that is, three accidents will only occur 22% of the time. Similarly, he can see that the chance of five NSAC accidents occurring is even smaller ($p(5) = 0.1322$).

The manager cannot actually determine probabilities using the Poisson law for the target values as the law only allows for whole numbers, but he can make an estimate from the plots of the cumulative probabilities on a graph. Figure 4 shows how the cumulative probabilities for the SAC and NSAC accidents differ for $k = 1, 2, \dots, 10$. He can now see at a glance that the chance of up to two SAC accidents occurring is roughly the same as up to three NSAC accidents. Furthermore, the

probability of less than 2.5 NSAC accidents occurring is approximately 0.4 while the probability of less than 1.3 SAC accidents occurring can be estimated at approximately 0.3.

Using the cumulative probabilities in Table 5, the manager sees that the probability of meeting these targets every month is low. However, the manager understands that the graph and table are representative of the current situation. For the targets to be met the manager realises that different safety measures have to be implemented so that the average number of accidents is reduced.

By using different values of the average number of accidents per month (μ_{SAC} and μ_{NSAC}) the manager can calculate probabilities and plot graphs to see which average rate will give him a reasonable chance of meeting the targets. Figure 5 shows how the probabilities of each number of events change as the average number of accidents changes. The cumulative probabilities are plotted for $\mu = 1, 2, 5$ and 10 , and the plot shows that as μ increases, so the chance of getting a small number of accidents every month decreases.

From this plot the manager can see that if he can reduce the average number of accidents to one per month then he will only fail to meet the NSAC target of 2.5 on a handful of occasions. If he can reduce the average number of accidents to two per month then he should meet the NSAC target approximately 75% of the time, while the SAC target would still only be met on approximately half the occasions. While this may not be too encouraging, he can use these probabilities as evidence that the targets may have been set too low

and that it is unrealistic to expect them to be met on every occasion.

Summary

Statistical analysis should be considered as a toolbox containing a wide variety of implements with which a statistician can determine the proper result. Just as a carpenter requires the correct tools to produce a piece of furniture, so a statistician needs the proper tools to perform the best analysis possible. Understanding the role and limitations of these tools is of paramount importance, as misuse only leads to misinterpretation of the situation.

It has been stressed that some of these tools require assumptions to be met. However, it is rare that all assumptions will be satisfied completely. The t-tests described in Case Study 1 are very robust and will produce a valid result in circumstances where some assumptions are not met. However, other methods are particularly delicate and highly susceptible to the slightest deviation from the assumptions underpinning the method. In the instances where assumptions are not fully met, it is the responsibility of the statistician to decide how suitable the tools at his/her disposal are to the situation.

Furthermore, while it may be possible to produce results from any set of data, to make a full and reasoned analysis, the statistician requires knowledge of the problem and the data. Unless the statistician knows the reason for the analysis and the exact source of the data it is unlikely that he/she can produce a fully enlightened summary.

Finally, if the statistician has collected data and analysed it correctly, then the conclusions that he/she draws are representative of the situation investigated. The conclusions may not be those that were hoped for; however, it is also the duty of the statistician to present the results as they are found, and in the correct context, so as not to mislead the audience.

The aim of this article has been to provide an insight into the benefits that can be gained from statistics. To use such analytical techniques in management is straightforward, but they can often be misused. It is hoped that those managers who already use some statistics will be encouraged to develop their understanding of the subject, and that those who are unfamiliar with the intricacies of the subject will be tempted to go beyond managing by the usual measures and investigate the power of statistical analysis.

There are several very good introductory textbooks and popular science books that discuss statistics in a straightforward and enlightening manner, some of these are given in the bibliography. A BT intranet web site (<http://odta.aom.bt.co.uk/acp/projpub.htm>) is available which addresses statistical issues commonly faced by managers. The author and his team are also available to answer further questions and will be happy to discuss any subsequent issues.

References

- 1 WONNACOTT, R. J.; and WONNACOTT, T. H. *Introductory Statistics*. Wiley, New York, 1985.
- 2 NELSON, D. *Dictionary of Mathematics*, Penguin, London, 1998.
- 3 MORONEY, M. J. *Facts From Figures*. Penguin, London, 1990.
- 4 MCCLAVE, J. T.; and BENSON, P. G. *Statistics for Business and Economics*. Macmillan, San Francisco, 1988.
- 5 GILBERT, M.; and CLIFT, D. *Root Cause Analysis of Field Productivity Barriers and Recommendations for BT*. BT Document No. 145LL/T/99/6001, 1999.
- 6 OAKLAND, J. S. *Statistical Process Control*. Butterworth Heinemann, Oxford, 1996.

- 7 WHEELER, D. J. *Understanding Variation*. SPC Press, Knoxville, 1993.

Bibliography

LARSON, H. J. *Introduction to Probability Theory and Statistical Inference*. Wiley, New York, 1982.

KANKI, G. K. *100 Statistical Tests*. Sage, London, 1993

LEWIS, H.W. *Why Flip A Coin?* Wiley, New York, 1997.

HUFF, D. *How To Lie With Statistics*. Penguin, London, 1991.

<http://odta.aom.bt.co.uk/acp/projpub.htm> (BT only)

Biography



Jonathan Malpass
Networks and
Information Services,
BT UK

Jonathan Malpass graduated from Portsmouth Polytechnic with an honours degree in Mathematical Sciences. He spent several years investigating improved mathematical and statistical techniques used in quantitative drug design, and was awarded a Ph.D. in Mathematics by the University of Portsmouth in 1994. During this time he lectured several undergraduate courses in statistics to students studying for both mathematics and other science degrees. On joining BT in 1998 as part of the Process Engineering Group, he has been working on the application of statistical methods to several aspects of resource management, including forecasting and statistical process control.

Jonathan Malpass can be contacted at jonathan.malpass@bt.com.

Network Electrical Protection

External sources of electromagnetic energy, both man-made and naturally occurring, can disrupt the transmission of telecommunications signals and create hazards to equipment and users. This article gives an overview of recent work in the field of network electrical protection. Developments in lightning protection, line isolation units, and noise suppression are described.

Introduction

Network electrical protection is as old as the network itself. From the very beginnings of telecommunications, attempts have been made to preserve the integrity of the electrical signal propagated over the copper network. Over the years this task has become more demanding as the network has been extended to all corners of the globe and the electromagnetic environment in which it operates has become more and more intense.

Today the network is expected to function normally in the presence of everything from electric fence pulses to lightning surges, from electrified railways to high-voltage power lines. This has occurred against a background of greater exploitation of the vast copper network infrastructure. The frontiers in terms of the type of services the network is capable of carrying have been pushed back and back from its humble beginnings as a carrier of telephony.

Along the way optical fibres have been introduced and in the long-term they will no doubt provide the ultimate solution to all issues of electrical protection. However, in recent years the pendulum has swung back somewhat, and copper has been given a new lease of life with the launch of BT Highway and the asymmetrical digital subscriber line (ADSL) trials.

This article reports on the continuing work to maintain the safety and integrity of the copper network, while at the same time not restricting its potential for even greater exploitation.

Lightning Protection

The United Kingdom is not noted for a high incidence of lightning activity,

but even so it is not unusual for there to be 80 000 ground strikes on a single day.

Figure 1 shows the distribution of lightning activity in a typical storm earlier this year, and Figure 2 shows the amount of lightning activity during 1998. It would not be economic to attempt to protect BT's network against direct strikes, but induced surges resulting from proximity strikes can usually be nullified providing the appropriate protection devices have been fitted. Fortunately, direct hits to one of BT's 4 million poles is not that common, but with so many ground strikes you can be sure that wherever the lightning goes to ground the BT network will not be far away.

Proximity strikes give rise to induced surges in BT's network. Once a lightning surge enters the network it can cause havoc. Cable is robust and is able to carry surges over long distances without any noticeable damage to the sheath. The resistibility of customer's equipment and

Figure 1 – Lightning strikes on 29 May 1999

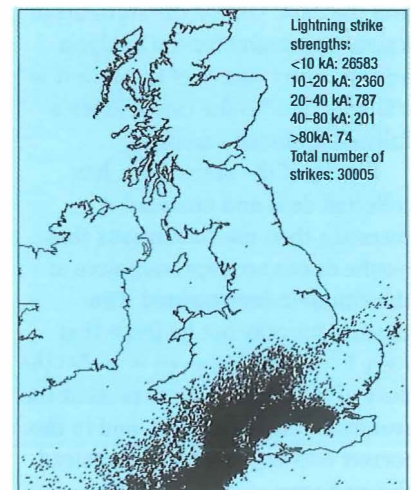



Figure 2—Lightning activity during 1998

exchange linecards is, however, far less and will frequently be damaged unless they are adequately protected.

Lightning protection has traditionally been achieved by gas discharge tubes (GDT). These devices consist of a spark gap in a tube filled with a rare gas mixture, which allows the electrical properties to be precisely controlled. When a lightning surge occurs, the high-voltage triggers conduction across the gap to a low-resistance earth connection. Once the lightning surge has been dissipated, conduction across the spark gap ceases and normal service continues. GDTs are high energy and robust devices capable of discharging surges up to 5 kiloamps. In addition, GDTs are virtually transparent to all network services, both those existing and those proposed, such as ADSL.

Over the years, GDTs have proven themselves to be reliable and effective lightning protection devices. However, in more recent years we have seen the robust electromechanical equipment used in exchanges and customers' premises replaced by sophisticated computer-controlled microelectronic systems. While being highly efficient these systems are, however, much more susceptible to induced overvoltage and can be damaged by fast-acting transient voltages. Protection devices using semiconductor technology are faster and more sensitive and better able to deal with fast-acting transients than the relatively slow GDT.

In recent years, semiconductor technology has increasingly been used in lightning protection devices. The device currently used in the network is a two-stage device. The first stage is a GDT, which provides the high-energy capability. The second stage consists of two fast-acting diodes, which are designed to provide protection against fast-acting transients. These protectors therefore combine the high-energy capability of the GDT with the fast-acting semiconductor. Together the two components provide comprehensive protection against induced surges of all types.




However, semiconductor devices have a higher capacitance than GDTs and while this is not an issue for existing services, which include ISDN2 and BT Highway, it may become an issue as bandwidths are extended further to include ADSL at 6 Mbit/s.

It is important that the deployment of services over the network is not restricted by the presence of protective devices, and to this end BT and its suppliers have been working on ways of reducing the capacitance of the semiconductor component of the protectors.

The current module design

The current PM6A design (Figure 3) is capable of withstanding a broad range of electrical surges such as high energy 8/20 μ s, 10/700 μ s to fast transient 5/50 ns waveforms. The module consists of a three-element gas tube and two transient voltage suppressors (TVS). The operation of the module is as follows, assuming the pulse is from line 'A' to GND. As the pulse ramps up, the voltage across the gas tube and TVS increases. Therefore, the gas tube begins to charge. When the voltage across the gas tube reaches the breakdown voltage of the TVS, the TVS sinks current and clamps the voltage across the gas tube at the

Figure 3—Current 6A module




TVS's breakdown voltage. The voltage across the gas tube is clamped to the TVS's breakdown voltage until the gas in the gas tube ionises and dissipates the pulse.

The typical response time of a semiconductor-type surge suppressor is in the picosecond range, while the gas-tube response time is in the microsecond range. The TVS device used, however, is limited in the energy that it can dissipate before being destroyed by the electrical surge. Further, these devices add significant levels of capacitance to the surge protection circuit. A typical gas tube has a capacitance of between 2 and 7 pF (picofarads), but the semiconductor used in conjunction with the gas tube increases the capacitance of the circuit to about 100 pF. The problem with such relatively high capacitance is that it limits the bandwidth and, therefore, the signal transmission rate of the transmission line to which the surge protector is connected. This is primarily due to the fact that the module capacitance is seen as a discrete discontinuity rather than a distributed circuit parameter.

The ultra-low-capacitance module (Figure 4) is capable of withstanding

Figure 4—Ultra-low capacitance module



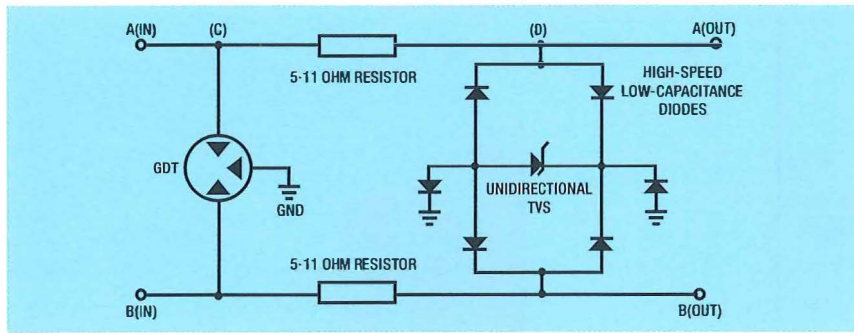


Figure 5—Balanced ultra-low-capacitance robust module

the same electrical surges as the 6A module. The module consists of a three-element gas tube and two transient suppressors (TVS) with two layers of high-speed rectifiers. The purpose of the parallel arrangement of diodes in series with the TVS semiconductor is to reduce the overall capacitance of the surge protector between line to ground and line to line. The parallel capacitance of the fast recovery diodes sum in series with the capacitance of the TVS semiconductor is:

$$C_t = \frac{1}{\frac{1}{C_{d1}} + \frac{1}{C_{d2}} + \frac{1}{C_a}}$$

where C_t is the capacitance of the clamping circuit between line 'A' to ground, C_{d1} is the capacitance of the parallel arrangement of the fast recovery diodes, C_{d2} is the capacitance of the additional arrangement of the diodes, and C_a is the capacitance of the clamping diode. Each fast recovery diode has a capacitance of approximately 10–15 pF and the clamping diode has a capacitance of approximately 180 pF. With these components, the circuit has a line-to-ground capacitance of 20–25 pF and a line-to-line capacitance of 11–13 pF. The parallel arrangement of the fast recovery diodes in a reverse-polarity configuration is used to allow the surge protector to operate bidirectionally.

Figure 5 shows the circuit design for the PM6B protector module. It not only provides for ultra-low capacitance, but also incorporates features for increased robustness and the added benefit of a balanced design. These features are discussed separately below:

Robustness

Although the solid-state TVS device exhibits extremely fast response

compared to a GDT it is unable to survive sustained high surge energy pulses. A 600 W device subjected to a 5 kV 10/700 ms pulse of 125 A will typically fail short circuit within 10–14 ms. In this situation the time taken for the gas within the GDT to ionise and initiate conduction is critical. At present it is necessary to specify the GDTs such that only the fastest are used in the protection module. This requirement to carefully match the GDT to the TVS may be avoided by the insertion of a series resistor as shown in Figure 5. Under conditions of a transient surge passing from left to right in this diagram, the voltage impressed onto the GDT is the sum of the voltage drop across the resistor between points 'C' and 'D' and the TVS clamping voltage of point D to GND. For a rising transient the increase of volt drop caused by the presence of the resistor will result in faster ionisation of the GDT and thereby limit the energy dissipated by the TVS.

Balanced design

The circuits shown in Figures 3 and 4 are unbalanced, meaning the capacitance between line and ground and line to line are different. The circuit shown in Figure 5 has the substantially low capacitance (< 25 pF) between lines and ground and either line to the other. In addition, however, in this case all the capacitances are within 5 pF of each other.

The balanced line has the advantage of providing improved noise immunity over an unbalanced line. Since transients are usually common-mode, it is important that the circuit operates in a balanced fashion, otherwise common-mode transients can cause differential mode disturbances, which will result in errors.




Figure 6—Internal construction of the protector module 6B

We shall return to the subject of line balance later in this article when considering noise suppression.

Summary

In summarising this section of the paper the requirement to reduce the capacitance of protection devices to less than 25 pF has been achieved using the balanced bridge configuration with series resistance. Configuring the diodes in a bridge arrangement allows the module to be in a balanced state for protection against both positive and negative transients and the TVS device need only be unidirectional. The addition of the series resistor ensures that even with relatively slow GDTs the GDT will ionise before the TVS is damaged. Figure 6 shows the internal construction of this protector module. This protector, which is suitable for use on the main distribution frames within the telephone exchange, will be introduced in November 1999. Similar protectors for use at the customer end of the network will be introduced at the same time.

Line Isolation Units

The second topic in this article concerns line isolation units (LIU). Line isolation is required on circuits that terminate on sites which are prone to a rise in earth potential. These sites are commonly referred to as *hot sites*. For the benefit of readers who may not be familiar with the concept of hot sites the following description is provided.

Hot sites

Electricity generating and transformer stations can experience a local rise of earth potential during certain types of fault on the electricity supply network. At some sites, particularly

those in areas of high soil resistivity, the rise in potential can reach several thousand volts relative to remote earth. This rise in earth potential will not normally present any hazard to people on the site since all metalwork will be bonded together and there will be no chance of exposure to different potentials simultaneously.

However, should a remote or true earth be introduced on to such a site then this would present a hazard to people and equipment. This is precisely the situation that occurs when BT lines are connected to hot sites. Copper-pair cables are connected to earth at the telephone exchange, and before they can be terminated on a hot site they must be isolated from the exchange earth. This is achieved by fitting an LIU at the customer end.

The electricity companies require telecommunications services to their sub-stations for telephony, telemetry and the control of their power switching and protective systems. They also have requirements for Kilostream and ISDN2 services, and separate LIUs have been designed for each of these services. A recent incident, however, has resulted in the requirement for yet another LIU, but before describing this new unit it is necessary to consider the assessment of the rise of earth potential.

Assessment of the rise in earth potential

Currents flowing in the ground as a result of a phase-to-earth fault in a high-voltage power line will cause a rise in earth potential, relative to remote earth in the regions where current enters and leaves the ground. To enable the precautions for safeguarding telecommunications plant against rise of earth potential to be determined, it is necessary to assess the magnitude of the potential rise of the station earth system under fault conditions. Figure 7 shows the conditions which arise when a fault occurs on a system whose neutral is earthed at the station.

An earth fault on a line fed from an electricity station causes a line fault current I_1 to flow. This current I_1 flows through the earth unless the power line is provided with a continuous earth wire or a metallic cable sheath bonded to the earth electrode system. In which case, a return current I_2 flows in that path and the current through the earth is $(I_1 - I_2)$. The currents $(I_1 - I_2)$ and I_2 flow back to the system neutral point over the paths shown. It is the earth current $(I_1 - I_2)$ flowing through the impedance Z of the station earth system which raises the potential of the system relative to that of the true earth. The rise of earth potential is thus $(I_1 - I_2)Z$ volts.

For a 400 kV power line, the total fault current can be as high as 60 kA flowing through the ground.


Soil resistivity

Soil resistivity is the resistance of the soil and the most common unit of measurement is the ohm-metre. This is defined as the resistance between opposite faces of a one-metre cube of earth. Typical values range from 2–3 ohm-metres for marshlands, 90–800 ohm-metres for sand to 500–10000 ohm-metres for rock. Thus the value of Z varies from location to location depending on the particular soil resistivity on which the electricity supply station stands. In areas of high earth resistivity, the rise in earth potential may reach 10 kV RMS during an earth fault on a power line.

Hot zones

The area over which the rise of earth potential extends is referred to as the *hot zone*. Here again this will vary according to local conditions but 100 m is a typical hot site radius. The end of the hot zone is defined as the point at which the rise of earth potential is down to 430 V relative to true earth. In the case of high reliability power lines (that is, where the associated circuit breakers operate within 0.2 seconds) a voltage of 650 V is used. The responsibility for defining the hot zone rests with the electricity industry but where they have not provided this information they have agreed that BT should assume it is 100 m. It is this last point that has given rise to difficulties recently. The LIUs supplied by BT are designed to be fitted at the generating station or sub-station and the electricity company concerned is charged accordingly. However, recently damage occurred to a customer's equipment as a result of a rise in earth potential where no hot zone boundary had been provided, and the customer concerned was some 300 m from the sub-station. Although this was an extreme situation where the earth resistivity was very high, this has caused some electricity companies to assess the

Figure 7—Effect of a fault in a high-voltage power line system



extent of the hot sites they control, and this work has identified further situations where the hot zone extends beyond the electricity supply site. Any part of a zone, which is outside the boundary of the electricity supply site, is called *off-site hot*.

An LIU for residential customers

As a result of the situation described above it has been necessary to design an LIU specifically for customers who have premises within an off-site hot boundary. The existing LIUs were designed for use by the electricity industry at their sites, whereas in the off-site hot situation isolation will be required at the premises of residential customers. This is, of course, an entirely different environment, and the design parameters need to reflect this. Otherwise the unit would, quite simply, not be used.

The objective therefore was to design a unit which would provide the necessary protection, but at the same time be less obtrusive than the current range of LIUs. For less obtrusive, read smaller. The LIU for the residential customer needed to be considerably smaller than that currently used to provide isolation on lines carrying a telephony service at electricity supply stations.

In general terms, the size of the current LIUs is proportional to their complexity and the required isolation. A rise in earth potential is normally centred on the high-voltage compound and the voltage/distance gradient is at its steepest close to the centre such that the voltage at the electricity supply site boundary is considerably down on the maximum value. The required isolation for off-site hot customers is therefore much less and the design parameters have taken account of this fact.

Apart from the size of the unit, another important consideration was, of course, cost. The politics of who pays for isolation provided to off-site hot customers has not been finally determined, and a discussion of those issues is not appropriate in this

article; suffice it to say that a cost reduction was necessary.

LIU 11A

The new LIU 11A is shown in Figure 8. It is housed in a case which has the same dimensions as a double-gang power outlet, and it is designed for wall mounting. The front plate provides the normal NTE5 customer interface.

The operation of the unit is based on the sensing of a local earth point with respect to the distant exchange earth (Figure 9). When the local earth potential rises above a set threshold, current flows to the line, triggering a pair of latching relays to open-circuit the line within 10 ms of going above this value. While in the open state, current is drawn from the exchange that feeds a timer circuit, the operation of which is indicated by a red LED on the front of the unit. After one minute, the timer attempts to re-close the circuit. If the voltage is below the threshold level, then service is restored.

One of the most significant differences between the LIU 11A and the existing range of LIUs is that the existing units maintain service before, during and after a rise of earth potential (ROEP), whereas the new unit provides service before and after but not during a ROEP. In fact, the interruption extends for 1 minute, which is considerably longer than the duration of the ROEP. This compromise was necessary to achieve the cost and size targets. However, it should be borne in mind that an ROEP only occurs under fault conditions and this is not an everyday occurrence. In addition, the customer will have a visual indication from the LED that

Figure 8—LIU 11A parts (front-plate, back-box and customer disconnect plug)



the line has been temporarily disconnected.

Most importantly though, the unit will protect the customer and his/her equipment in the event of an ROEP and is designed to be unobtrusive and inexpensive, which means that it will be used. The LIU 11A is currently subject to a field trial which started in September 1999.

The final topic in this article on electrical protection concerns low-frequency interference.

Interference

Transformers are extensively used in electrical engineering, and indeed it is difficult to imagine how we would manage without them. Transformers exploit a fundamental facet of physics, which is that whenever a current flows there is an associated magnetic field. This magnetic field will, if it links with another conductor, cause an induced EMF in the secondary circuit. Although desirable in the case of transformers, magnetic fields associated with current flow in power lines will cause an induced EMF in nearby parallel telecommunication lines. This unintentional induced voltage interferes with the normal operation of the telecommunication line (Figure 10).

Induced voltages on telecommunications circuits

The induced voltage appears on the twisted-pair telecommunications circuit as a common-mode (CM) voltage (Figure 11).

The CM voltage is the voltage measured with respect to earth. The

Figure 9—LIU 11A block diagram




Figure 10—Power line interference

differential mode (DM) is the voltage measured between each conductor of the circuit. The speech or data (telecommunications signal) is transmitted as a DM signal. Telecommunications cables are constructed in such a way as to minimise any DM noise signal. The cable balance is a measure of the extent to which CM noise is converted to DM noise by the parameters of the line. Poor insulation or high-resistance joints can lower the balance of a cable. Cables with a poor balance (less than 60 dB) will have a higher DM noise signal.

There are three main sources of interference in the BT network: power lines, electrified railways and electric fences.

Power lines

Power lines run together with telecommunications cables in almost every street in the country.

Underground power cables normally consist of three live conductors and one neutral conductor.

The resultant magnetic field is reduced because the current in each conductor is phase shifted by 120 degrees.

However, three-phase systems are never perfectly balanced and there is always a residual magnetic field which may cause interference to a telecommunications cable depending on other factors such as the length of the parallelism and the separation of the two cables.

In addition, there are situations especially in rural areas where overhead single-phase power supplies are used.

High levels of 50 Hz induction can create hazards to telecommunications users, whereas harmonics of the 50 Hz fundamental are frequently created by, for example, heavy industrial power users, and this gives rise to troublesome interference within the audio band.

Electrified railways

Running telecommunications cables parallel to overhead-electrified railways is to be avoided wherever

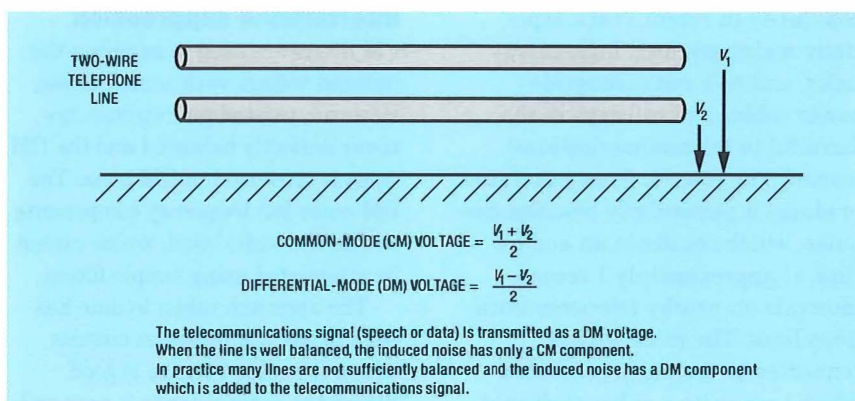



Figure 11—Common mode (CM) and differential mode (DM)


possible. This at first sight may not seem too onerous, but when one considers the zone of influence can extend to cover a corridor 6 km wide and includes both the east coast main line (Kings Cross to Edinburgh) and the west coast main line (Euston to Glasgow) and a number of other shorter lines, then it is apparent that this is not a trivial matter.

The overhead 25 kV system operated by Railtrack powers the locomotive from a contact wire suspended from the overhead catenary. The pantograph on the locomotive is designed to collect the

energy from the contact wire and the running rails are used as a return conductor. An accelerating locomotive may draw 500 A from the supply, which, being single phase, is totally unbalanced. This creates a huge inducing loop, which will link with nearby telecommunications cables and may result in induced voltages up to 30 V peak. Although the running rails are on ballast, current leaks into the earth and returns to the supply via this route.

Examples of CM and DM noise signatures from electrified railways are shown in Figure 12. The record-

Figure 12—Typical 50 Hz noise signatures



ings are from the west coast main line at Manchester. The sampling frequency was 44.1 kHz. The DM signal has been low-pass filtered to reject harmonics outside the voice frequency bandwidth.

Electric fences

Electric fences have grown in popularity in recent years, especially mains-powered high-energy units, and now rank alongside power cables and railways in the 'harmful to telecommunications league'. The electric fence generator produces a particularly troublesome pulse, which results in an audible click at approximately 1 second intervals on nearby telecommunications lines. The fence wire is connected to the pulse generator, which transmits a pulse, such that when an animal comes into contact with the fence wire it will complete the circuit, receive a short sharp but harmless shock, which it will not wish to repeat and so retreats back into the confines of the field. In practice, undergrowth which has not been cut back comes into contact with the fence wire and the generator is continually transmitting pulses into a low-impedance circuit.

As well as producing the clicking noise on telephony circuits, the

interference can also distort fax transmissions and produce errors on digital systems such as ISDN2.

A typical waveform for an electric fence pulse is shown in Figure 13; the pulse has a duration of approximately 1–2 ms.

Traditional methods of interference suppression

CM filters are used to suppress the induced voltage with some success. However, twisted-pair circuits are never perfectly balanced and the CM noise is converted to DM noise. The DM noise has frequency components within the audio band, which cannot be attenuated using simple filters.

The approach taken to date has been to use CM filters on circuits where the cable balance is good. Where the cable balance is poor and cannot be improved another technique is used. This is to convert the line to digital access carrier system (DACS). DACS uses digital transmission, which means the induced DM voltage can be rejected at the DACS remote unit where the digital signal is converted back to an analogue signal. The problem with this approach is that the DACS remote unit does not suppress the CM voltage, which passes through and may be converted to DM over the final drop to the customer.

Noise cancellation

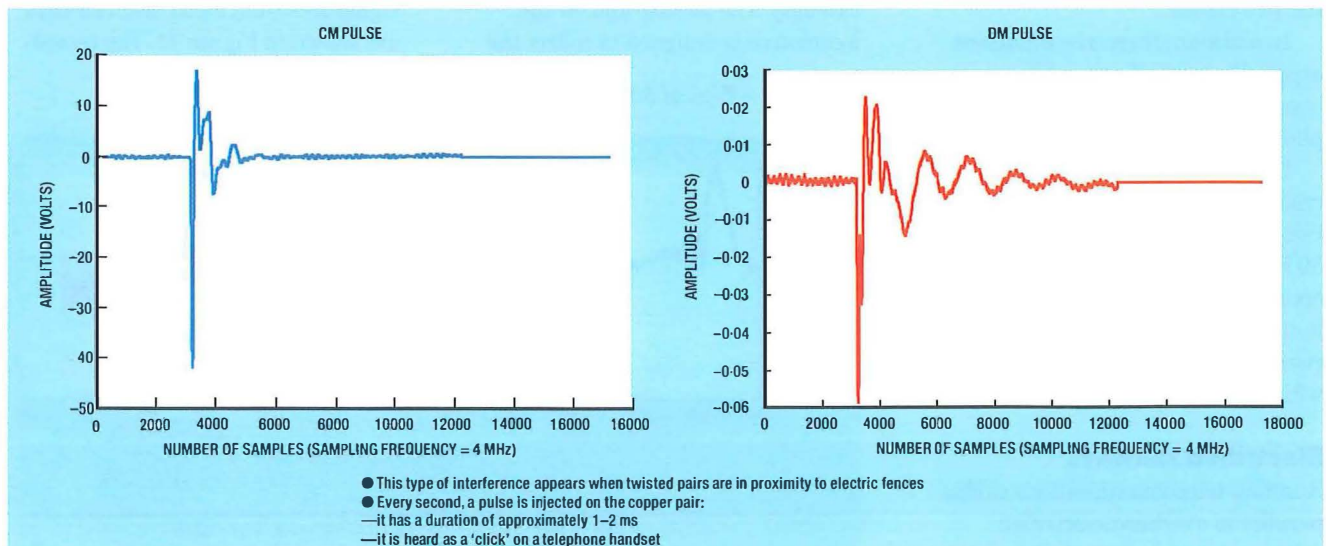
Although the so-called traditional methods have served us well over the years and produced significant noise suppression in many situations, customers' expectations in terms of a noise-free service have become ever more demanding, such that novel ways of dealing with the problem have been necessary. Noise from external sources has become more noticeable as the general background noise has all but disappeared following the exchange modernisation programme.

For the past three years, BT has been collaborating with the University of Essex on a joint project to develop an adaptive noise cancellation unit. Although the concept of noise cancellation is not new, the adaptive phased-locked buffer described here uses the periodic nature of the interfering signal to produce an accurate replica of the noise with only minimal computation.

The adaptive phase-locked buffer (APLB)

The principle of this technique is to use an anti-phase signal to cancel the noise interference. One period of the interference signal is stored in a buffer. The waveform in the buffer is synchronised to the incoming waveform and subtracted from it. The buffer is updated on every cycle of the noise to take

Figure 13—Electric fence noise



account of any changes in the noise waveform (Figure 14 refers).

Operation

The telephony/data signal (S) has superimposed on it a differential-mode noise signal (DM).

The APLB receives the reference signal (CM) and the residual/error signal (DE) and creates a replica of the DM noise waveform which is

Figure 14—Noise cancellation using the APLB




Figure 15—The APLB—principle of operation

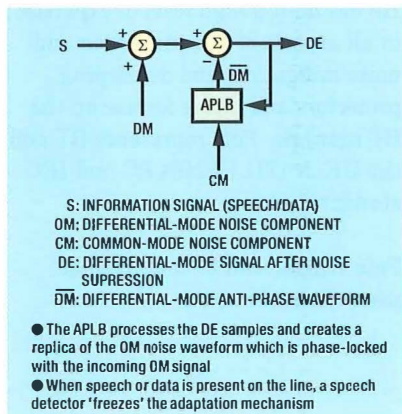



Figure 16—The APLB—cyclic buffer



phased-locked with the incoming DM signal (Figure 15 refers).

A cyclic buffer is used to hold one cycle of the periodic noise waveform. Once the anti-phase signal is synchronised with the incoming DM noise, then each location (tap) in the buffer corresponds to a specific part of the noise waveform (Figure 16).

All taps are updated during one noise cycle. If w_m is the value of the tap m at period n , then the buffer is updated as:

$$w_m(n+1) = w_m(n) + gDE_m(n),$$

where $m = 0, 1...N-1$, N is the number of taps and n is the current cycle.


The step size g controls the speed of adaptation and is set to 0 when speech or data is present on the line.

The noise signal is based on a fundamental frequency of 50 Hz and the presence of any non-50 Hz signal, such as speech or data, is detected. The process assumes that the noise signals do not change significantly from period to period. The cancelling process is suspended while there is speech or data on the line.

Synchronisation

Synchronisation between the incoming DM noise and the anti-phase signal is critical to the operation of the APLB. The contents of the buffer must be refreshed during one cycle of the noise signal such that the buffer waveform corresponds to one cycle of the noise signal. The updating frequency of the taps is the product of the number of taps and the fundamental frequency of the noise signal. A digital phase-locked loop is used to lock onto the noise signal (Figure 17).

Figure 17—Phase locking



The phase detector compares the zero crossings of the fundamental frequency noise signal with the buffer pointer and determines whether any adjustment to the phase of the cancelling signal is necessary. Any errors are smoothed by the loop filter and then applied to a delta sigma modulator (DSM). The DSM has an integer output, which determines the number of sample arrivals before the update of the next tap.

Next steps

The APLB approach relies on the periodic nature of the interference to operate effectively. Analysis of typical railway, electric fence and power cable noise signatures suggests that the APLB is a suitable approach for this type of interference and the results produced in laboratory conditions have been encouraging. The APLB method of noise cancellation will, unlike existing techniques, suppress both the CM and DM components of the noise without degrading the telecommunications signal, and requires relatively few computations which means it can be implemented on a single DSP chip. At the time of writing, prototype models are being prepared for a field trial beginning in October 1999.

The Future

It is clear that for the foreseeable future copper will continue to be the predominant medium for the transmission of telecommunications signals over the access network. It is also apparent that as more and more business is transacted electronically via the Internet the demands on that network will continue to grow. External sources of electromagnetic

energy will continue unabated and novel methods of suppressing these influences will be required. This article has described the most recent work in the field of network electrical protection.

Acknowledgements

The authors would like to acknowledge the following people who have made contributions to the work described: John Cook, Minesh Patel, David Croft and Paul Sadd of BT Advanced Communication Engineering, George Keratiotis and Professor Larry Lind of the University of Essex, Porta Systems, EA Technology.

Biographies



Bob Ruddock
Networks and
Information Services
BT UK

Bob Ruddock joined BT in 1964 as an apprentice. On completion of his apprenticeship, he worked as a precision test officer in West London. This involved using specialist equipment to locate faults in the external network. In 1972, he was successful in the Limited Competition and took up an appointment as Assistant Executive Engineer with responsibility for planning the access network in part of Northwest London. He subsequently held other planning posts in various parts of London. In 1984, he passed a promotion board and went to Telecommunications Headquarters to work on the interactive switched-star cable TV network. Since 1990, he has been responsible for network electrical protection. This involves specifying electrical protection equipment and associated practices, which are designed to protect the user and maintain the integrity of the telecommunications signal propagated over the network. He has an HNC in Electrical and Electronic Engineering and an Open University Degree.

Bob Ruddock can be contacted at bob.ruddock@bt.com.



Pete Whelan
Networks and
Information Services
BT UK

Pete Whelan joined BT in 1979 after graduating with a B.Sc.(Hons.) in Applied Physics from South Bank Polytechnic. He initially worked in the silicon fabrication area of BT Laboratories, managing the wafer testing of devices destined for use in submarine cable systems (40-Type transistors and ECL-40 integrated circuits), which included the development of automated test systems and test techniques. In 1989 he made the move from mV/mA to kV/kA measurements by joining the then Line Interface and Protection group managing the line protection aspects. He has built a high level of expertise in all aspects of line protection and noise mitigation, and developing protectors and filters for use on the BT network. Pete represents BT and the UK at ITU, CENELEC and IEC standards fora.

Pete Whelan can be contacted at pete.whelan@bt.com.

38th European Telecommunications Congress

Introduction

The 1999 Congress of the Federations of Telecommunications Engineers of the European Community (FITCE), held in Utrecht—the ‘Congress Capital’ of the Netherlands, in August 1999, created a valuable opportunity for the key players in the telecommunications industry to discuss major developments in the evolving area of telecommunications. The participants were operators, telecommunications systems and software producers, service providers and researchers.

IBTE contributed to the Congress by producing the Proceedings of the Congress, which was also published as the special August 1999 edition of the IBTE Journal.

Utrecht is a truly beautiful and enchanting city that is steeped in a wealth of history and tradition. At its centre stands the Domtoren, a clock tower which is all that remains of a medieval church. Another memorable feature of the city is the Central Museum which is located in a complex of historical buildings around a beautiful inner garden.

FITCE '99 had a truly international feel as it attracted delegates from 13 countries, as far afield as the United States of America. There were representatives from throughout Europe, including Austria, Belgium,

Czech Republic, France, Germany, Greece, Ireland, Italy, Netherlands, Spain, Switzerland and United Kingdom.

The Opening Ceremony

The opening ceremony of this year's Congress was held in the impressive auditorium of the Beatrix Theatre, where Dipl.-Ing. Guntram Kraus, President of FITCE, welcomed the delegates to Utrecht. He was followed by Ms Monique de Vries, Vice Minister of Transport, Public Works and Water Management of the Netherlands, who pointed out the importance of water management in a country where a large track of land was below sea level! The Minister then gave the audience some insight into the Dutch Government's forthcoming policy in respect of the telecommunications sector.

The other speeches at the opening ceremony were delivered by Herr Dieter Kaiser, Vice-Chairman of the Board of ETSI on ‘The Future of Normalisation in Telecommunications’, then by Mr Peter Scott of the European Commission, DGXIII, on ‘The Mutual Influence of Technology and Regulation’. Mr Patrick Morley, Chief Technology Officer and a member of the Board of KPN, gave an interesting talk on ‘Networking the Future—The Art of Deconstructing Networks’.

Now customary, the FITCE Overture ‘an Invitation to European Dialogue’ was performed by Musical Intermezzo con Corde during this session.

Technical Sessions

A total of 45 presentations over four days provided real substance. In addition various technical visits were arranged on the afternoon of the third day to Ericsson, KPN and Lucent Technologies sites in Holland. If these choices didn't whet your appetite the alternative was a workshop on Internet 2 comprising an introduction, various perspectives and the opportunity of discussing all the relevant issues.

The overall quality of presentation was excellent and the topics covered the whole spectrum of telecommunications. Each session comprised either three or four twenty-minute presentations followed by coffee or lunch giving ample opportunity to discuss the many issues raised during the talks. The presenters themselves were from all across Europe and beyond. Notably the first American presenters at a FITCE Congress, from the University of Southern California, opened the first session on the first day.

Topics ranged from nuts-and-bolts technology issues onto, much talked about, voice-data convergence through network planning and management, touching on legal and market issues before finishing the conference with a look at applications. This diversity provided an excellent overview of all aspects of telecommunications backed up by the variety of professions among the delegates.

Interoperability was frequently mentioned during the Congress, and may well be the key to ‘Networking the Future’. Convergence of voice and

Alison Brady and Simon Broom are from BT, and Felicia Holness is from Queen Mary and Westfield College, University of London

data networks has already started but still has a long way to go. A common theme among several presentations was that the current call quality of voice and other real-time services carried over Internet protocol (IP) networks is still the barrier to its success. Mechanisms to enable a guaranteed quality of service are needed before services such as voice-over-IP become a real threat to traditional telephony.

Several presentations covered mobile communications, which is currently one of the fastest growing market sectors in telecommunications. An evolutionary process towards future system generations, such as UMTS, will be the key to continued market growth, while the convergence of networks technologies and services will be the key to seamless mobility. This will help to open up the mass market in mobile communications although naming/numbering issues still need to be addressed.

Many of the challenges facing network operators in integrating multiple network technologies were discussed. These included the need to provide inter-domain network management to allow coordinated and integrated service management, customer care and billing. As the number of network technologies and network users increases, together with higher customer expectations, these challenges are sure to increase.

In the long-distance telecommunication market, liberalisation has turned out to be the main driver for the change in services. Supported by the rise of the Internet, it has created new commercial opportunities like the introduction of electronic bandwidth auctions and electronic commerce. Such global networks could make it as easy to do business with a company on the other side of the world as with one on the next street.

To support these new broadband networks and services the access network must also evolve. While the cost of bit transport in backbone networks continues to fall, the access network represents an increasing

part of a carrier's overall investment. Several presentations covered broadband access and its role in the future of networks.

Internet 2 Workshop

A workshop on Internet 2 was based on three presentations, namely, (a) 'Deutsche Telekom: Internet Revolution in Germany' by Herr Georg Böhm of Deutsche Telekom; (b) 'Advanced Internet Services Networks: A Status Update on Internet 2 and Related Projects' by Mr John Silvester, a Vice-Provost of University of Southern California and (c) 'GigaPort Network' by Mr Kees Neggers.

Congress Awards

Each year, awards are presented for outstanding papers and this year was no exception. Awards were made as follows:

- *Best paper*: 'ADSL Roll-Out in Belgacom' by Koen Berteloot of Belgacom, Belgium.
- *Most Promising Young Engineer*: Patrick Steemer for his paper 'Quantifying the Market Demand for In-Vehicle Telematics Internet Services', jointly written by Mark Hoogenboom and Albert Kuiper, all of Cap Gemini, the Netherlands.
- *Technical Committee Award*: David Simpson of Oftel, UK, for his paper 'Internet Service Providers—Anarchy in the UK'.

Dave Simpson, Oftel, receives his award from Technical Committee Chairman Vim van der Bijl



Cultural and Social Aspects

Aside from the technical content, one of the biggest components of the Congress is the cultural and social aspects. Everyone was keen to meet new people and spend time getting to know the other delegates. This general attitude makes FITCE a very pleasant forum for one's first international presentation.

The conference also catered for the partners of the delegates by having a separate partner's program, and a highlight was a visit to the famous Kröller-Müller museum situated in the national park.

Other cultural events included a reception at Sint Jans, an old but recently refurbished church near the centre of Utrecht, and a concert in the 'Nieuwe Kerk' (New Church) in The Hague. The young members of the Athenaeum Blazers (Wind) Consort of the Royal Dutch Conservatory at the Hague treated delegates to a stupendous brass and wind concert. The church was an auditorium worthy of the music played which included pieces by Dukas, Rossini, Strauss and Poulenc. In particular, Wagner and Sibelius will never be heard again without conjuring up the image of that setting.

A special evening in the Singer Museum at Laren enabled delegates to visit the unique collection of Belgian art as well as to listen to the ensemble The Broken Consort.

On the final evening, participants were treated to a stage performance by a local dance production company. The evening closed all too soon after a memorable and sensational laser display tracking the evolution of communication throughout the century.

Conclusion

All in all, it was a most enjoyable experience with many new friendships being formed. We are looking forward to a stimulating and attractive Congress in Shannonside, Ireland, in 2000.

Journal Awards for 1998/99

This journal is an important record by which the membership of the Institution of British Telecommunications Engineers (IBTE) and others can keep abreast of items of interest in telecommunications.

To encourage readers in furthering the role of the journal, and to give authors due recognition for outstanding contributions, the IBTE Board of Editors operates an annual award scheme. Prizes are awarded to the authors of articles, which, in the opinion of the Board, demonstrate excellence in content and presentation and which enhance the quality and range of contributions published.

For 1998/99, the Board awarded prizes for best article published in Volume 17 of *British Telecommunications Engineering*, together with three runners-up, and for best *Structured Information Programme* unit. The prizes were presented by IBTE President Chris Earnshaw, BT's Chief Technology Officer, at the IBTE Awards Dinner held on 1 October 1999 at the Mall Galleries in London.

Awards for Journal Articles

The prize for best article went to Andy Catchpole for his article 'Voice-Data Convergence and the Corporate

Voice-over IP Trial' published in the January 1999 edition. Andy received a crystal bowl inscribed with the IBTE logo, cash prize and certificate.

The authors of the three runner-up articles received IBTE crystal goblets, cash prizes and certificates. The first runner-up prize went to Chris Gibbings, Mark Barrett, Margarida Correia, Uma Kulkarni and Kevin Smith for 'Broadband Multicast on BT's Futures Testbed', published in the October 1998 edition. A second prize went to Andrew Grace, Richard Jacobs, Jon Cox and Angela Barrow for their article 'Intranet TV—Video Streaming for the World Wide Web', published in the April 1998 edition and a third prize went to Megan Brown for her article 'Customer Service—The Key to Satisfaction and Loyalty', published in the October 1998 edition.

Structured Information Programme Award

Victoria Hillier and Rhoda Holmes received the prize for the best unit from those published in issues 28–30 of the *Structured Information Programme* for their unit 'Forming and Managing Alliances in the Communications Value Chain', Chapter 16, Unit 8.

Article Assessment

The Board of Editors is assisted in its decisions for selecting articles for the awards by a number of readers who complete article assessments for each edition. Any reader interested in participating in this activity are invited to contact the Managing Editor on (0171) 843 7623.



Richard Jacobs and Andrew Grace with Chris Earnshaw



Megan Brown and Chris Earnshaw

Chris Earnshaw with Andy Catchpole, the Journal award winner



(l-r) Chris Gibbings, Uma Kulkarni, Mark Barrett, Chris Earnshaw and Margarida Correia



Peter Cochrane, Chief Technologist BT, continues his regular column in the Journal. Most people think of computer-telephony convergence, but here Peter is looking at the convergence of carbon and silicon.

Carbon-Silicon Convergence

For most people in the IT industry the word convergence conjures up a world of radios, TVs, hi-fis, personal computers, personal digital assistants, telephones, cameras and more, coming together to communicate and interact on a global network. They are right of course, but it seems to me that this is the easy and almost incidental part. Convergence will go much further to realise a far greater symbiosis between us and our machines—it is about the creation of networked things and intelligence to help us fit into a world changing faster than we have evolved to accommodate.

Consider the millions of people with pacemakers, cochlea implants, pain relief modules, and other forms of electronics imbedded in people. We no longer use IT to merely communicate, entertain, sell and trade, we use it to enhance human abilities and to sustain life. Only a fraction of us would be alive today without the intervention of technology. Imagine what would happen if we suddenly switched off all the computers and networks. In one fell swoop, a large fraction of humanity would die. We would be unable to create and distribute the food, clothing, water and energy required by the near 6 billion people on the planet.


Will our dependence on technology grow? Without doubt it will be difficult to refuse the future life enhancement and repair abilities on offer. At the frontier of today's research there are paraplegics with brain implants controlling computers by thinking, and as a result realising a communications channel for the first time. A handful of experiments with silicon retinal implants have been encouraging, as have direct interfaces to the human nervous system for the control of prosthetics.

Work on the use of silicon tracks to by-pass spinal and other gross nervous system damage is in its infancy but producing encouraging results. Closer to reality is the artificial pancreas, and the internal pharmacy. Imagine travelling the world in the sure knowledge you carry an internal drug store ready to electronically dispense the right antibiotic on demand—it is in the laboratory now.

A few years ago I suggested that chip implants would eradicate the need for keys, passports, driving licences, identity cards, money, bank accounts and central medical records. Two years later Professor Kevin Warwick at Reading University had a chip implant in his arm to automatically open doors, allow him access to secure buildings, and be tracked by his secretary. I now hear on the grapevine that some Italian diplomats are having electronic implants to counteract abduction. Obvious extensions would of course include tagging prisoners and criminals, but probably the greatest benefit would be the potential to rid ourselves of the inconvenient documentation and security paraphernalia of modern life.

Of course, nothing is certain or without competition. Genetically engineered solutions to our biological problems are also making great advances with the help of IT. Being able to model and fabricate piece parts, and repair elements, for our own bodies and systems is a worthy cause and a potentially huge market. The question is—who is going to get there first, and which technology will dominate? In the area of artificial hearts it looks as though the early efforts of the mechanical engineers will be eclipsed by carbon solutions using animals, and perhaps, genetic engineering.

Some find it comforting to think we are in control of our destiny. But I am not so sure we are, or indeed, if it would even be a viable or desirable proposition. The natural mode of our world is for chaos (in the mathematical sense) and evolution to go hand in hand. In such a world everything increasingly comes in



bursts—as is now visible with stocks, Internet activity, and e-commerce. As a very simple example, consider simple elements like a cup of coffee or a ladder bringing down mobile networks. At a large conference no one is using a mobile until 10.15 when coffee arrives. Two thousand people stream into the foyer and 100 mobile phones come out and overload the cell site in two minutes. On a motorway a ladder falls off a contractor's truck and brings down the 999 service. The mechanism? Every car that swerves to miss that ladder hits 999 and locks up the service until the ladder is removed.

In the future we will not think 'Intel Inside', but UMTS (3rd generation mobile) or Bluetooth (a new radio LAN technology) inside. Buy a microwave meal, and instead of a bar code, there is a printed chip transponder on the packet. The same will be true of products in the stores. All stock control and purchases will be automated without a checkout—each item will communicate with the chips you carry in your wristwatch, jewellery, and elsewhere. And back home that meal will be recognised by the microwave, the correct algorithm will be loaded, and the food cooked to perfection without your intervention. This may well alarm the cooks reading these words, but then again, where are the hunters among us? How many could now catch and skin a rabbit? Most would not appreciate the opportunity or enjoy the experience, so is food that cooks itself such a big step?


Historically we have always used technology to create more prime time for those things that we really want to do. But technological evolution will not provide any relief if we cannot

understand the mechanisms and apply some steer. Most aspects of modern life have made a subtle change from being dominated by atoms to being dominated by bits.

Well, it looks like part of this evolutionary path will see us converging in thought and in form with our own technology.

book reviews

Satellite Communications Systems (Third Edition) edited by B. G. Evans



The IEE's summer vacation school held each year at Surrey University has stood the test of time to become a reliable and valuable way to educate and introduce new workers to the field. As correctly stated in the preface, generations of students have benefited from the course which has always been well supported by the UK satellite industry.

Recognising the value of the material, the IEE has published the lecture notes in book form for sometime. This enables the work of the lecturers to reach a wider audience and provides a handy reference source all year round.

The third edition not only brings this fast moving subject up-to-date but is broader in scope and fully reflects the growing interest in satellites by a wider community. As a result there are several new and interesting chapters on growing areas of business such as personal communication systems, multimedia, military, navigation and positioning as well as the development of 'minisats'.

The volume still retains its comprehensive review on the basics of the

subject that provides a readable and useful reference source. As well as a thorough review of the history of the subject it introduces the reader into the fields of radio propagation, earth station and satellite design and communications channel engineering. The rapidly growing fields of digital broadcasting of video and audio material are introduced for the first time.

The fact that experts are drawn from satellite users gives a wider perspective than most similar satellite communications textbooks. Chapters covering satellite and network planning show how satellites play their part in an increasingly integrated field.

For those who intend to use satellites as part of their overall communications solutions there are chapters on the planning, networking and regulatory aspects of satellites. These areas are becoming increasingly important and will widen the appeal of the book.

The future is also explored. Here the increasingly important business aspects are considered in some detail and this helps the reader to understand the wide range of factors that have to be handled in the increasingly competitive and fast moving world. With the bewildering range of new systems being proposed and under development this is a very welcome addition to the traditional format.

Bringing together a wide range of expert authors from industry and universities this book will be of interest to anyone who seriously wants to appreciate the fundamentals of this rapidly developing and exciting field.


Published by the Institution of
Electrical Engineers
ISBN 0-85296-924-4
£75.00. xxvi + 727pp.

<http://www.iee.org.uk>

Reviewed by Russell Silk

Web-Based Network Management—Beyond the Browser)

by Curt Harler



An increasing number of network equipment vendors are providing a browser front end for the management of their equipment. *Web-based Network Management—Beyond the Browser* is a good practical guide to implementing a flexible web-based network access and management system.

The only criticism I have is that I do not like the inclusion of suppliers with their web pages. It turns the book into a catalogue and it will make the book out of date very quickly. Nevertheless, it does not detract from the practical value of this book, which is written in a clear easy to read style. It is not aimed at a reader who wants to build a web-based management system but more at an audience who want to decide if this type of system is for them.

There is a balanced view of the pros and cons of using web-based management tools. The types of tools available are described with full coverage of the benefits and the risks that will result. The discussion gives support to the implementation of web-based management systems. However, chapter three gives a good overview of the practice and methods required to maintain a secure system. It does not

book reviews

delve into the technicalities of network security. However, it gives clear guidance on the areas to be included in a security policy. This section also provides an overview of Java and its operating environment.

The final two chapters round off the book with the penultimate chapter describing real implementations of web-based management systems—how they were used to meet management goals focusing on the benefits that web-based management systems can provide.

The last chapter covers current and future relationship between web-based management systems and the current TCP/IP management components. As SNMP, RMON and RMON2 are some of the keys to knowing what is going on in a network it is useful to understand the relationship between these components and web-based management systems. The use of Neutral agents promise benefits in that it will be possible to use them to predict when a network could be running into problems.

Mr. Harler's book is a clear intelligent look at the products and issues associated with the implementation and integration of web-based network management products into a network management strategy. I enjoyed reading this book but I wish the tome were bigger with more information on the practical side of installing and commissioning of the web-based applications.

*Published by John Wiley and Sons
£22.95. xi + 244pp.
ISBN 0-471-32739-5
<http://www.wiley.com/compbooks/>
Reviewed by Philip Stevenson*


BT Demonstrates Gigabit Ethernet Over DWDM

BT has established the world's first Gigabit Ethernet link running over dense wave-division multiplexing (DWDM) equipment on its LEAnet high-speed Internet protocol (IP) testbed network.

The research project, achieved in collaboration with Lucent Technolo-

Optical filter design and analysis

by Christi K. Madsen and Jian H. Zhao



This is a highly technical book. A serious amount of mathematics is involved and hence a competent grounding in that specialism is a prerequisite. Indeed, the book is targeted at graduate students and researchers in the field. My own background is in the design and production of optical thin-film filters and hence I choose this chapter to 'dip' into first. The nomenclature of digital filter techniques is entirely different from that of optical thin films and I was glad to see both methods adequately described. The digital filter techniques (the Z-transform method) was described clearly but I found myself referring back to previous sections for further analysis and definitions. It became clear that this review was not a simple undertaking!

telecom focus

gies, is a major step in demonstrating high-bandwidth multimedia applications. It combines the strengths of DWDM, as the key to high transport capacity for future multi-service networks, with Gigabit Ethernet, an open networking standard based on a long line of Ethernet technology going back to the early 1970s. This heritage makes it attractive to the IT industry, as operators find it familiar.

The initial chapters set the scene. The plane wave model is derived shortly followed by the slab waveguide derivation. Technical sections on particular devices are followed (generally) by sections on the material properties and fabrication processes. For example, the short section on producing planar devices includes much useful information and is backed up by suitable references.

The text book describes a number of different filter configurations; for example, multi-port filters, lattice filters, thin-film filters, and Bragg-grating filters. The final chapters are taken up with optical measurements with respect to optical filter analysis and concludes with a brief future directions section.

Good clear diagrams are used throughout the book with a problem page at the end of each chapter (sadly, without answers).

In conclusion, I found the book a useful source of information for anyone embarking on the rapidly advancing field of optical filter design and analysis (the driving force here is the interest in dense wavelength division multiplexing in optical networks). Christi Madsen and Jian Zhao are to be congratulated on producing this book which will certainly find a place in my reference library.

*Published by John Wiley and Sons
ISBN 0-471-18373-3
£54.95. xiii + 408pp.
<http://www.wiley.com>
Reviewed by Ian Reid*

The network link operates between BT Adastral Park, Martlesham, and Cambridge. BT will be making the link available for experimental purposes to leading technology players in Cambridge, including members of the Cambridge Network and the university.

Stewart Davies, BT's director of advanced communications engineering, said: 'By creating advanced

networks and experimenting with them we are learning about their capabilities and implications for our business and our customers' businesses.'

'The project opens opportunities for operators to offer new services by using this proven and ubiquitous technology. Furthermore, it will be increasingly relevant for evolving networks in a campus environment with fibre-optic connectivity.'

Testing on the link started in July 1999. It uses Lucent WaveStar[®] OLS 80G DWDM and Gigabit Ethernet optical transponder units. DWDM nodes are located at Adastral Park and Cambridge, allowing the connection of a pair of gigabit routers over the 107 km link.

The line speed currently deployed is 1.25 Gbit/s, giving 1.0 Gbit/s of useful native Ethernet per channel on the DWDM system, providing a total useful data speed for the single fibre strand of 16 Gbit/s. While a major advance, this represents no more than a start on what will shortly be possible.

Stewart Davies said: 'In the near future, we expect to reach Ethernet speeds of 10 Gbit/s on each channel. Even with today's 16 channel DWDM that will give us data speeds of 160 Gbit/s per fibre, but such is the rate of development in multiplexing that in as little as a year from now 40 or even 80 channels will be available. Multiplying the two advances will take us well along the path to practical terabit networking.'

BT and Microsoft Agree to Create Mobile Internet Products and Services

BT and Microsoft have announced an agreement to combine efforts in developing world-leading mobile Internet and multimedia applications and equipment.

The worldwide agreement will lead rapidly to new types of service with broad consumer appeal.

The companies will join forces to:

- create a mobile multimedia service which allows customers to display

and access personalised information on hand-held wireless devices;

- develop hand-held interactive wireless devices based on the Microsoft Windows CE operating system; and
- establish an open industry forum for the development of third generation (3G) mobile Internet applications.

Both companies are to commit substantial resources to collaborative work in the belief that the market for mobile Internet and multimedia products and services represents a massive opportunity for rapid business growth, based on existing mobile networks as well as the forthcoming third generation (3G) networks, including the Universal Mobile Telecommunications System (UMTS).

Sir Peter Bonfield, BT's chief executive, said: 'By deepening our relationship with Microsoft we believe we can create world-class capabilities for our customers. Mobile access to the Internet is a fantastic growth market.'

Bill Gates, chairman and CEO of Microsoft Corp., said: 'Mobility is a key component of Microsoft's vision, and we are excited about working with a world-class network provider like BT to make that vision a reality.'

'Microsoft's platforms and services combined with BT's mobile networks will enable users to access rich, interactive information anytime, anywhere and on any device.'

The collaboration will give individual, personalised deliveries of Internet content and services to users of mobile telephones, laptop computers and a forthcoming new generation of interactive consumer devices.

This capability will complement Project Nomad, an existing collaborative development programme between BT and Microsoft, which lets users of mobile telephones read, send and receive e-mail and calendar information.

BT will develop additional consumer products, including games, music and other interactive services.

The two companies will work together to build a new generation of mobile handsets and computing devices which will exploit fully the power of BT's mobile networks to deliver multimedia content and interactive games.

They will also found a forum for the development of 3G and GPRS mobile applications. This will be an open industry forum based on Internet standards to encourage development of new multimedia applications for mobile devices, and the conversion of existing applications to work in the wireless mobile environment.

The first products and services from the new agreement are expected to become available to customers next year.

BT Steps Forward in Preparations for Next-Generation Internet

BT has announced an important step forward in unlocking the long-term potential of the Internet as a communication medium for its customers.

The company is in the first wave of European organisations to be awarded an initial tranche of commercial address block allocations for IPv6, the next-generation Internet working protocol which will push the commercial potential of the Internet beyond its present capabilities.

The allocation 'future-proofs' BT against IP-addressing limitations in the foreseeable future, and makes it possible for BT to respond rapidly to any new business potential arising from the expected formal launch of IPv6 early in the next decade.

Stewart Davies, BT's director of advanced communications engineering, said: 'IPv6 has huge potential to expand the scope and reach of Internet applications. It will open up exciting opportunities for new types of service in the next decade, including innovative network applications which exploit the greater scope for communication among machines and their users.'

'While new services for consumers are likely to be some years away, BT intends to maintain its position as a major contributor to IPv6 research and development.'

The planned commercial exploitation of the Internet has evolved very rapidly, taking the scope and scale of Internet applications far beyond those envisaged when the first-generation protocol, IPv4, was implemented. While many of the original constraints of IPv4 have been overcome in various ways, IPv6 will make the Internet a better environment for e-commerce.

Advantages offered by IPv6 include:

- Massively increased address space, so that it will be possible for the growth of the Internet to continue at explosive rates for the foreseeable future. IPv6 will be a building block in the 'networked society', in which many types of consumer and industrial devices, including mobile telephones, car navigation systems, home entertainment systems or even kitchen refrigerators could have permanent IP network addresses, and so be connected to the Internet.
- Improved potential for supporting different service qualities based on customer and application requirements. Advanced quality-of-service techniques for IP networks continues to be the subject of research and development, and IPv6 extends the potential to selectively recognise different traffic types and treat them appropriately in the network. It also provides the ability for the network to more efficiently recognise individual customer and applications flows, so that all the IP packets in an individual video call, for example, can be treated uniformly.
- Built-in support for authentication, data integrity and confidentiality, including encryption of packets and authentication of the sender of packets. These facilitate a new standard of security to Internet applications and could boost the prospects for e-business.

BT's preliminary plans include enabling IPv6 on its LEANet advanced IP testbed network, which has nodes in London, Cambridge and

Adastral Park, BT's research and development centre in Martlesham.

The company is also advanced in its work on mechanisms which will be available for a seamless transition of IP environments from IPv4 to IPv6, and the interoperability of the two protocols, as and when such transition becomes desirable.

BT has maintained an active interest in IPv6 over the past three years through its involvement in the global experimental IPv6 network, the 6bone. It has made substantial contributions to standards development work within the Internet Engineering Task Force (IETF), and is a founder member of the IPv6 Forum, a collaboration established to promote the development of IPv6. BT also participates in an EU funded project, 6INIT, with a brief to promote the introduction of IPv6 services in Europe.

Revolutionary Workstation Goes into Production

BT has announced that its Smartspace multimedia workstation of the future has been licensed for production by a British company, Incorporated Technologies Ltd.

The Smartspace concept was devised and demonstrated in prototype form at BT's Adastral Park near Ipswich. Created as BT's vision of a personal working environment for the on-line, multimedia age, it brings together all the benefits of a high powered workstation. It has video conferencing and Internet capability, full cinematic reproduction of audiovisual material, 3D sound and a user interface which uses touch, voice and biometrics to reduce effort and improve efficiency.

Smartspace offers a new type of virtual interface—a combination of desk and chair with a wrap around projection screen and 3D surround sound. Unlike virtual reality systems which involve helmets or headsets, it immerses the operator in the sights and sounds of the display without cutting him or her off from the outside environment.

Smartspace can be configured to display data in wide or split screen

format. The size of the display provides a workspace large enough to project many different data streams at once.

Its economical use of floor space makes it a practical proposition for limited office accommodation or even a home office.

As a videoconferencing access point, Smartspace provides all the benefits of personalised interaction, including the capability to control a remote camera through built-in chair sensors.

Smartspace can bring a distant specialist to a consulting environment in seconds. The consultant can speak face-to-face with a client hundreds of miles away, using the split-screen technology to conduct a live videoconferencing session on one side of the display and screen-based documents or other types of on-line information on the other.

Many forms of business consulting and professional services are expected to be important applications.

However, the technology also lends itself to medicine, counselling or even political applications, opening up the possibility of European MPs overcoming distance barriers by carrying out 'virtual' constituency surgeries.

Its design makes Smartspace ideal as a way of introducing people to virtual reality scenarios. These could include giving prospective house purchasers a very real flavour of the properties they are interested in, without having to travel from an estate agent's office.


Entertainment is also possible. Smartspace shows the latest DVD movies with full cinematic style reproduction and Dolby Digital AC3 sound reproduction. It is also a near-perfect environment for developing or demonstrating computer games software.

The Smartspace story began in April 1997 with a concept by BT engineers. The first prototype made its public debut in October 1997, since when it has been widely shown and attracted favourable comment at technology events in several countries including the United States, Japan, Belgium and the UK.

John Collins, a technology integration specialist who has been with the project from the beginning, said: 'Our

original unit provoked considerable interest from senior executives and technical specialists in major industries, including banking, entertainment, travel and industrial control.

'Our two years of development experience and feedback from potential users, coupled with considerable improvements in some of the available component technology, mean that today's Smartspace can be vastly more capable than the original, and far more practical to own.'



An artist's impression of a Smartspace workstation

BT to Offer Wireless Voice and Data Networks for Home and Office

BT is to introduce a range of wireless networking products enabling users to connect digital devices around the home or small office to the Internet or other on-line services. The networks, which will be available early in the year 2000, have been developed in BT-funded research by United States start-up Home Wireless Networks Inc. (HWN).

The system gives users a single, flexible integrated voice and data network at affordable prices. The modular design has four core components: a controller, data sockets, telephone sockets and handsets. Computers, peripherals and handsets are connected by wireless.

BT expects home networking to become a big growth sector and predicts that more than five million UK consumers will start using home networks within the next five years. At the same time, up to 10 million small businesses and teleworkers will find the flexibility and power of a wireless network an ideal fit for their expanding use of e-business. Wireless home networks will offer an ideal

complement to BT's ISDN and ADSL services, enabling wireless distribution of broadband services to multiple users in the home or office.

The two companies have been working together since June 1998 to develop products for the European market in response to the dramatic growth in Internet services and e-commerce. Sales of the products in the first year are expected to exceed £35 million (\$50 million) in the UK.

HWN's product line in North America combines the features of a wireless telephone system with the functionality of a wireless LAN. As a part of the BT-funded research, DECT and 802.11 compliant systems have been developed at HWN's UK facility in Cambridge. At the heart of the system lies a wireless 'intelligent gateway' to the global network, which will support all major international standards, including ISDN and ADSL.

BT will launch two-line and four-line systems with wireless smart handsets, wireless telephone sockets, and wireless data sockets. The systems are modular in design and feature expansion slots for the addition of a variety of additional products which will include ISDN access, Voice Mail, and 802.11 wireless Ethernet.

BT Outlines Plans for High-Speed Data Service

New connections will transform the Internet experience and provide the platform for a new wave of multimedia-rich services.

BT has outlined roll-out plans for new high speed data services which will turn an ordinary telephone line into a high-speed digital connection capable of carrying information at between 10 and 40 times the speed of a conventional modem.

It will significantly enhance existing services such as the Internet, videoconferencing, on-line education and information services, enabling users to download information, images, video or graphics almost instantly.

Getting information, education and entertainment on-line will soon

become as quick and easy as flicking through a magazine or changing TV channels, thanks to the new flat-rate 'always on' high-speed data connections for homes and businesses.

The technology, known as asymmetric digital subscriber line (ADSL), will be the catalyst for a wealth of new services and products, the new information services of the future. These could include:

- on-line shopping and electronic commerce;
- interactive games with players in different locations;
- teleworking;
- videoconferencing and video-phones; and
- education and distance learning.

ADSL is added to an existing telephone line by attaching advanced electronics to both ends: at the local exchange and at the office or home. The network is upgraded without digging up the streets and pavements.

Among the earliest parts of the UK to benefit will be London within the M25, Cardiff, Belfast, Coventry, Birmingham, Manchester, Leeds, Newcastle, Edinburgh, and Glasgow. 400 exchanges, covering almost six million households and businesses, will be upgraded by March 2000.

Sir Peter Bonfield, BT's chief executive, said: 'This roll-out plan is a major step towards making Britain a world leader in the information revolution. By committing to a large-scale roll-out, BT will play a vital part in stimulating a competitive mass-market. It will drive new information industries and services to the benefit of all.'

'Our intention is to be at the heart of the information society in the UK.'

These new services are part of BT's strategy to provide consumers and businesses with the facilities to enable them to operate effectively in the information age. In September 1998, BT launched the Highway family of products bringing digital data connections to the mass market. In

addition to its new high-speed data services, the company has substantial shareholdings in BiB and in BT Cellnet, which expects to introduce its high-speed mobile-data service (GPRS) early in the year 2000.

BT intends to provide ADSL connections wholesale to a wide range of service providers and other operators, enabling them to develop packages of digital content and digital connections to their own customers.

Corporate customers will be able to order high-speed data services for installation in the homes of teleworkers, giving them access to the full facilities of corporate networks.

One of many advantages of BT's high-speed data service is the 'always on' data connection which will give consumers convenient, fast access to their service provider.

Massive Expansion plan Internet Network in the UK

BT has announced a massive expansion plan for its Internet protocol (IP) backbone network in the UK. It will be 60 times faster and able to grow to 300 times the capacity of today's IP network.

By March 2000, BT will have expanded dramatically its IP backbone network from 14 to 103 points of presence (POPs), using its high-speed, high-capacity multi-service broadband platform. This will give BT the greatest number of access points of any UK IP network, opening the way to significantly faster and more efficient on-line communications for residential and business customers.

The announcement is part of BT's BeTaNet infrastructure—an advanced IP and multimedia network unveiled in February 1999. It will support the massive growth in IP traffic generated from Internet, e-commerce and multimedia services, expected to grow at around 400 per cent every year.

The latest multi-gigabit router technology and wave-division multiplexing (WDM) optical fibre systems, will give BT's IP backbone network the dramatic increases in speed and capacity.

BT also announced plans to deploy 100 000 extra IP dial ports (the modem gateway to the Internet network) by April 2000, with up to a further 100 000 planned by the following September.

Dave Hughes, BT's head of Internet and IP network services, said: 'The combination of these major network developments will position BT as the leading IP network services provider in the UK.'

'BT's IP network is fundamentally much better suited to carrying Internet traffic than the public telephone network, and does so at lower cost, making it possible to reduce prices for Internet service providers (ISPs) and their customers. These ISPs can be safe in the knowledge that however fast their networks grow, BT is committed to being at the forefront of Internet technology development and deployment.

'We envisage that, ultimately, this network will carry many tens of gigabits of backbone IP traffic and tens of millions of dial Internet calls.'

Currently, one in every four calls made across BT's networks is for IP. The IP dial access market alone is now forecast to grow to 35 billion call minutes by March 2000, and to 140 billion by March 2004.

IP backbone

By March 2000, following successful trials, BT will bring into service 14 gigabit routers at seven 'super node' sites across the UK.

These routers, connected by high-capacity WDM optical-fibre links, will support, initially, transmission speeds of up to five gigabits per optical fibre strand, with scope to increase bandwidth to 40 gigabits within two years. Each fibre strand can carry the equivalent of half a million simultaneous dialled Internet sessions rising to 20 million after two years.

By utilising BT's common multi-service broadband network, the IP backbone can handle large volumes of Internet traffic in a flexible, cost-effective and scaleable manner.

Dial Internet ports

The introduction of additional dial Internet ports will enable BT to

separate Internet calls from voice calls at its local exchanges, and stream the traffic on to its IP-backbone network rather than the core public telephone network. This is the most efficient network design for dealing with the mass adoption of Internet access in the UK.

BT Launches Smart Card Trial for Workers on the Move

BT has announced it is to launch a trial of smart card technology to give easy remote network access to employees working away from the office.

The trial, to be run with Microsoft, will involve up to 500 people and will use Microsoft's newly launched Windows Smart Card tool-kit.

Steve Brown, BT's head of smart cards, said: 'This pilot will be the first of its kind in the UK and will allow us to examine the smart cards in operation. BT intends to be a key player in smart card systems and this trial will be an important part in the development of our smart card strategy and technology.'

BT has 30 000 people who access the company's network and intranet using laptops and PCs while away from the office. Currently, they use a combination of passwords and randomly-generated security codes to gain access. These will be replaced with a simple PIN.

Employees taking part in the trial will be issued with a card reader and smart card which will contain a cryptographic application, and BT Trustwise certificates allowing fast and secure log-on to the network.

In the future, the cards could be developed to give a broader range of applications such as access to buildings, secure electronic corporate transactions, on-line banking, e-cash and customer loyalty programs.

Steve Brown said: 'Smart cards are going to be an important tool for businesses which want to exploit e-commerce and Internet opportunities. BT will become a key player in this market by partnering other companies to provide technology and solutions, as well as being a trusted name for customers to rely on.'

Nuisance and Hoax 999 Calls

BT welcomed representatives of other telephone companies and the emergency services to the inaugural forum of a working party to try to reduce the number of 999 hoax calls.

Lord Bassam, Home Office Minister, said in the opening address: 'The Home Office has given a great deal of attention to the problem of hoax 999 calls, working with the telephone operators and emergency services in the 999 Liaison Committee.'

'We deplore the making of hoax 999 or nuisance telephone calls. These are not children playing with the receiver or misdials but deliberate calls, made maliciously with intent to cause distress to the person receiving the call or to mobilise a fire appliance, police patrol car or ambulance.'

'For the emergency services, such calls pose a serious risk to life and waste valuable resources. They also reduce the efficiency of the service which is provided to genuine '999' callers. 'We have to work together to help eliminate these calls.'

Some 11 million false, nuisance and hoax calls are made to the emergency services each year out of 16 million 999 calls made from fixed lines and six million 999 calls made from mobile telephones.

Oftel Sets Out Ideas for Cheaper Internet Access

Oftel has set out a number of ideas and policy initiatives to reduce charges for interconnection to BT's network for Internet calls.

The initiatives aim to separate Internet charges from the current telephony charge structure, so that there is much greater flexibility in the way Internet calls are priced.

OFTEL's invited industry discussion on new approaches to interconnection charging, including:

- allowing Internet service providers greater control over the setting of retail prices;
- splitting the current pence per minute price into an initial set-up

charge and separate lower call duration charge to make long duration Internet calls cheaper;

- discounts on interconnection charges for long-duration calls; and
- options for providing wider unmetered access.

Most operators currently rely on BT to carry the first leg of the call, and the structure of the interconnection charges for this first leg is likely to impact on the range of retail prices that operators can offer their customers. BT presented to the Forum its new access services.

Oftel made clear at the meeting that it is for the telecommunications industry to agree any new pricing structures for these tariffs with BT. Oftel believes that more variety in BT's wholesale tariffs would encourage more innovative retail price packages for consumers. However, if new price structures could not be agreed, Oftel would be prepared to intervene to resolve any disputes.

Commenting on Oftel's suggestions, David Edmonds, Director General Telecommunications said: 'It is clear that the availability of a wider range of pricing options, including unmetered access, will encourage greater use of the Internet.'

'Our ideas presented to the industry should give greater flexibility and choice over charges for Internet access, resulting in even cheaper calls to the Internet, together with innovative pricing packages.'

'The industry should consider how they can make the most of these proposals to offer cheaper Internet access to their customers.'

Year 2000 International Calls

Telecommunications carriers from around the world expressed their confidence that calls between countries will get through over the millennium period.

The carriers announced the conclusion of a global testing programme coordinated by the International Telecommunication Union

(ITU) Year 2000 Taskforce at a meeting held at Inmarsat Headquarters in London which wrapped up its testing work. Many of the world's international telephone carriers have worked since March 1998 to simulate the rollover into the year 2000 and other critical dates. Telecommunications companies in North America, Africa, Asia Pacific, Europe, Middle East, Indian Sub-Continent and Central and South America linked international gateway voice and data switches, following Year 2000 testing standards agreed upon by the ITU Year 2000 Taskforce.

Over 140 global inter-carrier tests have been concluded across numerous combinations of international voice and data switches. All tests were successfully executed without any Year 2000 errors using the guidelines developed by the ITU Year 2000 Taskforce. This programme complemented the thousands of extensive tests conducted by telecommunications companies on their domestic interconnect and local networks.

This outcome confirms the industry's forecast that Year 2000-related concerns are unlikely to impact international calling in any major way during the century date change or leap year periods. Possible areas for concern have been highlighted in a number of well-publicised reports. Overall, however, industry testing indicates that there are unlikely to be any difficulties experienced in the vast majority of cases, although there are a number of countries that have either not shared their Y2K readiness plans or are experiencing difficulties in assessing their readiness.

Carriers demonstrated an unprecedented level of cooperation through participation in this voluntary, self-financed testing programme. The key benefits from this work have been to enable participating and non-participating carriers and customers around the world to assess the readiness of international telecommunications equipment and services. This work is in direct support of the continuity of their products, services and operations over the New Year transition.

Millennium Eve Demonstration to the Met

999 on 9/9/99!

Responding to continued concerns in the public arena as to whether the public switched telephone network—and 999 in particular—would be safe from the ‘millennium bug’ on New Year’s Eve (traditionally an extremely busy time for 999 calls), BT’s Y2k Programme set up a special end-to-end demonstration of the 999 system in conjunction with the Metropolitan Police.

The demonstration involved 999 calls being made from a payphone on the Adastral Park network—via the test PSTN, including BT OSS—and then extended by Megastream (to act as an ISDN 30 connection) to New Scotland Yard’s recently installed Ericsson ACP6000 switch. This simulated conditions on a typical 999 call. BT’s integration testing team at Adastral Park ensured that digital local exchange (DLE)/ digital main switching unit (DMSU) processor clocks were set so that on the day of the demonstration, 12 noon was equivalent to 12 midnight on 31 December 1999. In addition, the clock on the Scotland Yard Ericsson ACP6000 switch, and also the Met’s computer-aided

despatch system were also similarly rolled forward.

For good measure, to prove their confidence in their systems, the demonstration was held on 9/9/99.

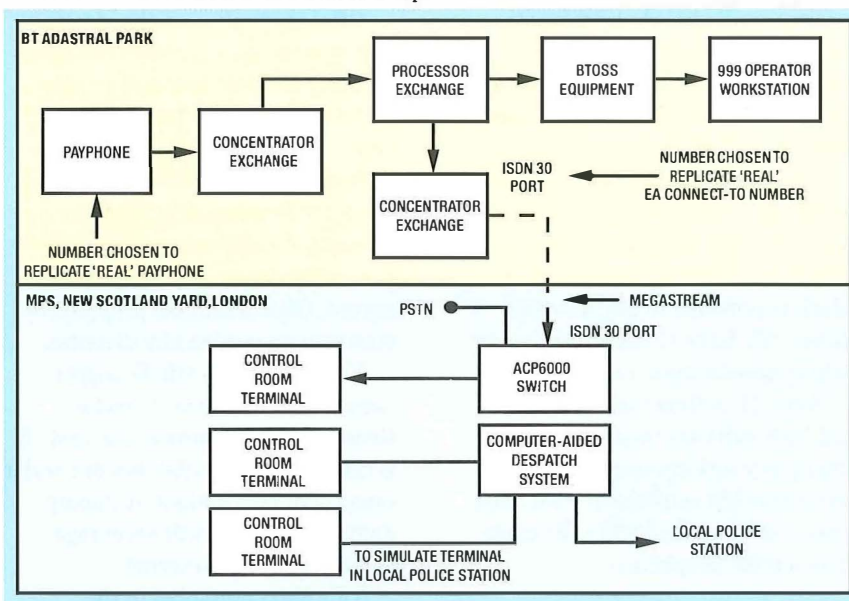
Calls were made (successfully) before, during, and after the ‘Millennium’ in the demonstration, to which a selection of senior players from Ofitel, the Home Office, the Inspectorate of Constabulary, and Action 2000 had been invited. BT Adastral Park and Scotland Yard were connected by video link.

The event was opened by Sir Paul Condon, the Metropolitan Police

Commissioner, and hosted by John Evans, Vice President, Association of Chief Police Officers, who subsequently confidently issued a press statement ‘999 System Ready for the Millennium—A demonstration by the Police Service and the telecommunications industry of the 999 emergency phone system shows that it will not be affected by millennium ‘bug’ problems, say police chiefs.’

Lizzie Beesley, BT’s Year 2000 Programme Director, added. ‘This was an important group to influence and they were delighted; it was a major bit of work, but well worth it’.

The 999 end-to-end demonstration set-up



forthcoming conferences

January 2000

Mobile IP The Hatton, London, 10–11 January 2000. SMI Conferences. Tel: (0171) 827 6000. Fax: (0171) 827 6001. <http://www.smi-online.co.uk/conferences/> E-mail: customer_services@smiconferences.co.uk

Broadcast at Internet 2000— Exploiting Convergence

Dorchester Hotel, London, 17–19 January 2000. IBC. Tel: (0171) 637 4383. Fax: (0171) 631 3214. <http://www.ibcusa.com>

Advanced Operating Systems— for Next Generation Mobile Information Devices

Radisson

Edwardian Marlborough, London, 17–18 January 2000. IIR Conferences. Tel: (020) 7915 5055. Fax: (020) 7915 5056. <http://www.iir-conferences.com/opsystems>.

E-mail: opsystems@iir-conferences.com

IP/PSTN Service Convergence

The Waldorf Meridian, London, 18–20 January 2000. IIR Conferences. Tel: (0171) 915 5055. Fax: (0171) 915 5056. <http://www.iir-conferences.com/ip-pstn>.

E-mail: ip-pstn@iir-conferences.com

xDSL 2000

Hotel Inter-Continental, Geneva, 20–21 January 2000. Access Confer-

ences International Ltd. Tel: (0171) 840 2700. Fax: (0171) 840 2701. <http://www.access-conf.com> E-mail: pippa@access-conf.com

Training & Development in

Telecoms The Kensington Hilton, London, 24–27 January 2000. IIR Conferences. Tel: (0171) 915 5055. Fax: (0171) 915 5056. <http://www.iir-conferences.com/ip-pstn>. E-mail: ip-pstn@iir-conferences.com

GPRS Implementation and Deployment.

London, 25–26 January 2000. ICM Conferences. Tel: (0171) 647 2331. Fax: (0171) 629 3533. E-mail: karineb@thguk.com

BRITISH TELECOMMUNICATIONS ENGINEERING

ISSN 0262-401X

Published in April, July, October and January by The Institution of British Telecommunications Engineers, Post Point 2D05, The Angel Centre, 403 St. John St., London, EC1V 4PL.

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 statement is made specifically by the Board.

© 1999: The Institution of British Telecommunications 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)—companies, universities, libraries and other bodies £60-00 UK, £65-00 overseas; private individuals £36-00 UK, £41-00 overseas.

Single copies and back numbers—companies etc. £15-00 including postage UK, £16-25 including postage overseas; private individuals £9-00 including postage UK, £10-25 including postage overseas.

Overseas customers paying by cheque should use a sterling cheque drawn on a London bank, a sterling Eurocheque or a sterling travellers cheque payable in London.

Orders can be paid by credit card by calling +44 171 843 7626.

Orders by post should be addressed to IBTE Office, PP: 2D05, The Angel Centre, 403 St. John St., London EC1V 4PL.

Cheques should be made payable to 'IBTE' and should be crossed '& Co'.

An index of recent back issues is on <http://www.ibte.org/>

IBTE Members requiring back issues should contact the IBTE Office on +44 171 843 7626.

Advertisements

All enquiries relating to advertisement space reservations should be addressed to The Advertisement Manager, IBTE Office, PP: 2D05, The Angel Centre, 403 St. John St., London EC1V 4PL. (Telephone: +44 171 843 7626. Fax: +44 171 843 7888.)

Communications

All communications should be addressed to the Editorial Office, IBTE Office, PP: 2D05, The Angel Centre, 403 St. John St., London EC1V 4PL. (Telephone: +44 171 843 7622. Fax: +44 171 843 7888.)

Copyright

The entire contents of *British Telecommunications Engineering* and the *Structured Information Programme* are covered by general copyright and the Board of Editors' permission is necessary for reprinting extracts.

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 published prior to 1978.

0262-401X/99 \$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	Carl Alderman	(0181) 666 5956
Martlesham Heath	Bryan Allman	(01473) 605696
Midlands	Dave Plumb	(01203) 492622
North East	Geoff Jenkinson	(01226) 248141
North West	David White (in Manchester)	(0161) 600 2468
Northern Home Counties	Keith Setchell	(01234) 277937
Northern Ireland	David White (in Belfast)	(01232) 215354
Paris	IBTE Office	+44 171 843 7622
Scotland	Cameron Scott	(0131) 668 5804
South West	Tim Jeanes	(0117) 920 5481
Southern Home Counties	John Dymott	(01202) 554001
Wales	Mike Thomas	(01492) 538103

IBTE HelpLine: 0800 028 0209

IBTE On-Line (Internet): <http://www.ibte.org/>

(BT intranet): http://networks.intra.bt.com/ibte_nat/

The Journal of the IBTE

Journal Subscription Order Form

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 2D05, The Angel Centre,
403 St. John St.,
London EC1V 4PL
United Kingdom

Annual subscription rates (four issues, including postage and packaging) are as follows:

Companies, universities, libraries and other bodies: **£60-00 UK; £65-00 Overseas**
Private individuals (to be paid for by personal cheque): **£36-00 UK; £41-00 Overseas**

Please supply four quarterly issues of the *Journal* of IBTE. I enclose a cheque/postal order to cover the yearly subscription as given above.

(Cheques and postal orders, payable to 'IBTE', 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. Customers may wish to use a local subscription agent. Customers wishing to pay by credit/debit card should call +44 171 843 7626.)

Title (Mr/Mrs/Ms etc) _____ Initials _____ Surname _____

Mailing Address _____


Postcode (Zip Code) _____ Country _____

Signature _____

Date _____

Please state the 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.*



The key to successful networks

IT ALL BEGAN WITH GUGLIELMO MARCONI.
WITHOUT HIS GENIUS, THE WORLD MAY
NOT HAVE HAD A VOICE. CERTAINLY,
THERE WOULD NOT BE COMMUNICATIONS
TECHNOLOGY AS WE NOW KNOW IT.

IT WAS HIS ABILITY TO INTEGRATE A
VAST SPECTRUM OF IDEAS, SKILLS AND
EQUIPMENT THAT OPENED THE DOOR TO THE
FUTURE.

AND THAT SAME, UNIQUE BREADTH OF
VISION DRIVES THE COMPANY THAT BEARS
HIS NAME TODAY.

WE HAVE THE PEOPLE, THE PRODUCTS,
THE EXPERTISE AND THE EXPERIENCE
THAT PROVIDE THE KEY TO A SUCCESSFUL
NETWORK.

WE CAN PLAN THAT NETWORK.

WE CAN BUILD IT.

WE CAN OPERATE IT FOR YOU.

WE CAN DO ANY ONE OF THESE THINGS.

OR WE CAN DO ALL OF THEM. TOGETHER.

WHEREVER YOU WANT TO GO, WE'LL GET
YOU THERE. FASTER. MORE SAFELY. MORE
RELIABLY. MORE ECONOMICALLY. HELPING YOU
OVERTAKE YOUR COMPETITORS.

THE WAY AHEAD HAS NEVER BEEN CLEARER.

MARCONI
COMMUNICATIONS

www.marconicomms.com

39th European Telecommunications Congress

FITCE 2000 Call for Papers

Telecommunications in the Third Millennium: New Dimensions, New Challenges.

23-26 August 2000, Shannonside, Ireland

Every year since 1962, the FITCE Congress has provided the European telecommunications community with an update on services, networks, operations and support systems. Over the years, ongoing change has added new dimensions to the challenges faced by FITCE members. The 39th European Telecommunications Congress in Shannonside offers FITCE an opportunity to explore these dimensions and share views on how the challenges can be tackled.

Contributors are invited to address the theme 'New Dimensions, New Challenges' by discussing changes, risks and opportunities in telecommunications and related industries. Presentations are expected to be of interest and value to fellow professionals in European telecommunications.

Sessions will be scheduled to fit the contributions received. Possibilities include:

- **The Technology Dimension** New solutions and their effects; evolution versus revolution in technology policy planning.
- **The Industry Dimension** New relationships in the equipment, network and service markets; privatisations, alliances and takeovers; increasing overlap with IT, media and other industries.
- **The User Dimension** How are end-user needs and expectations evolving? What efforts are the industry making to shape them? What usage patterns are emerging from new methods of service delivery?
- **The Regulatory Dimension** What is regulated, how does it affect the other dimensions? Are international standards bodies still worthwhile?
- **The Political Dimension** Is the vision of a unified global infrastructure still relevant? Does Europe have to compete to retain its voice?
- **The Professional Dimension** What are the career implications for FITCE members, and what should they and their employers do?
- **The FITCE Dimension** How should FITCE develop to maintain its value to members in a much-changed environment?
- **New Dimensions** What new dimensions will emerge in the Third Millennium?

Abstracts of 200 words or less are invited, with a deadline of **15 February 2000**. A brief biography of the author(s) should also be provided.

- Abstracts should be submitted in the English language.
- Papers should be unpublished.
- Abstracts should be sent before **15 February 2000** to: Paul Flanagan, FITCE Papers Coordinator, Room 5.8, Eircom Headquarters, 114-117 St Stevens Green West, Dublin 2, Ireland. (Telephone: +353 1 701 5515; Facsimile: +353 1 475 0140. E-mail: pflanagan@eircom.ie)
- The abstracts will be reviewed by the International Paper Selection Committee for relevance, technical content and originality. Authors will be informed whether their proposed paper has been selected for presentation by 31 March 2000.
- The full text of the selected papers, in English, is required by **1 June 2000**.
- Authors may wish to consult with their respective national representative of the FITCE Comité de Direction with regard to coordinating national contributions.



Contents

AT&T's 21st Century Network Architecture	76
The Market-Driven Network	85
The Full Monty: PSTN-Equivalent IP Telephony for Spain	92
Global Convergence—The Challenge of the 21st Century: Tele Danmark Experience	99
GPRS and the Mobile Data Revolution	107
UMTS Core Network Issues	113
Dense Wavelength-Division Multiplexing and Optical Networking	119
Interworking in the Future Broadband Access Network—A Question about IP and ATM	125
UK Digital Terrestrial Television	132
Managing IT Investment to Deliver Business Benefit	138
Extending the Boundaries of E-Business Transforming Back-Office Operations Between Telecommunications Operators	142
Advanced Global Electronic Messaging	149
Optimising Investment in the BT UK Access Network	154
Investment Decision and Control System for Access Planning Integrated Business and Network Analysis—A Case Study	160
Call Centres in BT UK Customer Service	165
Telepathy	174
Whither Video?—Between Sentiments and Signals	182
Active Networks—Making the Next-Generation Internet Flexible	188
Simulation Modelling in BT Operations	195
Statistics—An Indispensable Management Tool for Telecommunications	201
Network Electrical Protection	210

